The perceptual restoration of music in young children

Naomi E. Winstone

A thesis submitted for the degree of Doctor of Philosophy

2009

Abstract

Auditory perception often takes place in noisy environments. This can result in incomplete sensory input reaching the perceptual system, where fragments of the signal are missing or masked by other sounds. Adults can reconstruct missing sensory input in both speech and music perception using a process called perceptual restoration (e.g. DeWitt & Samuel, 1990; Warren, 1970), where missing sections of the signal that are replaced by noise can be reconstructed without disruption to continuity or intelligibility. Children can also reconstruct missing speech fragments with clear implications for their learning and development (Newman, 2004). The primary aim of the research reported here was to establish the operation of perceptual restoration in the musical domain for 4- to 6-year-old children. A series of eight related experiments investigated the effect of the duration and amplitude of the noise, the familiarity of the musical excerpt, and the presence or absence of lyrics on children's ability to identify familiar music using perceptual restoration. The role of melody and lyrics, and of pitch and rhythm in perceptual restoration were further investigated using a mismatch experiment. In addition, a developmental study compared these abilities to those of older children and adults. Taken together, the results demonstrate the operation of a musical perceptual restoration mechanism that is influenced by both the acoustic properties of the signal and higher-level contextual factors. The results further suggest that perceptual restoration in children's music perception is highly context-dependent and shows clear improvements with age. The present results show that music perception in children is an active process that involves a combination of incoming acoustic information, stored knowledge and expectation. It is argued that musical perceptual restoration in young children represents a fundamental adaption to environments that are frequently characterised by signal disruptions.

Declaration of originality

This thesis and the work to which it refers are the results of my own efforts. Any ideas, data, images or text resulting from the work of others (whether published or unpublished) are fully identified as such within the work and attributed to the originator in the text, bibliography or in footnotes. This thesis has not been submitted in whole or in part for any other academic degree or professional qualification. I agree that the University has the right to submit my work to the plagiarism detection service TurnitinUK for originality checks. Whether or not drafts have been so-assessed, the University reserves the right to require an electronic version of the final document (as submitted) for assessment as above.

Signed __________ *Date* 3/11/09

Acknowledgements

First, I would like to thank my supervisors Dr. Alyson Davis and Dr. Bart De Bruyn for encouraging me to pursue this research, and for their continued support throughout the process. I would also like to express my gratitude to the University of Surrey Research and Enterprise Committee for the award of a University Research Scholarship which has funded my research.

Second, I would like to thank my parents for giving me the inquisitive mind that led to my interest in research, and for arranging for me to go into their schools more than once. I am also very grateful to Clare Worsfold for her particular effort in securing participants. I thank all the teachers who made me feel so welcome in their schools, and for some very interesting discussions about the educational implications of this work. Particular thanks to the children from the following schools for making data collection so enjoyable:

Boxgrove Primary School, Guildford
Hinchley Wood Primary School, Esher
Loseley Fields Primary School, Godalming
Milford School, Godalming
St. Andrew's C. of E. Infant School, Farnham
St. Joseph's RC Primary School, Christchurch
St. Katharine's C. of E. Primary School, Bournemouth
St. Mary's C. of E. Primary School, Chiddingfold
St. Peter's C. of E. Primary School, Wrecclesham
Worplesdon Primary School, Guildford

Third, I am grateful to my fellow PhD students in the psychology department for providing a supportive and enjoyable environment in which to complete this project; in particular to Corinne Huntington for her encouragement and for many discussions about my research. Dr. Henriette Hogh provided me with some valuable comments on a draft of this work for which I am extremely grateful.

Finally, I would like to thank my husband Neil whose patience, endless support and belief in me has been such a vital part of the past three years.

CONTENTS

PART ONE - INTRODUCTION

Chapter 1: The perceptual restoration effect

Chapter 2: Theoretical and methodological approaches to perceptual restoration

PART TWO- EMPIRICAL STUDIES

Chapter 3: Experiment One

Does the perceptual restoration of music operate in children?

Chapter 4: Experiment Two

The effect of the duration of the missing fragments on the perceptual restoration of music in children

Chapter 5: Experiment Three

The effect of the amplitude of replacement noise on the perceptual restoration of music in children

Chapter 6: Experiment Four

Familiarity influences on the perceptual restoration of music in children: The perceptual restoration of initial and later phrases of familiar melodies

Chapter 7: Experiment Five

The perceptual restoration of songs in children

Chapter 8: Experiment Six

The contribution of lyrics and melody to the perceptual restoration of songs in children

Chapter 9: Experiment Seven

The role of pitch contour and rhythm in the perceptual restoration of melodies in children

Chapter 10: Experiment Eight
The effect of age on the perceptual restoration of music

Chapter 11:
An analysis of the effect of the musical characteristics of the stimuli on identification in experiments 1 to 8

Chapter 12: General Discussion

ILLUSTRATIONS

1. LIST OF FIGURES

2. LIST OF TABLES

PART ONE

INTRODUCTION

Overview

The primary aim of the first part of this thesis is to bring together evidence that forms the background to the research to be reported here. The first chapter begins with an explanation of the perceptual restoration effect and is followed by a review of the research conducted on perceptual restoration to date. This includes a consideration of how both the acoustic properties of the signal and listeners' prior knowledge and expectations influence the perceptual restoration process. This review highlights the fact that the majority of empirical work on perceptual restoration has focused on adult listeners and on restoration in the speech domain. This is followed by an analysis of the similarities and differences between speech and music perception, and an outline of children's musical development. It is argued that investigating the perceptual restoration of musical signals in young children represents an important empirical endeavour with clear implications for children's development. Chapter 2 presents a discussion of the theoretical implications of perceptual restoration and is followed by a description of the possible perceptual processes involved in the effect. An analysis of the methods used in previous research into perceptual restoration is then presented, which forms the basis of justification for the methods adopted in the series of experiments presented here. The discussion of these issues then leads to an outline of the aims of the present series of experiments: to obtain evidence that perceptual restoration in children operates in the musical domain and to determine some of the factors that influence its operation.

Chapter 1: The perceptual restoration effect

Introduction to perceptual restoration

Nearly all human behaviour takes place in environments that present redundant, irrelevant and disrupted sensory input. In order for accurate perception to take place it is important that the human perceptual system has ways of compensating for these disruptions, such that stability of perception can be maintained in even the most disruptive of environments.

Everyday environments are characteristically noisy. Noise, both transient and persistent, has detrimental effects on a variety of behaviours (Broadbent, 1958). In such environments, parts of auditory signals can be obliterated or masked by louder sounds, with the result that often sensory input is incomplete. The ability to fill in missing sensory input is a general property of sensory systems (Bremner, Johnson, Slater, Mason, Cheshire & Spring, 2007; King, 2007). This ability clearly represents an important adaptation to noisy and disruptive environments, where the ability to fill in missing sensory information means that for the most part it is still possible to make sense of what is seen, heard, felt, tasted and smelt. The perceptual system demonstrates a preference for coherence, and hence makes coherent percepts out of disrupted signals by "...filling in information and making best guesses" (Remijn, Nakajima & Tanaka, 2007, p.898).

The auditory perceptual system has many ways of compensating for disruptions to auditory signals in noisy listening conditions. One such mechanism, perceptual restoration, operates when noise masks or replaces missing sections of an auditory signal; the signal is perceived as continuous with no interruption and intelligibility is unaffected (Warren, 1970). Intelligibility is restored through a process of perceptual reconstruction of the missing sensory input, which is heard to be present despite the noise. The mechanism is often described as unconscious, occurring outside attentional control (e.g. Warren & Sherman, 1974). Noisy listening conditions are so common that perceptual restoration of missing or masked sounds is described as one of the "main tasks" of the auditory system (Remijn, Pérez, Nakajima & Ito, 2008, p.113). Disruptions to auditory signals often go unnoticed because mechanisms such as perceptual restoration allow listeners to focus on the content of the message and not the process by which it is received (Hawkins, in press).

The general effect of perceptual restoration is described in the literature under different labels: auditory induction; the continuity illusion; amodal completion; and auditory fill-in (Recanzone & Sutter, 2008). All refer to the ability of noise to restore coherence of a fragmented signal, whether this is a simple signal such as a steady tone, or a more complex signal such as speech or music. In contrast to a signal where missing fragments are replaced by noise, a signal where missing fragments are replaced by silence is not perceived as continuous. For example, when hearing a speech stream where a missing fragment is replaced by a louder cough, there is no evidence that the missing sound is not in fact present. The cough does not interrupt the continuity of the speech signal, making it difficult to detect that input is missing or incomplete. Conversely, a silent gap in a speech stream interrupts the perceived continuity of the signal; the fact that part of the sensory input is missing becomes particularly apparent. The perceived continuity of disrupted signals when missing information is masked or replaced by noise results from the sensitivity of the perceptual system to sounds that we might expect in a given context, and the resistance that the system demonstrates to potentially disruptive extraneous noise (Riecke, van Opstal & Formisano, 2008).

Perceptual restoration is reported as a specific case of auditory induction (Warren, 1999). Auditory induction refers to a group of effects where apparent continuity of a fragmented signal occurs when interrupted by a louder masking noise; missing sounds are heard to be present. There are two broad classes of temporal induction described by Warren (1984): contralateral induction, where a disrupted signal at one ear is restored when the other ear is presented with the intact signal; and temporal induction, where brief temporal gaps in a continuing signal are made to sound complete if replaced by a louder sound. There are, in turn, three forms of temporal induction: homophonic continuity, heterophonic continuity, and contextual catenation.

Homophonic and heterophonic continuity represent cases of simple temporal induction, where a steady-state sound is made to sound continuous through the insertion of a louder sound into any missing fragments. For homophonic continuity, the louder, replacement sound must be identical in its spectral characteristics to the signal it is replacing, such as where part of a tone is replaced by a louder tone. Heterophonic continuity is produced when a louder sound, with different spectral characteristics to the sound it is replacing, is inserted into gaps in the signal (e.g. where part of a tone is replaced by a louder cough). A similar form of continuity is described by Thurlow (1957) as an "auditory figure-ground effect" where, in a similar way to visual figure-ground effects, the more intense of two alternating sounds is heard as intermittent and the less intense is heard as continuous, representing the figure and ground, respectively. Thus, the ground is perceived to continue behind the figure. This only occurs when the sounds have similar frequencies and one is more intense than the other.

Both homophonic and heterophonic continuity require that the fainter sound is a steady-state sound, that is, identical before and after the interruption by a louder sound. In contrast, contextual catenation involves a more complex perceptual synthesis of missing information, in dynamic signals where the sound occurring before and after the missing fragments is not identical. The simplest form of contextual catenation would involve the apparent continuity of a rising and falling tonal glide[1] where missing parts are replaced by noise; a more complex example would be the perceptual synthesis of missing speech fragments. The temporal induction of speech is known as phonemic restoration, and is described as a "...special linguistic adaptation of temporal induction" (Warren, 1984, p.376).

Temporal induction involves a signal where all spectral regions are obliterated at a specific temporal point in a signal. This needs to be distinguished from spectral restoration, where some spectral regions of the signal at a single time point are removed, and replacing these missing spectral regions with noise results in restoration of the missing regions. There is much evidence that this kind of restoration occurs, for both simple tones and more complex speech signals (Bashford & Warren, 1987a; Bashford Warren & Lenz, 2005; Holloway, 1970; McDermott & Oxenham, 2008; Shriberg, 1992; Warren, Riener Hainsworth, Brubaker, Bashford & Healy, 1997). Whilst the distinction between temporal induction and spectral restoration is important on a conceptual level, the research reported here does not investigate spectral restoration, and instead concerns the temporal induction of musical signals; more broadly termed musical perceptual restoration.

---

[1] A tonal glide consists of a tone with gradual and continuous changes in frequency, either ascending or descending.

A further distinction is required between occasions where an auditory signal is complete but part is masked by a louder sound (such as a cough obscuring part of a musical performance), and the case in most experimental studies of perceptual restoration where part of the signal is actually removed and completely replaced by noise. In the former, perception of the obscured section when restored is veridical, since the signal persists through the noise. However, where part of the signal is removed and replaced by noise, perception of the restored section is actually illusory. As with many illusions, this represents an important approach to the study of normal perceptual processes that are not the product of an illusion (Ellis, 1999), and reveals insights into the organisation of perceptual information in the brain (Miller, Dibble & Hauser, 2001).

Importantly, Warren (1984) claims that regardless of whether part of an auditory signal is fully replaced or masked by noise, in both cases the signal temporarily ceases to exist. For this reason, similar processes are involved in restoring intelligibility to the listener (e.g. using context to interpret ambiguous input, making best guesses about the nature of disrupted information, and gathering evidence from what is present in the signal in order to evaluate these best guesses). Furthermore, by actually removing part of a signal, greater experimental control over what is actually presented to listeners can be achieved, and results in a conservative estimation of the operation of perceptual restoration. This is the method that will be adopted in the research to be reported here, following claims that this represents the most appropriate way to study the effect experimentally (Warren, 1983). Experimental work on perceptual restoration is thus a way of addressing an important real-world problem: that of stability of auditory perception in noisy environments.

There are two separate but related components of the perceptual restoration process (Bashford, Riener & Warren, 1992). First, if noise is able to hide the absence of missing fragments, then what is in reality a fragmented signal can be perceived to be continuous. This restoration of *continuity* is described as a low-level process (Drake & McAdams, 1999), where replacing missing sections of a signal with noise allows intact fragments to be treated as parts of the same auditory object (Plack & White, 2000). Second, for complex auditory signals such as speech and music, improvements in *intelligibility* as a result of replacing missing sections with noise occur in addition to the perception that the signal is continuous (Başkent, Eiler & Edwards, 2009). Thus, the content of what is missing is restored; this restoration of contextually-appropriate missing content requires higher-level processes, such as the influence of previous knowledge.

Even in optimal listening conditions, interpreting auditory signals such as speech and music involves a process of inference, with evidence for inferred content provided by the auditory input itself. This processing strategy represents an important adaptation to situations where input can be ambiguous or disrupted (Ellis, 1999), where inferences based upon listeners' prior knowledge of the signal content contribute to interpretation of the input. Importantly, such theories of how perception operates do not propose that perception is solely determined by higher-level processes, but that these higher levels interact with lower-level perceptual processing to determine what is perceived (Aslin & Smith, 1988). It is therefore not surprising that both top-down and bottom-up processes are considered to be involved in perceptual restoration[2]; top-down processes are involved in the restoration of the content of the missing fragments based on what is expected to be present, whilst the noise itself provides bottom-up confirmation of the restored signal by masking its absence (Samuel, 1981a). The closer the acoustic match between the extraneous noise and the properties of the missing segment, the stronger the perceptual restoration effect (Warren, 1999). This is because noise that is more similar to the sounds it is replacing is able to provide stronger bottom-up confirmation of expectations generated from higher levels.

The perceptual restoration of speech sounds, known as phonemic restoration, was first documented by Warren (1970; following brief reports by Cherry & Wiley, 1967, and Miller & Licklider, 1950). In Warren's study, listeners were presented with the sentence,

"*The state governors met with their respective legislatures convening in the capital city*".

The first *s* in the word *legislatures* was excised from the recording and replaced with a cough or tone. Listeners were given a typewritten copy of the sentence, and were asked

---

[2] Top-down processes in this context refer to the influence of knowledge on perception; bottom-up processes refer to extraction of the stimulus properties.

to locate the position of the cough or tone and identify whether a phoneme was missing. The results demonstrated the powerful effect of perceptual restoration; listeners were not able to complete either task. The missing phoneme was heard to be present, and was indistinguishable from sounds actually present in the sentence. Accurate localisation of the extraneous sound was difficult because it seemed to coexist alongside a seemingly intact sentence.

In contrast, when the same phoneme was excised and replaced with a silent gap, listeners were proficient at locating the gap in the sentence and identifying the missing sound. This indicates that the presence of noise, by masking the absence of the missing sound, allows the brain to fill in the missing segment, using the surrounding intact information to determine what is expected to be present. Context alone is not able to influence the restoration of missing phonemes, as perceptual restoration does not occur when there is no bottom-up confirmation of the expected content; this confirmation is not provided when silence replaces the missing phoneme.

It is important to note that there are limits to the intelligibility advantage provided by noise. Noise cannot enhance intelligibility of auditory signals if the noise is a constant, background noise capable of masking a stimulus completely (e.g. Elliott, Connors, Kille, Levin, Ball & Katz, 1979). Noise can only assist the restoration of brief temporal gaps in an auditory signal; the process of “filling in the gaps” requires the accurate perception of some of the intact contextual information surrounding the missing fragments.

This initial documentation of the phonemic restoration effect provided the impetus for many strands of research aiming to understand this auditory effect and the conditions under which it occurs. Most studies took one of two approaches to studying perceptual restoration, broadly corresponding to the distinction made above between bottom-up and top-down processes in perceptual restoration. In the context of perceptual restoration, top-down processes are related to the expectation part of the effect, and bottom-up processes relate to the confirmation part of the effect. Some research therefore has been primarily concerned with the stimulus-driven properties of the process and the bottom-up confirmation of restoration as a function of the nature of the extraneous noise, and has experimentally manipulated the nature of the noise and the incoming perceptual information. Other research has focused on the role of top-down processes in the effect, and how expectations generated by lexical knowledge or contextual factors influence the content of the restored portion of a signal. In order to investigate this, the contextual cues provided by the signal were manipulated. This distinction pertains to experimental methods; however, these two accounts are not mutually exclusive. The process of perceptual restoration itself is fundamentally dependent on the reciprocal interaction between top-down and bottom-up influences; Samuel (1981a) clearly identified the nature of this interaction:

> Phonemic restoration occurs when the bottom-up signal (e.g., white noise) is sufficient to *confirm* the presence of a schema that is *expected* on the basis of context (e.g., the rest of the word). There is a trade-off between these two factors; when expectations are strong, less confirmation (i.e., an acoustic signal less like the actual phoneme) is needed for restoration, and vice versa. (Samuel, 1981a, p.481)

Many descriptions of perceptual restoration include contributions from both top-down and bottom-up processes. For example, Verschuure and Brocaar (1983) explain perceptual restoration in terms of an automatic interpolation mechanism that uses contextual cues to infer the content of any missing information in an auditory signal. This mechanism, however, is only able to operate when it is not possible to perceive that any information is missing, such as when noise replaces any missing fragments. This account acknowledges the importance of the interaction between bottom-up (acoustic) and top-down (contextual) influences. Research evidence on the effect of acoustic and contextual influences on perceptual restoration in adult listeners will now be discussed in turn.

The effect of acoustic factors on perceptual restoration

The operation of perceptual restoration is affected by the acoustic properties of the input presented to the perceptual system. Experimental research into acoustic influences on perceptual restoration has largely involved testing the effect of manipulations based on the auditory conditions required for perceptual restoration to occur. One such condition requires that the amplitude of the replacement noise must be greater than the amplitude of the signal it is replacing, so that the absence of the missing sound can be masked (e.g. Kashino, 2006). Furthermore, the closer the match between the spectral characteristics of the noise and of the signal, the stronger the effect of perceptual restoration. The role of bottom-up factors in perceptual restoration is to provide confirmation of the expected content of missing information by masking its absence, hence bottom-up, acoustic manipulations vary the degree to which noise provides this confirmation. The greater the confirmation, the stronger the perceptual restoration.

It is quite plausible that noise is better able to provide bottom-up confirmation of a shorter missing fragment in an auditory signal than a longer interruption. Indeed, there is evidence in the literature that the duration of the missing information influences perceived continuity of interrupted tones, with shorter durations of missing fragments resulting in a stronger perception of continuity (Elfner & Caskey, 1965; Elfner & Homick, 1966). For example, Thurlow and Elfner (1959) presented listeners with a tone that was periodically alternated with a louder tone. The softer tone was perceived by listeners to continue through the louder tone, and this perception was affected by the duration of the interruption. Continuity effects were reduced by increasing the duration of the interruption. Similarly, Dannenbring (1976) reported how a tone was perceived by listeners to continue through a noise burst, but only where the noise was a few hundred milliseconds in duration. Beyond this value, the tone was perceived as discontinuous.

Listeners in a study by Shahin, Bishop and Miller (2009) were presented with words and pseudowords where one phoneme had either been removed and replaced by noise, or had noise superimposed upon it. The task was to report the words presented as either continuous or interrupted. The duration of the noise burst in each word was varied, but by a proportional amount in relation to the duration of the missing phoneme in each individual word, rather than by some fixed temporal value. When the noise burst replaced a small proportion of the word, the stimulus was more likely to be reported as continuous, whilst stimuli were reported as interrupted if a longer proportion of the word was replaced by noise. This suggests that stimuli with shorter missing fragments replaced by noise are more likely to be perceived as continuous than stimuli with longer missing fragments replaced by noise.

Further evidence of gap duration effects on perceptual restoration was obtained from a study conducted by Warren and Obusek (1971). The method used was similar to an earlier study by Warren (1970), where a phoneme or phonemes were removed from a sentence and replaced by a cough. Listeners were required to locate the position of a cough and to report whether they thought that the cough replaced the phoneme(s), that is, whether they perceived the signal as continuous or fragmented. Warren and Obusek found that listeners who were presented with the sentence where a cough replaced just one phoneme (the first *s* in the word *legislatures*) made more errors in localising the sound than listeners who were presented with the sentence where a cough replaced three phonemes, that is, an entire syllable of the word (*gis* in the word *legislatures*). This suggests that the cough restored the single phoneme to a stronger degree than the syllable, thus making it harder for listeners to locate the cough accurately. Warren and Obusek found that the duration of the interruption had no clear effect, however, on listeners' reports of the continuity of the signal.

These studies suggest that the shorter the interruption, the stronger the perception of continuity of a disrupted signal. Further research demonstrated that the effect, at least with reference to speech signals, might be determined by the size of a meaningful unit within speech, rather than a fixed temporal duration. Bashford and Warren (1987b) removed periodic fragments from spoken sentences of everyday speech and replaced them with either white noise or silence. They tested listeners' perception of continuity of disrupted speech by instructing them to adjust the rate of alternation between intact and disrupted sections to the longest duration of interruption at which the signal appeared to be continuous. Thus, the use of this method made it possible to determine the maximum duration of interruption that could be made to sound continuous through replacing it with white noise or silence.

When missing fragments were replaced by white noise, speech signals were perceived as continuous if the duration of missing fragments was no longer than 304 ms. When the gaps were longer than this value, listeners could detect the gaps in the signal and continuity was not perceived. Bashford and Warren noticed that this duration of interruption was very close to the average duration of the words in their sentences. In stark contrast, when the same missing fragments were replaced by silence, gaps were easily detected if they were any longer than 52 ms. The clear difference in duration thresholds in the noise and silence conditions clearly demonstrates the advantage that noise has over silence in replacing missing information; noise is able to make sound continuous a duration of interruption that would clearly sound fragmented if replaced by silence.

When isolated word lists were presented instead of coherent sentences, the threshold for gap detection was much lower for the noise condition, with any gaps longer than 161 ms being easy to detect. Interestingly, there was little difference to the threshold for gap detection in the silence condition; for isolated word lists, silent gaps were not detected by listeners if shorter than 61 ms. This indicates that the additional semantic context provided by a sentence increases the ability of noise to restore the continuity of fragmented signals; the effect on signals with silent gaps is much less pronounced.

Bashford, Meyers, Brubaker and Warren (1988) used the same threshold adjustment procedure that Bashford and Warren (1987b) employed. They aimed to investigate further the close relationship observed by Bashford and Warren (1987b) between the maximum gap duration at which signals were perceived to be continuous when replaced by noise and the average word length in the sentences presented. To this end, they tested listeners' continuity thresholds for the sentences spoken at a natural rate, and speeded up or slowed down by 15%. If a fixed temporal value is important in gap duration effects, then similar thresholds for the perception of continuity would be expected in all cases. If instead it is the unit size (i.e. the average duration of a word) that is important, then the threshold for continuity should change in accordance with the change in average word duration. This is exactly what Bashford et al. found; when the speech rate was increased by 15%, the average word length was reduced and so was the threshold for continuity (the threshold reduced by 18% in the noise condition and by 7% in the silence condition). Similarly, when the speech rate was decreased by 15%, the threshold for continuity increased in accordance with the increased word duration (an increase of 12% in the noise condition and of 15% in the silence condition was evident). These results suggest that in the case of speech perception, when the duration of a missing gap in a speech signal exceeds the average duration of one complete word, the ability of noise to restore the continuity of the signal is significantly reduced.

The effect of gap duration on the restored *intelligibility*, as well as perceived *continuity*, of fragmented speech signals has also been investigated. Powers and Wilcox (1977) presented listeners with sentences of everyday speech that were interrupted by gaps of either noise or silence. The rate of interruption was varied (gaps were between 130 and 650 ms in duration), but at all times the ratio between what was intact and what was removed remained the same (i.e. regardless of the gap duration, the missing sections were the same length as the intact sections). Listeners were instructed to write down as much of the sentences as they could. Regardless of the rate of interruption, noise as a gap replacement produced enhanced intelligibility of the sentences compared to silence, yet the effect of gap duration on the size of this advantage in intelligibility was not consistently linear in nature. Importantly, noise-filled gaps provided the greatest advantage over silence-filled gaps for a gap duration of approximately 333 ms. This closely corresponded to the average word duration in the sentences presented, hence when an entire word was missing it was very difficult to understand sentences when missing sections were replaced by silence. This finding, together with the results reported by Bashford and Warren (1987b) and Bashford et al. (1988), appears to suggest that the relationship between the duration of missing fragments and the unit size in speech is an important determinant of the effect of gap duration on the perceptual restoration of speech signals.

Taken together, these results suggest that the duration of what is replaced by noise influences the strength of the restoration of both continuity and intelligibility of disrupted signals. Noise is better able to provide bottom-up confirmation of the expected content of a short fragment; plausibly, the longer the duration of the missing information the harder it is for the noise to hide the absence of the missing information. A further explanation for gap duration effects on perceptual restoration relates to the fact that long gaps serve an ecological purpose in our perception of sounds such as speech, by separating syllables and words. Therefore, if these long gaps were completed by the perceptual system, this would, in some cases, represent an erroneous and inappropriate assumption (Remijn et al., 2007).

The duration of missing fragments is not the only acoustic factor that has been found to exert an influence on perceptual restoration. Research has demonstrated that an important condition for restoration to occur is that the replacement sound (e.g. white noise) be larger in amplitude than the signal it is replacing in order for it to act as a potential masker and obscure the absence of the missing fragment. The "masking potential" of the noise (Kashino, 2006) is described as a critical factor in perceptual restoration. Simple continuity experiments have confirmed the role of the masking potential of replacement noise; perceived continuity of fragmented signals increases in strength with masking level (e.g. Riecke et al, 2008; Riecke, van Opstal, Goebel & Formisano, 2007; Thurlow & Elfner, 1959). This effect plausibly occurs because, by masking the absence of missing fragments, noise provides strong bottom-up confirmation of the expected content of the missing sections.

The effect of replacement noise amplitude on perceptual restoration has been demonstrated further in experiments where the amplitude of the replacement noise in relation to the amplitude of the signal has been manipulated. Warren and Obusek (1971) investigated the perceptual restoration of a single excised phoneme where the replacement sound was either a silent gap, a cough, a loud tone or a loud buzz (8 dB above the peak intensity of the speech signal), or a soft tone or a soft buzz (of equal intensity to the speech signal). Listeners were asked to locate the replacement sound and report whether any speech sounds were missing (i.e. whether the signal was continuous or not). For all replacement sounds, perceptual restoration of the missing phoneme was stronger than when the missing phoneme was replaced by a silent gap, as shown by increased error in localisation of the extraneous sound. With reference to the individual sounds, the loud tone produced the strongest perceptual restoration as measured by listeners' initial judgements of whether or not the signal was continuous.

In total, the listeners were required to make this judgement four times. Whilst the loud tone always produced stronger restoration than the quieter tone across judgements, the effect of the amplitude of the buzz was inconsistent across judgements. Warren and Obusek concluded that no reliable intensity differences were observed. Moreover, both noise amplitudes were sufficiently loud to produce strong restoration (Samuel, 1981b), hence making it difficult to interpret the results of this amplitude manipulation. In addition, it is possible that, in hearing the signals four times, listeners misinterpreted the task demands and believed it unlikely that four trials would all require the same response. Being sensitive to listeners' expectations of a seemingly ambiguous task might represent an important consideration in the design of perceptual restoration experiments.

Further evidence that the amplitude of the replacement noise affects the perceptual restoration of speech signals was obtained by Powers and Wilcox (1977), who manipulated the intensity of the replacement noise in relation to the intensity of the speech signal to test its effect on the restored intelligibility of periodically interrupted sentences. Whatever the signal to noise ratio, the insertion of noise into missing gaps resulted in enhanced intelligibility of the speech when compared to intelligibility of speech with silent gaps. Furthermore, the level of intelligibility increased as a function of the intensity of the noise, where noise of greater intensity increased the improvement in intelligibility provided by the replacement noise, relative to silent gaps.

Similar results were reported by Eimas, Tajchman, Nygaard and Marcus (1996), who also investigated the effect of the amplitude of the replacement noise on phonemic restoration. As well as presenting signals with silent gaps, Eimas et al. compared the ability of standard white noise (higher in amplitude than the signal) and low amplitude noise (lower in amplitude than the signal) to restore missing sections of fragmented speech signals. Intelligibility of signals where missing fragments were replaced with low amplitude noise was poor, and showed little advantage over signals with silent gaps. In contrast, the standard white noise restored the missing fragments to the listeners, with the result that the speech sounded continuous.

It therefore appears that the duration of the interruption and the relationship between the amplitude of the noise and the amplitude of the signal both affect the ability of noise to restore continuity and intelligibility of auditory signals. Many researchers have also manipulated the nature of the noise used in a perceptual restoration experiment, to investigate if some types of noise produce stronger perceptual restoration than others. The crucial consideration here is that the acoustic properties of the noise should be as close a match as possible to those of the signal fragment the noise is replacing, in order to restore continuity and intelligibility to the highest degree, by providing strong bottom-up confirmation of contextually-generated expectations.

Whilst there are many examples of the extraneous sounds that could lead to perceptual restoration in everyday listening conditions (e.g. coughs, louder speech, machinery, environmental sounds, etc.), there are also many sounds that could replace missing sections of auditory signals in experimental situations. The effect does not only occur where white noise is used as a replacement for missing sections of a signal. Warren and Obusek (1971) investigated the ability of three types of replacement noise (tone, buzz and cough) to restore a missing phoneme. The measures of restoration were the degree of mislocalisation of the extraneous noise and listeners' reports of whether or not the noise replaced a phoneme (i.e. a continuity judgement). The listeners made four successive judgements of the continuity of the signals. All sounds replacing the phoneme led to stronger perceptual restoration than silent gaps (as demonstrated by increased error in localising the noise and fewer reports that the phoneme was missing). However, there were no obvious differences between the ability of these three different sounds to restore the missing phoneme, and there was little consistency across the four judgements. With reference only to listeners' first judgements, the cough and the tone produced stronger restoration than the buzz, suggesting a closer match between the speech and the cough and the tone than between the speech and the buzz. Nevertheless, without a consistent trend across judgements it is difficult to draw any firm conclusions from this study.

The potential problem with asking listeners to make four successive judgements in this study was outlined earlier. Furthermore, Samuel (1981b) offers a potential explanation for inconsistent effects of noise type on perceptual restoration. As only a single phoneme was removed, a high amount of context was present in the remainder of the sentence; thus, the measures of perceptual restoration were near ceiling for all noise types. This made it difficult to detect any differences in perceptual restoration caused by different types of noise. It is also highly likely that the effect is complex, and the ability of different sounds to induce perceptual restoration is dependent on the specific acoustic and contextual conditions present.

Layton (1975) replicated Warren and Obusek's study, but changed the method slightly, so that listeners were not aware that they would be hearing the same sentence four times in succession, as was the case in the original study. Layton suggested that as listeners in the original study were aware that they would hear four identical presentations of the same sentence, they might have anticipated where the critical missing phoneme would occur and directed their attention accordingly. This could have influenced their judgements and hence the results of the study. Layton compared the ability of a cough, tone and silent gap to restore a missing phoneme, and found evidence of superior perceptual restoration when a cough was the replacement noise when compared to a tone. A cough is acoustically the most similar to the speech, hence this supports the proposition that a close acoustic match leads to stronger perceptual restoration. Layton also offers another potential explanation; it may not only be the acoustic similarity between the noise and the signal that determines the strength of perceptual restoration, but also listeners' experience of a particular sound as a masker. In everyday listening conditions, a cough is frequently heard as a masker, whereas the tone is an artificial laboratory sound.

Similarly, Samuel (1996) reported differences between the perceptual restoration of phonemes when a tone was used as the replacement sound when compared to the usual white noise typically used in such experiments. It was found that the tone was generally less effective than the white noise in restoring intelligibility of the missing phoneme. Samuel concluded that the ability of replacement noise to restore a missing fragment of an auditory signal is not simple; instead, the strength of perceptual restoration is dependent on a complex interaction between the characteristics of the replacement sound and the replaced sound. This conclusion was supported by findings from Bashford, Warren and Brown (1996). In an experiment investigating the ability of noise to restore periodic interruptions in sentence stimuli, they found that speech-modulated noise (where the amplitude envelope of the replacement noise matched that of the speech) produced greater improvements in intelligibility than the typical white noise. Bashford et al. suggested that this was because speech-modulated noise provided important clues as to the identity of the missing speech. Furthermore, manipulating the centre frequency of the replacement noise has also been found to have an effect on perceptual restoration. Bashford et al. (1992) found that restored intelligibility for periodically interrupted sentences was greatest when the replacement noise (pink noise) had the same centre frequency as the speech signal. These experiments demonstrate that the match between the replacement and replaced sounds is crucial in inducing strong perceptual restoration.

This effect of a close acoustic match between replacement and replaced sounds on perceptual restoration has also been demonstrated not by manipulating the nature of the replacement noise, but by manipulating the nature of the replaced excised fragment to create the strongest match. For example, Samuel (1996) investigated the ability of white noise to restore missing phonemes; the strength of restoration was dependent upon the type of phoneme being restored. Perceptual restoration was strongest for fricatives and stops, which are argued to be the most acoustically similar to white noise. Similar results had been obtained in a study by Samuel (1981a). Here, both a tone and white noise were used as replacement noise. An interaction between the type of phoneme that was missing and the type of replacement noise was observed; vowels were better restored than fricatives when the tone was the replacement noise, and fricatives were better restored than vowels when the white noise was the replacement sound. This might explain earlier contradictory findings relating to the ability of different sounds to induce perceptual restoration; it is not only the relationship between the noise and speech in general that determines the strength of restoration, but the specific match between the individual sound to be restored and the replacement noise.

A similar investigation was reported by Samuel (1981b), and the results of the previous experiment were replicated. The ability of white noise to restore intelligibility of different phonemes was compared to that of a tone as the replacement sound. When white noise was the replacement sound, similar results were found to previous investigations, with fricatives and stops being restored to a greater degree than vowels. Conversely, when a tone was used as the replacement sound, vowels were restored to a greater degree than fricatives. The closer the match between the replacement noise and the phone class of the missing sound, the stronger the perceptual restoration.

Trout and Poser (1990) investigated the strength of perceptual restoration according to the type of missing consonant. It was predicted that voiced stops (/b, d, g/) would be restored by white noise to a lesser degree than voiceless stops (/p, t, k/), since voiceless stops and white noise share the characteristic of aperiodicity. Results supporting this hypothesis were obtained: the voiceless consonants were restored to a greater degree than the voiced consonants. It appears that only when the replacement sound is similar in nature to the sound it replaces, is it able to provide strong bottom-up confirmation of the expectations generated at a higher level. The effect of such high-level representations on perceptual restoration in adult listeners will now be discussed.

The effect of contextual factors on perceptual restoration

In an auditory signal such as speech, clues to the identity of any missing fragments can be provided by the context of the intact sections of the signal. Context can generate lexically-consistent expectations as to the most likely content of any disrupted fragments. One further source of information about the identity of missing phonemes in speech comes from the cues provided by coarticulation, which provide redundancy in speech perception. Therefore, even when a phoneme or portion of a word is removed, the phonemic cues surrounding the excised fragment can provide strong clues to the identity of the fragment. However, this phonemic encapsulation can be overridden by listener's expectations, as Warren (1970) and Warren and Warren (1970) report that even when these cues are removed along with the excised fragment, perceptual restoration still occurs. Furthermore, Warren and Sherman (1974) demonstrated that even when a phoneme is mispronounced prior to deletion (e.g. a stimulus word is produced as *commu/t/icating*), listeners restore the correct phoneme (/n/) despite the surrounding phonetic cues still present indicating that this is not the correct restoration. This suggests that other, stronger top-down contextual factors must exert an influence on perceptual restoration.

There is evidence in the literature that speech perception is more complex than a simple acoustic analysis of the present elements of speech. Top-down processes are also thought to exert an influence. According to some theorists (e.g. Marslen-Wilson & Warren, 1994), perception is not merely the result of a bottom-up flow of information from one level to an increasingly more abstract level. It is posited that higher-level representations, such as lexically-generated expectations, can exert an influence on lower-level perceptual analysis. The operation of top-down support for speech perception provided by previous knowledge is particularly true of noisy listening conditions (Hannemann, Obleser & Eulitz, 2007). Top-down processes in phonemic restoration are purported to influence the content of what is restored, based on a lexical expectation generated by the context within which the missing fragment occurs. By imposing contextual constraints on the content of what is restored, top-down processes ensure that only contextually-appropriate content is restored, allowing the operation of perceptual restoration to be not just adaptive but also selective.

Srinivasan and Wang (2005) propose a schema-based model of phonemic restoration, where the context in which a missing phoneme occurs is crucial in the process of its restoration. Intact information in the acoustic signal activates lexical representations, which are used to reconstruct the missing phoneme. Thus, when a missing phoneme is heard as present, this is not so much a reflection of its own attributes as the attributes of the context in which it occurs (Hawkins, in press).

Crucially, the effect of linguistic context on the perceptual restoration of speech is not restricted to the contextual information that precedes the missing fragment. Warren and Warren (1970) report how subsequent context can also influence the restoration of missing phonemes. Listeners were presented with the sentence, "*It was found that the *eel was on the* ____", where the symbol * represents a cough that replaced the missing fragment of the word. This incomplete word could be *wheel*, *heel*, *peel* or *meal* depending on whether the word at the end of the sentence was *axle*, *shoe*, *orange*, or *table*, respectively. The contextually correct word was restored by listeners, even though the disambiguating context did not occur until later in the sentence. Warren and Warren (1970) and Warren and Sherman (1974) proposed that the partially complete input is stored until the context allows the correct phoneme to be restored.

The role of top-down contextual influences in the perceptual restoration of speech was investigated in a series of experiments reported by Samuel (1981a). In Samuel's first study, participants were presented with a version of a two-, three-, or four-syllable word of high or low lexical frequency in which a phoneme had been excised and replaced with noise, or a version of the same word where the phoneme had not been removed, but the same noise was superimposed upon it. Participants were required to report whether or not the stimulus was intact (noise superimposed on the missing phoneme; "added" items) or not (noise replaces missing phoneme; "replaced" items). This paradigm made an important contribution to the literature on perceptual restoration. For the first time, through the application of signal detection analysis to perceptual restoration, it was possible to obtain a measure of perceptual restoration that was not influenced by a postperceptual bias[3]. This method considers both miss rates (reporting a stimulus with a missing phoneme replaced by noise as intact, suggestive of perceptual restoration of the missing phoneme) and false alarms (reporting an intact item as containing a missing phoneme replaced by noise).

The rationale of this method is as follows. If perceptual restoration is occurring when the phoneme is replaced with noise, then the resulting percept should sound identical to the version where noise is added to the stimulus, and as such is likely to be

---

[3] A postperceptual bias in this context represents listeners' tendency to report items as intact even if they did not perceive them to be so, as the result of a postperceptual decision process. Having both a miss rate and a false alarm rate means that this bias can be factored out in the analysis. A bias towards reporting items as intact makes it unlikely that the effect is truly perceptual in nature.

reported as intact. The stimulus with noise added to the critical phoneme is objectively intact, whereas the stimulus with noise replacing the critical phoneme is subjectively intact (Samuel and Ressler, 1986). If the phoneme is not being restored in the replaced version of the stimulus, it should sound quite different to an intact version of the stimulus with noise added, and hence will most likely not be reported as intact. The effect of experimental manipulations can be separated into a discriminability measure, which represents how similar the added and replaced stimuli sounded, and a measure of response bias representing a postperceptual decision to report stimuli as intact. Perceptual restoration is thus demonstrated by a low value of the discriminability measure, indicating that the restored and "real" phonemes sounded alike.

If lexical factors do influence perceptual restoration, then there should be greater evidence of perceptual restoration for high frequency words, which have a stronger lexical representation than low frequency words. Although small, an effect of word frequency was obtained in the predicted direction, where a missing phoneme from a high-frequency word was more strongly restored than a missing phoneme from a low frequency word. This suggests that the stronger top-down influences operating for the high frequency words helped to restore the missing phoneme, in turn suggesting that there are top-down influences on perceptual restoration. Furthermore, greater evidence of perceptual restoration was obtained for missing phonemes from the four-syllable words than from the shorter words, supporting the idea that the greater the amount of intact contextual cues available, the stronger the perceptual restoration. This effect of word length was also replicated by Samuel (1996), and is consistent with evidence of stronger lexical activation with increasing word length (Pitt & Samuel, 2006).

The position of the removed phoneme within the word was also manipulated by Samuel (1981a), with the expectation that a phoneme located at the end of the word would be restored to a greater degree than a phoneme in a beginning or middle word position, as there is more preceding context to influence perceptual restoration. What Samuel actually found was that phonemes occurring in the middle of a word were the most strongly restored, most probably due to an interaction between the amount of preceding context and the detrimental effects of backward masking. Furthermore, the phoneme position measure did not have its greatest effect on the measure of discriminability, but instead on the measure of postperceptual bias. It therefore appears that lexical factors alone cannot account for the effects of this manipulation.

In a second experiment reported by Samuel (1981a), using the same discrimination paradigm, the perceptual restoration of a missing phoneme from words was compared to the perceptual restoration of a missing phoneme from phonologically-legal pseudowords. The purpose of comparing the restoration of words and pseudowords is that only the words have a lexical representation. If top-down lexical factors do influence phonemic restoration, there is reason to believe that missing phonemes from words should be restored to a greater degree than those from pseudowords, which do not have an entry in the mental lexicon. Moreover, this manipulation potentially rules out the possibility that lexical effects on perceptual restoration are simply due to knowledge of phoneme sequences, since legitimate phoneme sequences are present in both words and phonologically-legal pseudowords, but only words have a lexical entry. Samuel found evidence of greater perceptual restoration of missing phonemes from words (e.g. *basis*) as opposed to pseudowords (e.g. *pafis*), across differences in word frequency and word length. This supports a strong top-down influence of lexicality on restoration. It is also important to note that this was a truly perceptual effect; the measure of postperceptual bias was not affected by these factors.

Samuel's third experiment took these manipulations a stage further to examine the greater effect of top-down influences on perceptual restoration provided in a sentential context. One important modification was made to the procedure in order to investigate the effect of context that a sentence provides. With the single word experiments previously reported, the restored phoneme would always lead to a lexically unique word; that is, there was only one possible word that could be produced through restoring the missing phoneme. In the sentence experiment, restoration could produce at least two different possibilities as to the word. For example, the words *battle* and *batter* constituted a word pair; the context of the sentence presented meant that only one of the two words was semantically correct. The final syllable was removed, leaving the fragment *bat**. The sentence either referred to a soldier and a *battle* or a pitcher and a *batter*.

As well as being required to judge whether each stimulus was intact (noise superimposed on top of a phoneme) or not (noise replaces a missing phoneme), listeners were also required to make a forced-choice judgement reporting which word they heard to be present (either *battle* or *batter*). This was an ingenious way to test whether the context was powerful enough to influence the content of the restored fragment even in the case of lexical ambiguity as to the correct identity of the missing fragment. In order to make this a test of the effect of predictability on restoration, listeners were also presented with sentences where the word to be restored was not predictable given the context: these sentences referred to a soldier and a *batter* or a pitcher and a *battle*.

The results obtained, on first inspection, appear counterintuitive. The discriminability measure indicated that it was easier to distinguish between added and replaced versions of the stimuli in predictable sentences, thus indicating a lesser degree of perceptual restoration of the missing fragments with the increased predictability provided by congruent sentential context. Furthermore, Samuel reports that the contextually-appropriate word was reported in the forced-choice word judgements. In addition, a response bias was evident; words with missing fragments replaced by noise were more likely to be reported as intact if the word was predictable based upon the context.

Whilst this appears to suggest that when higher-level contextual information is provided in sentential form, it does not provide sufficient top-down information to lead to perceptual restoration (as evidenced by the discriminability measure), Samuel offers a plausible explanation for this anomalous finding. It is proposed that, since a high degree of contextual information is provided, this leads to extremely high predictability for the word containing the missing fragment. This high predictability reduces the load on the perceptual system, leaving a larger amount of processing resources available to determine whether the presented stimulus is an added or replaced item.

The role of predictability in the perceptual restoration of a word within a sentential context remains ambiguous for a number of further reasons. First, this effect would be better studied using a stimulus with multiple interruptions rather than a single excised fragment, since the amount of context is not equated with the amount of missing information. It is possible that a silent gap in the same place as the replaced/added fragment would have led to the same degree of comprehension, since the large amount of context facilitates understanding.

Second, it is also important to determine the advantage of replacement noise over silence in enhancing intelligibility of interrupted sentences. The ability to discriminate between added and replaced versions of a stimulus (from which it would be inferred that restoration was not occurring) does not imply that the replaced version would not restore intelligibility relative to the same stimulus with a silent gap. Indeed, using repetition accuracy as a measure of perceptual restoration, Bashford et al. (1992) found greater evidence of perceptual restoration for high predictability sentences than low predictability sentences, when missing fragments were replaced with noise rather than silence. Third, in only requiring a simple judgement about a single phoneme to be made, the replaced/added discrimination paradigm does not measure intelligibility increases produced by perceptual restoration directly (because no measure of intelligibility is made), and as shown here, the discriminability measure can be affected by processing factors that result in ambiguous conclusions.

Further evidence for top-down contextual influences on perceptual restoration comes from findings that the strength of perceptual restoration is reduced in the case of missing sections from lists of isolated words with no semantic connections between them (e.g. Bashford et al., 1992; Bashford & Warren, 1987b; Bashford et al., 1996). This suggests that the context in which a missing fragment occurs is crucial in generating expectations as to what would most likely be the content of the missing section. Thus, expectations appear to be a fundamental part of the perceptual restoration process.

In a further experiment, Samuel (1987) was interested in whether the top-down influences exerted on perceptual restoration varied according to the number of items that could legitimately be restored given the contextual information. This was investigated by comparing the restoration of missing phonemes from lexically unique words (only one possible restoration; e.g. **esion* can only be *lesion*) versus lexically ambiguous words (several possible restorations; e.g. **egion* could be *legion* or *region*), using the added/replaced discrimination paradigm. There are two competing hypotheses as to how lexical uniqueness could influence perceptual restoration. With a lexically ambiguous item, it is plausible to suppose that perceptual restoration should be stronger because, as there are multiple candidates for the restored item, it is more likely that a lexical representation will be activated. Conversely, it is equally plausible that the perceptual restoration mechanism is designed to produce a unique product of restoration, and so will be detrimentally affected when there are multiple possibilities as to the content of the missing section.

What Samuel actually found was that perceptual restoration was strongest when stimuli were lexically ambiguous, thus supporting the former hypothesis. Lexically unique stimuli were found to result in a postperceptual bias towards reporting the stimuli as intact. However, here the reduced perceptual restoration of lexically unique stimuli was inferred from a greater value of the discriminability measure between replaced and added items. Therefore, this can be explained with reference to processing resources; unique items do not require processing resources to disambiguate between possible candidates for restoration, thus leaving greater resources for the process of discrimination. In addition, with only one possible lexical match for a lexically unique item, it is more likely that a match will not be made, than in the case where there are multiple possibilities for the identity of the missing fragment. Hence, without access to a lexical representation, perceptual restoration does not occur for lexically unique items to the same degree as lexically ambiguous items.

It is important to note that one of these explanations refers to increased discrimination and one to reduced perceptual restoration. The argument for increased discrimination is primarily related to processing factors and hence it could be argued that it has little to do with perceptual restoration. Norris (1995) describes the misleading nature of the application of signal detection theory to perceptual influences, and even Samuel himself notes the problem with applying signal detection theory to perceptual restoration:

> In applying signal detection techniques to such high-level domains (rather than to simple signal detection), there are many ways that apparently minor procedural details can affect the outcome.
> (Samuel, 1996; p.48)

Samuel (1987) also reports a second study looking at the effects of another type of lexical uniqueness, namely, temporal uniqueness. At the beginning of hearing a word, there are multiple candidates for its identity but as the word proceeds, the number of candidates reduces until only the unique lexical representation remains. This is referred to as temporal uniqueness, as it is dependent on the point in time of the word being presented. To investigate whether perceptual restoration is sensitive to temporal uniqueness, Samuel compared the perceptual restoration of words with common or uncommon first syllables. Greater restoration was found for missing phonemes from words that were unique as a result of the first syllable ("early-unique" words; e.g. *veg*____, *vamp*____) as opposed to "late-unique" words that only became unique as a result of the presentation of a later syllable (e.g. *com*_____, *dis*____). In both of these experiments, the critical uniqueness effects were found only for words and not for matched pseudowords, and the method used affords the conclusion that these are true perceptual effects. The results imply that there is a critical role for the lexicon in both speech perception in general and perceptual restoration in particular.

Further evidence for the role of lexical factors in phonemic restoration was obtained from an extensive study by Samuel (1996), using the added/replaced discrimination paradigm. Again, the focus of these experiments was to investigate lexical influences on perceptual restoration by comparing the perceptual restoration of missing phonemes from words and pseudowords. Under some conditions, the predicted lexical effect was found (greater perceptual restoration for words than pseudowords), but this effect was only weak under other conditions (depending on the type of phoneme that was missing and the position of the phoneme within the word, together with the word length). Samuel nevertheless argues that lexical influences on phonemic restoration are real, and suggests a possible reason for the inconsistent findings. The method of measuring the discrimination between added and replaced items, and testing the effect of lexicality (words versus pseudoword) on the size of this discrimination, assumes that increased lexical activation from words as opposed to pseudowords affects the items where noise replaces a missing phoneme exclusively. Even though the item where noise is superimposed on an intact phoneme does not have a missing phoneme, masking noise is still present. Thus, lexical activation might affect the perceived intactness of both added and replaced items, meaning that a discrimination difference as a result of lexical manipulations would not be observed.

To control for this, Samuel made a small change to the method, where listeners were asked to rate how similar either the replaced or the added stimulus was to an intact version of the word. This allowed for the effect of lexical activation on manipulated (noise added or noise replaced) versus intact words to be tested. With this modification, more reliable lexical influences were demonstrated. There was greater evidence for perceptual restoration of missing phonemes in words versus pseudowords; words were more likely to be rated as intact than pseudowords. Interestingly, this was not the case when the initial phoneme was a fricative or a stop, with white noise as the replacement or added noise. Here, the match between the phoneme and the noise was so close, that the bottom-up confirmation component of the effect was immensely strong. Due to the trade-off between top-down and bottom-up influences in perceptual restoration, it was not possible to detect the effects of top-down influences, since the bottom-up component was accounting for almost full perceptual restoration. This result highlights the importance of considering both aspects of the effect when interpreting research findings.

Samuel (2001) claimed that the top-down flow of information in perceptual restoration, in activating a lexical representation, means that phonemes can be identified when the acoustic evidence presented to the auditory system is incomplete. Extensive research evidence indicates that restored phonemes sound like real phonemes (to the extent that the two are often indistinguishable from each other). In order to determine whether lexical activation is producing a “robust phonemic percept” (Samuel, 1997, p.98) in the process of perceptual restoration, experiments have combined perceptual restoration studies with investigation of other perceptual phenomena. This is to see if restored phonemes “behave” in the same way as intact phonemes in perceptual tasks, where no explicit phonemic decision is required. In using a secondary test, it is possible to see if top-down influences are the same with a restored stimulus as an intact stimulus where there is direct acoustic stimulation.

Samuel (1997) reports a study that investigated the ability of words with a phoneme removed and replaced with either noise or silence to produce adaptation effects (where repeatedly presenting a sound affects the perception of similar sounds). Only the stimulus with a noise-replaced phoneme produced adaptation effects, which mirrors the behaviour of intact words; crucially, where perceptual restoration did not occur, neither were adaptation effects observed. This is strong evidence for the role of top-down lexical influences on perceptual restoration; the phonemic percept produced through the process of perceptual restoration functions in the same way as a “real” phoneme would in perceptual processing.

Similarly, Obusek and Warren (1973) used perceptual restoration stimuli in the place of intact stimuli to see if secondary perceptual effects occur, here using the verbal transformation effect (where repeated presentation of a stimulus changes how it is perceived). Verbal transformations occurred more for noise-replaced stimuli, where a phoneme was replaced by noise, than silence-replaced stimuli, where the phoneme was replaced by silence. This is further evidence that top-down processes operate in the perceptual restoration effect and that the phonemic percepts that these top-down processes produce not only sound like real phonemes but also function in a similar way.

Samuel and Ressler (1986) propose that the influence exerted on phonemic restoration by top-down factors is due to the fact that attention is misdirected to the global level (the context), in an analogous way to visual illusions such as the Müller-Lyer illusion. They attempted to redirect attention to the phonemic (local) level using a cueing paradigm. By redirecting attention to this level, it might be possible to shut-off the top-down flow of expectations, and thus reduce the restoration of missing phonemes. If this is possible, it suggests that phonemic restoration is an auditory illusion created by the top-down flow of lexical expectations. Results to this effect were obtained: when listeners were given an attentional cue directing them to the location and identity of the critical phoneme that was either replaced by noise or had noise superimposed upon it, it was easier to discriminate between added and replaced versions of the stimulus. This suggests that perceptual restoration was not occurring to the same degree as when no attentional cue was given, or only a single cue was provided. Information about both the location and the identity of the critical phoneme was required. It is suggested that in the phonemic restoration effect, expectations of what might be present in the stimulus are generated on the lexical level, and insufficient processing is directed at the local level to disconfirm this expectation, and process that the expected sound might not in fact be present.

It is important to consider that, in using a discrimination paradigm, attentional cues would plausibly reduce processing demands on the system with the result that greater processing resources are available to make the discrimination between added and replaced items. Furthermore, this explanation is specific in reference to phonemic restoration, and hence may not apply to perceptual restoration in other domains.

Nevertheless, further evidence of the role of attention in phonemic restoration was obtained in a study conducted by Samuel (1991). Samuel presented listeners with thousands of trials where they were required to make the added/replaced judgement. Mere exposure did not lead to any increase in the measure of discriminability between added and replaced items, for either words or pseudowords. This suggests that perceptual restoration is an automatic process, and that repeated exposure does not allow listeners to override the powerful effect. It appears that more explicit training is needed, and that exposure alone is not sufficient to allow the reallocation of attention from a global to local level. Listeners need to be told how to reallocate attention, and to where.

In the next experiment reported by Samuel (1991), a variety of cues, increasing in degree of explication, were provided to direct attention to the local level and the critical phoneme prior to listeners' judgements of whether the critical phoneme was replaced by noise or had noise superimposed upon it. In the baseline condition, a series of dashes were presented prior to the auditory stimuli; the number of dashes corresponding to the number of letters in the word. No explicit attentional information was provided in this condition. In other conditions, the syllable in which the critical phoneme occurred was shown to listeners, the syllable and the position of the critical phoneme within the syllable were shown to listeners, or the identity of the word and the position of the critical phoneme within the word were revealed to listeners. The greater the level of information provided in the cue, the greater the discrimination of added and replaced items relative to the baseline condition. However, only providing information about the identity of the word and the position of the critical phoneme within the word resulted in a value of the discriminability measure that was significantly greater than baseline discrimination.

These experiments suggest that it is possible to overcome the effect of lexical activation in speech processing and phonemic restoration, but only if the information provided is highly explicit. The inability of other, less informative cues to override the effect of phonemic restoration indicates that it is a very strong effect that occurs automatically unless strong and explicit cues are provided to redirect attention.

Some experimenters have investigated audiovisual influences on restoration, and whether a concurrent visual presentation displaying the face of the speaker alongside the auditory stimulus provides extra top-down information and hence leads to stronger perceptual restoration. Trout and Poser (1990) used the added/replaced discrimination method as a measure of perceptual restoration in either an auditory only condition (where only the auditory stimulus was presented), or an auditory and visual condition (where the face of the speaker was presented concurrently). This was to test if the addition of visual information led to the generation of stronger expectations about the identity of the missing phoneme. A tone and a flash of light were used as the masker in the auditory and visual conditions, respectively. There is evidence in the literature that visual information enhances speech perception (Kim & Davis, 2003; Sumby & Pollack, 1954); yet Trout and Poser did not find that the addition of visual information affected perceptual restoration. Similarly, Repp, Frost and Zsiga (1992) found no extra utility from concurrent visual cues in the detection of speech in noise. Taken together, these findings suggest that visual influences on perceptual restoration are weak, and that the mechanism does not need concurrent visual information in order to operate.

The research presented in this review indicates that there are a number of factors that influence the operation of perceptual restoration, which can be broadly categorised as acoustic and contextual influences. The observed effects of these influences are not always consistent, and under some experimental conditions are only represented as weak effects. A key consideration that has emerged from discussion of these influences is that interpretation of research findings needs to consider the possible interaction and trade-off between the bottom-up and top-down processes that influence perceptual restoration. More detailed understanding of exactly how perceptual restoration operates has developed as new methods of study have been identified. The use of neurophysiological methods represents one way in which this has been achieved. A further strand of research, aiming to understand the neurophysiological basis of perceptual restoration, will now be described.

The neurophysiological basis of perceptual restoration

More recent research into perceptual restoration has attempted to elucidate the neural mechanisms underlying the phenomenon. This endeavour can provide further insights into how perceptual restoration operates, and how it differs from veridical perception. Such research is motivated by the ecological importance of perceptual restoration in reducing the detrimental effects of noisy environments on communication (Husain, Lozito, Ulloa & Horwitz, 2005).

In an investigation of the neural basis of the restoration of continuity of fragmented speech using functional magnetic resonance imaging (fMRI), Heinrich, Carlyon, Davis and Johnsrude (2008) tested for listeners' perceived continuity of words where missing vowels were replaced by white noise. Activity was evident in the posterior middle temporal gyrus and superior temporal sulcus, areas described as, "an anatomically distinct, objective physiological correlate of the continuity illusion in human listeners" (p.1737).

Petkov, O'Connor and Sutter (2007) attempted to find neural evidence for the perception of illusory continuity where missing sections of a stimulus are replaced with noise. The authors were looking for a neural representation of perceptual restoration that involved three separate neural responses that represent the perception of a continuous signal through a missing gap, where the noise hides the absence of the missing section. First, in order for neural evidence for the perceptual of illusory continuity to be displayed, sustained responders (those neurons that would show a sustained response to the signal) must fire continuously during the missing sections, such that noise provides continued neural activation during the gap. Second, offset responders must not detect the end of the signal at the beginning of the gap and third, onset responders must not detect that the signal begins again at the end of the gap. If these three responses are displayed, then this is evidence that the signal is perceived to continue through the gap, and that these processes play a role in the perceived continuity of fragmented signals.

The response of the primary auditory cortex in monkeys was used in order to test the neural basis of illusory continuity. It was found that neurons in the primary auditory cortex that responded to the signal continued to respond through the noise, behaving as if the signal were present. This is evidence that in order for restoration to occur, the neural activity present during the missing section should be indistinguishable from the neural activity that would have occurred if the signal were continuous. Offset and onset responders also behaved as if the signal were continuous; offset responders did not respond to the start of the gap, and onset responders did not respond to the end of the gap.

Whilst this experiment shows that perceptual restoration is well represented by the response of neurons in the primary auditory cortex, the exact mechanisms responsible for perceptual restoration are not revealed. Petkov et al. (2007) do suggest that the perceptual restoration mechanism is highly adaptive:

> the ecological pressure to maintain stable representations of interrupted sounds is important enough that it has been selected upon...through multiple neuronal encoding mechanisms. (Petkov et al., 2007, p.161)

This in turn suggests that it is unlikely that a single brain region or cellular mechanism might be responsible for perceptual restoration. Furthermore, in searching for the neural mechanisms underlying perceptual restoration, it is important to remember that the neural activity corresponding to what is actually perceived is not likely to be related to a simple acoustic analysis of the incoming signal. The interactive nature of perception as demonstrated by perceptual restoration suggests that the neural activity underlying this kind of active perception is most likely to be in areas that combine incoming sensory information with top-down facilitation from previous experience (Frith & Dolan, 1997).

Event Related Potential (ERP) methods have also been utilised in the attempt to uncover the neurophysiological basis of the perceptual restoration process. Micheyl, Carlyon, Shtyrov, Hauk, Dodson and Pullvermüller (2003) demonstrated that participants showed evidence of the perception of continuity of a tone with missing sections replaced by noise. This was reflected in the Mismatch Negativity (MMN) component of the ERP, thought to reflect "the response of an automatic, preattentive change-detection system." (p.748). The MMN therefore reflects whether stimuli are perceived as similar or different. An oddball paradigm was used, where a standard stimulus was presented, being either a tone containing a silent gap (condition 1) or an intact tone (condition 2). The deviant stimulus in both conditions was the tone with a gap replaced by noise. Micheyl et al. also manipulated the frequency of the replacement noise to either match or differ to the frequency of the tone.

In condition 1 where the standard stimulus was a tone containing a silent gap, the MMN response was larger where the frequency of the replacement noise in the deviant stimulus matched that of the tone. Under these conditions, the noise-replaced stimulus should sound continuous and hence very different to the silent gap standard stimulus, and this was reflected by the size of the MMN response. Conversely, when the frequency of the noise did not match that of the tone, the noise-replaced stimuli were not perceived as continuous and hence did not appear to be that different to the silent gap stimuli, as shown by a smaller MMN response.

In condition 2 where the standard stimulus was the continuous tone, the opposite profile of results was observed. When the frequency of the noise in the deviant stimulus matched that of the tone, the continuous-sounding percept was similar to the continuous tone, as reflected in a small MMN. Conversely, when noise of a different frequency was used to replace the gap, the signal did not sound continuous and so was distinguishable from the continuous tone and was registered by the MMN. In a psychophysical experiment Micheyl et al. ran alongside the ERP study, listeners' reports of the perceived continuity of the different stimuli corroborated the ERP data. The MMN is believed to be generated preattentively, suggesting that the illusion of continuity is an automatic process outside of conscious control.

Sivonen, Maess, Lattner, and Friederici (2006) measured both behavioural responses (reaction times) and ERP responses to recognition of words within spoken sentences; these words were either intact or had a phoneme replaced by a cough. In addition, the word was either highly expected or less expected given the sentence context. The ERP component of primary interest in this experiment was the N400 component, which reflects semantic, top-down processing and the difficulty of integration with the semantic context. Enhanced amplitude of the N400 reflects increased processing demands when integrating the word with the semantic context.

The N400 response to noise-replaced words was later, but not larger in amplitude, than the response to the intact words. This suggests that the words could still be integrated into the semantic context even though part of the word was missing. Strikingly, intact, less expected words elicited a more negative response than highly-expected noise-replaced words. The response of this component suggested that it is easier to integrate highly expected words into the semantic context, even if part of the word is missing, than a less expected word. Because this suggests that semantic processing was not detrimentally affected by the noise-replaced phoneme, the authors argue that the missing phoneme had been reconstructed and that phonemic restoration involves a “top-down repair process” (p.177), where word recognition under these conditions utilises predictions based on the semantic and phonetic information presented in the intact portions of the signal.

In a second experiment (Sivonen, Maess and Friederici, 2006), recognition of words with a missing phoneme replaced by a silent gap was compared to recognition of intact words in the same sentence context. Even the words containing a silent gap could be recognised and integrated within the semantic context (again the N400 response was later to silence-replaced than intact words but was not larger). The authors argue that word recognition, when sufficient contextual information is present, does not require all of the phonemes within the word to be intact.

The behavioural data across both experiments, which consisted of listeners’ reaction times for word recognition, indicated that words with a phoneme replaced by noise were recognised in a similar time frame to intact words; recognition of silence-replaced words was much slower. Crucially, in neither experiment was there a behavioural test of the perception of continuity of the different stimuli; that is, whether the listeners actually thought that the missing phoneme was present when replaced by noise. Thus, perceptual restoration of both the continuity and intelligibility of the signal might not have been occurring.

Furthermore, the two separate experiments do not provide a direct answer to the perceptual restoration question. It is not critical to know how the activation from intact words differs from manipulated words. This would be expected to be different because the amount of present phonetic information is different. The crucial difference is that between noise and silence conditions, since here the amount of intact phonetic information is equal; the only difference being the added auditory input provided by the noise. This comparison would reveal a great deal about the actual brain activation underlying the perceptual restoration process.

Many of the studies investigating the neurophysiological basis of perceptual restoration have investigated only the continuity illusion part of the effect, and as such have used simple tones as stimuli rather than more ecologically relevant signals such as speech. Shahin et al. (2009), however, attempted to investigate the brain regions involved in the perceptual restoration of fragmented speech, using an event-related fMRI design where participants were required to report stimuli where missing phonemes were either replaced by noise or had noise added to them as continuous or interrupted. Brain activity was compared across listeners' responses to the two stimuli, and, based upon the findings, Shahin et al. emphasised the role of two processes in perceptual restoration: a continuity process and a repair process.

They propose a model of perceptual restoration where a "continuity network" (left posterior angular gyrus, superior temporal sulcus) first assesses the representation of the signal provided by low-level processing mechanisms. This assessment analyses the coherence of the signal, to determine whether or not it is fully intact. If part of the signal is detected as missing (but replaced with noise), the "repair process" is recruited (Broca's area, bilateral interior insula, pre-supplementary motor area). This process restores the missing information, before reporting back to the continuity system. If the repair process has adequately restored the missing content, a "continuous" response is registered. Alternatively, if the content has not been restored, an "interrupted" response is registered. The brain regions responsible are not believed to be specific to speech, but domain-general, providing that the sound of the domain operates with internal templates of stored representations.

With research findings identifying the possible neural mechanisms responsible for perceptual restoration in the human brain, it is important to consider that the general perceptual process of making coherent percepts out of disrupted input might not be a process specific to human perceptual systems. The discussion now turns to reports of a similar mechanism operating in the communication systems of non-human animals.

Perceptual restoration in animals

The basic process of making fragmented auditory input coherent by filling in missing sections is not necessarily restricted in its operation to humans. There is evidence in the literature of the perceptual completion of missing fragments in auditory signals by non-human animals. For example, Sugita (1997) demonstrated temporal induction in the cat. Cats were trained to discriminate between an intact tonal glide and a tonal glide with a gap. When noise was added to the gap, discrimination between intact and interrupted glides was possible when the noise was of lower intensity than the tone. When the noise was of higher intensity than the missing part of the tonal glide, this discrimination became very difficult, with very similar patterns of responses seen between humans and cats. The response of cells in the cat primary auditory cortex that respond to tonal frequency glides was stronger when missing portions were replaced by noise compared to a much weaker response when replaced by silent gaps.

The operation of a sensory filling-in process has also been demonstrated in birds. Braaten and Leary (1999) found that the addition of noise into gaps in birdsong induced perceptual restoration. Starlings responded with significantly greater frequency to species-specific starling song than budgerigar song when gaps were replaced with noise rather than silence. Braaten and Leary suggested that perception of song in birds does rely on top-down perceptual processes, and that this finding highlights important analogies between human language and animal language in this case of birdsong. More recently, Seeba and Klump (2009) demonstrated that in the process of perceptual restoration of birdsong segments, starlings are influenced by familiarity and prior experience. Thus, this result demonstrates further parallels between the operation of perceptual restoration in humans and non-human animals.

Furthermore, Brumm (2006) found that birds possess sensory mechanisms to deal with the potential interference from noise to birdsong communication, by adapting the acoustic properties of their song to be heard in noisy environments. Similar adaptation has been demonstrated by Slabbekoorn and Peet (2003) and Slabbekoorn and den Boer-Visser (2006). This demonstrates a similar flexibility in perception to that in humans by suggesting that where masking noise is constant, not transient, birds demonstrate a behavioural plasticity by producing a song that can be heard. This also suggests auditory perception in noisy conditions to be a reciprocal process, where both the producer and perceiver have strategies to prevent communication from being disrupted in noisy conditions.

Miller et al. (2001) found evidence of perceptual restoration in cotton-top tamarins, who display a particular response to their species-specific call. It was found that these primates responded equally to intact calls and calls containing gaps filled with noise, but to a lesser degree to calls containing silent gaps. Miller et al. claim that these findings indicate that the neural mechanism responsible for perceptual completion may have evolved over 40 million years ago in a common ancestor. This is purportedly due to the overlap in evolutionary history between humans and primates. Furthermore, the findings suggest that both human and nonhuman primates may organise sensory information in a similar way. In a similar study conducted by Petkov, O'Connor, and Sutter (2003), Rhesus monkeys were trained to tell the difference between complete stimuli, stimuli containing silent gaps, and stimuli containing gaps filled with noise. The addition of noise into gaps made the detection of differences with intact stimuli difficult; stimuli containing silent gaps were easily distinguishable from intact stimuli. These results suggest that there are similar sensory mechanisms in both non-human primates and humans that result in perceptual completion of gaps in auditory signals when noise replaces the missing gaps.

Evidence of an auditory perceptual completion process in non-human animals contributes to the emerging picture of perceptual restoration by highlighting the operation of such a mechanism as an example of a very general perceptual filling-in process. Thus, perceptual restoration cannot be viewed as a strictly human mechanism operating only in complex forms of communication such as speech. Further evidence for the operation of very general sensory filling-in processes comes from investigations of filling-in in other domains in human perception. Such evidence for filling-in missing sensory input in the visual domain will now be presented.

Perceptual restoration in the visual domain

Visual perception, like auditory perception, often takes place in environments where input can be ambiguous and disrupted. Parts of objects can be occluded by other objects in a visual array such that visual input can be incomplete. There is evidence in the literature that the process of perceptual completion occurs in the visual domain as well as the auditory domain (e.g. Johnson, 2004; Pessoa, Thompson & Noë, 1998), supporting a constructivist view of perception where perception involves more than is present in the sensory input (Johnson, 2004). Craton (1996) demonstrated how, even in infancy, a partially occluded object is perceived as intact. At 6.5 months of age, the object is perceived as intact, whereas at 8 months of age the identity of the object is also perceived. Occulder width (presumably analogous to gap duration in auditory signals) has an influence on visual perceptual completion in infancy (Johnson, 2004); with increasing age comes the ability to fill in wider gaps in perception. Visual perceptual completion of occluded fragments in newborn infants was demonstrated by Valenza, Leo, Gava and Simion (2006), suggesting perceptual completion in the visual domain to be a processing bias that operates from early in life, but also questioning the role that visual experience with objects plays in the process. Nevertheless, the operation of visual perceptual completion in infancy is taken as evidence that perception is not passive, but active (Johnson, 2004).

Further, more direct evidence of visual perceptual restoration in adults was obtained by Schultz-Westre (1985). When a portion of a videotaped sentence in American Sign Language was removed and replaced with video white (analogous to white noise), deaf viewers used contextual cues to restore the missing information; intelligibility was unaffected by the disruption. This suggests that similar perceptual mechanisms are used for the perception of language, whether spoken or signed. However, without knowing whether the viewers perceived that the sentence was continuous, it is difficult to conclude that this is strong evidence of visual perceptual restoration.

Jordan, Thomas and Scott-Brown (1999) described the perception of continuity in visual word recognition as “graphemic restoration”, or the “illusory letters phenomenon”. Where missing parts of words were visually occluded, no letters were perceived to be missing and intelligibility was unaffected. This result suggests that visual word recognition can still operate with incomplete input where the absence of missing letters is hidden. Thus, research on a similar form of perceptual filling-in in the visual domain represents this process as a general-purpose mechanism where perceptual systems in domains other than the auditory system attempt to make coherent percepts through filling in missing information. Even within the auditory domain, there is evidence that such a process is not restricted to speech perception, but operates within other auditory domains such as music perception.

Perceptual restoration in the musical domain

Many environments where music perception takes place can be noisy. For example, parts of a musical performance can be obscured by someone in the audience coughing. Consequently, perceptual restoration operating in the musical domain would represent an important adaptation to such environments. The first documentation of musical perceptual restoration in the literature was a report by Sasaki (1980). One note of a familiar melody was removed and replaced by white noise, and adult listeners were asked to locate the missing note. The results indicated that perceptual restoration of the missing note was occurring, since localisation of the missing note was poor. An effect of familiarity was also obtained, where the note that was restored was congruent with the musical context in which it occurred, and the role of bottom-up influences was demonstrated, where the duration of the replacement sound exerted an effect on temporal localisation. A decrease in error with increasing duration of replacement noise was evident in the results, representing reduced perceptual restoration. Sasaki also investigated perceptual restoration of speech using the same method, and obtained similar results. This early investigation of musical perceptual restoration demonstrated that the phenomenon exists, and indicates important similarities between the perceptual processing of music and language.

DeWitt and Samuel (1990) reported the only systematic study of musical perceptual restoration currently present in the literature. They tested adult listeners for the perceptual restoration of a single note in a musical context, with a variety of manipulations. Of primary interest was the effect of familiarity with the musical stimulus on perceptual restoration. In the first experiment, the perceptual restoration of a single note in nursery rhymes and popular melodies was tested using the added/replaced discrimination paradigm reported in earlier studies of phonemic restoration (e.g. Samuel, 1981a). A series of speech stimuli were also presented for comparison. As expected, strong evidence of perceptual restoration was obtained for the speech stimuli. Furthermore, greater evidence of perceptual restoration was evident for a missing note from familiar melodies than for a note presented in isolation, thus suggesting that top-down influences from increased contextual information do operate on the perceptual restoration of music.

In a second experiment, this familiarity effect was tested more explicitly by comparing the perceptual restoration of a single note in a familiar melody, with a single note in an unfamiliar melody. The characteristics of familiar and unfamiliar melodies were matched for rhythm and implied harmonic structure[4]. Contrary to expectation, missing notes from familiar melodies were restored to a lesser degree than those from

---

[4] The sequence of notes in a melody corresponds to a series of accompanying chords that imply the harmonic progressions underlying the melody. Such harmonic structure is important in generating expectations about likely musical events; the rules of Western music determine which harmonic progressions are most predictable given the preceding context.

unfamiliar melodies. However, this result shows some important parallels with the phonemic restoration study reported by Samuel (1981a), where the high degree of context provided by a sentence led to increased discrimination between added and replaced versions of the stimulus, leading to the inference of reduced perceptual restoration. On a structural level, both melodies and sentences represent large holistic units. Therefore, the two results plausibly occurred for the same reason, that increased familiarity and contextual cues free more processing resources to make the discrimination between added and replaced items. These studies together therefore highlight important parallels between the operation of perceptual restoration in music and speech perception, at least as demonstrated by this discrimination measure.

In the next experiment, DeWitt and Samuel tested for the effect of predictability on musical perceptual restoration. Instead of comparing familiar and unfamiliar melodies, they compared the perceptual restoration of a missing note from melodies with different levels of predictability. They varied the tonality and rhythm to produce different levels of predictability[5], and presented the intact melody to listeners before the added and replaced versions in order to establish a representation of the melody to compare with the distorted version. Surprisingly, greater predictability resulted in a higher value of the discriminability measure, leading to the inference of reduced perceptual restoration. As with a similar finding relating to the effect of predictability on perceptual restoration in the speech domain, processing factors can account for this seemingly anomalous finding.

The perceptual restoration of smaller musical units, chords and scales, was tested in two further experiments. Here an effect of key-based expectation was found: greater perceptual restoration was demonstrated for notes in a scale that conformed to an established key context. This suggests that the perceptual restoration of missing notes in musical scales is influenced by expectation, and that contextual factors play a role in musical perceptual restoration. In a similar way, the perceptual restoration of missing notes from chords was based on key-based expectation. What was restored depended on what was expected, based on the given musical context. Taken together, the results of

---

[5] In a given musical context, some rhythmic and harmonic progression are more predictable than others. For example, constant rhythmic patterns are more predictable than changing rhythmic patterns, and implied harmonic patterns that correspond to the musical key are more predictable than those that constitute a random pattern of notes.

DeWitt and Samuel's experiments appear to suggest that low-level representations, such as chords and scales in music, and phonemes in speech, play a greater role in perception than higher-level representations such as melodies and sentences.

Whilst it is appealing to accept these results, especially as they corroborate Samuel's earlier experiments with phonemic restoration, there is a need to investigate the factors that influence the perceptual restoration of music through the use of different methods. Some of the issues surrounding the use of the added/replaced discrimination method were outlined earlier, particularly with reference to interpreting the effect that experimental manipulations have on the discriminability measure. Measuring discrimination between a stimulus in which part has actually been replaced by noise, and a stimulus where noise is merely superimposed upon the stimulus at the same point, says nothing about the ability of the noise to influence the intelligibility of the signal.

When the ecological value of a perceptual restoration mechanism is considered, it becomes clear that the purpose of such a mechanism is to compensate for signal disruptions by restoring perceived continuity and intelligibility of the signal. In experimental terms, this is best achieved by comparing the ability of replacement noise to enhance intelligibility of a fragmented stimulus when compared to intelligibility of the same stimulus with silent gaps. Thus, it could be argued that the discrimination paradigm does not test the real ecological nature of perceptual restoration, since no reference is made to the case where a signal contains silent gaps. Whilst increased familiarity can lead to better discrimination between added and replaced stimuli due to processing load factors, the effect of familiarity may result in an improvement in intelligibility of stimuli where noise replaces missing fragments, relative to silence. Therefore, whilst the discrimination paradigm is extremely useful for isolating the true perceptual nature of the perceptual restoration effect from postperceptual decision bias, it may not be able to detect the effect that different factors, such as familiarity and predictability, have on the intelligibility of a disrupted signal. These factors need to be investigated using other, more sensitive methods that measure comprehension directly, such as repetition accuracy or stimulus identification.

There is evidence that, in both speech and music perception, what is actually perceived in an auditory signal is not directly analogous to the content of the signal itself. Kaminska and Mayer (1993) report that this dissociation was previously taken as a unique function of speech perception, and what separated speech perception from the perception of other, non-speech sounds. Their finding of the perceptual restoration of musical signals is taken as evidence that such active perception also operates in the music domain, thus suggesting that perceptual restoration studies can reveal deeper parallels between music and speech processing than previously believed to exist. Kaminska and Mayer deleted single notes from six-bar music sequences, and replaced them with white noise. The measure of perceptual restoration was subjective; participants were asked to rate on a scale of 1 to 7 how certain they were that they could hear the music beneath the noise. Kaminska and Mayer tested for the effects of familiarity and predictability; and found that only predictability influenced restoration. However, this is not a reliable measure of perceptual restoration; at which point on a rating scale is it correct to place the cut-off point for evidence of perceptual restoration?

In summary, experimental evidence of the operation of musical perceptual restoration, while sparse, appears to suggest that a similar mechanism to that identified in speech perception operates to restore missing content of musical input given sufficient contextual support. Therefore, musical perceptual restoration might represent a fundamental adaptation to noisy acoustic environments.

The research presented on acoustic and contextual influences on perceptual restoration, and on the neurophysiology of perceptual restoration and perceptual restoration in the visual and auditory domains, all report research on adult participants. From an ecological perspective, the auditory perceptual systems of children would also benefit from compensatory mechanisms, such as perceptual restoration, to support perception in noisy listening conditions. This issue is fundamental to the purpose of this thesis, and will be discussed in the following section.

Perceptual restoration in children

Investigating perceptual restoration in children is an important part of furthering understanding of the phenomenon. To date, the perceptual restoration of speech, phonemic restoration, has been studied in children, but other forms of perceptual restoration require experimental study. If children show evidence of comparable perceptual restoration to adults, this might suggest there to be some similarities between perceptual processing in adults and children.

From an early age, children demonstrate the ability to separate sounds from different sources in noisy environments. Furthermore, over the first year of life, infants develop the ability to focus attention on their own name in the context of background speech (the "cocktail party effect"; Newman, 2005), and to separate foreground from background sounds ("auditory figure-ground perception"; Marsh, 1973). Whilst these abilities demonstrate processes that children possess to assist auditory perception in noisy listening conditions, more complex mechanisms, such as perceptual restoration, are required in order to ensure that comprehension of disrupted input is not detrimentally affected in a noisy environment.

Perceptual restoration is a particularly important issue to be considered in the context of children's auditory perception, following reports of the detrimental effects that everyday noise can have on their development (e.g. Marsh, 1973; Mills, 1975). Noise levels in children's learning environments often exceed recommended limits (noise levels below these limits minimise the detrimental effects of noise on children's development, e.g. Mills, 1975). Therefore, the ability to compensate for these sounds becomes a critical concern. The importance of perceptual restoration for children becomes even more apparent following findings that high noise levels in learning environments can detrimentally affect the acquisition of skills important for reading development (Maxwell & Evans, 2000). Furthermore, it is at the stage of speech and language acquisition that noise can have its most serious disruptive consequences on children's development (Manlove, Frank & Vernon-Feagans, 2001).

It has also been found that the redundancy in speech communication that assists the perception of disrupted input in adults might be reduced in children (e.g. Manlove et al., 2001; Mills, 1975), and that cues in redundant speech that can be used by children are easily masked by background noise such as the babble produced by multiple simultaneous talkers (Elliott et al., 1979). There is also evidence that the ability to utilise the redundancy in speech to interpret ambiguous words might be related to intelligence levels in children, (Marsh, 1973), with the consequence that the ability to organise auditory information in a noisy environment might be harder for children than adults. Crucially, these considerations are not restricted to classroom learning environments; children's home environments are often sufficiently noisy that potentially important signals, such as speech and music, face competition from many sources of noise

(Newman, 2005). In short, the potential importance of a perceptual restoration mechanism for children's learning and development must not be underestimated.

The first experimental evidence of the perceptual restoration of missing phonemes in children was reported by Ackroff (1981). Children were presented with a sentence where a phoneme had been removed from a word and replaced by noise. Children were asked to locate the position of the noise within the sentence. According to this measure of the effect, Ackroff found that 8-year-old children showed greater evidence of perceptual restoration than 6-year-old children (i.e. greater localisation error), thus illustrating developmental differences in perceptual restoration ability.

In a further study of perceptual restoration in children, Walley (1988) compared the recognition of individual words, in conditions where a phoneme had been replaced by noise or had noise added to it, in 5- to 6-year-old children and adults. Whilst children did show evidence of perceptual restoration by correctly recognising the words with phonemes replaced by noise, there was a larger difference between recognition of noise-added and noise-replaced words in children when compared to adults. This result suggests that the disruptions to perception caused by incomplete input are greater for children than adults. This is possibly due to weaker lexical knowledge or weaker contextually-based expectations in children.

Nevertheless, since no comparison is made to recognition of the words with a phoneme replaced by a silent gap, there is no measure of whether replacing a phoneme with noise improves recognition when compared to silence. Word recognition in this context could have been based purely on the intact sections of the word. Walley also found that both adults and children were affected by acoustic similarity effects, where fricatives (that are spectrally more similar to the replacement noise) were restored to a greater degree than nonfricatives. Whilst this suggests that perceptual restoration in adults and children is affected by the same acoustic factors, Ackroff (1981) found no evidence of a similar acoustic similarity effect in children's perceptual restoration. Therefore, the effect of acoustic factors on perceptual restoration in children is not clear.

Koroleva, Kashina, Sakhnovskaya, & Shurgaya (1991) investigated the perceptual restoration of a missing phoneme replaced by noise in a group of 5- to 6-year-old Russian children and a group of adults. Participants were asked to locate the missing phoneme and report whether or not the stimulus was intact. Koroleva et al. found that the children showed stronger perceptual restoration than a sample of adults according to these measures, and that, in contrast to adults, the position within the word and the type of phoneme that was excised had little effect on perceptual restoration in children. Since the words presented were chosen to be familiar to the children, the operation of perceptual restoration suggests that top-down factors were exerting an influence on the perceptual process in children, and possibly to a greater degree than the adults. However, it is not made clear whether the words presented would be more familiar to children than adults since English translations are not provided. Furthermore, this result could simply represent a response bias in children to report stimuli as intact; no measure of false alarms is provided. It is therefore not possible on the basis of these results to make any firm conclusions about the factors influencing the operation of perceptual restoration in young children.

A further study by Koroleva, Shurgaya, Kashina and Sakhnovskaya (1996) investigated perceptual restoration in young children by comparing boys with girls. This was motivated by evidence in the literature of gender differences in verbal abilities in young children. Perceptual restoration of a single phoneme was studied using the same method as Koroleva et al. (1991), and girls showed evidence of stronger perceptual restoration for consonants than vowels, a difference not evident in boys. Furthermore, girls showed stronger perceptual restoration for consonants in middle word positions rather than initial word positions. For the reason that the replacement noise used was more similar to consonants than vowels, Koroleva et al. (1996) concluded that girls are better at using acoustic information (bottom-up cues) for perceptual restoration than boys. In addition, the finding relating to the position of the critical phoneme within the word suggests that girls are also better able to use top-down cues from context in perceptual restoration. It appears that the ability to use the interaction between top-down and bottom-up information is more fully developed in girls than boys at the age of 5 to 6 years, and, in using this interaction, that processing in girls is more similar to that of adults than processing in boys. These differences are suggested to be due to gender differences in the rate of development of verbal systems, but refer exclusively to the perceptual restoration of speech signals.

One paper that is fundamental to the research reported in this thesis is a more recent study by Newman (2004), which investigated the perceptual restoration of sentences with multiple periodic gaps replaced with either noise or silence, in 5-year-old children and adults. Because alternating fragments were removed, the amount of intact information was equal to the amount of missing information. Children demonstrated improved repetition accuracy of spoken sentences when the gaps were replaced with noise than with silence, suggesting that they were able to use the context provided by the intact portions of the sentence to perceptually restore the missing fragments. Newman argues that, in using context in this way, children can use previous stored knowledge in the process of speech perception. By its very nature, perceptual restoration requires background knowledge to influence perceptual processing; knowing what should be present in any missing sections is a critical part of the effect.

Whilst the relative size of the noise-filled gaps over silence-filled gaps advantage was approximately equal between adults and children, overall children's repetition accuracy was more seriously affected by disruptions to the signal. It may be that young children do have access to lexical information that can influence perceptual restoration, but that the ability to use this information is less developed in young children when compared to adults. Nevertheless, the noise-filled gaps over silence-filled gaps intelligibility advantage in young children does suggest that from a young age children are able to combine top-down and bottom-up information in the process of perceptual restoration. Indeed, if the measure of perceptual restoration here is taken to be the size of the intelligibility advantage of stimuli with noise-filled gaps over stimuli with silence-filled gaps, then the correct conclusion would be that children can perceptually restore missing segments of speech to the same degree as adults. The overall trend for children's repetition scores to be lower than those of adults may be in part due to task factors, a possibility that is outlined below.

First, and most importantly, direct comparisons are made between data collected from adults and children, despite the fact that these data were collected using different methods. Adult participants, on hearing each sentence, were required to type what they heard into a computer. Conversely, children repeated the sentences they heard into a microphone. Newman herself admits that these different methods would likely result in different numbers of overall errors between adults and children, although argues that the crucial comparison between intact, noise-filled gaps and silence-filled gaps stimuli would be unaffected. Second, the testing session was extremely long (96 trials; for children, divided into 8 blocks of 12 trials), hence differences in sustained attention and concentration levels could partially account for differences between adults and children.

Studying perceptual restoration in children younger than 5 years of age requires different methods to research conducted with older children and adults. Repetition accuracy and discrimination between added and replaced versions of a stimulus are not feasible with toddlers, due to difficulty in describing the task in simple terms. This problem was addressed by Newman (2006), who used a preferential looking paradigm to investigate the perceptual restoration of the middle phoneme of a word from a pair (*cat* or *dog*; *lobster* or *dinosaur*) in 2-year-old toddlers. The measure of word identification employed was the amount of time spent looking at a corresponding image of the word presented, in conditions where this word was complete, contained a silent gap, or contained a gap filled with noise. Newman was testing to see if the presence of noise in the gap enhanced the intelligibility of the word (as measured by a greater length of time fixating a match between spoken word and visual referent) when compared to a silent gap. Despite manipulating the word presented, the design (within- or between-subjects), and the presentation format (words presented in isolation or in a sentential context), Newman failed to find evidence of perceptual restoration in toddlers. In fact, the visual preference for the image corresponding to the word presented was lower when the gap was filled with noise when compared to a silent gap.

So why might toddlers not show evidence of perceptual restoration? Newman suggests a variety of plausible explanations. First, toddlers have less prior lexical knowledge, and that which they do have is not robust enough to contribute to perceptual restoration. Second, children at this age are very open to new words; hence, they might have treated a noise- or silence-replaced word as a new word. Third, the noise used might not have been similar enough to the excised phoneme to induce perceptual restoration. Finally, less efficient processing speed in young children may prevent lexical knowledge contributing to on-line auditory perception. Whilst these explanations are quite feasible, it is also important to consider the implications that the method employed may have for explanation of the null result.

When making inferences from children's looking behaviour, a degree of caution is required. Whilst it is well documented that children prefer to fixate a match between a word and its visual referent, this measure can be difficult to interpret (Houston-Price and

Nakai, 2004). Infants come to such an experimental situation with pre-existing preferences, and the match with cross-modal information is not the only factor that can affect preferential visual attention. The relative salience of the stimuli for individual infants can also affect children's looking behaviour (Burnham & Dodd, 1999). A further problem with this study is that the words presented could be as easily identified in the silence condition as the intact condition in some of the experiments, as measured by infants' looking time towards the image that matched the spoken word. This means that the addition of noise was unlikely to improve intelligibility relative to its absence.

It appears that children do show evidence of the ability to compensate for signal disruptions through perceptual restoration of missing fragments, but this has only been reliably demonstrated from the age of 5 years. The difficulties in testing for perceptual restoration in young children has resulted in the application of a variety of experimental paradigms being employed to address the same problem, hence it is difficult to make solid comparisons between different studies conducted in this area.

Whilst there is evidence that perceptual restoration operates in the speech domain for children, there is no such evidence that it also operates in children's music perception. The possible operation of perceptual restoration in children's music perception represents the key question to be addressed in this thesis. This review will now explore relationships between speech and music perception, and the development of children's music perception, as a background to making hypotheses about the operation of a perceptual restoration mechanism in the music domain in child listeners.

The relationship between speech and music perception

There are many similarities between speech and music and the way in which they are processed which suggest that perceptual restoration might operate in a similar way in the two domains in children. From an evolutionary perspective, there are many accounts that music and speech do not have separable evolutionary histories, and may even be common parts of a shared precursor. For example, the first basic function of both language and music was to express emotion through the use of rhythm and prosody (Besson & Schön, 2003). Two similar theories about the evolutionary relationship between music and language are the "musilanguage" model of Brown (2001) and the "Hmmmmm" model of Mithen (2005). Brown argues that both music and language have a shared precursor in the form of a communication system where sound was used to express reference and emotion. The presence of this common ancestor is cited as the reason for the common structural properties of music and language. Similarly, Mithen (2005) presents evidence to suggest that both language and music evolved out of "Hmmmmm", a communication system that was holistic, manipulative, multi-modal, musical and mimetic.

A similar stance is taken by Masataka (2007), who argues that the evolutionary role of music may have been as an intermediate stage in the evolution of language. This is supported by observations that even very early in development, infants perceive speech sounds as music, and attend to the melodic and rhythmic aspects of speech. Masataka claims that because infants are predisposed to attend to the musical aspects of speech, they are well equipped to acquire spoken language.

Such accounts of a similar evolutionary role for music and language are strengthened by propositions that music and language share many features, such as parallel neural and cognitive mechanisms (Brown, Merker & Wallin, 2001), and that both are universal human traits (Mithen, 2005; Peretz & Hyde, 2003). Structurally, both music and language possess a phonetic, syntactic and semantic level (Aiello, 1994), and according to Zatorre, Belin and Penhune (2002), share the same structural organisation on the basis of these levels (in speech: phoneme – word – sentence; in music: tone – melody – song). Both music and language follow specific and fixed developmental time courses, do not need specific training for accurate perception (Zatorre et al., 2002), and children learn both music and language by acquiring the rules that govern the combinations of individual elements through exposure (Aiello, 1994).

There is much evidence that there are overlapping neural resources for music and language processing in both adults and children (Koelsch, Gunter, von Cramon, Zysset, Lohmann & Friederici, 2002; Maess, Koelsch, Gunter & Friederici, 2001; Patel, Gibson, Ratner, Besson & Holcomb, 1998; Zatorre et al., 2002). In addition, there is evidence of common perceptual mechanisms in speech and music perception; categorical perception is a prime example (Aiello, 1994). It is therefore possible that perceptual restoration is also a common perceptual mechanism working to preserve communication in different auditory domains in children.

A key aspect of the perceptual restoration process is the role that the listeners' expectations play in the process of filling in missing input. Music perception is viewed as an active process, where expectations are generated and updated as a piece of music proceeds (Bigand, 1997). The process of expectation is defined as, "...the anticipation of upcoming information based on past and current information" (Schmuckler, 1997, p.292). Expectations operate in music perception in the same way as lexical activation generates expectations in speech perception (Besson & Schön, 2003). A distinction with critical implications for perceptual restoration is that between schematic expectations (processing biases related to passively assimilated musical regularities) and veridical expectations (related to our representation of a specific piece of music). Schematic expectations based on likely musical events are generated by schemas, whereas veridical expectations based on specific and familiar musical events are generated by memory representations (Bharucha, 1994; Marmel, Tillmann & Dowling, 2008).

DeWitt and Samuel (1990) demonstrated the effect of schematic expectations on the perceptual restoration of music; yet the role of veridical expectations in the perceptual restoration of music is not well understood. In speech perception, expectations about specific speech events are based on activation of entries in the speech lexicon. A similar "musical lexicon" has been proposed to operate in the musical domain, being a representational system containing representations of all the music one has been exposed to (Peretz & Coltheart, 2003). Thus, the operation of veridical expectations in musical perceptual restoration, based on activation of the musical lexicon, is quite possible, but as yet has not been studied experimentally.

There are many similarities between speech and music perception that might influence the comparable operation of perceptual restoration in the two domains, and there is evidence to suggest that in adults perceptual restoration operates in both speech and music perception. However, it would be premature to assume that the same might be true for children. There are also many differences between the ways in which speech and music are processed that preclude a direct assumption that perceptual restoration operates, and operates in the same way, in the two domains in children.

First, speech and other auditory domains differ in the complexity of the auditory signal, with speech generally seen to be the most complex of auditory signals (Saberi & Perrott, 1999). Stimulus complexity could influence the operation of perceptual restoration in children in ways not apparent in adults, due to differences in perceptual processing. Second, the size of a critical unit in speech and music differs. A phoneme in speech cannot be directly compared to a note in music, especially in terms of the familiarity of an individual unit. Whilst a phoneme is familiar to the extent that it can be named on hearing, a musical note can only be aurally identified by those with absolute (perfect) pitch. Familiarity effects on perceptual restoration in children are likely to be quite different to the same effects on adults' perceptual restoration.

Third, whilst there is a developmental increase in the desire and ability to make perceptual input meaningful, and thus use expectations to interpret ambiguous input (Solley & Murphy, 1960), it is highly plausible that the role of expectation in music perception would also differ to that in speech perception, particularly in children. Moreover, perceptual restoration requires previous knowledge and expectations to contribute to the perceptual process. It is possible that children's expectations in music perception will not operate to the extent where they can contribute to the restoration of missing input. Moreover, differences in abilities such as working memory, segmentation, speech analysis and spectral analysis could all plausibly lead to age differences in the perceptual restoration of speech, but the factors that could contribute to musical perceptual restoration are not well known. Finally, children's musical development could also interact with these factors to influence the operation of musical perceptual restoration. An outline of important developments in children's music perception will now be given.

The development of music perception in children

Music perception, on some level, operates from early in life (McDermott, 2009; Trehub, 2003). This early ability to process music and make kinaesthetic responses to music affords experiential opportunities for learning about culturally specific musical regularities and regularities of musical structure.

Whilst many children receive considerable musical exposure during the preschool period, it is the onset of formal schooling that has the most significant influence on musical development (Hargreaves, 1986). At school, children are exposed to songs on a daily basis that typify the music of their culture. This allows children to develop knowledge of the rules that typically govern musical sequences, as well as develop stored representations of particular pieces of music that form the beginnings of their musical lexicon. Such exposure is crucial; from a young age and throughout development music is interpreted in the context of culturally-specific tonal scale frameworks (Dowling, 1999). Continued early exposure means that by the age of 6 or 7 years, children have all the skills necessary for accurate music perception (Hargreaves, 1986).

From a very young age, the rhythmic aspects of music elicit kinaesthetic reactions in children in the form of bodily movements (Bentley, 1966). This early response to rhythm is consolidated, such that these early rhythmic abilities form the basis for later developments including reading, motor activity, and the organisation and recall of stored information (Gardner, 1971), and the development of turn-taking in speech (Mithen, 2005). By the age of 5 years, children have quite a sophisticated way of organising rhythmic information (Hargreaves, 1986) which plays a key role in the perception of musical melody. Even by this young age, song organisation involves an interaction between rhythm and pitch, and from the age of 18 months, children can remember familiar melodies, indicating that they have stored representations of familiar rhythms and melodies by this age (Dowling, 1999). In both music perception and performance, melodies are primarily organised in terms of pitch contour (the global pattern of rises and falls in the pitch of notes), with research evidence suggesting that 4- to 6-year-old children discriminate melodies on the basis of their pitch contours (Morrongiello, Trehub, Thorpe & Capodilupo, 1985).

It is proposed by Hargreaves (1986) that children's music development shows evident parallels with their cognitive development. When considering children's music development, Serafine (1989) emphasises the importance of considering the difference between perception as a passive process, and cognition as an active process. Serafine suggests that when investigating children's ability to process music, it is important to look beyond merely what children are able to do at different ages, but to investigate the relative contribution of perception and cognition. Studying the perceptual restoration of music in children can provide insights into active processes of cognition in young children's music perception.

Perceptual restoration has crucial implications for children's learning and development. It is critical that children can make sense of perceptual experience so that information can be assimilated into their growing store of knowledge. In order to gain maximum benefit from explicit classroom instruction and informal learning experiences provided by caregivers, it is imperative that the extraneous sounds that could mask information (such as scraping chairs, coughs and the general everyday tumult of sounds) do not disrupt perception. Whilst the utility of a perceptual restoration mechanism for speech communication is quite clear, what might be the benefits of perceptual restoration for music perception?

There have been many claims that music is crucial for children's development (e.g. Racette & Peretz, 2007), and there has been a plethora of research into the possible transfer effects of music experience to other areas of development. For example, Lamb and Gregory (1993) found that music processing is related to reading ability, and that music training can help language development in children's early school years. They also explain how musical pitch discrimination abilities are important for phoneme discrimination. Similarly, early music perception has been shown to play a role in reading development (Anvari, Trainor, Woodside & Levy, 2002). The benefits of music training for IQ and cognitive abilities (Schellenberg, 2004, 2005), pitch processing in both music and language (Schön, Magne & Besson, 2004), spatial reasoning abilities (Bilhartz, Bruhn & Olson, 2000), and decoding the emotional content of speech prosody (Thompson, Schellenberg & Husain, 2004) have all been demonstrated in children. Finally, melody can aid memory for text by functioning as a mnemonic device (Rainey & Larsen, 2002).

There is a well-established link between melody in music and prosody in speech, with evidence that skill in perception of one is related to perception of the other (Patel & Iversen, 2007). The benefits of the melodic features of Infant Directed Speech (IDS) in making speech appealing and meaningful to infants, and for aiding speech discrimination early in life, are well documented (e.g. Bergeson & Trehub, 2007; Bosch & Sebastian-Galles, 1997; Fernald, 1989, 1991). Moreover, there is also evidence that there is a relationship between music perception and perception of speech prosody throughout development, with musical experience strengthening the relationship (Schellenberg, 2004).

This research evidence suggests that music perception plays a key role in development. In addition, current educational practice encourages the use of music as a learning device in many areas of the UK curriculum for young children, including language and communication, linking sounds and letters, and learning rhyming sounds (Department for Children, Schools and Families, 2008). Therefore, if music perception in children was active in the same way as speech perception (e.g. Newman, 2004), where through perceptual restoration missing musical input could be reconstructed, children's learning and development would not be detrimentally affected by everyday noisy listening conditions. The operation of perceptual restoration in children's music perception would also have important theoretical implications in terms of understanding the relationship between perceptual and cognitive processing in the music domain.

Summary

The review of the literature that has been presented here has identified perceptual restoration to be a mechanism that makes disrupted sensory input coherent not only in auditory and visual domains in adults and children, but also in the communication systems of animals. However, the mechanisms of perceptual restoration are still poorly understood (McDermott & Oxenham, 2008; Seeba & Klump, 2009). The chapter that follows considers the implications of perceptual restoration for theories of perception, and outlines some of the perceptual processes that might be involved in the perceptual restoration mechanism. The methods that have been used in previous studies of perceptual restoration are then examined in more detail in order to establish the most appropriate method to be adopted in the experiments reported in this thesis, to investigate the perceptual restoration of music in young children.

Chapter 2: Theoretical and methodological approaches to perceptual restoration

Theoretical Approaches to Perceptual Restoration

The operation of perceptual restoration has implications for theories of the nature of perception and of the structure of the perceptual system. For example, the operation of perceptual restoration questions whether perception is data-driven (where all the information required for perception is present within the sensory input) or schema-driven (where the knowledge and expectations of the perceiver interact with the content of the sensory input in the process of perception). In the context of auditory perception, bottom-up processes involve the perceptual system directly extracting information about features of the stimulus, whereas top-down processes include the influence of knowledge on perception (Sekuler & Blake, 1994). The two processes interact, such that knowledge provides the context in which stimulus features are interpreted.

Similarly, in demonstrating that top-down processes exert an influence on perception, the perceptual restoration effect has implications for the debate surrounding the modularity of the perceptual system. As a theoretical framework, the perceptual restoration of speech is consistent with interactive activation models such as the TRACE and COHORT models of speech perception (e.g. Elman & McClelland, 1984, 1988; McClelland & Elman, 1986), and perceptual restoration in all domains appears to closely follow rules of Auditory Scene Analysis (Bregman, 1990). In order to present a theoretical basis for the research to be reported here, the following section will discuss the theoretical implications of perceptual restoration.

Direct versus constructivist approaches to perception

The direct approach to perception claims that inner mental processes have no role in perception, with all the information required for accurate perception being present in the sensory input. There is no intermediate stage in the perceptual process where knowledge and expectations contribute to the final percept; instead, there is a direct link between sensory stimulation and what is perceived. This approach is also termed the "ecological approach" in the visual domain, after J. Gibson's (1979) claims that visual perception is well adapted to our sensory environments, and that the perceptual environment affords all the information that is necessary for accurate perception. Perception is a natural process that has evolved to deal with the world as it is in its natural form, where perception involves an interaction between the perceiver and the natural environment. According to this approach, the perceptual environment provides rich sensory input and perceptual processing is driven solely by this input, with no influence from higher-level cognitive processes.

In order for perceptual restoration to be explained as a purely data-driven, bottom-up process, the information provided within the auditory input itself must be sufficient to allow missing input to be reconstructed. This is quite possible for the restoration of simple sounds such as steady tones where restoration involves a simple continuation of the signal through a disruption. In contrast, the perceptual restoration of more complex, dynamic signals such as speech is difficult to explain without reference to listeners' knowledge; for example, lexical knowledge ensures that contextually-appropriate speech sounds are restored (e.g. Samuel, 1981a; Warren & Warren, 1970). In both speech and music, it is the learned rules particular to the construction of sound sequences in that domain that are used to infer the content of any missing input (Warren, 1999). This role of previous knowledge in perception cannot fully be accounted for by the acoustic information contained within the stimulus. Further, given everyday perceptual conditions, where we have to contend with the effects of occlusion, disrupted input, and irrelevant material, sensory systems must provide us with meaningful as well as contextually-appropriate information. It is therefore plausible that top-down information plays an important role in adaptive perceptual mechanisms such as perceptual restoration.

In contrast to the direct approach to perception, the defining feature of the constructivist approach to perception is that the process of perception involves more than a passive analysis of incoming sensory data. The perceiver themselves takes an important role in perception, bringing their previous knowledge and expectations to the process, and the perceiver is involved in "actively constructing a coherent percept" (Hawkins, in press, p.11). According to this view, perception is not merely a mechanistic process, but the perceiver is an important mediator between the stimulus and the percept.

The stimulus itself provides only the raw material for perception, with meaning being added in the mind of the perceiver (Rookes & Willson, 2000).

Support for constructivist approaches to perception comes from situations where the sensory input reaching the senses is incomplete, disrupted, or misleading. There are many situations where there is no direct link between the actual content of sensory input and what is perceived to be the content of the stimulus, such as in the case of perceptual illusions. The meaning applied to such sensory input is the product of mental processes. Perceptual restoration appears to violate any theory of a direct match between internal percepts and the state of the external world (Vallar, 2006), and as such, is often cited as evidence for constructive approaches to perception. This is for the reason that sensory input not present in the acoustic signal can be reconstructed through the influence of knowledge and expectations. The importance of active cognition in perceptual restoration has been continually emphasised: "What we hear does not simply correspond to the physical properties of the sounds reaching our ears." (Kashino, 2003, p.18).

Gregory's theory of perception (1970) was developed to account for visual perception yet many concepts are relevant to the study of auditory perception. Perception, according to Gregory, is a product of the interaction between sensory data, stored knowledge, and inferential cognitive processes. Perception involves the generation of hypotheses based on the sensory data, with knowledge and inference contributing to the formation of these hypotheses. Gregory claims that the operation of visual illusions demonstrates how top-down knowledge can override the bottom-up information presented to the senses. Not only does perception involve more than mere sensory input according to this view, but also behaviour can continue even when part of the sensory input is missing. This suggests a crucial role for experience, knowledge and inference in the process of perception. However, if stored information is an intrinsic part of the perceptual process, this theory needs to account for how this stored information is acquired, if perception is not based on sensory data alone (Schiffman, 2001). It is possible that perceptual primitives afford experiential opportunities to build up a store of knowledge that can contribute to the process of perception.

The Gestalt psychologists viewed perception as dynamic, and rejected direct, data-driven approaches to perception (Schiffman, 2001). They claimed that by considering perception in terms of sensations alone, perception of individual stimuli can be described, but not the perception of the relationship between stimuli. Principles of perceptual organisation must be given precedence in describing perception according to this approach, where we perceive globally coherent forms rather than individual elements of stimuli. For example, a familiar melody is not identified on the basis of individual frequency patterns that define individual notes. Instead, it is the global melodic contour that derives from the relationship between individual notes that enables melody identification (Schiffman, 2001).

There is an important caveat in the application of constructive, schema-driven accounts of perception to the explanation of perceptual restoration. Top-down facilitation from stored knowledge and expectations cannot lead to perceptual restoration when the bottom-up acoustic evidence does not support this interpretation, such as where missing information is replaced by a silent gap (Bregman & Dannenbring, 1977). As such, stimulus-driven acoustic evidence is a key part of the perceptual restoration process. Such evidence includes replacement noise being louder than the signal it replaces such that it acts as a potential masker of the signal, and that neural activity responding to the signal does not stop when the replacement noise begins. Moreover, the continuation of the signal must reappear immediately after the noise burst, so that the intact sections occurring before and after the gap are perceived to be parts of the same auditory object (Plack & White, 2000). Whilst some evidence for continuity effects have been demonstrated for cases of "partial anchoring" where the signal does not continue after the noise burst, the duration of perceived continuity is extremely short (Aronoff, 2006). Finally, there should be no sharp drop in amplitude evident where the signal stops and the noise burst begins (Bregman & Dannenbring, 1977). The content of restoration, what is filled in, depends on what is predicted to be there, whether that is speech, music, a steady tone or a rising/falling tonal glide. The evidence to support this prediction as the most appropriate interpretation of the signal content comes from the available acoustic evidence.

Top-down processes are likely to play a key role in perception when signals are disrupted, such as when they occur in noise (Magnuson, McMurray, Tanenhaus & Aslin, 2003). In this situation, knowledge and context constrain interpretation of what is an ambiguous signal. According to Shinn-Cunningham and Wang (2008):

> ...our robust ability to understand noisy signals...relies on integrating bottom-up sensory information with top-down knowledge of the likely source content. (Shinn-Cunningham & Wang, 2008, p.300).

Thus, both what we expect to be present and the acoustic evidence for this expectation appear to be critical in the process of perceptual restoration.

In speech perception, top-down processes are as critical for perception as bottom-up processing (Davis & Johnsrude, 2007; Hawkins, in press). Lexical constraints exert a large influence on speech perception, especially when the input is ambiguous. Such higher-level knowledge interacts with bottom-up perceptual processing during speech perception. The COHORT model of speech perception (Elman & McClelland, 1984) assumes interactive processing and the model has been found to behave in a similar way to humans in response to perceptual restoration stimuli, and as such has been used as a theoretical framework in the context of phonemic restoration.

According to Engle, Fries and Singer (2001), cognitive neuroscience has recently seen a paradigm shift from passive to constructive approaches to perception, which consider both hypothesis-driven (top-down) processing and stimulus-driven (bottom-up) processing. Whilst active, constructive processes provide us with hypotheses about the likely content of sensory signals, confirmation of these hypotheses is required from the stimulus itself. In perceptual restoration, the noise that replaces missing fragments of the signal provides critical confirmation of the expected content of the missing fragments. In Gregory's (1970) description of the role of hypotheses in perception, he stated that probable hypotheses are favoured over improbable hypotheses. Replacement noise makes our expectations of the content of missing fragments probable hypotheses, as there is no evidence that fragments are actually missing. If perception were purely a bottom-up operation, then perceptual restoration would be unlikely to occur, since this approach claims that only information present in the stimulus is extracted. Equally, bottom-up processing is critical for perceptual restoration to occur; it operates to check that expectations might be correct. Too much reliance on top-down processing is likely to lead to errors in perception.

The interaction between top-down and bottom-up perceptual processes in the context of perceptual restoration has largely been discussed with reference to adult listeners. Newman (2004) suggests that the operation of phonemic restoration in children demonstrates such an interaction, but whether a similar interaction occurs in perceptual restoration in the musical domain for children is an important consideration.

Perceptual restoration and modularity

Experiments on perceptual restoration make an important contribution to the debate surrounding the architecture of the perceptual system; specifically, whether the perceptual system is strictly modular or interactive (Samuel 1996). A modular perceptual system would be one that is not influenced by higher-level, more abstract representations and the knowledge that the perceiver has (Fodor, 1983). This approach gives a small role to expectations in perception, because knowledge used to generate expectations is unavailable to the perceptual module. Only information from lower-levels of processing are available to the module; for example, knowledge that both lines in the well-known Müller-Lyer illusion are equal in length does not allow the perceiver to overcome susceptibility to the illusion.

Conversely, an interactive perceptual system would involve the influence of more abstract representations on a lower level of processing. The evidence that has been presented here indicating that lexical factors can influence lower-level processing of phonemes suggests that an interactive system might be involved in perceptual restoration. Since there is evidence that expectations operate in a top-down flow to the lower levels of the perceptual system and bottom-up perceptual information provides confirmation of these expectations (Samuel & Ressler, 1986), it appears that the system is fully interactive, with a complex interaction between top-down and bottom-up influences on perception (Bowers & Davis, 2004; Frith & Dolan, 1997; McClelland, Mirman & Holt, 2006). The TRACE model of speech perception (McClelland & Elman, 1986) is consistent with this interactive view; lower levels of processing can be influenced by higher levels of processing, and there is a bi-directional flow of information combining contextual and phonetic cues in speech perception.

Where lower-level processes can be influenced by higher levels of processing, it is difficult to see how perception can be informationally encapsulated in a module. Phonemic restoration in particular has been cited as evidence against the modular view, since the effect demonstrates that phoneme perception is influenced by the lexical context (Bowers & Davis, 2004), and demonstrates an interaction between syntax, semantics and phonology which could not happen in a strictly modular system (Prinz, 2005). In music perception, the schema-based theory of music perception (Krumhansl & Castellano, 1983) is based on an interactive system, where perception is determined by an interaction between knowledge of musical structure and incoming musical input. Thus, music perception is also dependent on prior musical experience.

There are exceptions to this debate. It has been suggested that evidence of the influence of top-down processes on perception does not necessarily negate the possibility that the system is modular (Coltheart, 1999). A module itself can have internal structure, where information flow can be both top-down and bottom-up. According to Fodor:

> The claim that input systems are informationally encapsulated must be very carefully distinguished from the claim that there is top-down information flow within these systems...phoneme restoration provides considerable prima-facie evidence that phoneme identification has access to what the subject knows about the lexical inventory of his language. If this interpretation is correct, then phoneme restoration illustrates top-down information flow in speech perception. It does not, however, illustrate the cognitive impenetrability of the language input system. To show that the system is penetrable (hence informationally un-encapsulated), you would have to show that its processes have access to information that is not specified at any of the levels of representation that the language input system computes...If...the 'background information' deployed in phoneme restoration is simply the hearer's knowledge of the words in his language, then that counts as top-down flow within the language module. (Fodor, 1983, p.76)

This quote suggests that, contrary to the belief that perceptual restoration contradicts proposals of a modular perceptual system (Bowers & Davis, 2004), the way that top-down influences operate on perception in the restoration process is not necessarily evidence against modularity. It is therefore plausible that top-down flow of information within the musical module, as is likely to occur in musical perceptual restoration, is also not a violation of modularity.

Auditory Scene Analysis and perceptual restoration

In everyday listening situations, a complex mixture of sounds reaches our ears. The auditory perceptual system must take this complex input and find the best way to represent reality. Auditory Scene Analysis (Bregman, 1990) involves breaking down the mixture of sounds constituting the auditory input and grouping sounds according to constituent parts. The ways in which sounds are grouped together follow the Gestalt principles of proximity, closure and similarity. Within a mixture of sounds, sounds of similar timbres are grouped together (the principle of similarity) and sounds closer together in duration or pitch are grouped together (the principle of proximity). These Gestalt principles of auditory organisation are believed to influence auditory scene analysis from infancy onwards (Demany, 1982), with these processes being the means by which auditory input becomes meaningful cognitive experience. As well as accounting for auditory scene analysis, Gestalt theories of auditory organisation, such as the principle of closure, can explain the phenomenon of illusory continuity in perceptual restoration effects. Missing parts of a fragmented signal are restored so that the percept satisfies the principle of closure. The system deals with fragmented input by connecting the fragments in the most plausible way, using top-down cues as a guide.

In fact, Bregman claims that perceptual restoration occurs directly because of these processes of auditory scene analysis, describing perceptual restoration as, "another oblique glance of the auditory scene analysis process in action." (Bregman, 1990, p.28). Spectral components of the noise are taken and reallocated to the signal during auditory organisation because the most plausible organisation is that the missing section is in fact present, as the noise is hiding its absence. According to this viewpoint, perceptual restoration and the continuity illusion occur because the auditory scene analysis process is functioning normally, not as a result of the brain being somehow tricked, as the term "illusion" can imply.

Bregman, Calantonio and Ahad (1999) describe auditory scene analysis and illusory continuity of disrupted signals as the result of a common process; that of linking together parts of a sequence that have similar frequencies. Both represent problems encountered by the auditory system in everyday environments, and have similar theoretical implications. By linking together intact fragments of auditory signals, the restoration of continuity allows the intact fragments to be treated as part of the same auditory object (Plack & White, 2000). Similarly, according to Shinn-Cunningham and Wang (2008):

> our robust ability to understand noisy signals...relies on integrating bottom-up sensory information with top-down knowledge..., taking into account all kinds of evidence that a particular sound feature belongs to a particular object. (Shinn-Cunningham & Wang, 2008, p.300)

Bregman (1990) outlines rules that determine whether or not continuity of an auditory signal with missing fragments replaced by noise is perceived as continuous. These rules are illustrated with reference to two intact portions of a continuous sound, occurring before and after a gap in the signal (A1 and A2, respectively), which is replaced by noise (B):

1) "No discontinuity in A" rule: there should be no evidence that the missing part of the signal is not present in the auditory input.
2) "Sufficiency of evidence" rule: some of the neural activity corresponding to the replacement noise must be identical to that which would have occurred if the missing sound had continued through the gap. This explains why perceptual restoration is strongest if the missing sound and the replacement noise have the same or a similar frequency.
3) "A1-A2 grouping" rule: there must be evidence that the sections of intact signal either side of the missing portion came from the same sound source. This is because perceptual restoration can only fill in the gaps between already formed auditory streams (Tougas & Bregman, 1990).
4) "A is not B" rule: The replacement noise (B) must not be interpreted by auditory scene analysis as part of the signal (A).

These rules provide important insights into the perceptual processing involved in perceptual restoration, and the neural mechanisms that might be responsible. The section that follows outlines some explanations of how perceptual restoration might operate, and some of the possible perceptual processes involved. These processes are not mutually exclusive; a combination of different perceptual processes is likely to operate in the process of perceptual restoration.

Explanations of perceptual restoration

Four explanations of the possible perceptual processes involved in perceptual restoration are outlined in this section. The "neural reallocation" account represents an attempt to explain how sounds missing in reality are heard to be present through the process of perceptual restoration. The "neural perseveration hypothesis" and the "no evidence for discontinuity" account explain perceptual restoration with reference to the evidence available to the perceptual system. Finally, the explanation that considers perceptual restoration as a hypothesis-generating and evidence-gathering process takes into account both listeners' expectations of the likely content of missing information and the evidence available in the auditory signal to support these expectations.

The "neural reallocation" account

When missing sounds are heard to be present through the operation of perceptual restoration, this implies that sounds not present in reality are reconstructed by the perceptual system. In the process of perceptual restoration, it has been suggested that some of the energy of the replacement sound (e.g. white noise, cough) is used up in order to physically restore the missing sound. Thus, the ability of the replacement noise to restore missing fragments is due to a perceptual synthesis of the missing input by the replacement sound (Bashford & Warren, 1979; Warren, 1984; Warren, Obusek & Ackroff, 1972). It is further suggested that in the process of perceptual synthesis, a portion of the neural input of the replacement sound is reallocated to restore the sound, resulting in a residue of part of the replacement sound that is not used for this purpose (the "neural reallocation" account; Repp, 1992; Warren, 1984; Warren, Obusek & Ackroff, 1972).

The neural reallocation account of perceptual restoration was originally applied to the explanation of nonverbal temporal induction (e.g. of simple tones), but its application to the perceptual restoration of speech was tested experimentally by Repp (1992). This experiment aimed to determine where perceptually restored speech sounds come from when not actually present in reality, and Repp considered two alternative possibilities. First, it is possible that perceptual restoration is in reality an illusion, and that the restored phoneme is merely a perceptual experience of sounds expected to be present. The role of the replacement noise is simply to confirm that this might be the case; it does not directly contribute to the reconstruction of missing sounds. Repp called this possibility the "top-down completion hypothesis". Alternatively, the "segregation hypothesis" proposes that part of the high frequency energy of the replacement noise is perceptually subtracted and used to reconstruct the missing fragment. The result is that the replacement noise changes in perceived pitch (and plausibly also intensity) as a result of the subtraction of part of its composition, leaving a residue that sounds different to the noise if it were presented in isolation.

Repp asked participants to subjectively rate the pitch of the noise replacing part of a sentence, and compared this to ratings of the pitch of the noise when presented in isolation. The noise replacing the missing phoneme was rated as lower in pitch than the control noise probe presented in isolation, but only in two out of the five experiments conducted. Thus, no conclusive evidence to support either hypothesis was obtained in this experiment.

It is possible that listeners' judgements were biased by experimental factors. For example, are listeners really in a position to make a reliable judgement about the nature of the replacement noise when previous work (e.g. Warren 1970) has shown that listeners cannot even localise the noise accurately? Clearly, the very nature and strength of perceptual restoration make it difficult to perceive the replacement sound accurately. Further, there was no behavioural test of whether perceptual restoration was actually occurring when the noise replaced the missing fragment.

More reliable evidence for the neural reallocation account of perceptual restoration was obtained by Warren, Bashford and Healy (1992) and Warren, Bashford, Healy and Brubaker (1994). In these experiments, the noise replacing a missing fragment of a continuous tone or speech signal was perceived by listeners to be lower in volume, as a result of restoration of the missing fragment, than the noise heard in isolation. In some cases, the timbre of the sound was also perceived to change by listeners. This supports a perceptual subtraction account of perceptual restoration, in which neural reallocation of part of the replacement sound to perceptually reconstruct the missing signal content is responsible for the restoration of the missing fragment.

However, even if phonemic restoration appears to operate in a similar way to nonverbal temporal induction, as the neural reallocation account suggests, this does not mean that there is no schema-driven component to the effect. All forms of auditory induction might be schema-driven, to the extent that they utilise contextual cues. The schema in operation could vary in complexity, with a more complex schema operating on a speech signal than a simple sound sequence or tone. As many experiments have found powerful top-down influences on perceptual restoration, and have proposed that the relationship between bottom-up and top-down influences in the process of perceptual restoration is one of a complex interaction in nature, it is unlikely that one account of its operation can explain all aspects of the effect.

The "neural perseveration hypothesis"

When a missing fragment in an auditory signal is replaced by noise, if the noise provides continued neural activation during the gap, where the same neural units that would have been stimulated by the signal are stimulated by the noise, then the signal can be perceived to continue through the noise (Elfner & Caskey, 1965). This explanation is termed the "neural perseveration hypothesis" (Bregman & Dannenbring, 1977), and is similar to the "sufficiency of evidence" requirement for perceptual restoration reported by Bregman (1990) and outlined earlier in this chapter. Whilst this provides a good explanation for how simple tones can be perceived to continue through a noise burst, the difficulty that this explanation has in explaining more complex effects such as the perceptual restoration of a missing phoneme has been identified (Bregman & Dannenbring, 1977). For example, it provides no account of how contextually-appropriate sounds are restored.

Nevertheless, Warren et al. (1994) report that by only inferring continuation of a signal through an interruption where there is neural evidence of continued activation through the interruption, inappropriate perceptual restoration of input does not occur, if this is not the most logical interpretation of the sensory input. Thus, according to this explanation, the restoration of continuity of a fragmented signal is determined by the evidence available to the perceptual system. If the acoustic evidence suggests continuity to be the most appropriate interpretation of the input, the fragmented signal is perceived to be continuous. Yet if the acoustic evidence does not support such an interpretation, the signal is not perceived to continue through noise. A further account of perceptual restoration based on available acoustic evidence will now be described.

The “no evidence for discontinuity” account

What is, or is not, present within an auditory signal can be an important determinant of the operation of perceptual restoration. According to the “no evidence for discontinuity” account of perceptual restoration, a fragmented signal is perceived to be continuous providing there is no evidence available to the perceptual system indicating that anything is actually missing from the signal. In contrast, perceptual restoration does not occur where gaps in the signal can be detected (Kashino, 2003; Riecke, Mendelsohn, Schreiner & Formisano, 2009). It is important to note that this explanation is most relevant to explaining the perceived continuity of fragmented signals, rather than perceptual restoration of the content of missing information. Therefore, this explanation is consistent with a data-driven approach where the stimulus provides the necessary information for the perception of continuity. This explanation also accounts well for the perceived difference between stimuli where missing fragments are replaced by noise and by silence. A silent gap provides clear evidence for the discontinuity of the signal, whereas noise, if it satisfies acoustic requirements to hide the gap, provides no evidence for discontinuity.

The no evidence for discontinuity explanation of the restoration of continuity of a fragmented signal is often described as a simple explanation of the effect, where decomposition of the signal into the restored portion and the noise residue, as proposed by Repp (1992), is not necessary for continuity to be perceived. The simple determinant is whether or not the auditory system registers a gap: “...continuity of an object will be perceived unless the discontinuity of the object is explicitly represented in the auditory system” (Kobayashi, Osada & Kashino, 2007, p.398). Furthermore, continuing neural activity in the gap is not necessarily required; the important factor is whether signal offsets and onsets around the signal can be detected.

Auditory edges, such as offsets and onsets of a signal around a gap, define auditory objects. If edges are detected around the missing gap when replaced by noise, then the signal and noise are treated as separate auditory objects, not the continuation of one through the other (Nelken, 2004). The role of the masker (noise) is to prevent energy changes in the frequency channel of the signal at offsets and onsets (Riecke et al., 2007); thus, offsets and onsets are not detectable (Başkent et al., 2009).

Bregman and Dannenbring (1977) found that edge detection was more important in the perception of continuity of a steady tone containing a gap filled with noise than the energy of the noise relative to the tone. Because noise masks the sudden drop in intensity of the tone when it stops, thus providing no evidence for discontinuity, "the auditory system projects the sound into the noise" (Bregman & Dannenbring, 1977, p.155).

With reference to visual perceptual completion, Pessoa et al. (1998) raised the important theoretical question of whether such perceptions of continuity when missing sections are occluded or masked in visual or auditory perception result from the perceptual system ignoring an absence, or detecting a presence. The no evidence for discontinuity explanation implies that the perception of continuity of an auditory signal results from the auditory system detecting a presence in missing fragments, such that the absence of the signal is not detected. However, Kobayashi et al. (2007) found that the perception of continuity of fragmented signals where noise did not replace the missing fragments could occur if attention was directed away from the gap, by directing it towards a flashing visual stimulus. This seems to suggest that being able to ignore an absence also results in the continuity illusion.

Whilst no evidence for discontinuity is a plausible explanation, it is clear that experimentally it has been used to account for the perception of continuity of a simple signal rather than the restoration of the content of missing fragments of a more complex signal, such as speech or music. Both the neural perseveration and no evidence for discontinuity accounts are better able to account for the restoration of continuity than the restoration of intelligibility. In order to explain the restoration of intelligibility of fragmented signals, a further explanation taking the context surrounding the missing fragment into account is required.

Perceptual restoration as a hypothesis-generating and evidence-gathering process

Many theories of auditory perception in ordinary listening conditions take a prediction-driven approach to the analysis of sound (Pickering & Garrod, 2007), and continuous inferences based on expectations are seen as the basis of speech perception (Ellis, 1999; Hawkins, in press), and a key part of music perception (McAdams &

Bregman, 1979). These expectations are constantly updated on-line with new incoming information; this forms the basis of predictions about what is likely to come next in the auditory sequence (Engle et al., 2001). This process of inference is also described as a key part of the perceptual restoration process (e.g. Bregman, 1990; Koroleva et al., 1991; Shinn-Cunningham, 2008), where what is filled in is based on what is expected to be present (Kashino, 2003). Indeed, such a prediction-based approach to perception is described as even more important in cases where auditory input can be ambiguous or disrupted (Pickering & Garrod, 2007).

Thus, perceptual restoration can be represented as a process where expectations and contextual information generate hypotheses about the most likely content of any missing or ambiguous auditory input (the "hypothesis-generating" process). Crucially, such hypotheses must be evaluated for plausibility against the available acoustic evidence that can be gathered from the auditory input (the "evidence-gathering" process. According to Kashino (2003):

> The brain continuously produces hypotheses about acoustic events and looks for evidence for the hypotheses in the sounds coming into our ears. The world we perceive is nothing but the most plausible interpretation of the input signals. (Kashino, 2003, p.23)

Thus, what is (or is not) present within the acoustic signal provides corroborating (or contradictory) evidence for the hypotheses generated on the basis of expectation. Noise, by hiding the absence of missing information, provides evidence that contextually-generated hypotheses of the probable missing content are likely to be correct, whereas silence instantly disconfirms such hypotheses by making it quite clear that the expected content is in fact missing. Similarly, top-down expectations control a matching process, where bottom-up features consistent with expectations are selected, and those inconsistent are inhibited (Grossberg, 2003). This matching process occurs during perceptual restoration; what is perceived is determined by what is expected in the context, and a disrupted stimulus is mapped to an internal percept (Aronoff, 2008). Where there is no evidence that the stimulus does not match the percept, this process functions normally. However, if the stimulus clearly does not match the percept (such as a stimulus with silent gaps) this process cannot occur.

This explanation of perceptual restoration is able to account for the role of higher-level process involved in the perceptual restoration of more complex auditory input such as speech. Indeed, this interaction between processes of expectation and corroboration is consistent with many theories of speech perception (e.g. Rumelhart, 1977). The ability of the expectation/corroboration interaction to account for the perceptual restoration of speech in adults suggests that it may also be able to account for the perceptual restoration of music in children, and represents an issue that will be addressed in this thesis.

Summary

Simple, data-driven accounts of perceptual restoration are adequate to account for simple perceptions of continuity of fragmented signals, but higher-level representations are required to account for the highly context-dependent restoration of missing input from more complex signals such as speech and music. Where context is important in perceptual restoration, this suggests that the effect is unlikely to be totally data-driven. Therefore, it is most likely that perceptual restoration does not depend entirely on top-down or bottom-up processes, but an interaction between the two where, through the processes of expectation and corroboration respectively, sense can be made of disrupted auditory input.

Approaches to measuring perceptual restoration

The experiments to be reported in this thesis require the application of the perceptual restoration effect to a previously unresearched area: children's music perception. Therefore, consideration of the methods used in previous perceptual restoration experiments is fundamental to informing the design of the present research. Since the first experimental reports of auditory continuity effects and perceptual restoration, researchers have taken many different approaches to measuring perceptual restoration. These different approaches have been developed in order to be most appropriate for the aspect of the effect under investigation, and the nature of the participants tested. This section outlines the major methods used in previous perceptual restoration studies (see Table 2.1), and then describes how the methods to be employed in the present series of experiments reflect this aspect of previous research, in also choosing the most appropriate way of investigating the perceptual restoration of music in young children. Justification for these methods, based on an analysis of previous methods, will outline exactly why these methods are appropriate.

Early studies of perceptual restoration used a localisation test of restoration (see Table 2.1 for a description). This method would be difficult to use with very young children; asking children to locate a noise and report whether it replaced a musical note would almost certainly appear a very arbitrary task. When children perceive task demands to be arbitrary, they are often unsure how to respond to the task, sometimes believing the task to be more complicated than it actually is. Furthermore, extensive research cautions against the use of task instructions involving unfamiliar words (e.g. Waterman, Blades & Spencer, 2000), and task instructions that require some further information or knowledge to interpret correctly (Donaldson, 1978). Where children's interpretation of adults' questions depends upon their knowledge of language, it is essential that the task demands can be explained using as simple language as possible.

Many previous studies of perceptual restoration with adult participants involved laboratory tasks requiring complex perceptual decisions to be made. In many cases, such as the localisation, discrimination and subjective rating methods described in Table 2.1, it would be very difficult to explain the task to a young child using language with which they are familiar. Furthermore, many of the tasks could appear quite arbitrary. Adults, particularly the sample often used for such laboratory studies, are used to being asked to carry out what might appear to be strange tasks. In contrast, children are not and can easily misinterpret a task they do not understand which can lead to misleading results. Whilst Ackroff (1981) used a noise localisation measure with children, these children were older than those to be tested in the current research. Furthermore, this measure fails to address the ability of perceptual restoration to restore intelligibility of a fragmented signal directly. This ecological aspect of perceptual restoration represents a key concern of this thesis. The repetition accuracy measure used by Newman (2004) with child participants was a more appropriate task to present to children, where the task instructions could be clearly understood without requiring further knowledge. Furthermore, the task was something children are frequently used to doing in an educational setting (listening to and repeating speech). Moreover, this method does directly consider the perceptual restoration of intelligibility of fragmented signals.

Table 2.1. Description and analysis of the methods used in previous perceptual restoration experiments

Type of Method	Description	Advantages	Disadvantages	Previous studies
Localisation	Listeners are presented with a signal where a single phoneme or musical note is excised and replaced with noise or silence. Listeners are asked to locate the position of the noise or silent gap and report whether the noise or silent gap replaces the phoneme or musical note. The degree of mislocalisation is used as a measure of perceptual restoration.	- Simple task for listeners - Can be conducted with unfamiliar stimuli	- Tests the perception of continuity without really addressing intelligibility - Does not take false alarms into account - Difficult to interpret: does mislocalisation by 6 phonemes really indicate increased restoration than mislocalisation by 3 phonemes? (Samuel, 1981a)	Ackroff (1981); Koroleva et al. (1991); Layton (1975); Sasaki (1980); Warren (1970); Warren & Obusek (1971); Warren et al. (1972)
Added/Replaced discrimination	Listeners are presented with a version of the stimulus where a single phoneme or musical note has been removed and replaced with noise, and a version of the same stimulus where the phoneme or musical note has not been removed, but noise is superimposed on top of it. Listeners are asked to report whether the stimulus is intact or not. If perceptual restoration is occurring when the phoneme or note is actually missing, the two stimuli should sound identical. Signal detection analysis is used to determine an index of discriminability as well as a measure of postperceptual bias.	- Takes false alarms into account - Can be conducted with unfamiliar stimuli	- The task is far removed from the way we use perceptual restoration in our everyday lives. It fails to address the intelligibility of the stimuli - It assumes that restoration is complete, that the process restores stimuli to the intelligibility of intact stimuli - In failing to make the comparison between stimuli with noise-filled gaps and stimuli with silence-filled gaps, the ecological nature of restoration is not addressed - Signal detection theory, when applied to such high-level domains, can be complicated by even small procedural factors (Samuel, 1996)	DeWitt& Samuel (1990); Samuel (1981a); Samuel (1981b); Samuel (1991); Samuel (1996); Samuel & Ressler (1986); Trout & Poser (1990)
Subjective rating	Listeners are asked to make a judgement about some aspect of the stimulus, such as determining the threshold at which the number of interruptions per second are detectable, or rate how sure one is that a signal continues beneath noise.	- Allows more detailed aspects of the restoration process to be tested, such as the perceptual subtraction account - Stimuli do not need to be familiar	- Relies purely on subjective reports - Does not reflect the use of perceptual restoration in everyday life	Bashford & Warren (1987); Bashford et al. (1988); Kaminska & Mayer (1993); Repp (1992); Warren et al. (1994)
Repetition accuracy	A stimulus is presented to the listener where fragments have been removed and replaced with noise or silence. In some cases a control condition is also presented where the intact version of the stimulus is presented for the listener to repeat.	- Directly addresses the ecological purpose of perceptual restoration - Objective measure - Avoids the problems associated with hits/ misses and false alarms - Takes relative differences between conditions into account	- A universal coding system must be applied to score participants' responses	Bashford et al. (1992); Bashford et al. (1996); Koroleva et al. (1996); Newman (2004); Powers & Wilcox (1977); Verschuure & Brocaar (1983)

In the review of relevant literature presented in the preceding chapter, many problems surrounding interpretation of the added/replaced discrimination paradigm (e.g. Samuel, 1981a; see Table 2.1) have been identified. Furthermore, the difference between a missing sound replaced by noise and a sound with noise added to it is very subtle, and the method requires a large number of trials to obtain a reliable difference between the two stimulus types (Newman, 2006). This method also fails to address the issue of intelligibility directly. For these reasons, this method is not the most appropriate for use with young children in the present context.

When just a single phoneme or musical note is removed from an auditory signal, listeners are asked to locate the missing information, or make some sort of judgement about the sounds that they hear. However, making a single excision from an auditory signal can be considered as inappropriate for studying the effect of perceptual restoration on the intelligibility of disrupted signals. With a single excision, the amount of intact contextual information is not equal to the amount of information that is removed. Thus, even intelligibility with a single silent gap would be unlikely to be detrimentally affected. Furthermore, Warren (1983) explains how deleting a single complete sound is not ecologically valid since in real life noisy environments it is more likely that several sections of a signal would be disrupted. Therefore, in order to study the effect of signal disruptions on the intelligibility of auditory signals, the amount of intact information must be equated with the amount of information removed. This is best achieved using multiple alternating interruptions, where for every missing fragment of a given duration, an intact section of the same duration follows.

The present series of experiments is primarily concerned with addressing the ecological purpose of perceptual restoration; that is, how noise is able to improve intelligibility of fragmented signals relative to its absence. The analysis presented in Table 2.1 suggests that repetition accuracy of a signal with multiple alternating interruptions addresses this question well. This is because versions of a signal where missing fragments are replaced by noise and where they are replaced by silence are identical in all respects apart from the bottom-up acoustic stimulation throughout the missing fragment facilitated by the replacement noise. However, there are potential difficulties in using a repetition accuracy measure with young children and musical signals. Whilst by the age of 5 years children can accurately repeat speech, children vary in the degree to which their music perception and production skills follow the same trajectory (Atterbury, 1985). Furthermore, correct repetition of speech is based on content only, not rhythm and prosody. In order for a child to repeat a musical phrase accurately, they must correctly reproduce both the pitch and the rhythm of the notes. Children also differ in their comfortable pitch range for singing (Welch, Himonides, Saunders & Papageorgi, 2008). As such, repetition accuracy of musical signals is not likely to be an accurate reflection of children's perception of the music, and thus, of the operation of perceptual restoration in this context.

Methods to be employed in the current research

It is clear that previous methods used to test for the operation of perceptual restoration are inappropriate for investigating the perceptual restoration of music in young children. The difficulty with applying previously used methods to the context of children's music perception will be overcome in the present series of experiments by employing identification accuracy as the dependent measure of perceptual restoration. Identification of a familiar melody involves a process of mapping what is present in the sensory input to a stored representation of the melody (Hébert & Peretz, 1997). This mapping procedure is also involved in perceptual restoration (Aronoff, 2008). Therefore, identification of disrupted input can still occur provided the evidence that input is missing is limited, and that the sensory input can still be mapped to the stored representation. Identification judgements are described as the "best indicator of the achievement of the mapping procedure between the stimulus and a stored memory representation" (Hébert & Peretz, 1997, p.520).

The melodies to be identified in the experiments presented in this thesis are taken from children's nursery rhymes and play songs, in an attempt to present ecologically relevant stimuli. These stimuli are important in children's development; knowledge of nursery rhymes is related to the development of phonological skills, where the rhyming pattern of the lyrics is important in the development of rhyming skills and sound categorisation. These skills are crucial to later reading development (Maclean, Bryant & Bradley, 1987). They also teach concepts such as letter names and counting (Trehub & Trainor, 1998). Nursery rhymes and play songs are universally familiar and as such are often used in tests of adults' music perception (e.g. DeWitt & Samuel, 1990).

Furthermore, familiar melodies such as those from nursery rhymes and play songs are remembered from the age of 18 months and are thus seen as stable percepts in a child's environment (Dowling, 1999).

In addressing the perceptual restoration of missing signal fragments when replaced by noise, Newman (2004) makes a critical comparison between the case where missing fragments are replaced by noise, and where the same fragments are replaced by silence. Crucially, the same amount of intact information available for identification is present in both cases. The only difference is that noise hides the absence of the missing information whereas silence does not. Particularly with a highly familiar stimulus such as a nursery rhyme, it is possible that some stimuli with silent gaps can be identified. Therefore, a good test of the operation of perceptual restoration is whether noise can provide an intelligibility advantage for signals over and above intelligibility of the intact sections alone.

Newman (2004) was also the first researcher to include an intact control condition, where the repetition accuracy of intact sentences was tested alongside stimuli with noise-filled gaps and silence-filled gaps. This was an important addition to the design since it allowed an analysis of the effect of disrupted signals in general on children's speech perception, as well as providing a measure of optimum performance. In most cases, however, a critical comparison between intact and noise-filled gaps conditions as a measure of perceptual restoration is not entirely appropriate. First, this requires that in order for perceptual restoration to be demonstrated, there should not be a significant difference between the intelligibility of intact signals and those with noise-filled gaps. This assumes that restoration should be complete; yet it is more likely that noise provides an intelligibility advantage relative to its absence without necessarily restoring intelligibility to that of an intact stimulus.

Furthermore, whilst noise-filled gaps and silence-filled gaps stimuli are identical except for the addition of noise into missing sections, a comparison between an intact stimulus and one with noise-filled gaps means that the effect of noise in general exists aside from any intelligibility differences. Hearing bursts of white noise would be an unusual experience for the majority of children (Walley, 1988) and could lead to effects quite aside from perceptual restoration if the comparison between intact and noise-filled gaps stimuli were taken as the critical measure of perceptual restoration. Comparing stimuli with noise-filled gaps to those with silence-filled gaps as a measure of perceptual restoration means that all stimulus characteristics are equal apart from the critical presence or absence of noise.

Whilst stimulus identification provides the most appropriate measure of perceptual restoration in the context of the perceptual restoration of music in young children, it is not the only informative measure of perceptual restoration that can be obtained through an identification task. Expectations are a key part of the perceptual restoration process, and it has been suggested that the closer the match between listeners' expectations about how something should sound, and how close the actual stimulus is to this expectation, the faster the stimulus will be identified (Mills, 1980). For this reason, the time taken to identify stimuli will be analysed to see if this measure corroborates identification scores.

In the first part of this thesis, previous research into the perceptual restoration effect has been presented, together with a consideration of similarities and differences in the perception of music and speech, and an outline of important developments in children's music perception. The theoretical issues involved in perceptual restoration, together with possible explanations of how the effect occurs, have been presented. Finally, in the last section, different methods used to investigate perceptual restoration have been outlined, together with a justification for the methods to be employed in the current series of experiments and why they are the most appropriate way of assessing the perceptual restoration of music in young children. The section that follows summarises the aims of the research to be presented here.

Aims of the current research

There were five main aims of the present research, which represent the research questions to be addressed through the empirical studies. First, the primary aim of the present research is to investigate whether perceptual restoration in young children is restricted in its operation to speech perception, or whether it also operates in another auditory domain, that of music perception. The primary focus will be on the ecological importance of perceptual restoration. This aspect of perceptual restoration has been emphasised in recent work (e.g. Shahin et al., 2009). It is possible that the perceptual restoration of music in children has not previously been investigated because the ecological value of a perceptual restoration mechanism in music perception is not seen as being as important as such a mechanism in speech perception. It will be argued here that for children, perceptual restoration in music perception is equally, if not more important than perceptual restoration in speech perception, given the major role of music perception in children's development and the way that young children learn through music.

Second, the perceptual restoration of music has only been studied in an adult population. Investigating the perceptual restoration of music in children is important for understanding the perceptual processes involved in the effect, and thus represents an important aim of this research. For example, if musical perceptual restoration does not occur in children, then this raises the question of what differences exist between auditory perceptual processing in adults and children. Equally, if perceptual restoration in children is restricted to the speech domain, it will be necessary to consider possible differences between children's speech and music perception.

Third, if evidence of the perceptual restoration of music in children is demonstrated, a further aim of this research is to establish the role of expectation (top-down influences of familiarity and context) and corroboration (bottom-up acoustic influences) in the effect. Manipulations in previous studies of musical perceptual restoration have largely been restricted to top-down factors such as familiarity and predictability. Understanding the role of bottom-up acoustic factors represents an equally important endeavour, and is a key part of understanding the effect in music perception. In addition, the perceptual restoration of music has previously only been demonstrated for a single musical note in a longer melody. Whilst Samuel reports a personal experience of hearing perceptual restoration of multiple periodic fragments in music (DeWitt & Samuel, 1990), this requires experimental investigation. With reference to the expectation part of the process, research has shown that young children do have stored representations of familiar music, and expectations operate in music perception from a young age (Dowling, 1999). Whether these representations are established enough to be used to assist perception when sensory input is disrupted represents an important question to be addressed.

Fourth, familiarity is cited in the adult perceptual restoration literature as a key component of the perceptual restoration process. Nevertheless, the concept of familiarity requires further clarification; in the context of perceptual restoration, familiarity might involve being familiar with one component of a signal (such as just the tune, or the words) or it might involve an exact match with a previously stored representation (such as a complete song). In previous research into the perceptual restoration of music, "music" has in most cases been taken to constitute a single melodic line, and research has not considered the different role that, for example, pitch and rhythm play in the effect. An important aim of the present research is to investigate which parts of a musical signal are best supported by perceptual restoration.

Fifth, this research aims to test for any developmental changes in the perceptual restoration of music, based on findings that such developments in perceptual restoration ability occur in the speech domain.

Perceptual restoration is clearly an ecologically important perceptual process. Yet little is known about the perceptual restoration of music in comparison to the perceptual restoration of speech. If the perceptual restoration of music does occur in children, this would reveal important insights into auditory perceptual processing of music at a young age. It is also important to investigate how the effect operates in children, and whether similar factors to those that influence the perceptual restoration of speech also influence the perceptual restoration of music in children. The chapters that follow report a series of eight related experiments designed to address these issues.

PART TWO: EMPIRICAL STUDIES

Overview

Clear implications, both theoretical and educational, of the operation of perceptual restoration in children's music perception have been outlined in the preceding chapters. These implications formed the basis of the aims of this research: to investigate musical perceptual restoration in young children, and to establish some of the factors that influence its operation. With these aims in mind, what follows is the report of eight related experiments that address these issues.

Experiment 1 establishes the experimental methods to be employed in this series of experiments and was designed in order to test for the operation of musical perceptual restoration in 4- to 6-year-old children. Experiments 2 and 3 aimed to determine the influence of acoustic factors on musical perceptual restoration, and to this end involved manipulations of the duration of missing fragments in the signal and the amplitude of the noise replacing missing fragments. Contextual influences on perceptual restoration are investigated in Experiments 4 and 5, through the manipulation of the typicality of musical excerpt presented and through investigation of the perceptual restoration of vocal music. Experiments 6 and 7 aimed to determine more explicitly the role of familiarity in musical perceptual restoration, through investigation of the components of a musical signal that are critical for perceptual restoration. Finally, Experiment 8 addresses developmental improvements in musical perceptual restoration. Consideration is then given, through a by-material analysis, as to whether the musical characteristics of the excerpts presented across experiments could influence perceptual restoration.

Chapter 3: Experiment One

Does the perceptual restoration of music operate in children?

Introduction

Children are faced with noisy environments on a daily basis. The primary aim of the present study was to investigate young children's ability to compensate for disruptions to musical signals in noisy listening conditions using perceptual restoration, and to evaluate what this might reveal about the nature of music perception in children. The first issue to be addressed was thus whether perceptual restoration operates in the music domain for child listeners.

Perceptual restoration and auditory perceptual processes in children

Perceptual restoration in adults has often been described as providing evidence that auditory perception is highly interactive (e.g. McClelland et al., 2006), where expectations that have developed from passive exposure, and knowledge-based contextual constraints, influence perception of stimulus features (Ellis, 1999). Perceptual restoration also indicates that auditory perception is constructive; in restoring missing input, sounds not present in reality are heard to be intact (Riecke et al., 2007). In contrast, the view of a perceptual system that simply processes the incoming sensory input in terms of what it does or does not contain leaves no role for the perceiver in the perceptual process; perception is passive not active, and primary attention is given to the stimulus and not the perceiver in the perceptual process (Vallar, 2006). In the process of perceptual restoration, the perceiver also plays a key role:

> [mechanisms such as perceptual restoration] allow listeners to repair degraded signals and lead them to believe they are perceiving an intact auditory pattern rather than fragments joined together with patches they themselves had created. (Warren, 1999, p.134)

In a similar way, evidence of the operation of perceptual restoration in children could provide insights into the nature of perceptual processing in young listeners. According to Aslin and Smith (1988), the perceptual systems of young children are heavily influenced by experience, and the passive perceptual learning that contributes to the knowledge base that forms the basis of this experience occurs early in life (Bharucha, 1994). Investigating the operation of the perceptual restoration of music in children can provide insights into the nature of music perception, and more generally auditory perception, in young children. Of interest is whether children can use their previous musical knowledge to help them interpret ambiguous or incomplete sensory input, and hence whether music perception in children is fully interactive and involves a bi-directional flow of information, where higher-level representations can influence lower-level perceptual processing.

Perceptual restoration demonstrates global information processing of disrupted information (Remijn et al., 2007). Where noise replaces missing sections of an auditory signal, and hence the global perceptual coherence of the signal as continuous remains unaffected, perceptual restoration operates with the result that local disruptions (i.e. where missing information is actually replaced by noise) are not registered. When missing sections remain silent, this global coherence is lost and hence local disruptions are particularly salient. Children are also believed to be global information processors, as shown, for example, by their greater susceptibility to visual illusions when compared to adults (Vernon, 1966). In the auditory domain, there is evidence that children's music processing demonstrates a similar global preference (Trehub, 2001). Therefore, the operation of perceptual restoration in children can provide further support to accounts that children's perceptual systems, particularly in the music domain, also favour global configural features rather than local cues.

Perceptual restoration in adults demonstrates the ability to make and evaluate hypotheses about the most likely content of disrupted information. Evidence for these hypotheses is gathered from the information contained within the auditory signal. This hypothesis-generating and evidence-gathering process (Bregman, 1990; Koroleva et al., 1991; Shinn-Cunningham, 2008) is an important part of the ability to compensate for missing or disrupted sensory input. For example, if a familiar melody is heard where some sections are obscured by a louder sound, such as a cough, the listener would have expectations about what would most likely be heard in the disrupted sections based upon their stored template of the melody in their musical lexicon. These hypotheses are then evaluated against the evidence provided by the auditory signal. If the extraneous noise is a potential masker of the missing information, there is no evidence to suggest that the expected sections are not in fact present. In this situation, where the input is perceived to be complete, a match will be registered between the sensory input and the stored template of the melody. However, if the missing sections are not masked, the evidence suggests that the hypotheses about the content of the missing information are incorrect. As a result, the sensory input does not match the stored template of the melody, or only constitutes a partial match, which could detrimentally affect melody identification.

If the perceptual restoration of music does operate for children, it suggests that this hypothesis-generating and evidence-gathering ability operates in the musical domain for children, which would represent an important adaptation to environments where disrupted or ambiguous sensory input can be common.

Research evidence on perceptual restoration in children

Children's ability to fill in missing parts of their perceptual environment has been well documented in the visual domain (e.g. Craton, 1996). Similar evidence has been obtained for the auditory domain; there are some reports in the literature that the perceptual restoration of speech signals does occur in children.

The first documentation of perceptual restoration effects with child participants was reported by Ackroff (1981), where children between the ages of 6 and 8 years showed some evidence of the perceptual restoration of missing phonemes, with Walley (1988) finding similar evidence for 5- to 6-year-old children. Koroleva et al. (1991) obtained evidence of the perceptual restoration of speech in 5- to 6-year-old Russian children, with a further study (Koroleva et al., 1996) revealing gender differences in this ability. These differences suggested that girls' responses were more like those of adults than boys' responses, and girls also showed a more advanced ability than boys to combine top-down and bottom-up information in the process of perception. This ability is equally relevant to musical perceptual restoration as it is to speech restoration, hence in the present experiment it will be important to investigate whether similar gender differences operate on the perceptual restoration of music.

These previous studies of perceptual restoration in children tested for the perceptual restoration of a single phoneme. Newman (2004) obtained evidence of perceptual restoration in 5-year-old children by comparing the repetition accuracy of sentences where multiple alternating fragments were removed and either replaced by white noise or silence. This type of stimulus was seen as a better indicator of children's ability to use the context and their previous knowledge to assist the perceptual restoration process. Newman found that children's repetition accuracy of the sentences was far superior when the missing sections were replaced by noise than by silence, and interpreted these findings as evidence that children were using their knowledge and the context to interpret what would most likely be the content of the missing sections. This in turn suggests that children can combine top-down and bottom-up information in the process of perceptual restoration, and, more generally, in speech perception.

Research findings on the perceptual restoration of speech in children are relatively sparse when compared to the amount of research on the effect in adults, which is somewhat surprising given the clear implications of perceptual restoration for children's development. Furthermore, there is as yet no evidence for a comparable mechanism operating in children's music perception. Such a mechanism would have obvious advantages, given the importance of music to children's development (e.g. Racette & Peretz, 2007), and the large amount of young children's learning that is song-based. Music has often been described as a universal human ability (Cross, 2003), and something that plays a pivotal role in children's cognitive and social development. However, many of children's everyday environments, especially classrooms where music-based learning might take place, are particularly noisy (Manlove et al., 2001). The ecological importance rather than just the theoretical interest of perceptual restoration as a laboratory phenomenon is something that is becoming increasingly apparent, and has been emphasised in recent work (e.g. Riecke et al., 2009; Shahin et al., 2009).

Adults possess the ability to reconstruct missing sections of musical signals using perceptual restoration, in a comparable way to the same ability in the speech domain. Sasaki (1980), in finding such evidence, suggested that this demonstrated the influence of higher levels of processing on the perception of musical signals. DeWitt and Samuel (1990) carried out a systematic study of the perceptual restoration of music, based on a schema-based theory of music perception (Krumhansl & Castellano, 1983) where previous knowledge was seen as an important influence on music perception. Their findings demonstrated such an influence, with the implication that music in adults is not processed passively. Kaminska and Mayer (1993) also claimed that music processing is more likely to be schema-driven that stimulus-driven, based on data demonstrating that the perceptual restoration of music does operate in adults, and does so in a very similar way to the perceptual restoration of speech. Smith (2005) also emphasised the parallels between the operation of perceptual restoration in music and speech perception, describing it as a domain-general phenomenon.

Given that there are believed to be similar developmental underpinnings for speech and music (McMullen & Saffran, 2004), there is reason to expect that perceptual restoration will operate in the music domain for children. Similarly, there are believed to be common processing mechanisms that operate in both music and speech perception (Saffran, 2003), of which perceptual restoration might be an exemplar. The common function of perceptual restoration can be better understood by considering whether in humans the operation of the mechanism is specific to speech communication, or a more general mechanism preserving communication in different auditory domains.

The present experiment investigates the perceptual restoration of music by asking children to identify excerpts from nursery rhymes and play songs in intact or disrupted forms. Nursery rhymes and play songs are important for children's development and are the musical genre children are most frequently exposed to (Fitch, 2006; Trehub & Trainor, 1998). Although identification provides an indirect measure of perceptual restoration, it is nonetheless ecologically valid as it represents the use of perceptual restoration in the kind of situation in which it might be put to use in everyday conditions (i.e. being able to identify a melody in a noisy environment). Furthermore, a method requiring judgements about stimuli where noise is added to or replaces a phoneme, for example, is not really appropriate for testing children. Asking children to identify nursery rhymes simplifies the task demands.

If perceptual restoration does operate for musical signals in children, melodies will be easier to identify when missing sections are replaced by noise, than when the same sections are replaced by silence. Exactly the same amount of intact information is present in both cases; the only difference being the bottom-up confirmation of expectations provided by replacement noise.

Method

Participants

Sixty-six children (27 girls) participated in this experiment. The children were from the Reception and Year One classes of a primary school in a predominantly white, middle-class area of Surrey, UK. The children were between the ages of 4 years 7 months and 6 years 8 months, and had a mean age of 5 years 9 months (69.03 months; *SD* = 7.06). This age group was selected to allow comparisons to be made with the phonemic restoration results reported by Newman (2004), and because children of this age can reasonably be expected to understand the task instructions and task demands. Consent for participation was gained both from parents and from the school Headteacher. All children in the sample were native English speakers, and had normal sensory development as reported by the school. Further details about the participants in each condition are shown in Table 3.1.

Table 3.1. The gender split and mean age of the children in each of the three experimental conditions.

Experimental Condition[a]	**Gender**	**Mean age in months (*SD*)**
Intact	M=15; F=7	68.41 (7.01)
Noise-filled gaps	M=10; F=12	69.55 (6.76)
Silence-filled gaps	M=14; F=8	69.14 (7.67)

[a] *n* = 22 for each condition

Design

Participants were randomly allocated to one of three independent conditions. Each participant was presented with the same 15 nursery rhyme and play song melodies[6] (see Table 3.2), but these were modified in different ways depending on the experimental condition.

---

[6] Hereafter, the term "nursery rhyme" is used to refer to both nursery rhymes and play songs.

Table 3.2. The name, duration, time and key signatures of the fifteen stimuli presented.

Nursery rhyme	**Excerpt duration (s)**	**Time signature**	**Key signature**
Baa Baa Black Sheep	5.90	4/4	C major
The Grand Old Duke of York	5.25	4/4	C major
Twinkle Twinkle Little Star	5.75	4/4	C major
Happy Birthday to You	4.75	3/4	C major
The Wheels on the Bus	5.32	4/4	C major
The Hokey Cokey	6.50	2/4	C major
Row, Row, Row Your Boat	4.25	4/4	C major
Old Macdonald had a Farm	4.50	4/4	C major
London Bridge	4.50	4/4	C major
Jack and Jill	4.39	2/4	C major
Incy Wincy Spider	5.00	4/4	C major
Humpty Dumpty	8.00	2/4	C major
Hickory Dickory Dock	5.50	4/4	C major
If You're Happy and You Know it Clap Your Hands	6.75	4/4	C major
Heads, Shoulders, Knees and Toes	3.25	4/4	G major/ C major

In the noise-filled gaps condition, alternating 100 ms fragments of the nursery rhyme melody were removed and replaced with white noise. McDermott and Oxenham (2008) argue that bursts of white noise in experimental situations perform exactly the same function as a cough or a clap, for example, would perform in everyday listening conditions in replacing missing sensory input. Furthermore, white noise as a gap replacement has most frequently been used in previous perceptual restoration experiments and thus, by following this trend, the present findings should be easier to interpret in the context of previous work in the area of perceptual restoration.

In the silence-filled gaps condition, exactly the same fragments were removed and replaced with silence. Finally, in the intact condition, the nursery rhymes were presented intact with no fragments removed. This control condition is extremely important as it provides a measure of optimal performance against which to compare identification of disrupted signals (Newman, 2004). Intact in this context describes a totally unmanipulated signal, not where noise is superimposed upon an intact stimulus, as in experiments by Samuel and colleagues (e.g. DeWitt & Samuel, 1990; Samuel 1981a).

The duration of missing gaps was set at 100 ms in line with the multiple phonemic restoration stimuli reported by Warren (1999). By removing alternating fragments, the amount of intact information is equal to the amount of missing information. This equates the amount of contextual support for identification with the amount of information to be restored (see Figure 3.1). This method actually leads to a conservative estimate of perceptual restoration for the following reason. If only a single note is removed, then perceptual organisation is likely to revolve around the more dominant intact sections. When as much of the signal is removed as is present, there should be no such dominant organisation in favour of an intact stimulus (Drake & McAdams, 1999).

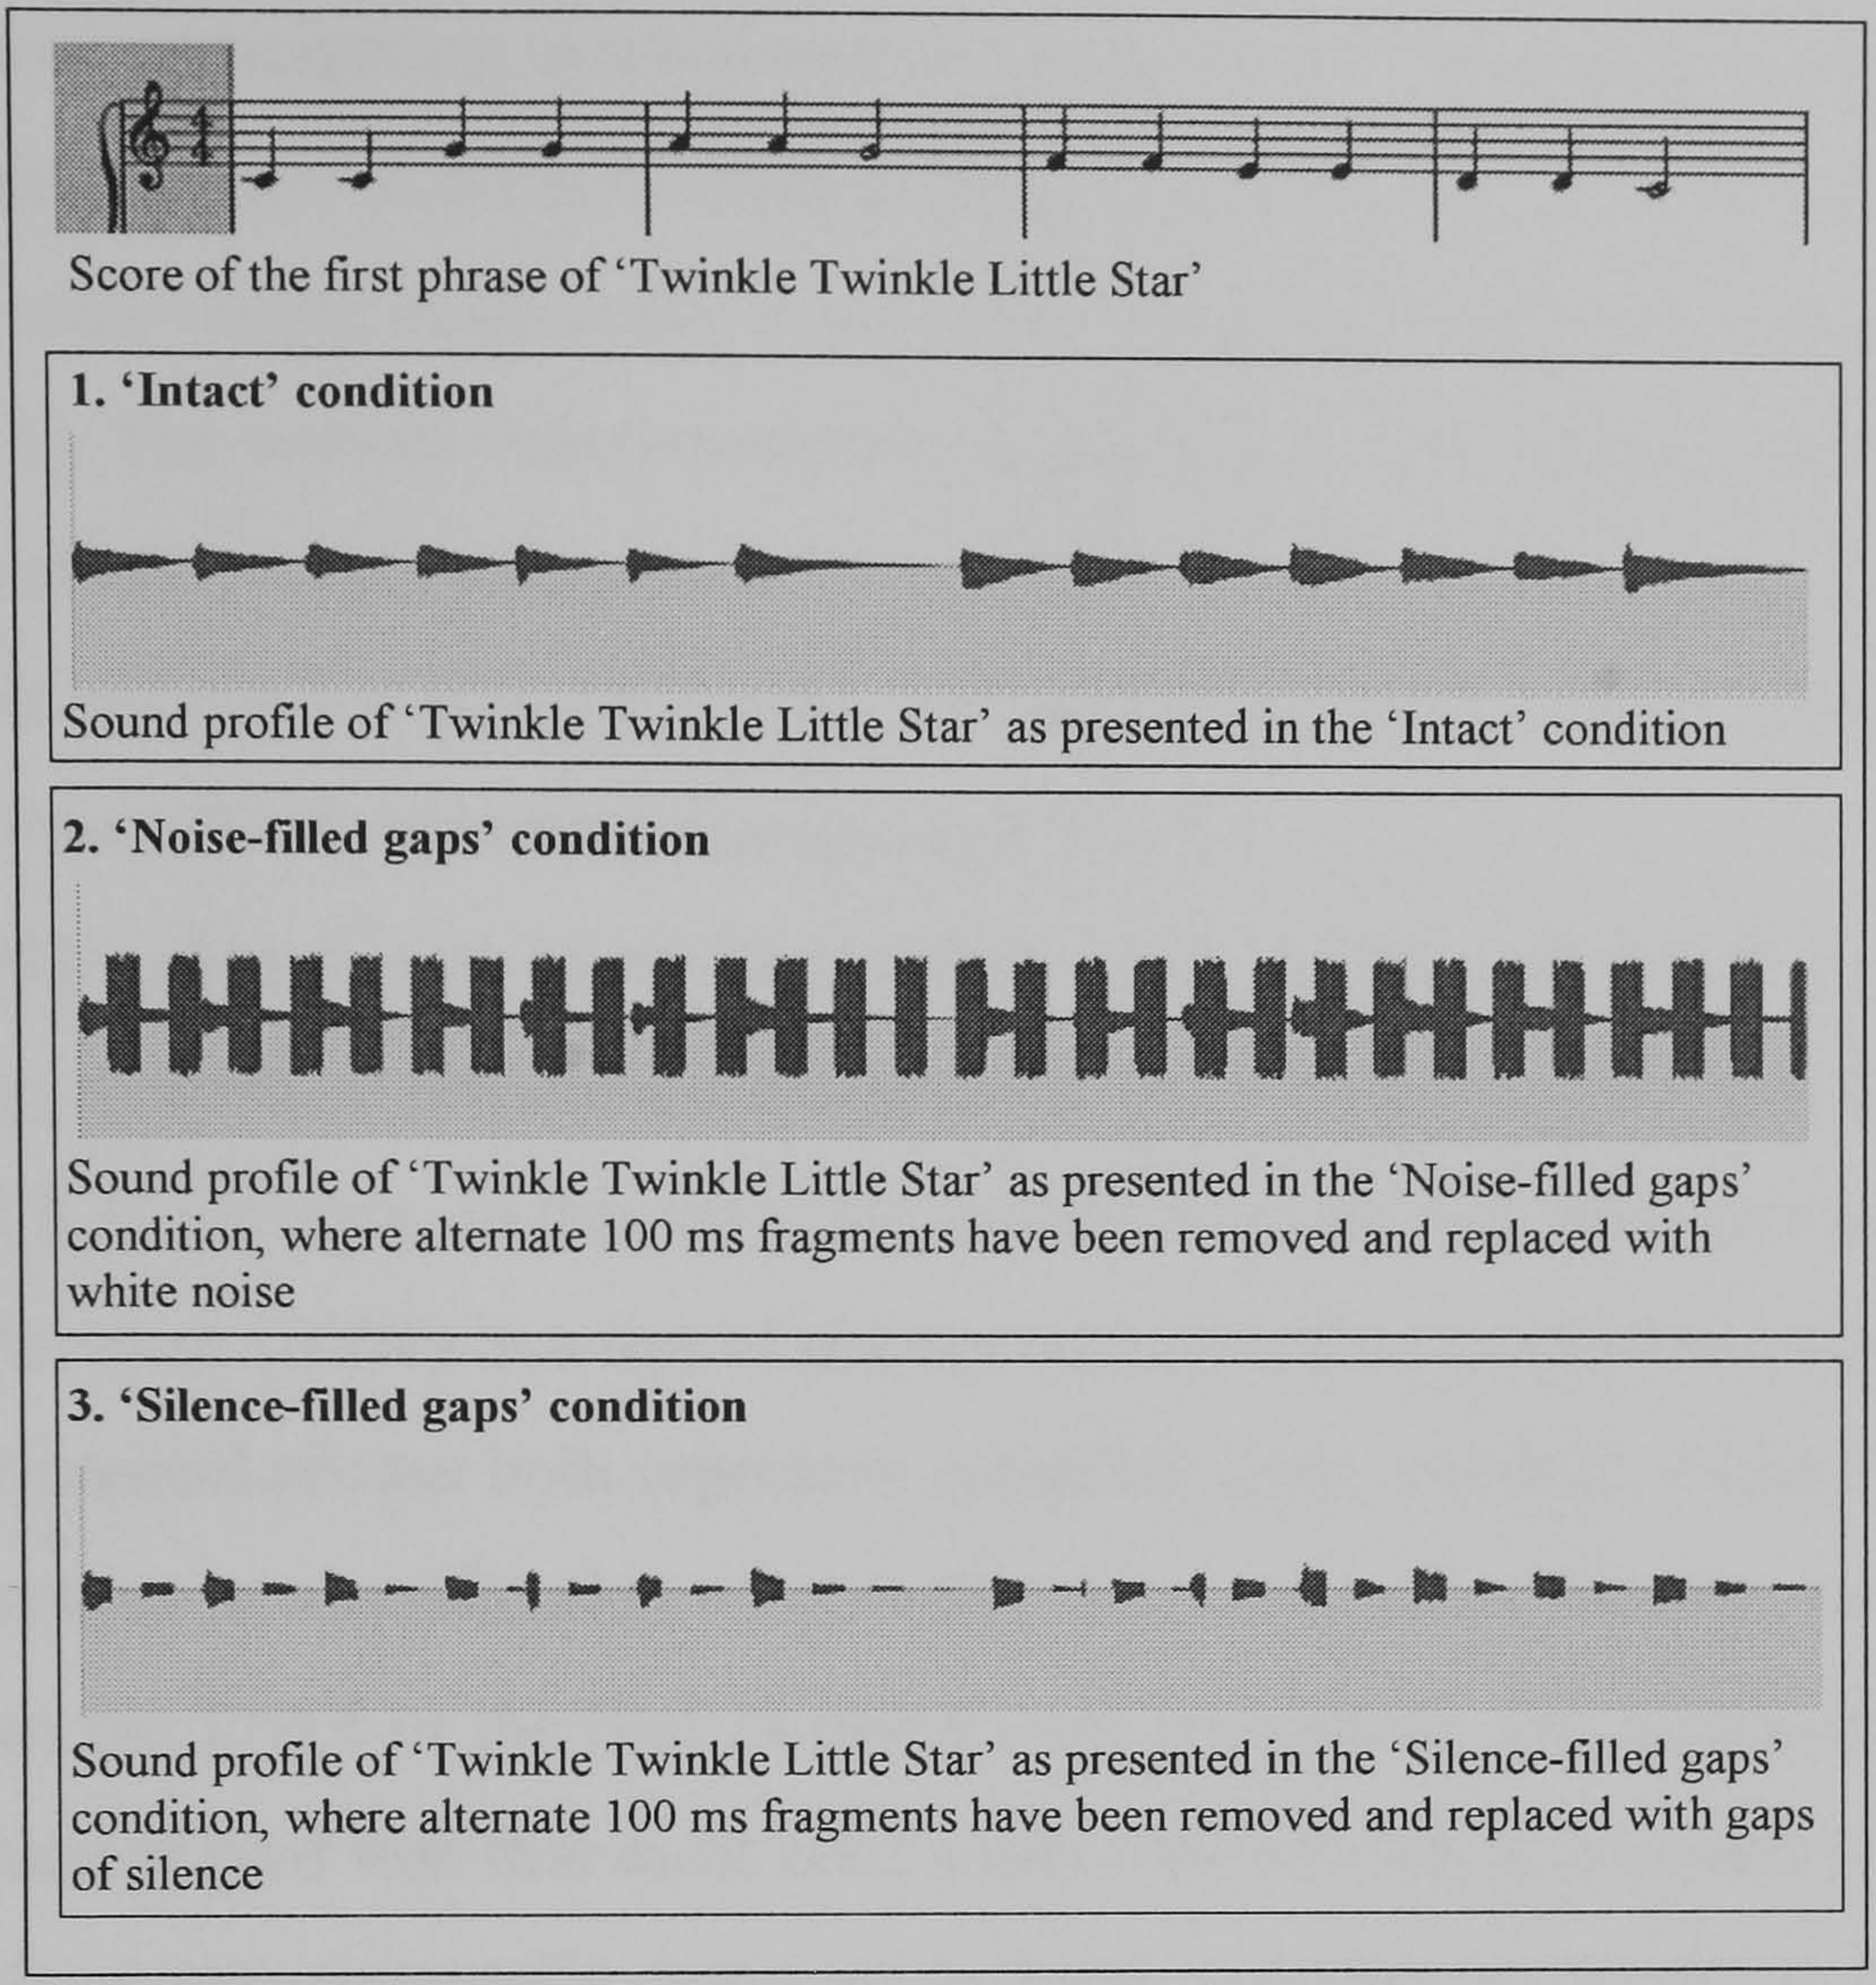

Score of the first phrase of 'Twinkle Twinkle Little Star'

1. 'Intact' condition

Sound profile of 'Twinkle Twinkle Little Star' as presented in the 'Intact' condition

2. 'Noise-filled gaps' condition

Sound profile of 'Twinkle Twinkle Little Star' as presented in the 'Noise-filled gaps' condition, where alternate 100 ms fragments have been removed and replaced with white noise

3. 'Silence-filled gaps' condition

Sound profile of 'Twinkle Twinkle Little Star' as presented in the 'Silence-filled gaps' condition, where alternate 100 ms fragments have been removed and replaced with gaps of silence

Figure 3.1. Example sound profiles of one of the stimuli presented: score of *Twinkle Twinkle Little Star* and sound profiles of the intact, noise-filled gaps and silence-filled gaps versions of the stimuli.

Whilst the case here where sections of the signal are removed and completely replaced by noise does not generalise well to real life where it is more likely that a louder sound will mask, rather than replace, sections of the signal, experimental situations like the present study need to rule out the effects of incomplete masking (Warren, 1984). In addition, complete replacement of missing sections gives a conservative estimate of restoration, and hence by using this method it is more likely to underestimate the ability of perceptual restoration to restore intelligibility of disrupted signals than to overestimate it (Vallar, 2006).

Materials

The school participating in the experiment was sent a familiarity questionnaire prior to the testing session in order to confirm that the nursery rhymes that would be used as stimuli were taught in the early years of the school and sung on a regular basis[7] (see Appendix 1). The stimuli were produced by playing the monophonic melody line of the first phrase of each nursery rhyme on an Oxygen8 MIDI keyboard (piano sound selected; default settings for volume, reverb, and pan). This keyboard was connected to a Macintosh G4 PowerBook operating Garageband V3.0.4.

A complete phrase[8] was taken from each nursery rhyme rather than an excerpt of a fixed duration, in order to present a coherent unit (see Table 3.2); the stimuli had a mean length of 5.31 seconds. In this way, the stimuli are similar to the sentences presented by Newman (2004) in a test of phonemic restoration in children, since a sentence and a musical phrase both represent complete units. Further, White (1960) found no differences in identification of melodies as a result of the number of notes presented (either the first 6 or the first 24 notes were presented). The tempo at which each stimulus was played was that most commonly represented on recordings of nursery rhymes; it has been experimentally demonstrated that melodies are harder to identify if played at a non-typical tempo (Warren, Gardner, Brubaker, & Bashford, 1991).

Noise or silence was inserted into gaps in the stimuli using Soundstudio V2.2.4. The peak amplitude of the stimuli was equated using a normaliser to 20 dB below the maximum bandwidth of the file, and the amplitude of the replacement noise was kept constant, at a level 12 dB below the maximum bandwidth of the file. A programme written in Matlab (V5.2.1) presented the stimuli to the child in a random order, and recorded the child's response and identification time.

---

[7] Across all experiments presented in this thesis, schools confirmed that the proposed stimuli were taught in school and sung regularly; thus, all stimuli met the conditions for inclusion in the set.

[8] A musical phrase is a section of music that has a sense of completeness; natural breaks in the melody usually signify the division between each phrase.

Procedure

Each child was tested individually in a quiet room in the school. The experimenter (author) told the child that they were going to play a music game, where they had to work out what some children's tunes were:

> "I need your help to work out what some children's tunes are. They are sounding a bit silly on my computer and I don't know them very well."

The child listened to the stimuli through headphones and was presented with two demonstration trials (always *Baa Baa Black Sheep* and *The Grand Old Duke of York*), followed by 13 experimental trials. If the child did not identify the nursery rhyme from the demonstration trial, no further explanation was given, and the next trial followed immediately. On presentation of each experimental trial, the child's task was to name the nursery rhyme. The experimenter told the child,

> "Listen to each tune carefully. As soon as you know what the tune is, say "stop!" and you can tell me the name of the tune. You do not need to wait until it has finished playing."

No feedback was given as to whether the child's response was correct or incorrect; the next trial followed immediately. If a child gave no response, or could not identify the stimulus, this was scored as incorrect. No time limit was imposed upon children giving a response. After all of the trials had been presented, the child was praised and thanked for his/her participation.

Results

Correct identification scores

Each child was given a score out of 13 that corresponded to the number of melodies they identified correctly in the experimental trials. Data were examined for differences in identification scores as a result of gender (see Appendix 2). There was no significant effect of gender, so data were collapsed across this factor. Figure 3.2 shows the mean number of correct identification responses for children in each condition.

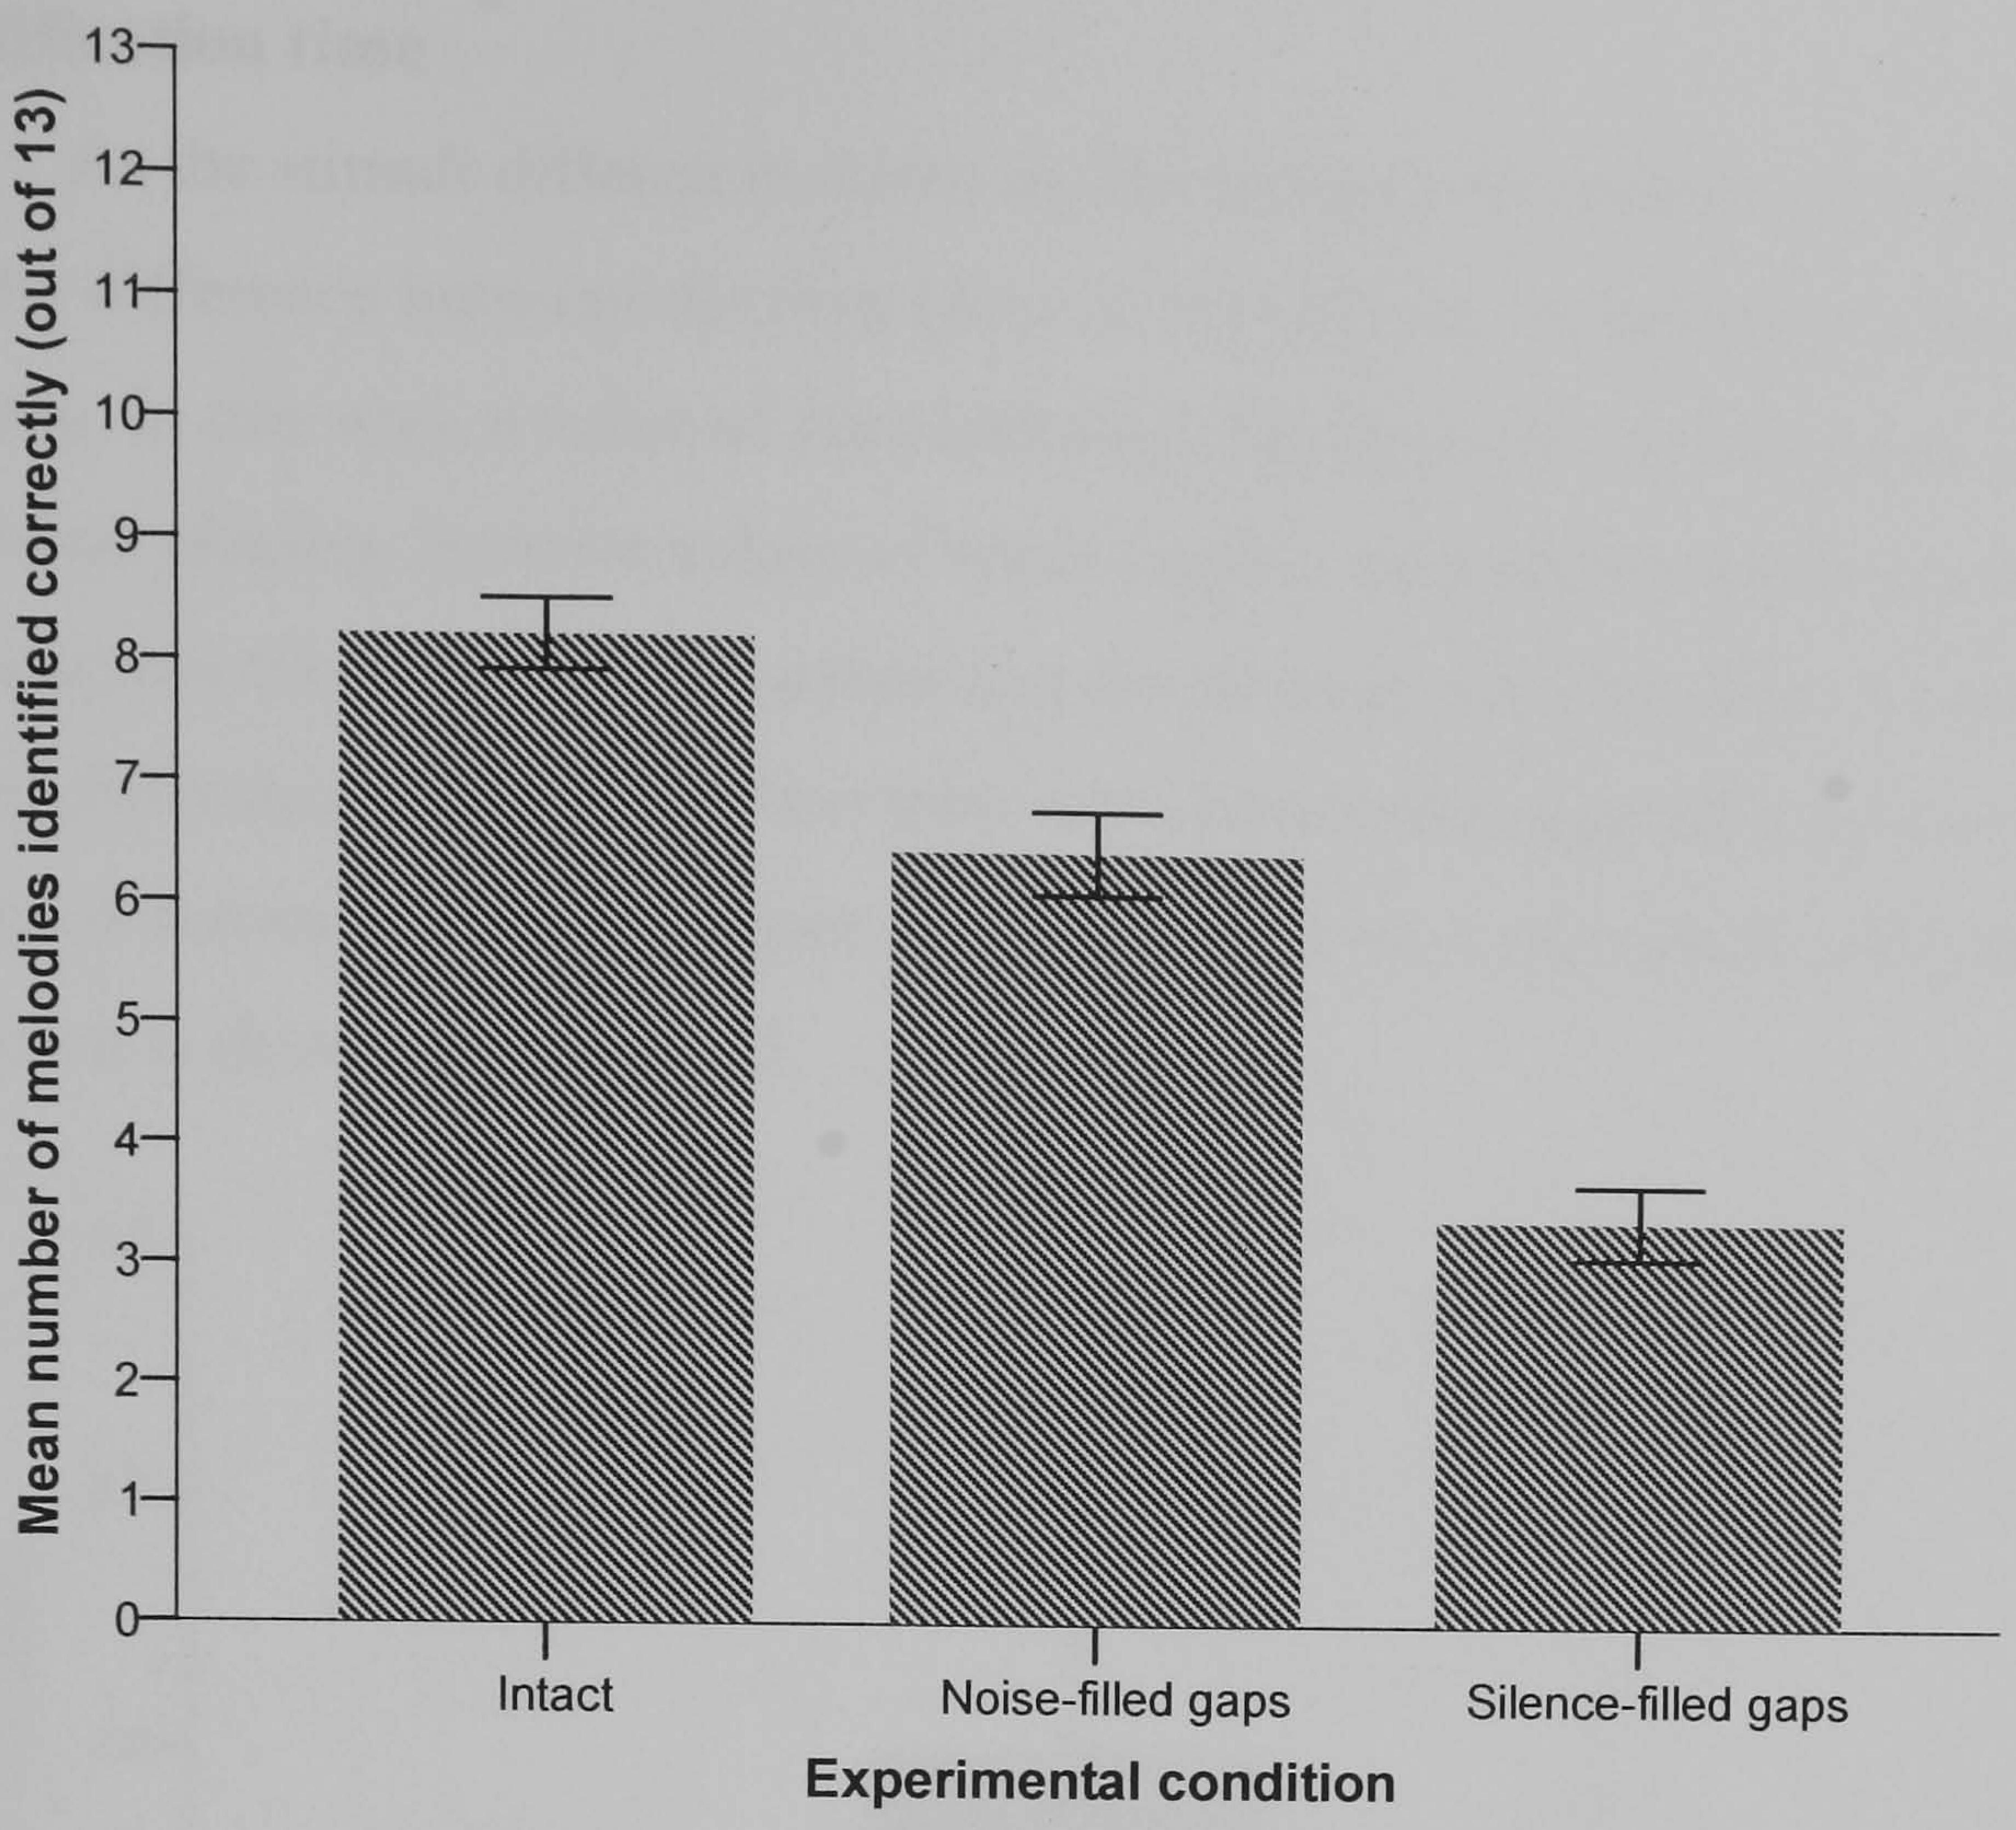

Figure 3.2. Mean number of melodies identified correctly (± *SE*) by children in the three experimental conditions (intact, noise-filled gaps and silence-filled gaps).

A one-way ANOVA with the three-level factor of experimental condition (intact, noise-filled gaps or silence-filled gaps) revealed a significant condition effect on correct identification scores, $F(2, 63) = 59.37, p < .001, \eta^2 = .65$. Post hoc tests[9] (Tukey HSD) comparing each pair of conditions revealed significant differences in all cases (intact versus noise-filled gaps, $p < .01$, $d = 1.18$; intact versus silence-filled gaps, $p < .001$, $d = 2.43$; noise-filled gaps versus silence-filled gaps, $p < .001$, $d = 1.29$). Figure 3.2 shows that overall, the intact melodies were significantly better identified than both the stimuli with noise-filled and silence-filled gaps. However, the key result is the finding that there was also a significant advantage of noise-filled gaps stimuli compared with silence-filled gaps stimuli on the number of melodies identified.

---

[9] For between-subjects factors reported across all experiments, Tukey HSD is used as the post hoc test where group variances are homogenous. Where group variances are heterogenous, Games-Howell post hoc tests are used.

Identification time

As the stimuli differed in duration, the measure of identification time employed was the difference between the time taken to identify the stimulus and the stimulus duration. In this way, a value of zero indicates that the stimulus was identified exactly as it finished playing. Positive values of identification time represent the number of seconds after the excerpt had finished that the melody was identified correctly. No gender differences in identification time were found (see Appendix 2) hence data were collapsed across gender. The mean identification time in seconds for children in each condition is shown in Figure 3.3.

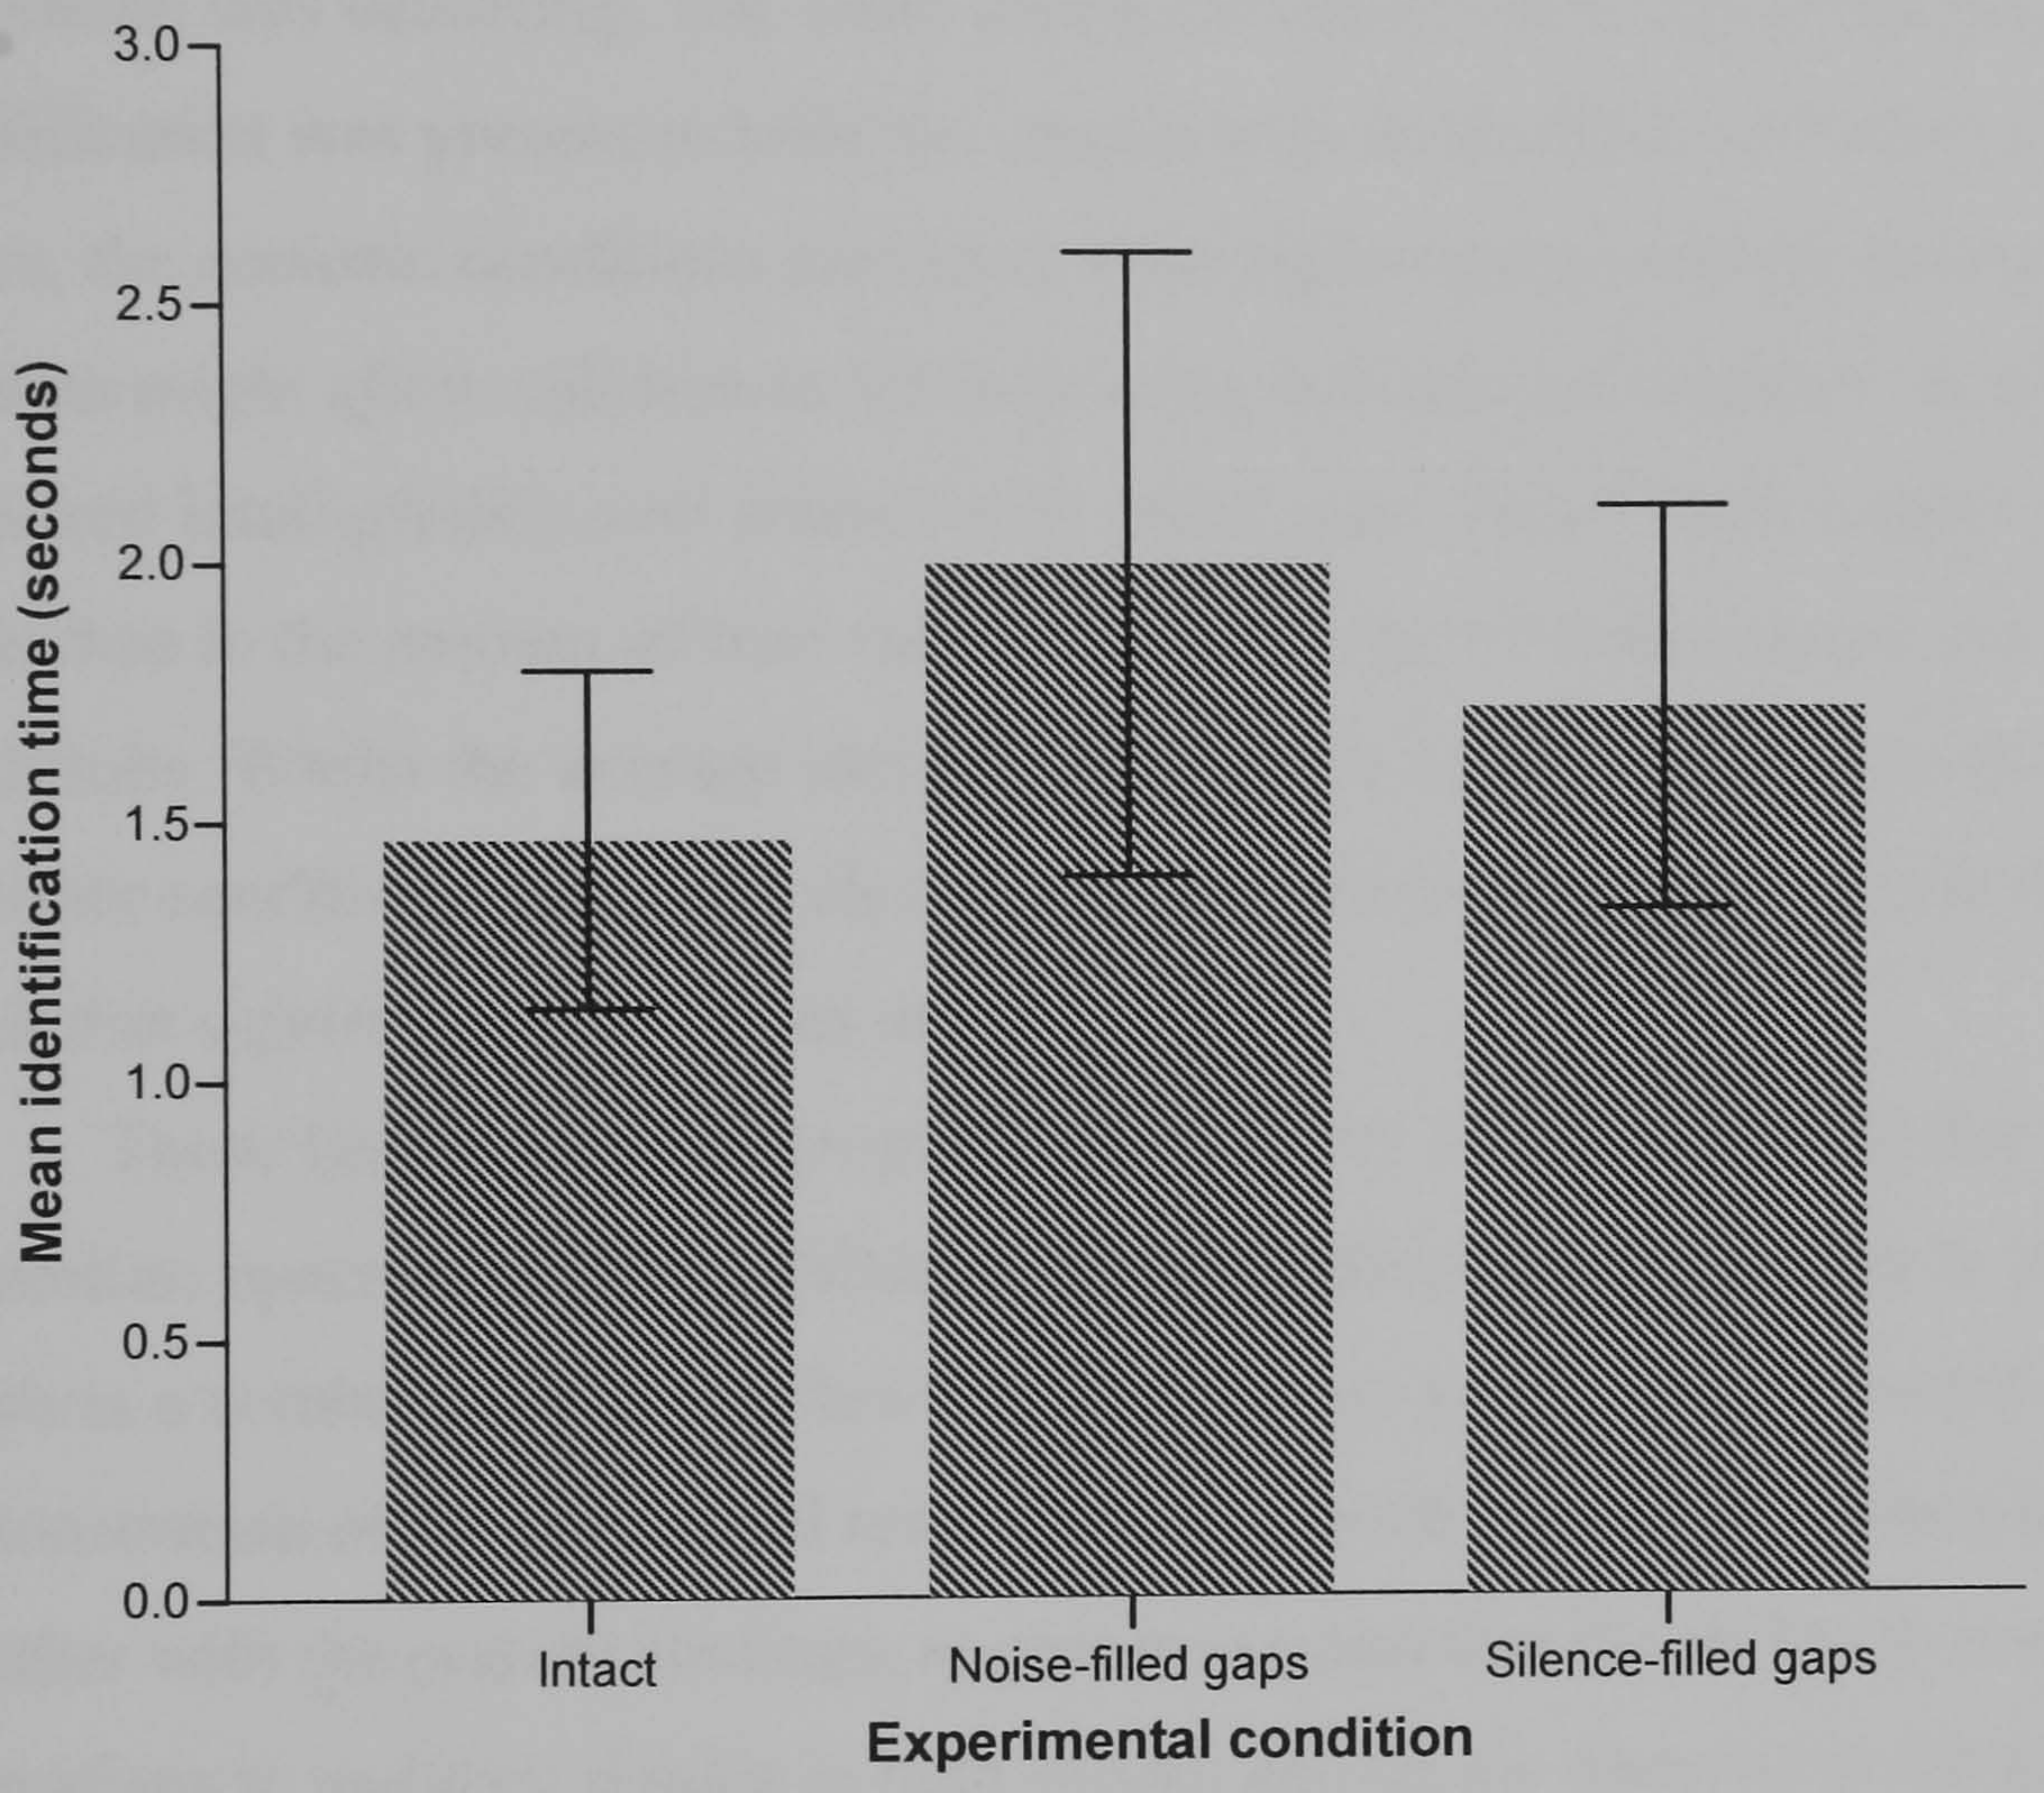

Figure 3.3. Mean identification time (identification time – stimulus duration; ± *SE*) in seconds for children in the three experimental conditions (intact, noise-filled gaps and silence-filled gaps).

These data were analysed using a one-way ANOVA with the three-level factor of experimental condition (intact, noise-filled gaps or silence-filled gaps). There was no significant main effect of experimental condition on identification time, $F(2, 63) = .15, p > .05$.

Discussion

The main aim of this first experiment was to investigate young children's ability to perceptually restore missing fragments of a non-linguistic signal, namely musical melody. The results were clear cut. First, there were no effects of gender. Second, there was a clear advantage in identification of melodies with noise-filled gaps when compared to the same melodies with silence-filled gaps. This indicates that perceptual restoration was occurring. The same amount of intact information that might afford identification was present in both the stimuli with noise-filled and silence-filled gaps; hence, the acoustic conditions provided when replacement noise is inserted into missing sections might allow children to fill in missing acoustic information, resulting in enhanced intelligibility over stimuli with silent gaps. Third, there was no significant difference in the amount of time taken to identify the melodies across the three conditions. Whilst the average identification time was faster in the intact condition than the other conditions, there was clearly a large amount of variation in the data with the result that significant differences did not emerge.

These findings extend propositions made by Newman (2004) that perceptual restoration operates in young children, demonstrating that perception in young children involves a combination of top-down and bottom-up processing. Newman's (2004) demonstration of the perceptual restoration of speech signals in 5-year-old children, together with the present findings, appear to suggest that the ability to compensate for disruptions to auditory signals in both speech and music domains is not restricted to adults but is well developed in early childhood.

The lack of any gender differences in musical perceptual restoration in children stands in contrast to the findings of Koroleva et al. (1996) in the speech domain. However, Koroleva et al. interpreted their gender differences in terms of differences in rates of verbal development between girls and boys. It may therefore be the case that such different rates of development do not exist in music perception, or if they do, they occur in areas that have no impact on musical perceptual restoration. Petzold (1969) reports finding no evidence of gender differences in children's music perception, hence the former possibility might well be correct. Furthermore, factors other than gender may result in individual differences in children's musical development.

Examination of identification scores indicates that although there was a significant noise-filled gaps over silence-filled gaps advantage in melody identification, there was also a significant difference in terms of identification scores between intact and noise-filled gaps conditions. This suggests that whilst perceptual restoration was occurring, this process was not able to restore intelligibility to the status of an intact stimulus. It might be that in children, signal disruptions do have a detrimental effect on music perception, but crucially the presence of noise still provides support to perception relative to its absence.

In addition, the results also suggest that identification of stimuli in the silence-filled gaps condition was not at floor level. With a highly familiar stimulus such as a nursery rhyme, it might be possible to identify a melody without the accompanying perception that it is complete. In the silence-filled gaps condition, there is clear evidence that sections of the melody are missing. In the noise-filled gaps condition, where identification was approximately twice as good as that in the silence-filled gaps condition, it is possible that children were unaware that sections were actually missing.

Consideration of some of the likely perceptual processes involved in the perceptual restoration of music might offer insight into exactly what it is about noise that improves intelligibility so measurably over silence. The findings from the present experiment suggest that children's music perception does display a preference for global coherence. Whilst stimuli in both the noise-filled gaps and silence-filled gaps conditions contain disruptions at a local level, only by replacing missing sections with noise does the stimulus satisfy the requirements of a perceptually coherent whole.

With reference to the proposed hypothesis-generating and evidence-gathering process involved in perceptual restoration, the present findings suggest that not only are both aspects important for perceptual restoration, their interaction is one of the key components of the perceptual restoration process. If perceptual restoration were based purely on hypothesis generation (analogous to top-down influences on perception) then one would expect no difference in identification scores for noise-filled gaps and silence-filled gaps conditions, as the same amount of intact information that would lead to the generation of these hypotheses is present. The evidence-gathering process is where the noise-filled gaps condition gains its advantage; it is through this process that corroborating evidence for the hypothesis is generated. This also suggests that not only are children able to use previous knowledge in the process of perception, they can also combine this previous knowledge with incoming acoustic information in the process of perception.

These findings suggest that music perception in children is not passive; their knowledge and expectations play a key role in the perceptual process, and it appears that there is sufficient information within a musical signal to allow children to generate hypotheses about the likely content of missing fragments. Furthermore, stored representations of familiar nursery rhyme melodies in 4- to 6-year-old children appear to be established enough to support perception of disrupted input. Thus, in this way, these findings build on those of Kaminska and Mayer (1993) in suggesting that music processing not just in adults, but also in children, appears to be schema-driven, and not purely stimulus-driven.

Taken together, these findings indicate that perceptual restoration in children is not restricted to the linguistic domain, and provide further support for accounts that music and speech might share similar processing resources (Saffran, 2003). Not only does this result suggest that music perception at an early age involves more than a simple analysis of the present acoustic elements in a signal, but when combined with similar findings on the perceptual restoration of speech, suggests that general auditory perception also operates in this way. The ability to combine incoming perceptual information with stored knowledge is critical in everyday listening conditions.

The kind of active music perception demonstrated through perceptual restoration, which is not detrimentally affected by extraneous noise, affords the potential to gain maximum benefit from musical learning experiences at an early stage of life. Both speech and music might well share some cognitive and perceptual principles in development, given that a mechanism that works to prevent communication being disrupted by everyday listening conditions operates in parallel in both auditory domains.

The first aim of this series of experiments was to establish the operation of perceptual restoration in children's music perception, and the results of this experiment clearly demonstrate that perceptual restoration operates for children in the music domain. However, it is not yet possible to determine whether perceptual restoration of music in children operates in the same way as perceptual restoration in adults, and responds in a similar way to manipulating the acoustic characteristics of the stimulus. Research on the perceptual restoration of speech in adults has found that the acoustic properties of the signal, and of the relationship between the signal and the replacement noise, influence the strength of perceptual restoration and, in some cases, whether the effect actually occurs. The next experiment was designed to test for the effect of acoustic factors on the perceptual restoration of music in children. This first experiment has indicated that perceptual restoration does operate in children for musical signals; it is now important to investigate the conditions under which this crucial mechanism can operate.

Chapter 4: Experiment Two

The effect of the duration of missing fragments on the perceptual restoration of music in children

Introduction

The ecological importance of a perceptual restoration mechanism has been so widely acknowledged that much research has endeavoured to determine the acoustic parameters under which the effect does or does not occur (Heinrich et al., 2008), and the relationship between the physical properties of a stimulus and the perceptual response to the stimulus. The particular parameter of interest in the present experiment is the duration of the missing fragments in the auditory signals (the "gap duration"), and the resulting effect on perceptual restoration. In testing for this relationship in the context of the perceptual restoration of music in children, it is also important to compare the results with previous work in this area in order to see if this relationship is similar in the speech and music domains, and between adult and child listeners.

The effect of gap duration on the perceptual restoration of speech

Perceptual restoration provides a good illustration of how perception involves an interaction between top-down and bottom-up processes. In perceptual restoration, top-down processes restore the content of the missing signal based on what is expected to be present, whilst the noise itself provides bottom-up confirmation of the restored signal by masking its absence.

When investigating the perceptual restoration of speech signals, researchers have manipulated the duration of the missing fragment(s) in order to test whether this factor affects the efficiency of the perceptual restoration mechanism. Some research has only considered the effect of gap duration on the perception of illusory continuity (the perception that a fragmented signal is continuous when gaps are replaced with noise), whereas some research has investigated the effect of gap duration on the perceptual restoration of intelligibility of fragmented speech. Most studies investigating the effect of gap duration on perceptual restoration have employed signals where multiple interruptions are made, rather than removing, for example, a single phoneme or a single musical note. Regardless of whether single or multiple interruptions are made, research has demonstrated that perceptual restoration is stronger for shorter gap durations, in both speech (e.g. Warren & Obusek, 1971) and in music (Sasaki, 1980). Similar results have also been obtained for visual perceptual completion, where the width of an occluder affects infants' perception of trajectory continuity through the occluder (Bremner et al., 2007). Furthermore, increasing age is related to the ability to fill in wider gaps in visual perception (Johnson, 2004).

The ability of noise to restore the continuity of fragmented auditory signals is affected by the duration of the interruption (Elfner & Caskey, 1965; Elfner & Homick, 1966; Shahin et al., 2009; Thurlow & Elfner, 1959). Furthermore, when asked to locate the position of a cough that replaced either a single phoneme or an entire syllable (consisting of three phonemes) in a spoken sentence, participants in a study by Warren & Obusek (1971) made more error in localisation of the cough for a single phoneme than the longer syllable. Since larger localisation error indicates stronger perceptual restoration of the missing sound, this experiment provides evidence of stronger perceptual restoration of a shorter interruption than a longer interruption.

Research on the perceived continuity of speech signals where missing input is replaced by noise has revealed a context-dependent effect of gap duration (Bashford et al., 1988; Bashford & Warren, 1987b). Bashford and Warren (1987b) asked listeners to adjust the duration of interruptions in speech sentences to the maximum duration at which they thought the speech sounded continuous. The maximum duration at which speech was perceived to continue through noise (304 ms) closely corresponded to the average word duration in the sentences presented. The maximum gap duration for the perception of continuity was higher for the spoken sentences than isolated word lists, confirming the role of contextual information in the perceptual restoration of fragmented speech.

The most interesting finding came from a second experiment (Bashford et al., 1988) where the maximum gap duration at which continuity was perceived when sentences were speeded up or slowed down by 15% was found to remain close to the average word duration in each case, rather than a fixed temporal value. This indicated that the unit size and its relationship to perceptual organisation are crucial to gap duration effects operating on the perceptual restoration of speech signals.

Powers and Wilcox (1977) were interested in the effect of gap duration on the intelligibility of fragmented speech. They tested for the perceptual restoration of everyday speech sentences with multiple alternating interruptions of between 130 ms and 650 ms. They found that for all gap durations, replacing the missing sections with noise resulted in increased intelligibility of the speech when compared to replacing the same sections with silence, thus demonstrating the operation of perceptual restoration in all cases. Interestingly, Powers and Wilcox compared the size of the noise-filled gaps versus silence-filled gaps advantage for each gap duration and found that noise had the greatest advantage over silence for a gap duration of 333 ms. In a similar way to the results of Bashford & Warren (1987b), this gap duration often resulted in an entire word being missing, and this meant that intelligibility in the silence condition was extremely poor.

In her investigation of the perceptual restoration of speech in children, Newman (2004) also included a gap duration manipulation. This was included as a way to avoid floor and ceiling effects in her child and adult participants who were tested using the same stimuli, but nevertheless the results are consistent with other findings regarding the effect of gap duration on the intelligibility of fragmented speech. The children and adults in Newman's study were tested on their repetition accuracy of everyday sentences where alternating gaps of either 100, 150, 200 or 250 ms were removed and replaced with noise or silence. Regardless of whether gaps were filled with noise or silence, Newman found that intelligibility was best for stimuli with 100 ms gaps, with intelligibility of sentences with 250 ms gaps being the worst. However, Newman only reports the effect of the duration manipulation pooled across noise-filled gaps and silence-filled gaps conditions, without reporting how gap duration affects intact stimuli, stimuli with noise-filled gaps, and stimuli with silence-filled gaps individually. Hence, there is no way of determining whether it is the noise-filled gaps or silence-filled gaps condition that is primarily driving this duration effect, or a combination of effects in both conditions.

These results highlight an important question for further research; whilst analysis of gap duration effects in terms of the noise-filled gaps over silence-filled gaps advantage in intelligibility is an important measure, this effect needs to be considered in terms of what is driving the effect. The inconsistent size of the difference between the two conditions with variation in the gap duration indicates that the effect of gap duration might be greater for one condition than the other. The findings of Powers and Wilcox (1977) appear to suggest that it is an increased detriment to intelligibility with silent gaps with a longer gap duration that leads to an increased noise-filled gaps over silence-filled gaps difference, but this has not been tested explicitly.

To summarise previous research, two main findings have emerged. First, the perception of continuity and the perceptual restoration of intelligibility of fragmented speech seem to be greatest when the duration of missing information is brief. Second, the size of the advantage noise-filled gaps have over silence-filled gaps in terms of intelligibility is also affected by the duration of the interruption, and might represent a separate effect of gap duration.

It is important to remember that where multiple alternating fragments are removed from a signal, although the duration of the gaps is altered, the ratio between what is removed and what is intact remains constant. Where alternating fragments of a given duration are removed from a signal, essentially half of the signal is removed, and half remains intact, regardless of the actual duration of the removed fragments. When gaps are, for example, 100 ms in duration, for every 100 ms gap removed, the next 100 ms of the signal are left intact, and so on. This means that the same amount of intact information affording identification of the signal is present regardless of the gap duration, but it is likely that perceptual organisation of the intact information might differ. It is also important to reiterate that the same amount of intact identifiable information is present in the signal regardless of whether the gaps are filled with noise or with silence.

There seem to be some limits to the effect of gap duration on perceptual restoration, at least in terms of the perception of continuity. Riecke et al. (2008) suggest that where the duration of gaps in auditory signals is less than 100 ms, a conceptual problem in terms of auditory grouping occurs. Because it is very difficult to perceive sections as missing when they are this brief, this can lead to false alarms in the perception of continuity. This therefore suggests that when testing for gap duration effects on perceptual restoration, 100 ms should be the shortest gap duration tested. There also appears to be an upper limit at which noise can restore the continuity of a missing section; when testing the perceived continuity of tones, Warren (1999) reports how noise cannot restore the continuity of signals with gaps of 300 ms or longer. Other research suggests that the effect decreases in strength for gap durations beyond this limit (e.g. Riecke et al., 2008).

The perceptual restoration of music, and in particular the effect of gap duration on musical perceptual restoration, has not received as much experimental attention as similar manipulations on the perceptual restoration of speech. Sasaki (1980) reported how the duration of noise replacing a single excised note in a familiar piece of music influenced listeners' ability to localise the noise, with localisation being easier with a longer duration of replacement. Since this implies that perceptual restoration was weaker with longer gap durations, it appears that the same gap duration effect might operate on the perceptual restoration of music. However, this requires more explicit testing with multiple interruptions and a measure that takes into consideration intelligibility with noise-filled gaps when compared to gaps filled with silence.

The possible adaptive function of gap duration effects on perceptual restoration

It is possible that the gap duration effects described above operate in order to maintain the adaptive nature of perceptual restoration. In order for perceptual restoration to be adaptive, it needs to be flexible and only operate in the kind of acoustic environments where filling in missing information is the appropriate thing to do (Riecke et al., 2009). For this reason, the gap duration effect might operate to constrain the operation of perceptual restoration to the conditions under which it would satisfy this requirement. Short, transient interruptions to auditory signals are quite common, and often arise quite accidentally as a consequence of characteristically noisy everyday environments. On the contrary, longer accidental interruptions are relatively rare, and more importantly, long interruptions in speech discourse can signal important events, such as the end of a sentence or clause. Therefore, in many situations, long gaps are supposed to be present and thus filling in these gaps would not be adaptive. The reduced rate of perceptual restoration for longer gap durations may therefore arise because in many situations it is best to treat long interruptions as if they are supposed to be there (Remijn et al., 2007). Whether this also applies to musical perceptual restoration, and to perceptual restoration in children, are important considerations.

It is also possible that the gap duration effect arises as a result of the way that perceptual restoration operates. The role of noise in perceptual restoration is to hide the absence of missing information (Verschuure & Brocaar, 1983). It is quite plausible that it is easier for noise to hide the absence of a shorter missing fragment, in the same way as it is easier for the signal before the interruption to be perceived as coming from the same source as that continuing after the interruption if the interruption is transient (Bregman, 1990).

In the present experiment, there were two main aims. First, the present experiment was designed to test for the operation of a gap duration effect on the perceptual restoration of music in children. Second, as well as investigating overall intelligibility, the experiment aimed to investigate the effect of gap duration on the relative size of the advantage noise-filled gaps have over silence-filled gaps in terms of intelligibility, and to investigate what might be driving differences in this relationship.

Method

Participants

Eighty-eight children (38 girls) from a primary school in a predominantly white, middle-class area of Surrey, UK, participated in this experiment. None of these children had participated in Experiment 1. The children were taken from the Reception and Year One classes of the school, and were between the ages of 4 years 1 month and 6 years 8 months, with a mean age of 5 years 6 months (66.50 months; *SD* = 8.35). Information about children's general sensory development was obtained from the school and children did not participate in the experiment if they had any significant hearing problems. All children were native English speakers, and consent to participate was obtained from both the child's parent and the school Headteacher. Further details about the participants in each condition are shown in Table 4.1.

Table 4.1. The gender split and mean age of the children in each of the four experimental conditions.

Type of Gap Replacement	**Gap Duration**	**N**	**Gender**	**Mean age in months (*SD*)**
Noise-filled gaps	200 ms	22	M=11; F=11	59.05 (7.29)
	300 ms	22	M=14; F=8	70.91 (6.91)
Silence-filled gaps	200 ms	22	M=9; F=13	59.45 (6.72)
	300 ms	22	M=16; F=6	71.29 (7.51)

Design

Participants were randomly allocated to one of four independent conditions. Participants in each condition heard the same fifteen nursery rhyme melodies presented in Experiment 1 (see Table 3.2), with the melodies being modified in different ways. Participants in the noise 200 ms condition were presented with nursery rhyme melodies in which alternating 200 ms fragments were removed and replaced with white noise. The noise 300 ms condition was identical to the noise 200 ms condition except that alternating 300 ms fragments of each stimulus were removed and replaced with white noise. In the silence 200 ms condition, participants heard the melodies with 200 ms fragments removed and replaced with a silent gap; in the silence 300 ms condition alternating 300 ms fragments were removed from the melodies and replaced with silence. Thus, two variables were manipulated: the duration of the missing gaps (200 ms or 300 ms); and the type of gap replacement (noise-filled gaps or silence-filled gaps).

Materials

The stimuli were produced in exactly the same way as for Experiment 1, and the peak amplitude of the signals and the amplitude of the replacement noise were set to the same levels as in Experiment 1. The school participating in the present experiment also completed the same song familiarity questionnaire that was used in Experiment 1 as a way to check the familiarity of the proposed stimuli (Appendix 1). As for Experiment 1, a programme written in Matlab V5.2.1 presented the stimuli to the children in a random order and recorded their response and identification time.

Procedure

The procedure was identical to that of Experiment 1.

Results

The analysis here includes the data from the noise-filled gaps and silence-filled gaps conditions from Experiment 1 (with 100 ms gaps), as in all respects except for the duration of the missing fragments, the stimuli were identical. This allows comparisons to be made between identification of stimuli with three different gap durations (100 ms, 200 ms, & 300 ms).

Correct identification scores

Each child was given a score out of 13 corresponding to the number of melodies that they identified correctly in the experimental trials. Data were examined for gender differences (see Appendix 2), and since no significant differences were found, data were collapsed across gender. The mean number of melodies identified correctly by children in each condition is shown in Figure 4.1. The mean identification score from the intact condition of Experiment 1 is shown for comparison.

The data were analysed using a 2 (gap replacement condition: noise-filled gaps; silence-filled gaps) × 3 (gap duration: 100 ms; 200 ms; 300 ms) between-subjects ANOVA[10]. The main effect of gap replacement condition was significant, $F(1, 126) = 179.12, p < .001, \eta^2 = .59$. Figure 4.1 shows that for all gap durations, identification scores were superior when missing sections were filled with noise when compared to the same sections filled with silence. The main effect of gap duration was also significant, $F(2, 126) = 18.27, p < .001, \eta^2 = .23$. Post hoc tests (Games-Howell) revealed that identification of stimuli with 100 ms gaps was significantly superior to stimuli with 200 ms gaps ($p < .05$, $d = 0.50$) and to stimuli with 300 ms gaps ($p < .05$, $d = 0.73$); there was no significant difference between 200 ms and 300 ms gap conditions ($p > .05$). The interaction between gap replacement condition and gap duration was non-significant, $F(2, 126) = 2.19, p > .05$.

---

[10] This analysis was also run with age included as a covariate due to the lower age of participants in the 200 ms conditions compared to the other conditions. Neither main effects nor the interaction were affected by this. Examination of the adjusted means revealed that this age difference might inflate the identification scores of the 300 ms conditions relative to the others.

In summary, stimuli with short interruptions (100 ms) were easier to identify than those with longer interruptions, and replacing missing fragments of all durations with noise resulted in superior melody identification to replacing the same fragments with silence.

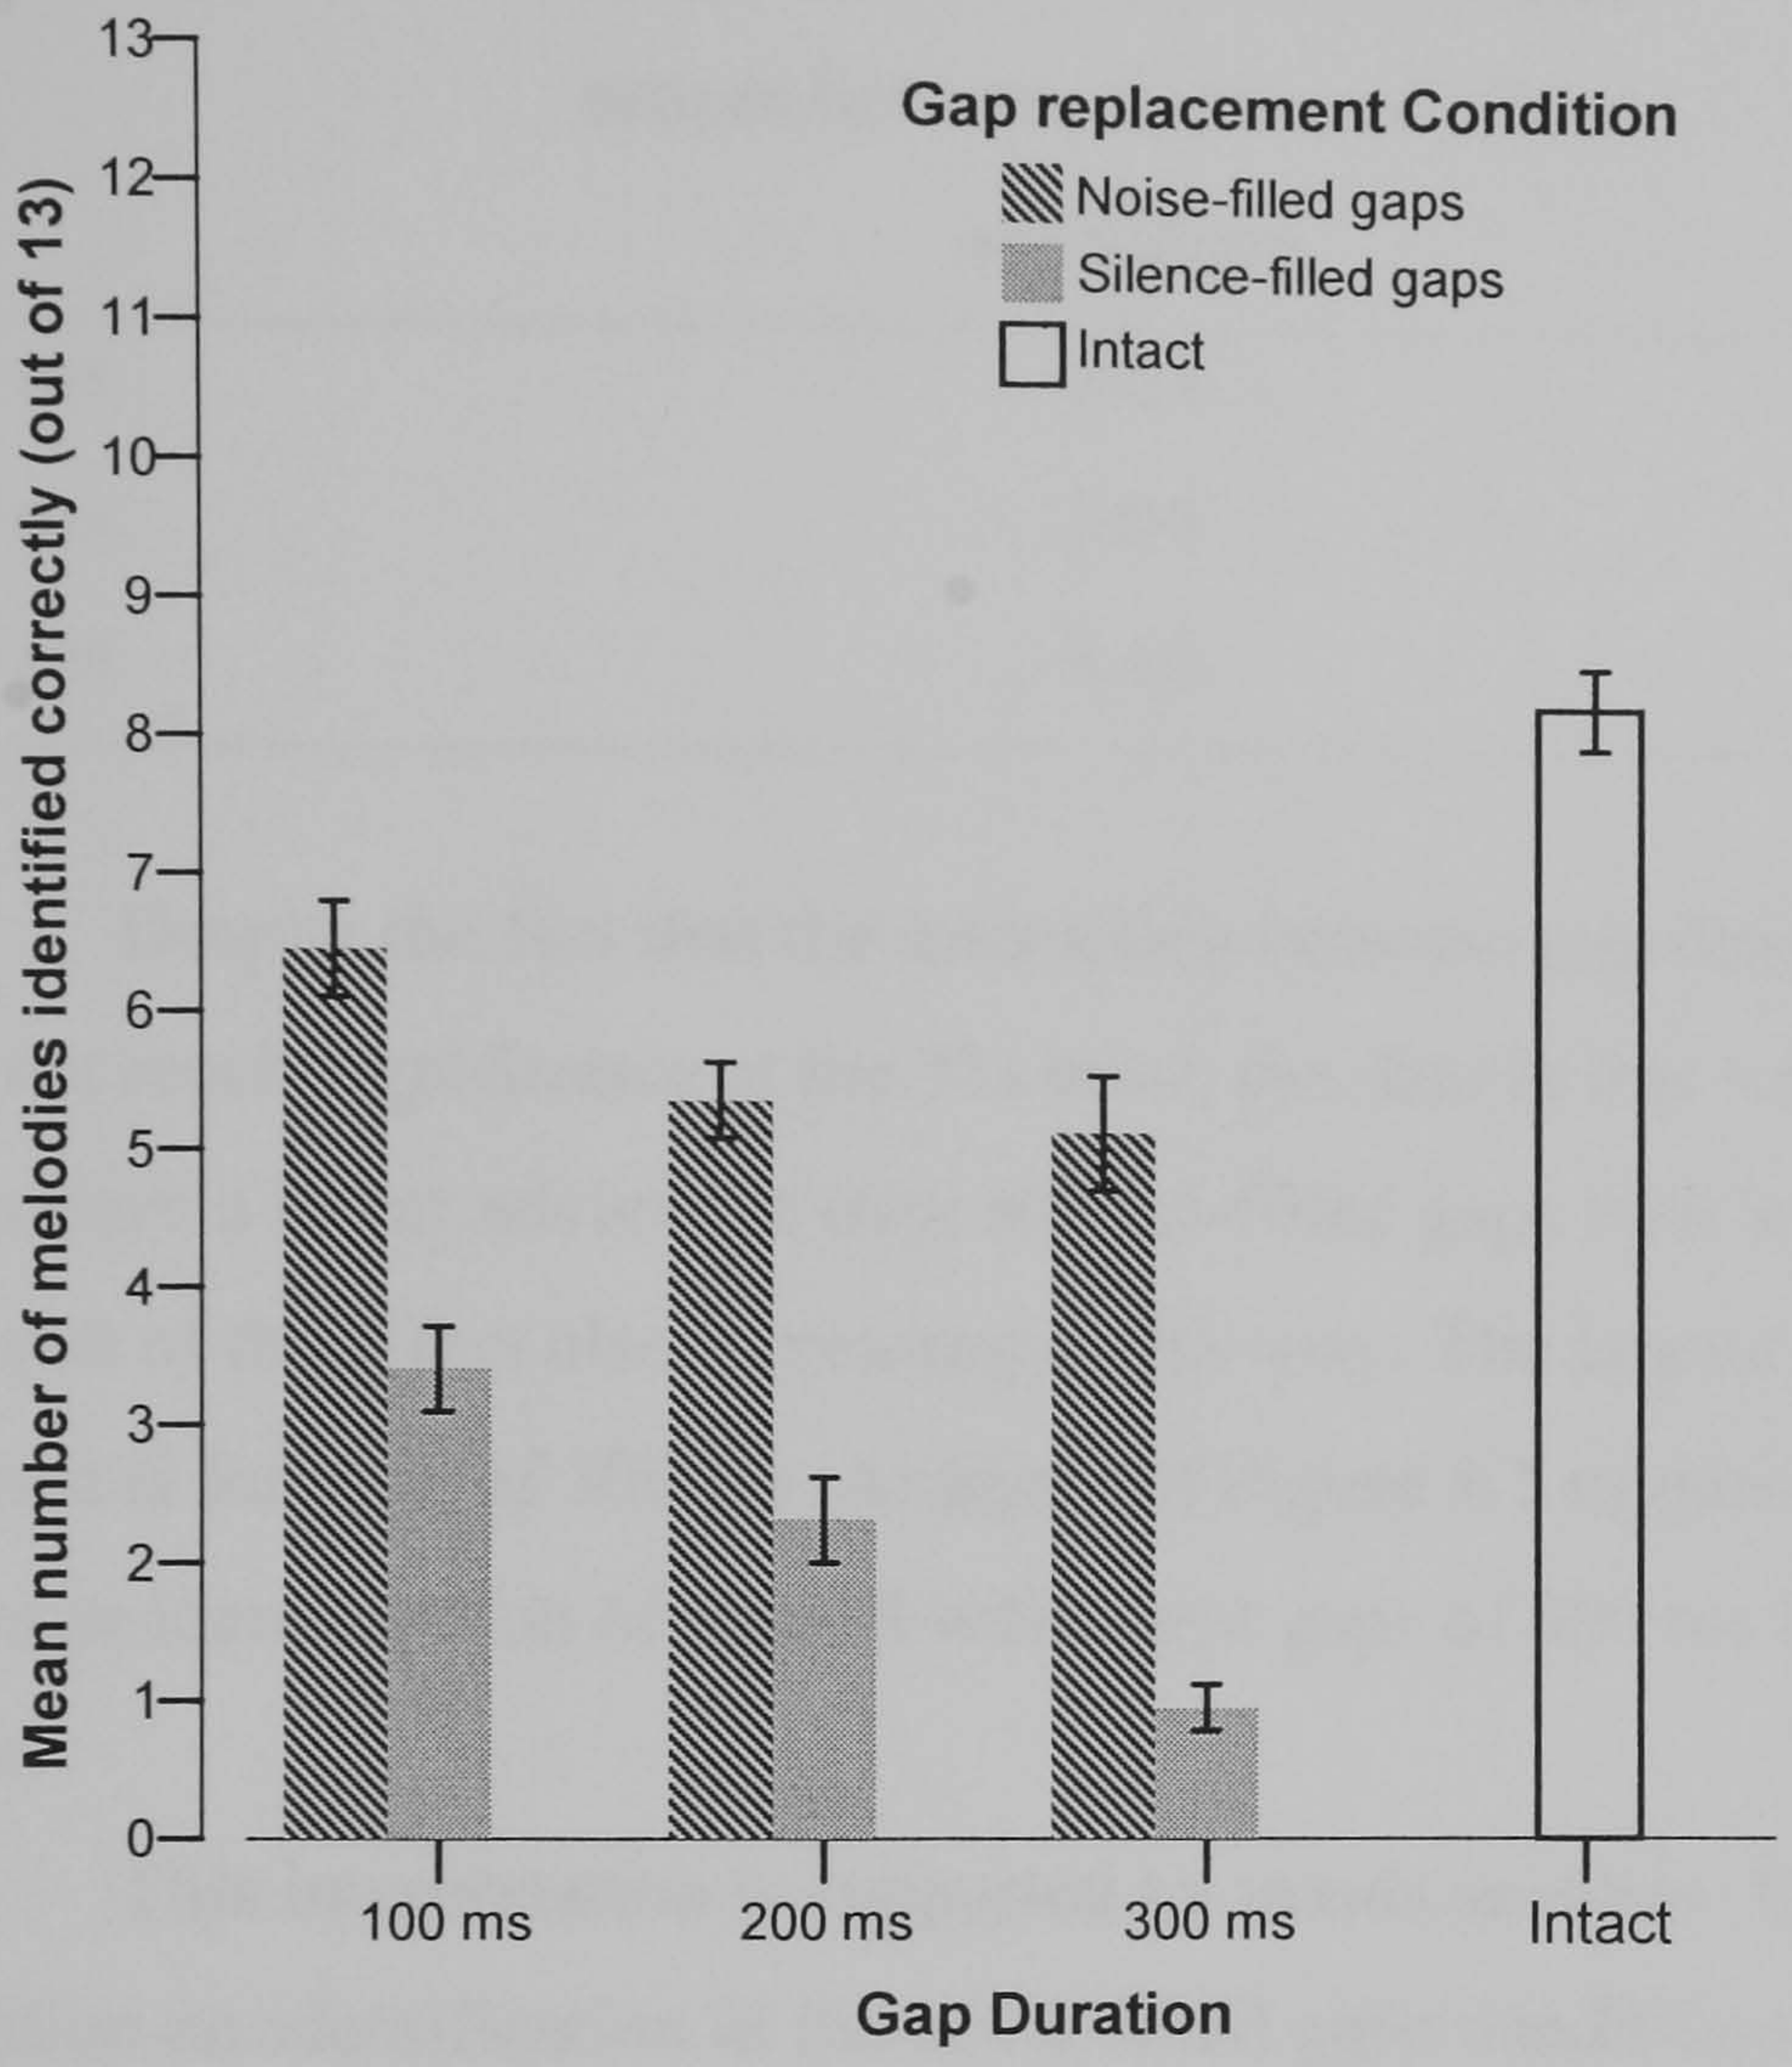

Figure 4.1. Mean number of melodies identified correctly (± *SE*) by children in each of the six experimental conditions (noise- and silence-filled gaps with 100, 200 & 300 ms gap durations; the intact condition from Experiment 1 is shown for comparison).

Analysis of the effect of gap duration on the size of the noise over silence advantage

One of the aims of the present experiment was to examine further the effect of gap duration on perceptual restoration by investigating the possible reasons for the emergence of the effect. As well as considering the effect of gap duration on identification scores, the effect on the size of the advantage noise-filled gaps have over silence-filled gaps is also considered. The difference in identification scores for noise-filled gaps and silence-filled gaps conditions for each gap duration is shown in Table 4.2, together with the effect size of the difference.

Table 4.2. The difference in means between noise-filled gaps and silence-filled gaps conditions at each gap duration.

Gap duration	**Difference in mean identification scores between noise and silence conditions**	**Effect size of the difference**
100 ms	3.04	$d = 1.87$
200 ms	3.04	$d = 2.08$
300 ms	4.19	$d = 2.14$

Despite the fact that the interaction between gap duration and gap replacement did not reach significance at the 5% level, the data in this table suggest that noise-filled gaps have a larger advantage over silence-filled gaps with increasing gap duration, with the size of the effect also increasing in this way. The largest advantage of noise over silence is for gaps of 300 ms. Analysis of Figure 4.1 suggests that this is primarily because identification of stimuli with silent gaps of 300 ms in duration is especially poor.

This interpretation is supported by trends analysis. Whilst the effect of gap duration on identification in the noise-filled gaps condition does show a significant linear trend, $F(1, 63) = 7.04, p < .05, \eta^2 = .10$, the significance of the linear trend was much greater for the silence-filled gaps condition, $F(1, 63) = 41.32, p < .001, \eta^2 = .40$. This suggests that gap duration has a more consistent effect on identification of stimuli with silence-filled gaps than those with noise-filled gaps.

Identification time

The average identification time for correctly identified melodies was calculated for each participant in the same way as for Experiment 1. The data from one child in the silence 200 ms condition and from seven children from the silence 300 ms condition were excluded from this analysis as they did not identify any melodies correctly. There was no significant difference in identification times as a result of gender (see Appendix 2), so data were collapsed across gender. The mean identification time in seconds for children in each condition is shown in Figure 4.2. Data from the intact condition of Experiment 1 are shown for comparison.

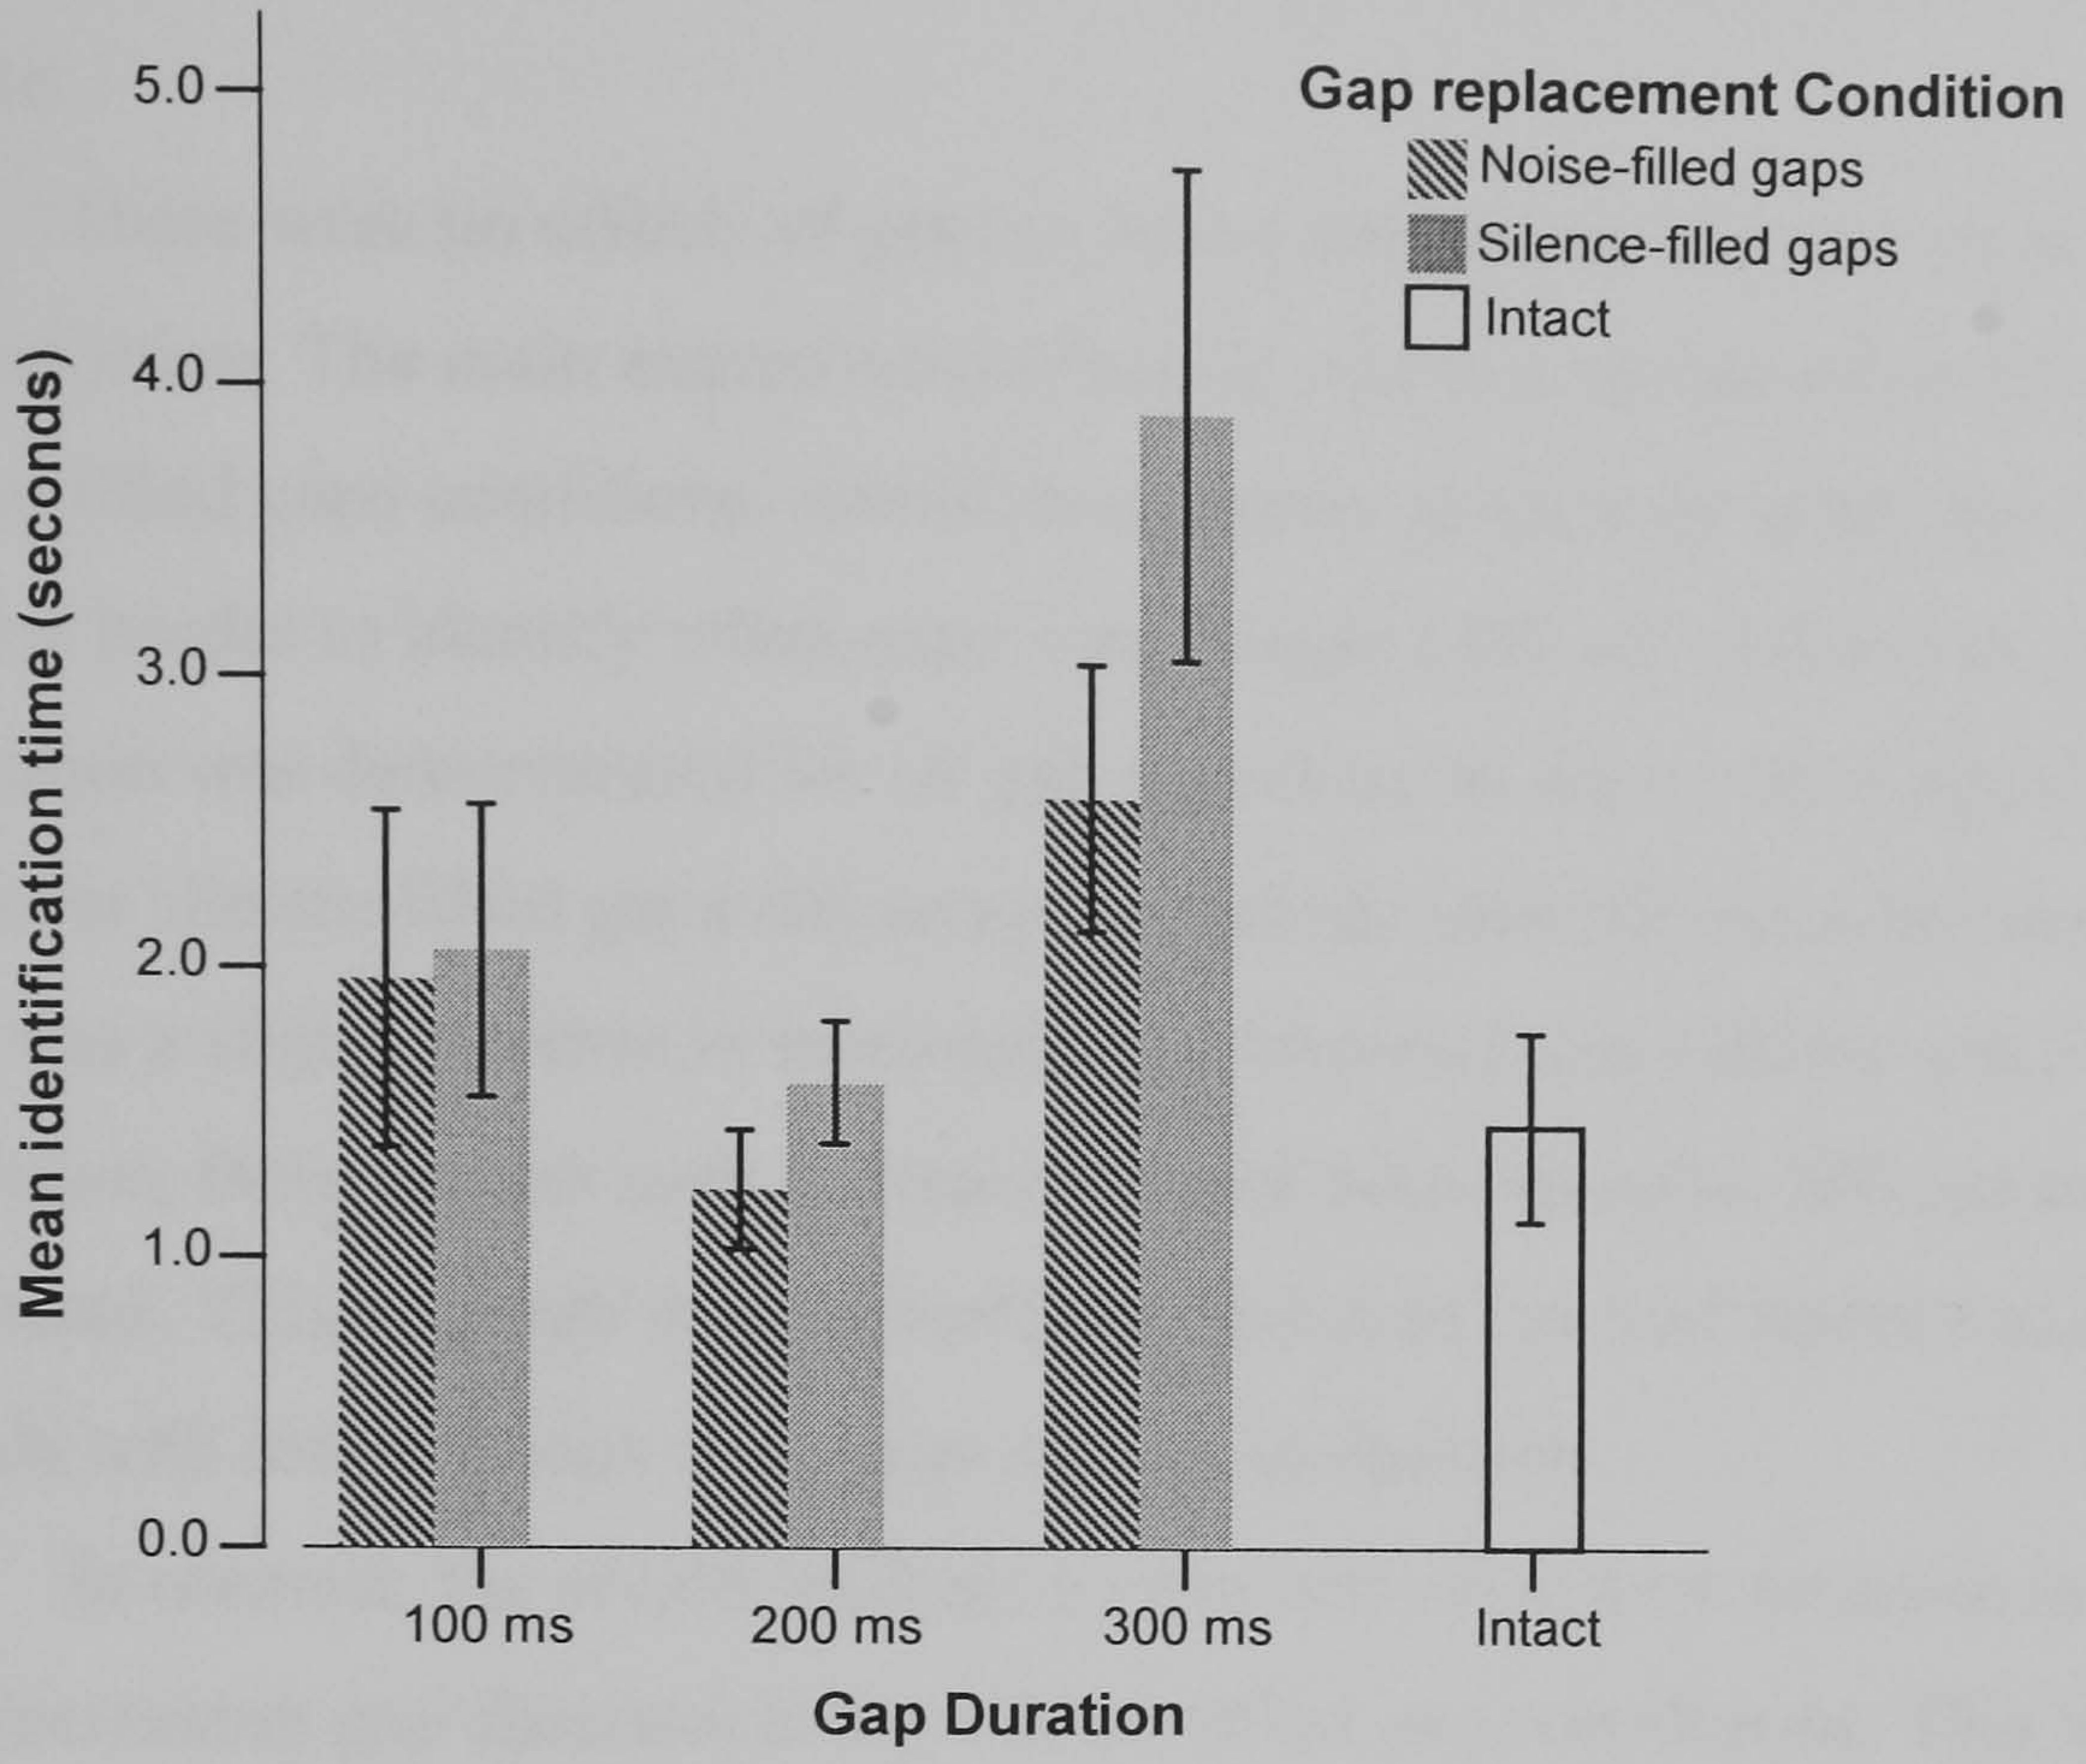

Figure 4.2. Mean identification time (identification time − stimulus duration; ± *SE*) in seconds for children in each of the six experimental conditions (noise- and silence-filled gaps with 100, 200 & 300 ms gap durations; the intact condition from Experiment 1 is shown for comparison).

The data were analysed using a 2 (gap replacement condition: noise-filled gaps; silence-filled gaps) × 3 (gap duration: 100 ms; 200 ms; 300 ms) between-subjects ANOVA. The main effect of gap replacement condition was non-significant, $F(1, 131) = .92, p > .05$, and there was also a non-significant main effect of gap duration on identification time, $F(2, 131) = 2.08, p > .05$. The interaction between gap replacement condition and gap duration was also non-significant, $F(2, 131) = .20, p > .05$.

Discussion

The main aim of the present experiment was to test for a comparable gap duration effect on the perceptual restoration of music in children to that reported in the literature for the perceptual restoration of speech, and to consider why this effect might operate.

There were no effects of gender on the number of stimuli correctly identified in any condition. The main experimental finding was that across noise-filled gaps and silence-filled gaps conditions, stimuli were easier to identify when gaps were brief (100 ms) and harder to identify when gaps were longer (300 ms). However, perceptual restoration was demonstrated for all gap durations, as shown by a significant noise-filled gaps over silence-filled gaps advantage in melody identification. Interestingly, although there was a slight decrease in intelligibility between noise 100 ms and noise 200 ms conditions, there was no such decrease evident between noise 200 ms and noise 300 ms conditions. This suggests that perceptual restoration can still operate in the music domain with interruptions as long as 300 ms in duration.

In contrast, the results indicate a clear and consistent decrease in intelligibility with increasing gap duration in the silence-filled gaps conditions. This was supported by a highly significant linear trend. It therefore becomes apparent that the increased noise-filled gaps over silence-filled gaps advantage for stimuli with longer gaps in the present study, and plausibly also in previous research (Newman, 2004; Powers & Wilcox, 1977), might be because gap duration has a greater effect on stimuli with silence-filled gaps than stimuli with noise-filled gaps.

This is likely to be the case because when short silent interruptions occur, the natural rhythm of the signal (speech or music) is not detrimentally affected to a large degree. However, with longer interruptions, the rhythmic continuity is more seriously disrupted. Replacing missing sections with noise restores this continuity; when interruptions are long, stimuli with silence-filled gaps are more seriously affected. This serves as an illustration of the adaptive nature of perceptual restoration; it shows its strongest advantage by continuing to operate under conditions where perception of stimuli with silent interruptions is especially poor.

Analysis of the identification time data revealed no significant difference between noise-filled gaps and silence-filled gaps conditions, but did suggest that the duration of the missing gaps has an effect on the time taken to identify the stimuli, even though this difference was not statistically significant. The data suggest that stimuli with 300 ms gaps took longer to identify than stimuli with 200 ms gaps and 100 ms gaps, thus suggesting that identification is slower when longer disruptions are present within a melody.

The present results have replicated findings in the speech domain. Newman (2004) reported that young children demonstrated superior repetition accuracy of spoken sentences with shorter gaps removed when compared to longer gaps. In addition, close parallels are evident between the findings from the present experiment and those of Powers and Wilcox (1977). In the experiment reported by Powers and Wilcox, evidence of the perceptual restoration of speech was obtained for all the gap durations they tested, as shown by a significant noise-filled gaps over silence-filled gaps advantage in intelligibility. In addition, the largest noise-filled gaps over silence-filled gaps advantage was demonstrated for a gap duration of approximately 333 ms, which can be compared to the largest advantage for gaps of 300 ms in the present experiment. Powers and Wilcox explained this effect in terms of the unit size; a gap duration of 333 ms closely approximated the average word duration in their speech stimuli, meaning that with a gap of this duration, whole words were often missing.

The average note duration of the melodic stimuli in the present experiment was 379 ms. Therefore, it is possible that for some of the melodies a similar explanation applies; whole notes might have been missing from the melodies when 300 ms gaps were removed. The removal of a whole note or of transitional cues between notes from a melody disrupts the melodic contour (the global pattern of pitch changes in a melody) to a greater extent than part of a note, and melodic contour is an important cue to melody identification from infancy (Trehub, Bull & Thorpe, 1984) through to adulthood (Dowling, 1978). When a whole note is missing and replaced by noise, in a highly familiar context such as this, perceptual restoration can still restore the missing note. The absence of whole notes and transitional cues is more noticeable when gaps are silent, thus offering a possible explanation for the increased duration effect with silent gaps when compared to noise-filled gaps in the context of music perception.

It is not simply the case that with longer gap durations, more of the signal is removed than for a shorter gap duration, because when making multiple alternating interruptions to a signal, half of the signal is present and half is removed regardless of the gap duration. In all noise-filled gaps and silence-filled gaps conditions tested, the same amount of intact information affording identification of the melody was present. The crucial difference, therefore, must be that the different versions of the melodies present different acoustic conditions; the strength of the perceptual restoration mechanism is affected by these conditions. For example, when 100 ms fragments are replaced by noise, it is easier to infer the continuation of the signal than when longer fragments are replaced by noise, and it is easier for noise to hide the absence of the missing information. Thus, noise might provide stronger acoustic evidence for, and hence stronger bottom-up corroboration of, the expected content of brief missing fragments.

When 300 ms fragments were removed, whilst the duration of intact sections was longer than when 100 or 200 ms fragments were removed, these intact sections occurred less frequently, because a gap of 300 ms was present before the next intact section was heard. Thus, it appears that if interruptions are brief, even if intact sections are also brief, hearing intact input regularly seems to improve melody identification. This may relate to the generation and confirmation of expectations that is so critical to perceptual restoration; confirmation of expectations occurs more regularly for brief interruptions.

In the speech domain, the upper durational limit at which continuity of a fragmented signal can be perceived has been suggested to be 300 ms. This could possibly function in an adaptive way to prevent meaningful gaps in speech being filled in when in fact they are supposed to be there. However, the present experiment demonstrated that perceptual restoration can still operate for a gap duration of 300 ms for musical signals, as shown by the advantage noise-filled gaps have over silence-filled gaps at this duration. There are several possible reasons for this discrepant finding. First, 300 ms is cited as the upper durational limit for the perception of continuity of speech signals, not the restoration of intelligibility. It is possible that in music, the perception that the stimulus is continuous could decrease significantly with increasing gap duration, whilst intelligibility, especially for a highly familiar stimulus, would be affected to a lesser degree.

Second, it might be easier to infer the continuation of a musical signal than a speech signal through a longer gap replaced with noise. Third, as previously explained, the average note length of the melodies presented in this experiment seems to be longer than the average length of a word in a spoken sentence. Thus, the cut-off point at which continuity can be perceived might represent a longer gap duration for music than in the case of speech. Without any measure of whether the children actually perceived the stimuli to be continuous or fragmented, it is difficult to distinguish between these possibilities. It will therefore be important in future experiments to include such a subjective report measure in order to make a more explicit distinction between the effect of manipulations on the perception of continuity and the perceptual restoration of intelligibility.

The aim of this experiment, in the context of the series of experiments presented here, was to determine whether acoustic factors have an influence on the perceptual restoration of music in children. The results indicate that they do; to summarise, perceptual restoration operates for gap durations of 100 ms, 200 ms and 300 ms, and intelligibility of melodies is increased when shorter rather than longer fragments are missing. The effect of gap duration on overall intelligibility and on the size of the noise-filled gaps over silence-filled gaps advantage in intelligibility shows clear parallels to the effect of similar manipulations on the perceptual restoration of speech (e.g. Powers & Wilcox, 1977).

This result provides evidence that the acoustic properties of the signal affect the perceptual restoration of music in children. However, the duration of missing fragments is not the only acoustic manipulation that has been tested on the perceptual restoration of speech. Much research has been conducted on the nature of the noise that replaces missing sections, and what criteria it needs to fulfil in order to allow perceptual restoration of the missing information. The most important finding has been that in order for perceptual restoration to occur, the amplitude of the replacement noise must be larger than the amplitude of the signal it is replacing, so that it can act as a potential masker of the signal and thus hide the absence of the missing information (Kashino, 2006). Only then can the noise provide bottom-up confirmation of expectations as to the content of the missing fragments. Testing for the operation of this constraint on the perceptual restoration of music in children forms the focus of the next experiment.

Chapter 5: Experiment Three

The effect of the amplitude of replacement noise on the perceptual restoration of music in children

Introduction

Previous research on the perceptual restoration of missing auditory input has shown that there are certain acoustic conditions under which perceptual restoration occurs. The role of replacement noise in perceptual restoration is crucial in providing bottom-up corroboration of expectations as to the most likely content of missing information. However, the presence of noise in missing sections of an auditory signal is only able to restore missing content if the noise is able to act as a potential masker of the signal (the "masking potential" rule; Kashino, 2006), and thus exceeds the peak amplitude of the signal it is replacing.

Masking and the perception of continuity

It has been proposed that there are two separate processes at work in perceptual restoration (Shahin et al., 2009). First, the presence of noise that hides the fact that input is missing produces illusory continuity of what is in reality a fragmented signal. Second, the listener's previous knowledge is recruited, using the surrounding context as a guide, to restore the missing content of the signal and hence intelligibility. Bashford et al. (1992) made a similar distinction between two separate processes operating to restore missing fragments of an auditory signal, where restored intelligibility is dependent on the perception of illusory continuity. It therefore appears that the masking potential of the noise might be a crucial part of the process; only if the amplitude of the replacement noise exceeds the peak amplitude of the signal can it act as a potential masker and restore the continuity of the signal. Consequently, the amplitude of replacement noise might affect the perceptual restoration of both continuity and intelligibility.

Verschuure & Brocaar (1983) describe perceptual restoration as the result of an "automatic interpolation mechanism", which fills in sections of an auditory stimulus which are absent in reality. However, this mechanism is only able to operate if it is not possible to tell that any sections are actually missing, that is, only if illusory continuity is perceived. This provides further support for the proposition that perceptual restoration involves two processes, and that intelligibility is best restored where illusory continuity is perceived. The importance of the ability of the replacement noise to act as a potential masker in order for illusory continuity to be perceived is also stressed by other researchers (e.g. Remijn et al., 2007; Shriberg, 1992). Only where noise can act as a masker for missing information can it provide perceptual closure of a missing fragment (Bregman, 1990), and mask the transitions between the offset of the signal and the onset of the gap (Nelken, 2004; Petkov et al., 2007). Noise that is a potential masker of the missing information also provides continued neural activity during the missing section; neurons firing during the presentation of the signal must also respond to the masker if the signal is to be perceived as continuous (Petkov et al., 2007). It is important to note that whilst noise that exceeds the peak amplitude of the signal it is replacing allows it to function as a potential masker of the signal, the term "potential" is critical to its interpretation since noise intensity is not the only factor that influences auditory masking.

The effect of replacement noise amplitude on perceptual restoration

The amplitude of the replacement noise in relation to the amplitude of the signal has been found to affect simple perceptions of continuity of interrupted signals such as tones (e.g. Riecke et al., 2007; Riecke et al., 2008; Thurlow & Elfner, 1959). Continuity is not perceived where the amplitude of the noise is lower than that of the signal (Riecke et al., 2008).

Further evidence that the amplitude of the replacement noise affects the perception of continuity of fragmented speech was obtained by Eimas et al. (1996). Listeners were asked to judge whether a word with a missing phoneme replaced by noise was complete, and so this response acted as a measure of whether the missing phoneme was perceptually restored. Judgements that the word was complete, indicative of perceptual restoration of the missing phoneme, were made on 93% of trials when noise greater in amplitude than the signal was replacing the missing phoneme. This percentage fell to just 17.4% when a low amplitude noise, lower in amplitude than the signal, was the gap replacement. Low amplitude noise was only marginally better than silence ("complete" judgement made on 5% of trials) in giving the illusion of continuity. These results clearly illustrate that the perception of continuity of a fragmented signal cannot be produced by inserting noise into the missing fragments unless the noise is able to act as a potential masker of the missing information.

There is evidence in the speech domain that the ability of replacement noise to restore intelligibility of a fragmented signal is dependent on the amplitude of the noise in relation to the amplitude of the signal. Warren and Obusek (1971) removed a single phoneme from a sentence, and asked listeners to locate the missing fragment and indicate whether the extraneous sound replaced the missing fragment. They used a variety of sounds as the replacement noise, including a cough, a buzz, or a tone. For two of the sounds (buzz or tone), Warren and Obusek varied the amplitude of the replacement sound, with the sound either being of the same amplitude as the signal, or 8 dB above the amplitude of the signal.

For both amplitudes, listeners were more likely to report that no sounds were replaced (thus indicating that restored phonemes were indistinguishable from those actually present) when the missing phoneme was replaced by noise than when replaced by silence, indicating that perceptual restoration was operating. In addition, the perceptual restoration produced by the louder sounds was slightly stronger than that produced by the softer sounds, but not to a significant degree, and was not consistent across a series of trials. However, as the quietest sound here was not below the peak amplitude of the signal, both levels of noise could act as potential maskers and hence lead to illusory continuity, even if the quieter sound was only a partial masker. Thus, this study did not make a comparison between noise above and below the peak amplitude of the signal.

Powers and Wilcox (1977) tested for perceptual restoration of speech signals with a wider variety of replacement noise levels. Whilst overall intelligibility of the signals increased as the intensity of the replacement noise increased, this increase was not consistently linear in nature. The strength of perceptual restoration as a function of replacement noise amplitude might be the result of the interaction of many acoustic parameters, for example the duration of the missing fragments and the nature of the noise in relation to the nature of the signal it is replacing.

To summarise previous research on replacement noise amplitude effects, it is evident that the amplitude of the replacement noise affects the degree to which fragmented auditory signals are perceived to continue through noise, and the degree to which noise restores intelligibility of missing speech. The effect of noise amplitude on the strength of perceptual restoration might represent an important constraint on its operation. The adaptive function of this constraint is to prevent perceptual restoration of gaps in perception in inappropriate circumstances. For example, if a missing sound could not have been masked by a louder sound, then the assumption that the missing sound continued behind the louder sound could in many circumstances be incorrect (Kashino, 2003).

In exploring the effect of the amplitude of the replacement noise on perceptual restoration, investigating the two proposed processes in perceptual restoration (the restoration of continuity and intelligibility, e.g. Shahin et al., 2009) is particularly important. Analysis of the data from Experiment 2 highlighted that having a measure of whether children believe the melodies to be fragmented or continuous might provide potentially important data in addition to a measure of the intelligibility of the melodies. Further, in examining cognitive processes, the importance of combining objective experimental data with subjective reports of cognitive experience is highlighted by Kingstone, Smilek and Eastwood (2008). It is argued that by measuring both aspects of the phenomenon under study, the consistency of experimental data with subjective reports can be examined. For these reasons, the design of the present experiment incorporates children's subjective reports of the continuity of the melodies as an attempt to measure their perceptual experience of the stimuli.

It is evident that in the perceptual restoration of speech, the amplitude of the replacement noise is crucial in determining whether perceptual restoration can occur. However, no such manipulation has been carried out on the perceptual restoration of music. When parts of a musical signal are missing, can noise of any amplitude restore intelligibility, or does the noise amplitude have to exceed the peak amplitude of the signal? The masking potential rule may not necessarily apply to all cases of perceptual restoration; it might somehow be specific to speech signals and simple tones. A further question concerns child listeners, as noise amplitude effects have never been tested on children's perceptual restoration. It may not necessarily follow that children will behave in a similar way to adults when presented with replacement noise of different intensities. Hearing white noise is likely to be an unusual experience for children. Furthermore, if replacement noise is too loud, it may distract children from the content of the intact sections of the signal, and so lead to *reduced* intelligibility.

The aim of the present experiment was to test for a similar effect of the amplitude of replacement noise on the perceptual restoration of music in children; in particular, whether noise only provides an intelligibility advantage over silence if the noise is greater in amplitude than the signal it is replacing. This was achieved by manipulating two variables: 1) the amplitude of the replacement noise, and 2) the duration of the gap the noise was replacing. Gap duration was included as a manipulation in this experiment in order to test whether the effect of replacement noise amplitude on musical perceptual restoration differs as a result of other acoustic parameters.

Method

Participants

One hundred and twenty-six children (61 girls) from two primary schools with predominantly white, middle-class catchment areas in Surrey and Dorset, UK, participated in this study. None of the children had participated in previous experiments. Children between the ages of 4 years 2 months and 6 years 10 months were tested, taken from the Reception and Year One classes of the school. The sample had a mean age of 5 years 7 months (67.82 months; *SD* = 8.56). All children were native English speakers, and consent for their participation was obtained from both the child's parent and the school Headteacher. The schools were asked about children's general sensory development and children did not participate in the experiment if they had any significant hearing problems. Further details about the age and gender of the participants in each condition are shown in Table 5.1.

Table 5.1. The gender split and mean age of the children in each of the six experimental conditions.

Noise Amplitude	**Gap Duration**	N	**Gender**	**Mean age in months (*SD*)**
High Amplitude Noise	100 ms	21	M=13; F=8	71.29 (7.14)
	300 ms	21	M=12; F=9	70.43 (8.64)
Low Amplitude Noise	100 ms	21	M=10; F=11	62.67 (8.61)
	300 ms	21	M=10; F=11	61.90 (7.01)
Very Low Amplitude Noise	100 ms	21	M=10; F=11	69.33 (9.17)
	300 ms	21	M=10; F=11	71.29 (7.51)

Design

Participants were randomly allocated to one of six independent conditions. Participants in each condition were presented with the same fifteen nursery rhyme melodies presented in previous experiments (see Table 3.2), but they were modified according to the experimental condition. Conditions varied according to the type of noise that was used to replace missing fragments in the melodies. In "high amplitude noise" conditions, the amplitude of the replacement noise exceeded the peak amplitude of the signal (the noise amplitude was 10 dB and the peak amplitude 20 dB below the maximum bandwidth of the file). In "low amplitude noise" conditions, the amplitude of the replacement noise was below that of the peak amplitude of the signal, (30 dB and 20 dB below the maximum bandwidth of the file, respectively).

Finally, the "very low amplitude noise" conditions were included as a comparison to the low amplitude noise conditions. The amplitude of the replacement noise in both cases was lower than that of the signal, but the amplitude of the replacement noise in the low amplitude noise conditions was higher than that of the very low amplitude noise conditions. Thus, when replacement noise is lower in amplitude than the signal it replaces, does the more intense replacement noise improve intelligibility of the signals, relative to less intense replacement noise, or are both equal in their effect on signal intelligibility? In the extra low amplitude noise conditions, the amplitude of the replacement noise was 45 dB below the maximum bandwidth of the file. This amplitude was selected based on pilot work; this was the lowest amplitude at which noise was audible. In this way, this condition is very similar to the low amplitude noise condition reported by Eimas et al. (1996). The aim of this experiment was not to look for linear effects of noise amplitude on perceptual restoration, since previous research had suggested that these effects do not occur. Instead, these conditions were included in order to test specific questions related to the effect of noise amplitude on perceptual restoration in child listeners.

In addition, and as a test of whether noise amplitude and gap duration interact in their effects on perceptual restoration, conditions also varied according to the duration of the missing fragments. The shortest and longest gap durations from Experiment 2 (100 ms and 300 ms) were used in the present experiment.

Materials

The stimuli were produced and presented in the same way as for Experiments 1 and 2. The schools participating in the present experiment were sent the same song familiarity questionnaire that was used in previous experiments (Appendix 1).

Procedure

The procedure was identical to that in Experiments 1 and 2, except for the addition of a measure of whether children in the different conditions thought the tunes were continuous or fragmented. At the end of the testing session each child was asked,

> "In the tunes we have just been listening to, do you think that little pieces of the tune were missing, or do you think that all of the tune was there?"

The order of the options within the question was counterbalanced to avoid a bias to respond with the first or second option within the question. The question was not asked separately for each trial to avoid children thinking that they should vary their answer from trial to trial; their most frequent experience would be that all questions do not have the same answer. The child's response was recorded by the experimenter.

Results

Correct identification scores

Each child was given a score out of 13 that represented the number of melodies they identified correctly in the experimental trials. Data were examined for gender differences and a significant difference emerged in the high amplitude noise 300 ms condition (see Appendix 2). For this reason, gender was included as a factor in further analyses to test for a possible interaction with other variables. The mean number of melodies identified correctly by children in each condition is shown in Figure 5.1.

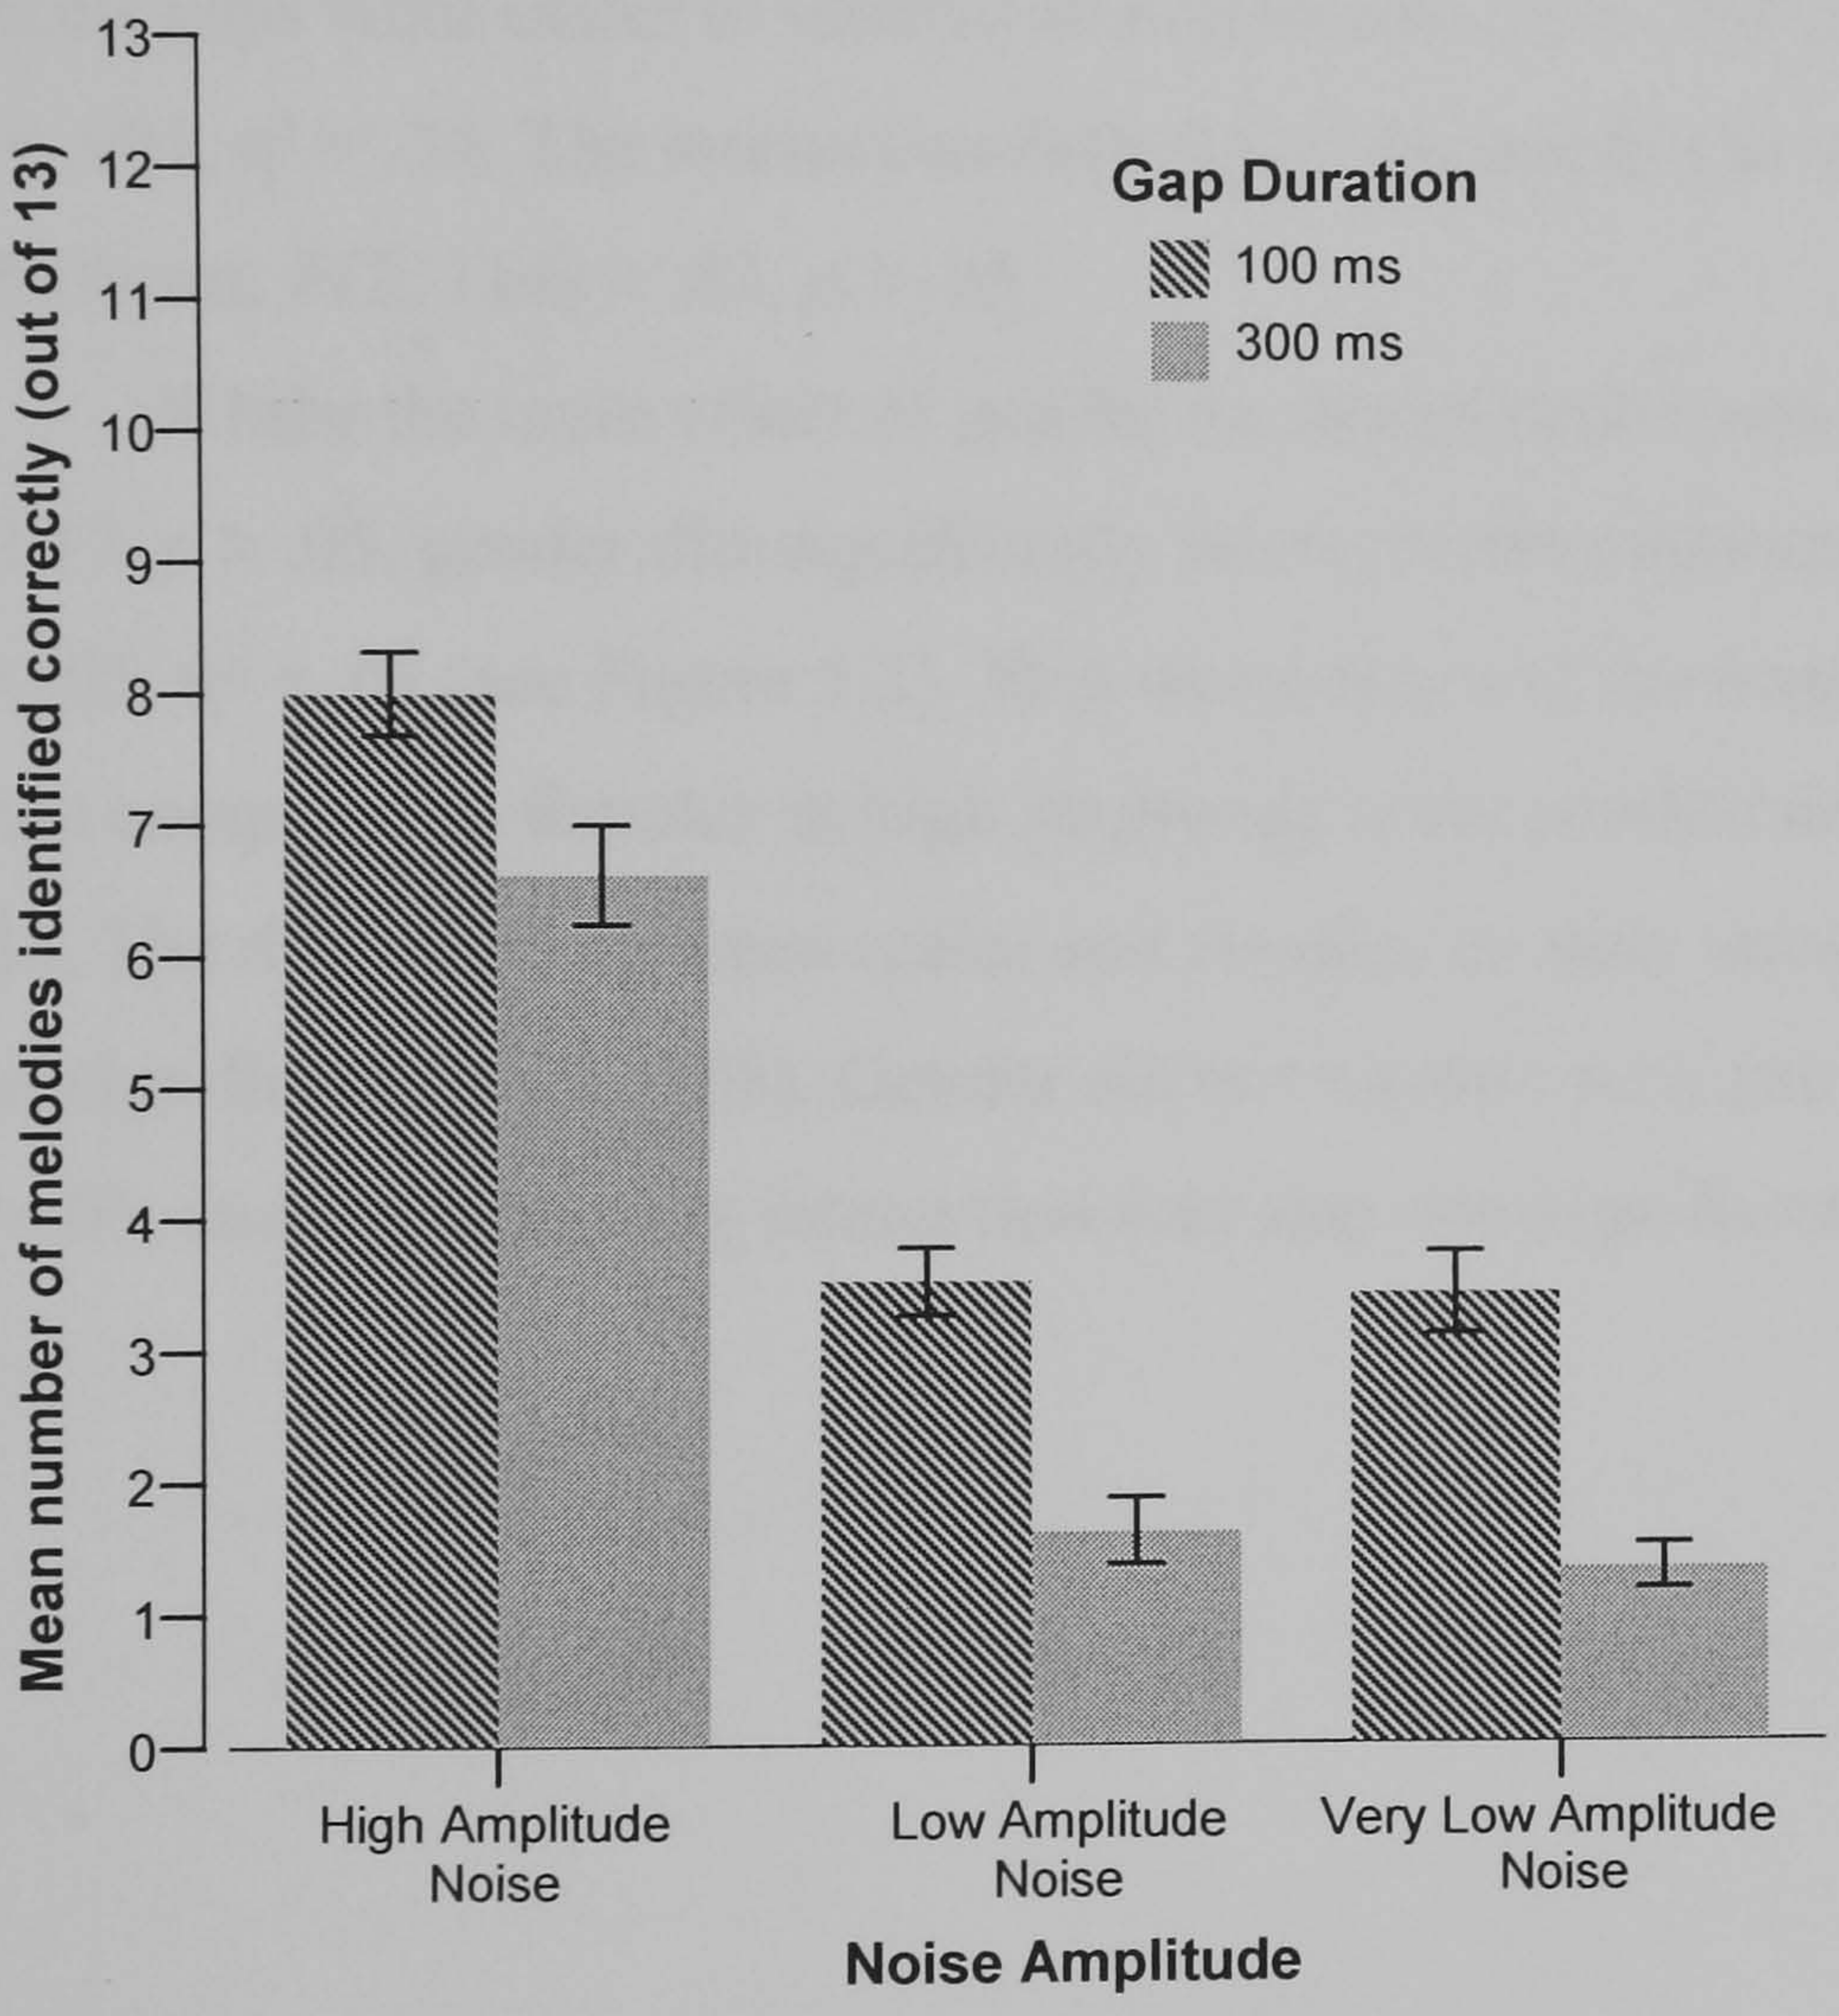

Figure 5.1. Mean number of melodies identified correctly (± *SE*) by children in each of the six experimental conditions (high amplitude noise, low amplitude noise, very low amplitude noise; 100 ms or 300 ms gaps).

These data were analysed using a 3 (noise amplitude: high amplitude noise; low amplitude noise; very low amplitude noise) × 2 (gap duration: 100 ms; 300 ms) × 2 (gender: male; female) between-subjects ANOVA[11]. The analysis revealed a significant main effect of noise amplitude on identification, $F(2, 114) = 181.92, p < .001, \eta^2 = .76$. Identification of melodies with gaps replaced by high amplitude noise was significantly superior to identification of melodies with gaps replaced by low amplitude noise and very low amplitude noise (Games-Howell: $p < .001$, $d = 2.74$ and $p < .001$, $d = 2.85$, respectively). The difference between low amplitude noise and very low amplitude noise conditions was non-significant (Games-Howell: $p > .05$).

The gap duration effect from Experiment 2 was replicated here; melodies with 100 ms gaps were easier to identify than melodies with 300 ms gaps, $F(1, 114) = 59.48$, $p < .001, \eta^2 = .34$. The interaction between noise amplitude and gap duration was non-significant, $F(2, 114) = .83, p > .05$.

Whilst the main effect of gender on identification was non-significant, $F(1, 114) = 3.17, p > .05$, gender did significantly interact with noise amplitude, $F(2, 114) = 3.94$, $p < .05, \eta^2 = .07$ (see Figure 5.2). This was a result of the superior performance of males when compared to females in high amplitude noise conditions, $t(40) = 2.80, p < .01, d = 0.85$. The difference between males and females in other noise amplitude conditions was non-significant (all *p*s >.05). Gender did not interact with gap duration, $F(1, 114) = 1.04$, $p > .05$, and the three-way interaction was also non-significant, $F(2, 114) = .28, p > .05$.

---

[11] Due to the slightly lower mean age of children in the low amplitude noise conditions relative to the other conditions, this analysis was also run with age included as a covariate. Neither the main effects nor the interactions were affected by this.

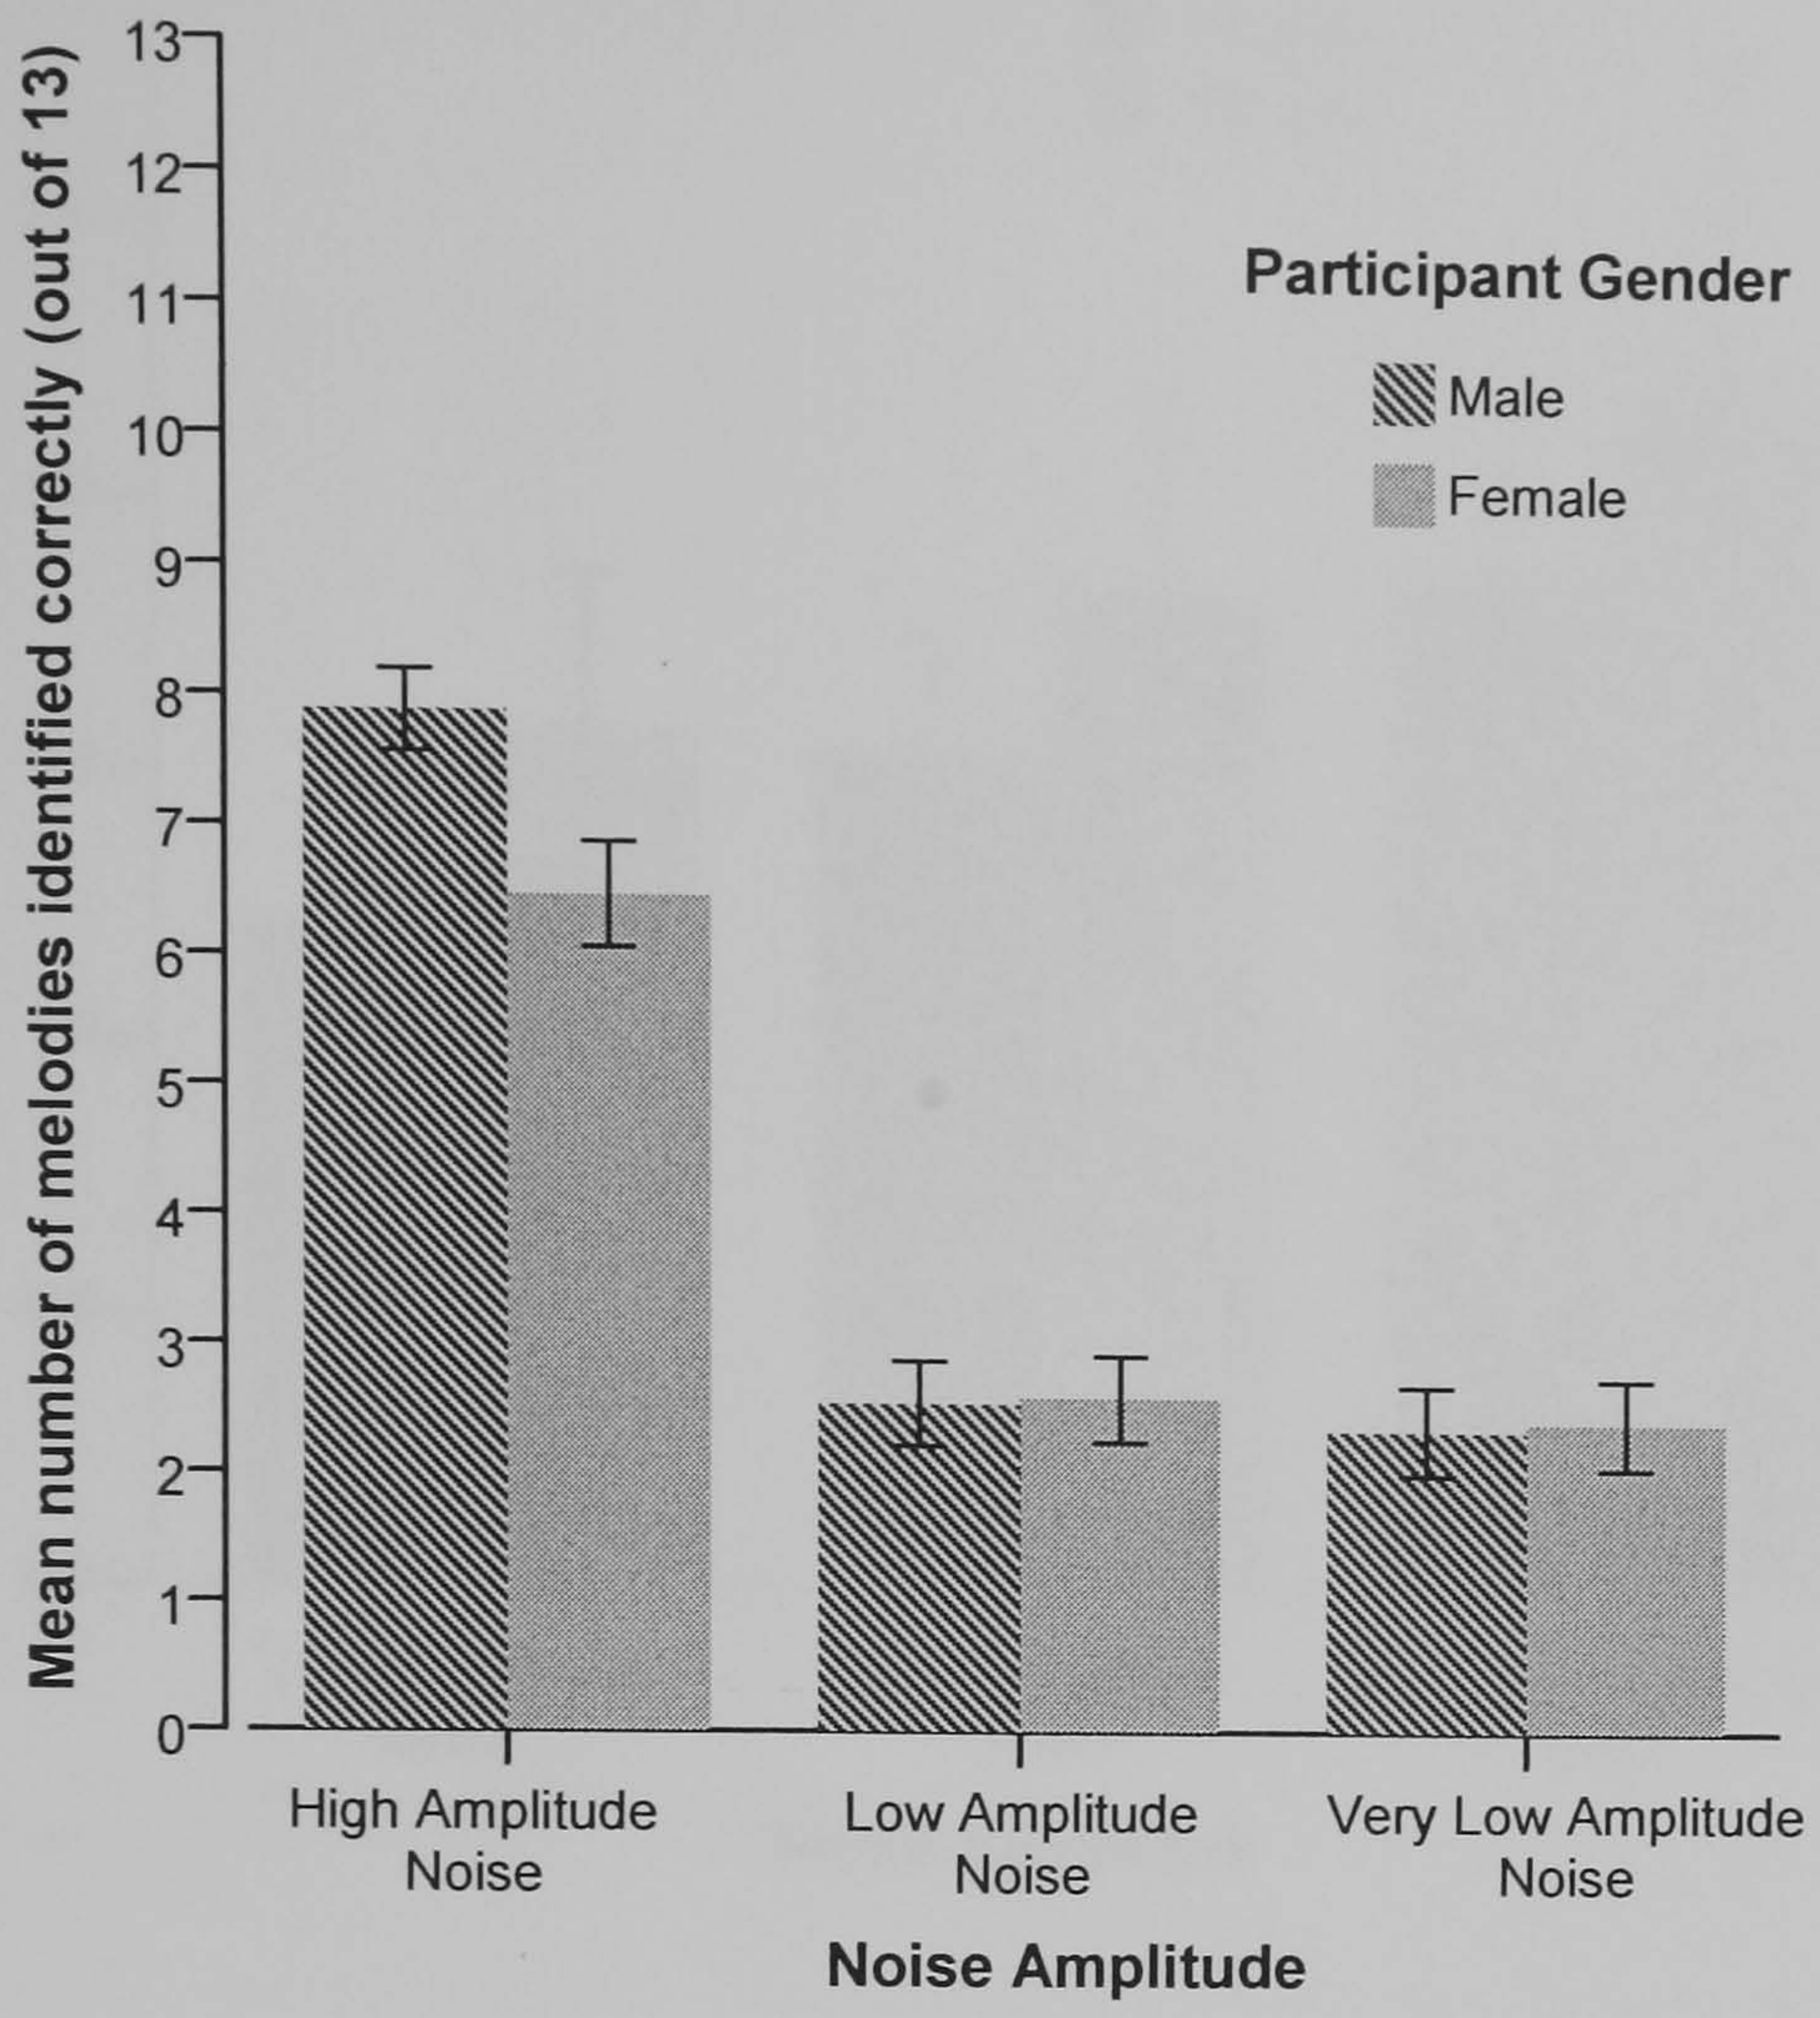

Figure 5.2. Mean number of melodies identified correctly (± *SE*) by males and females in each of the noise amplitude conditions (high amplitude noise, low amplitude noise and very low amplitude noise).

Identification time

The average identification time for correctly identified melodies was calculated for each participant in the same way as for previous experiments. This analysis excludes data from one child in the very low amplitude noise 100 ms condition, two children from the very low amplitude noise 300 ms condition, and four children from the low amplitude noise 300 ms condition as they did not identify any melodies correctly.

In the same way as for identification scores, the identification time data were examined for gender differences (see Appendix 2). There were no significant gender differences so data were collapsed across gender. The mean identification time (in seconds) is shown for children in each condition in Figure 5.3.

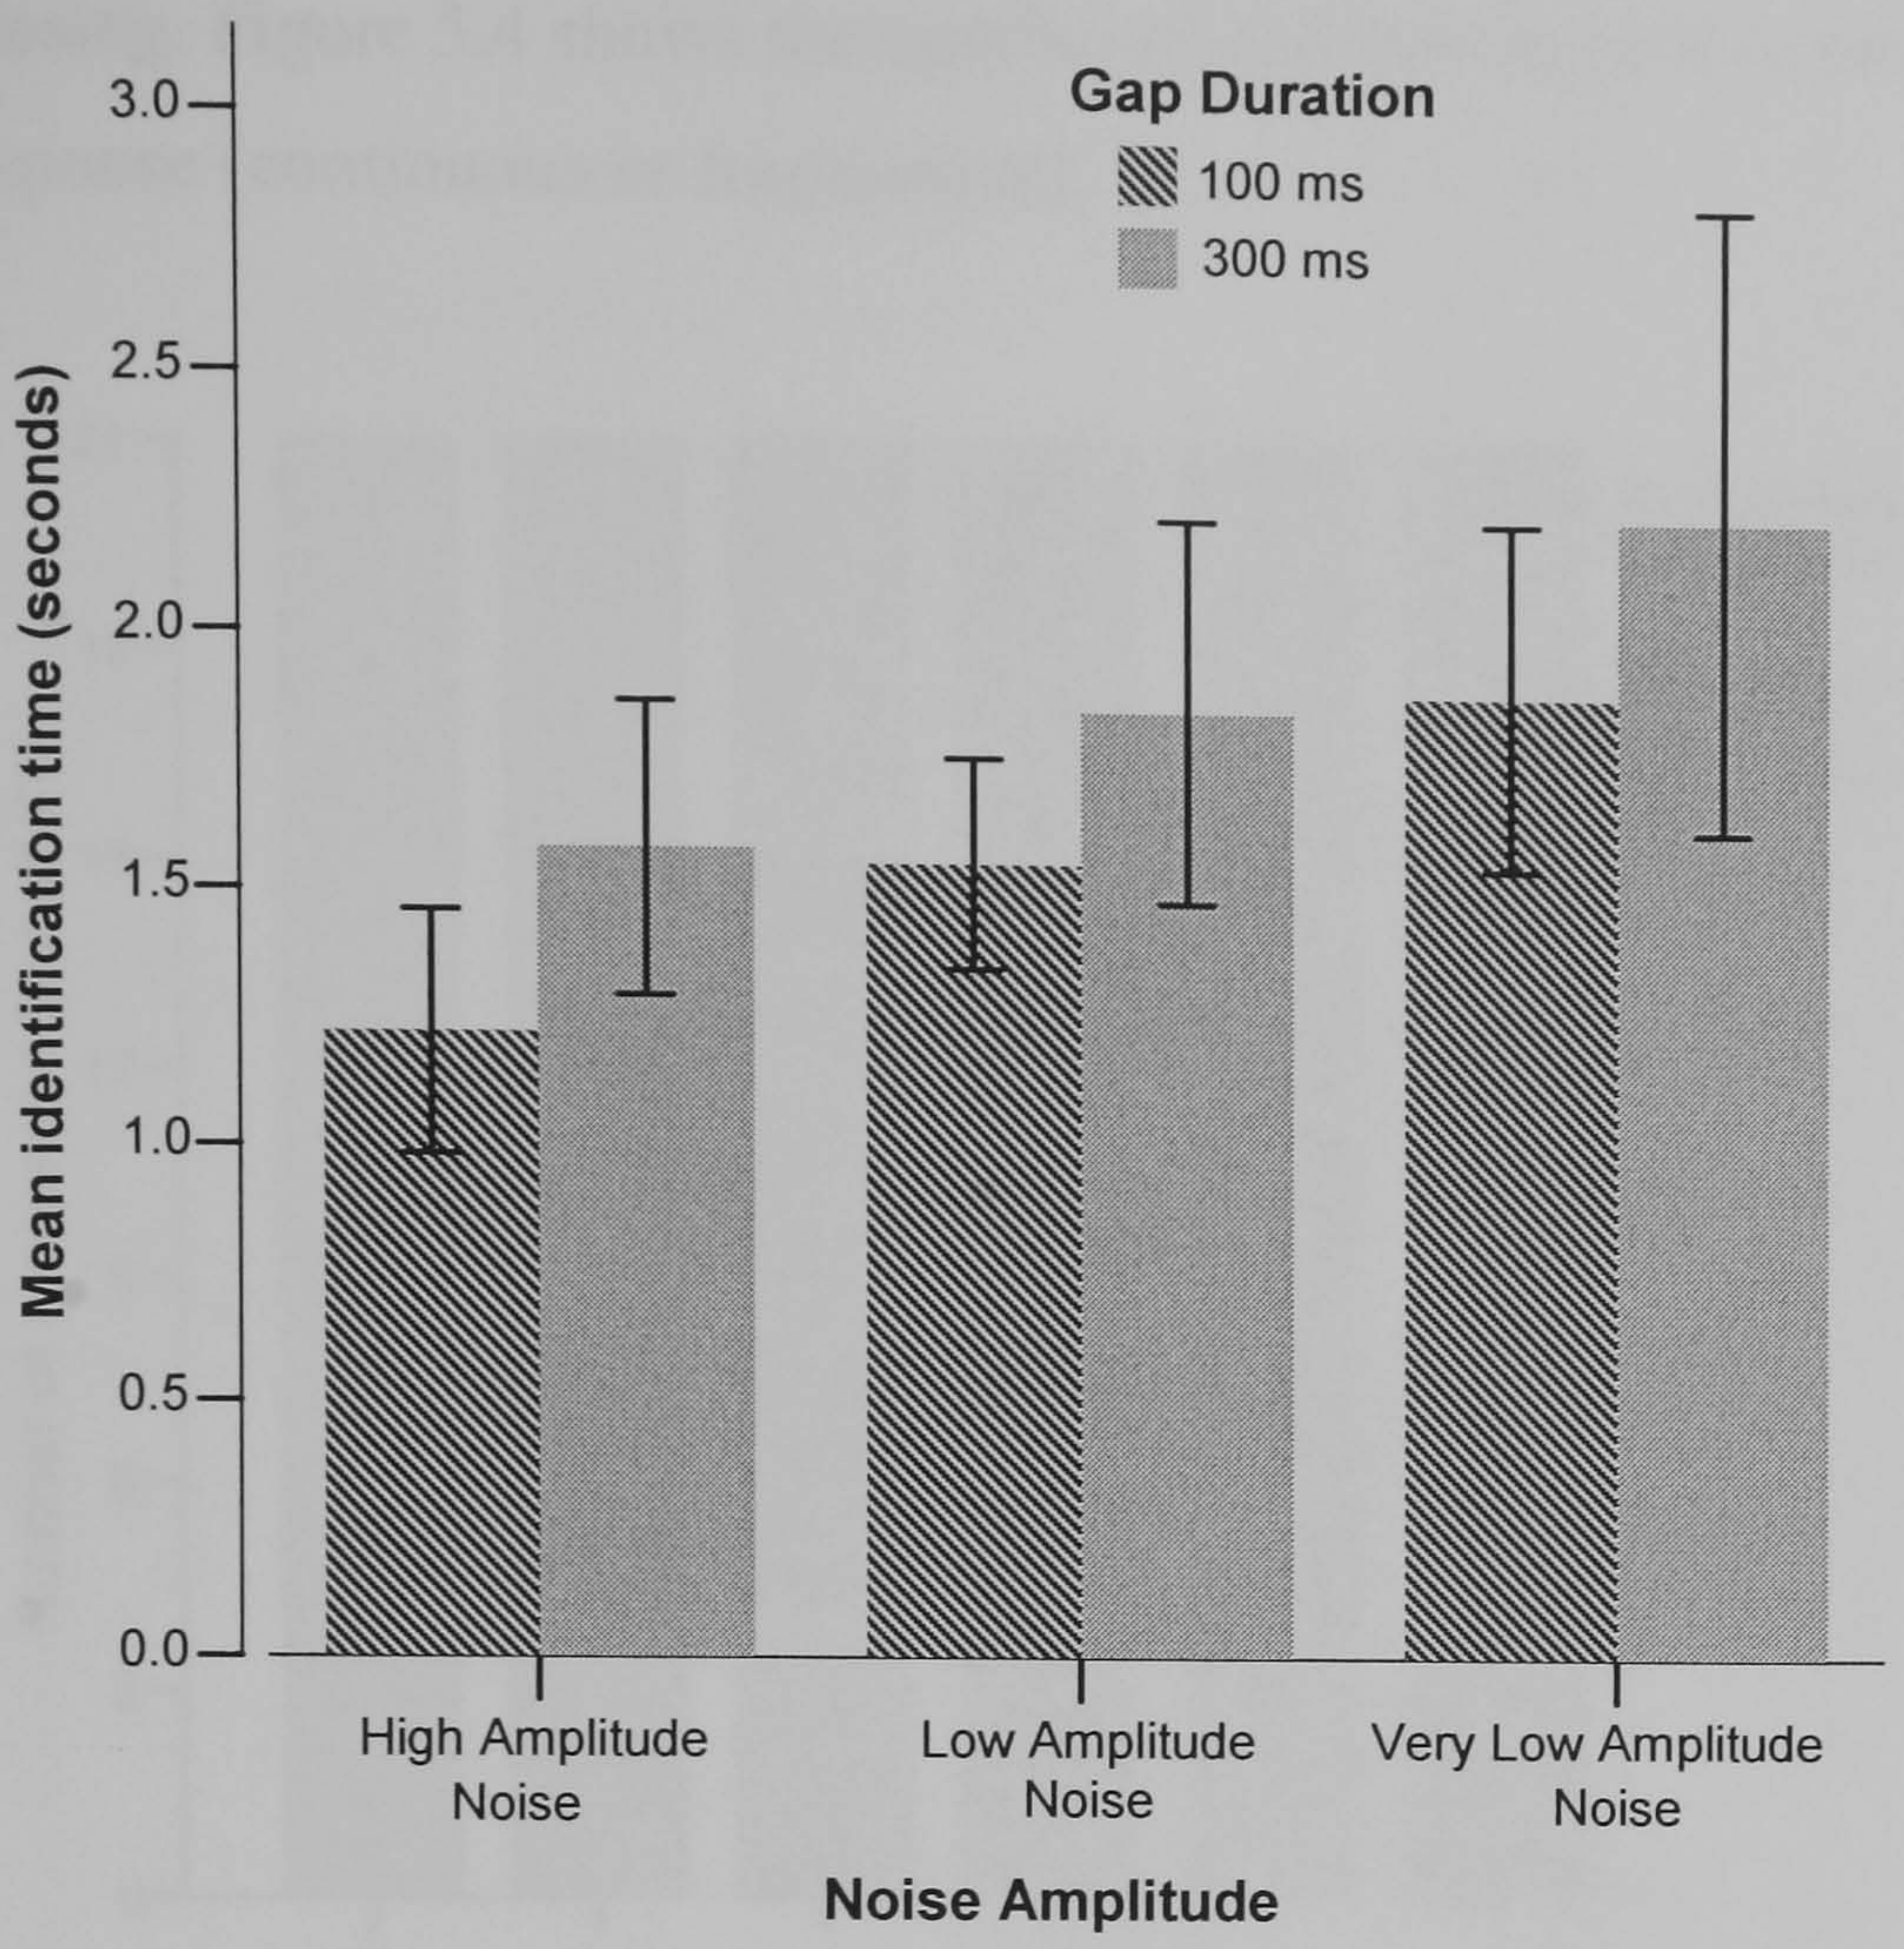

Figure 5.3. Mean identification time (identification time − stimulus duration; ± *SE*) in seconds for children in each of the six experimental conditions (high amplitude noise, low amplitude noise, very low amplitude noise; 100 ms or 300 ms gaps).

These data were analysed using a 3 (noise amplitude: high amplitude noise; low amplitude noise; very low amplitude noise) × 2 (gap duration: 100 ms; 300 ms) between-subjects ANOVA. The main effect of noise amplitude was non-significant, $F(2, 113) = 1.68, p > .05$ and furthermore, there was no significant effect of gap duration on identification time, $F(1, 113) = .84, p > .05$. The interaction between noise amplitude and gap duration was also non-significant, $F(2, 113) = .07, p > .05$.

Perception of Continuity

Each child was classified as either reporting that the signal was continuous, or reporting that the signal was fragmented, based on the response they gave when asked whether they thought that all of the tune was there or that little pieces of the tune were missing. Figure 5.4 shows the number of children in each condition who gave each response (continuous or fragmented).

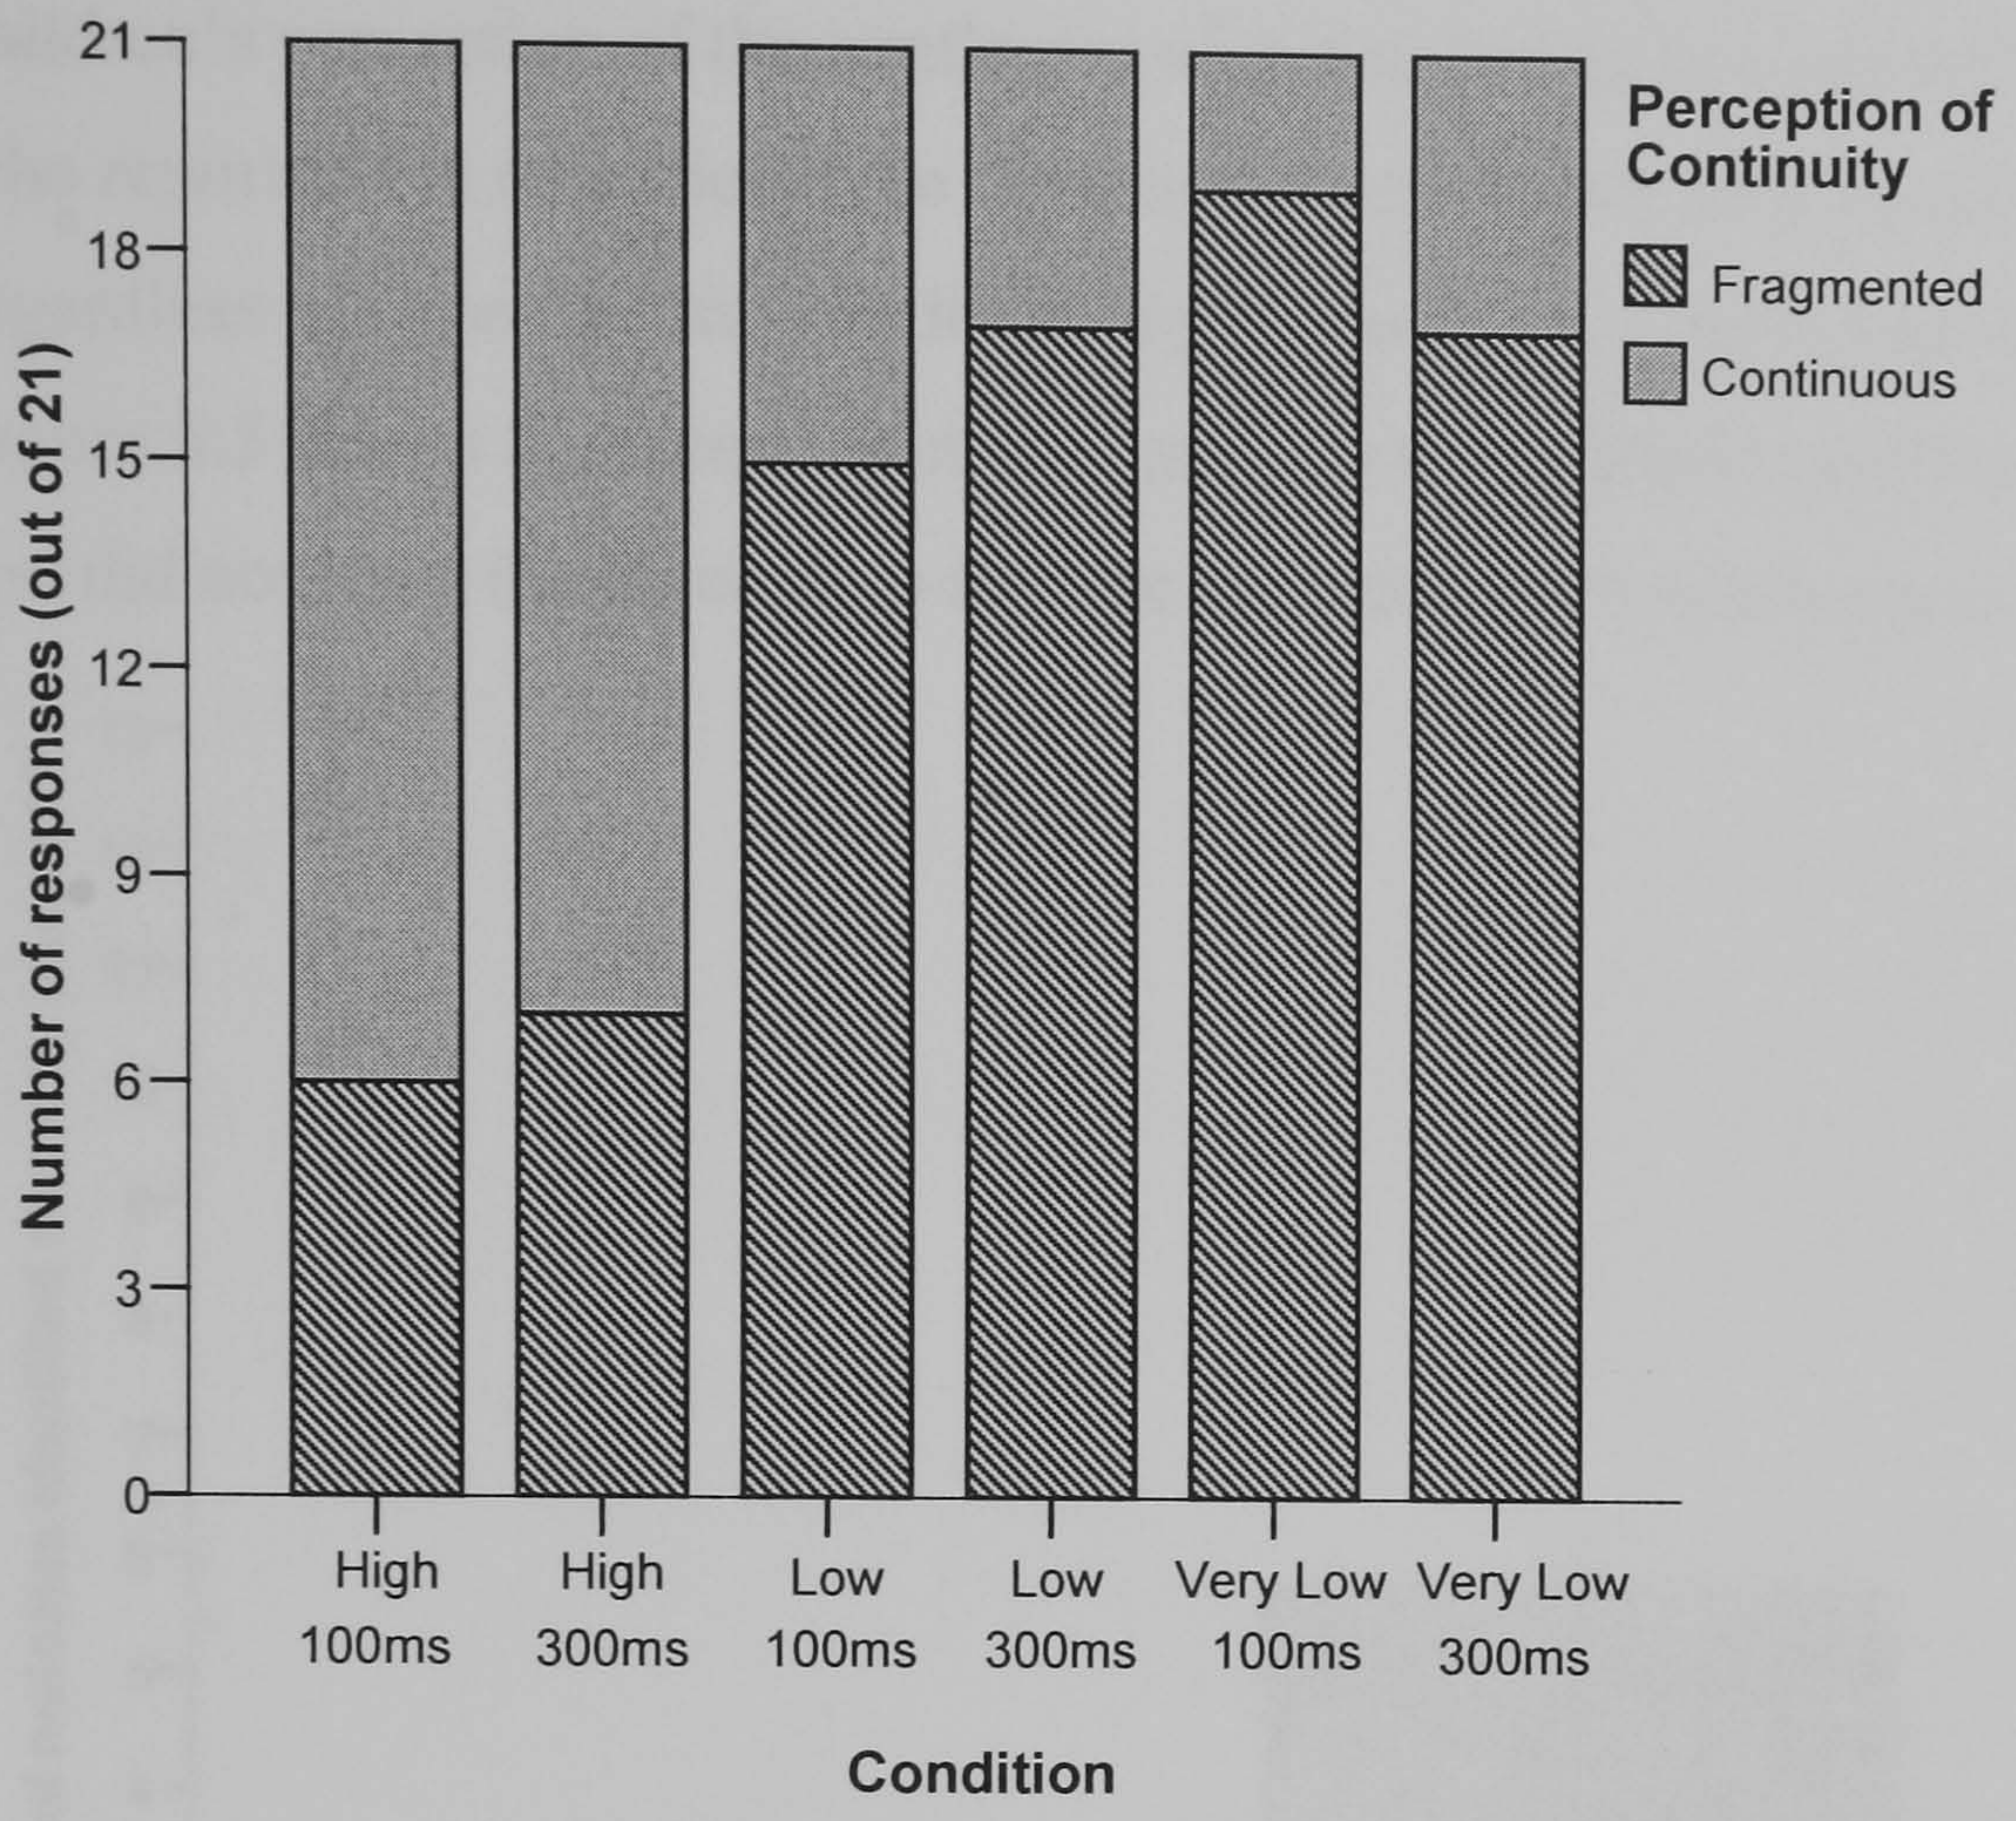

Figure 5.4. Total number of children in each condition (high amplitude noise, low amplitude noise, very low amplitude noise; 100 ms or 300 ms gaps) reporting the melodies as fragmented and reporting the melodies as continuous.

These data were analysed using a 6 (condition: high amplitude 100 ms; high amplitude 300 ms; low amplitude 100 ms; low amplitude 300 ms; very low amplitude 100 ms; very low amplitude 300 ms) × 2 (perception of continuity: continuous; fragmented) chi-square. Observed and expected frequencies for chi-square analyses across all experiments are reported in Appendix 3. The chi-square revealed a significant association between condition and the perception of continuity, χ^2 (5, $N = 126$) = 32.25, $p < .001$, Cramer's $V = .51$. As Figure 5.4 shows, the majority of children reported the melodies to be continuous when replacement noise was a potential masker of the missing information (high amplitude noise conditions), whereas most children thought the melodies were fragmented when replacement noise amplitude was lower than that of the signal (low and very low amplitude noise conditions).

Given that there is a relationship between replacement noise amplitude and children's perception of the continuity of the melodies, it is possible that the children who reported the melodies to be continuous were better able to identify the melodies regardless of experimental condition, and regardless of the actual status of the signal. Figure 5.5 shows the mean number of melodies identified correctly for children who did and did not have the perception that the melodies were continuous.

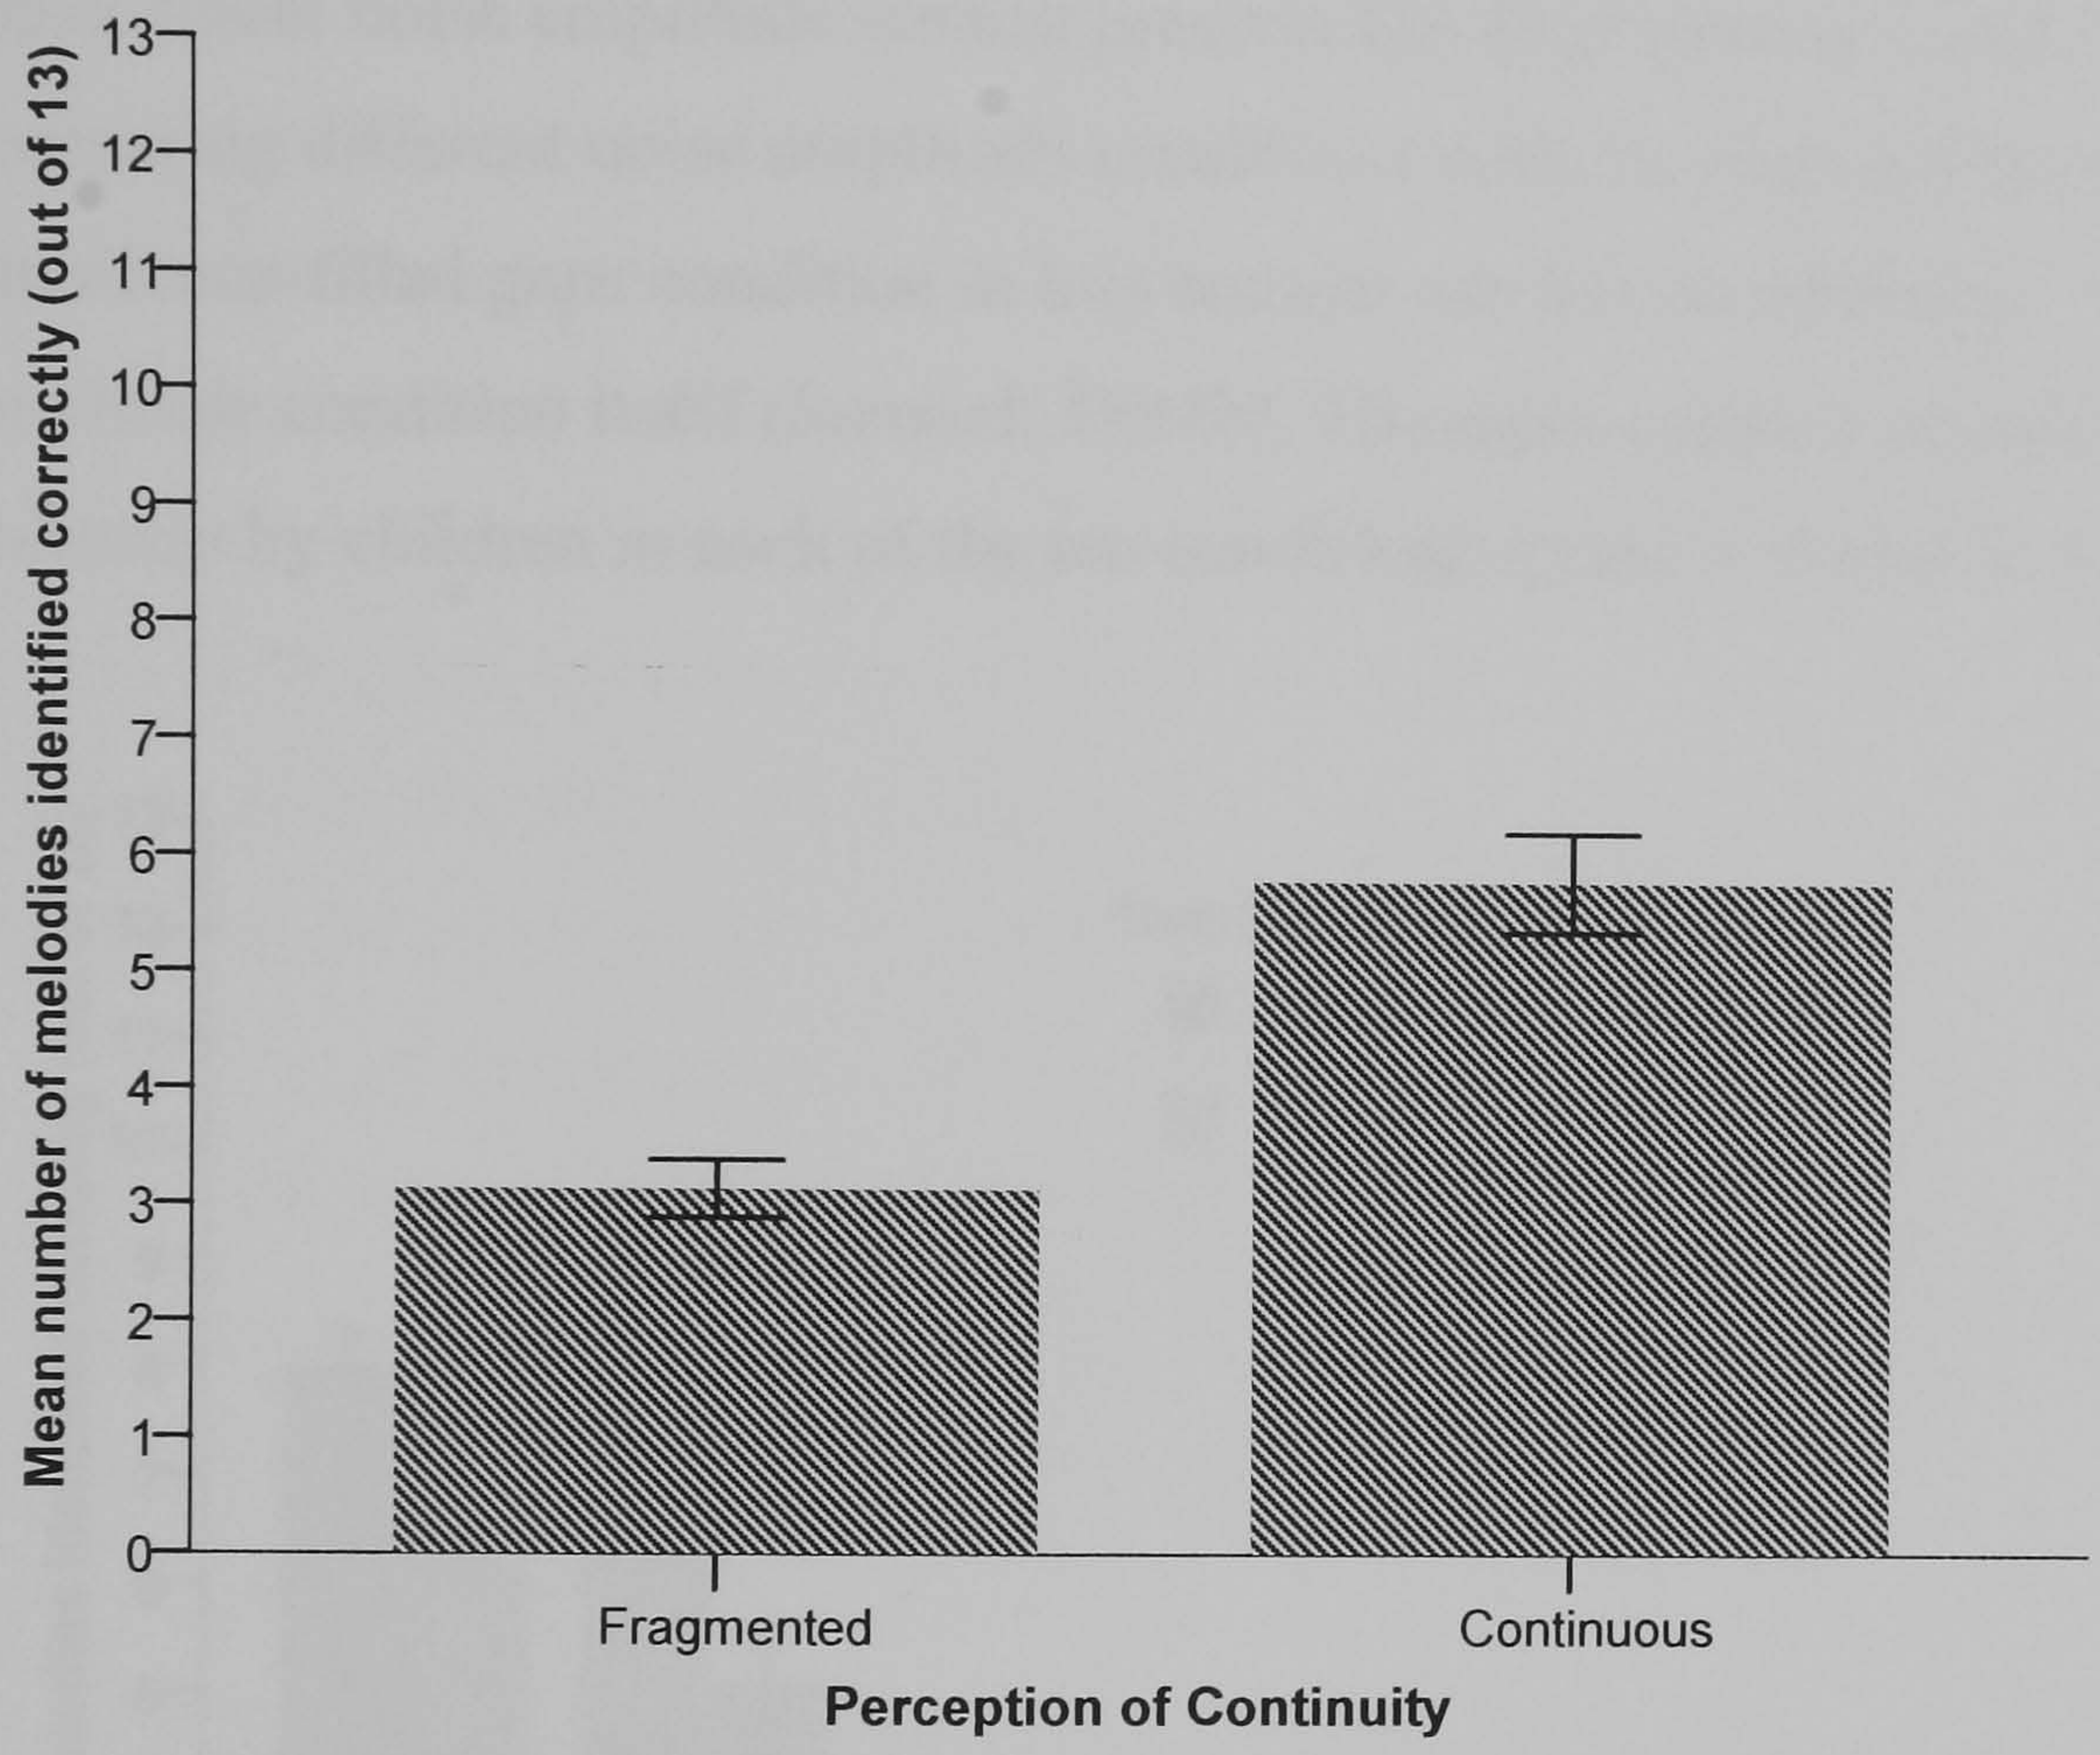

Figure 5.5. Mean number of melodies identified correctly (± *SE*) by children who reported the melodies to be fragmented and children who reported the melodies to be continuous.

An independent-samples *t* test showed that the mean identification score of those children who reported that the tunes sounded continuous was significantly higher than the mean identification score of children who reported that the tunes were fragmented, $t(74) = 5.38, p < .001, d = 0.93$.

Comparison between Experiment 3 and Experiment 2

Correct identification scores

In order to test whether replacement noise of different amplitudes leads to a significant intelligibility advantage over silent gaps, it is necessary to make a comparison between the data from the present experiment and the data from Experiment 2. This allows a noise-filled gaps over silence-filled gaps comparison to be made for all replacement noise amplitudes, and also comparisons can be made with the standard replacement noise amplitude stimuli presented in Experiments 1 and 2. As well as comparing different noise amplitude conditions with the silence-filled gaps condition, the silence-filled gaps condition in this context can be considered an "extreme" noise amplitude condition itself (Samuel, 1981b). The mean number of melodies identified correctly by children in each of the ten conditions tested is shown in Figure 5.6.

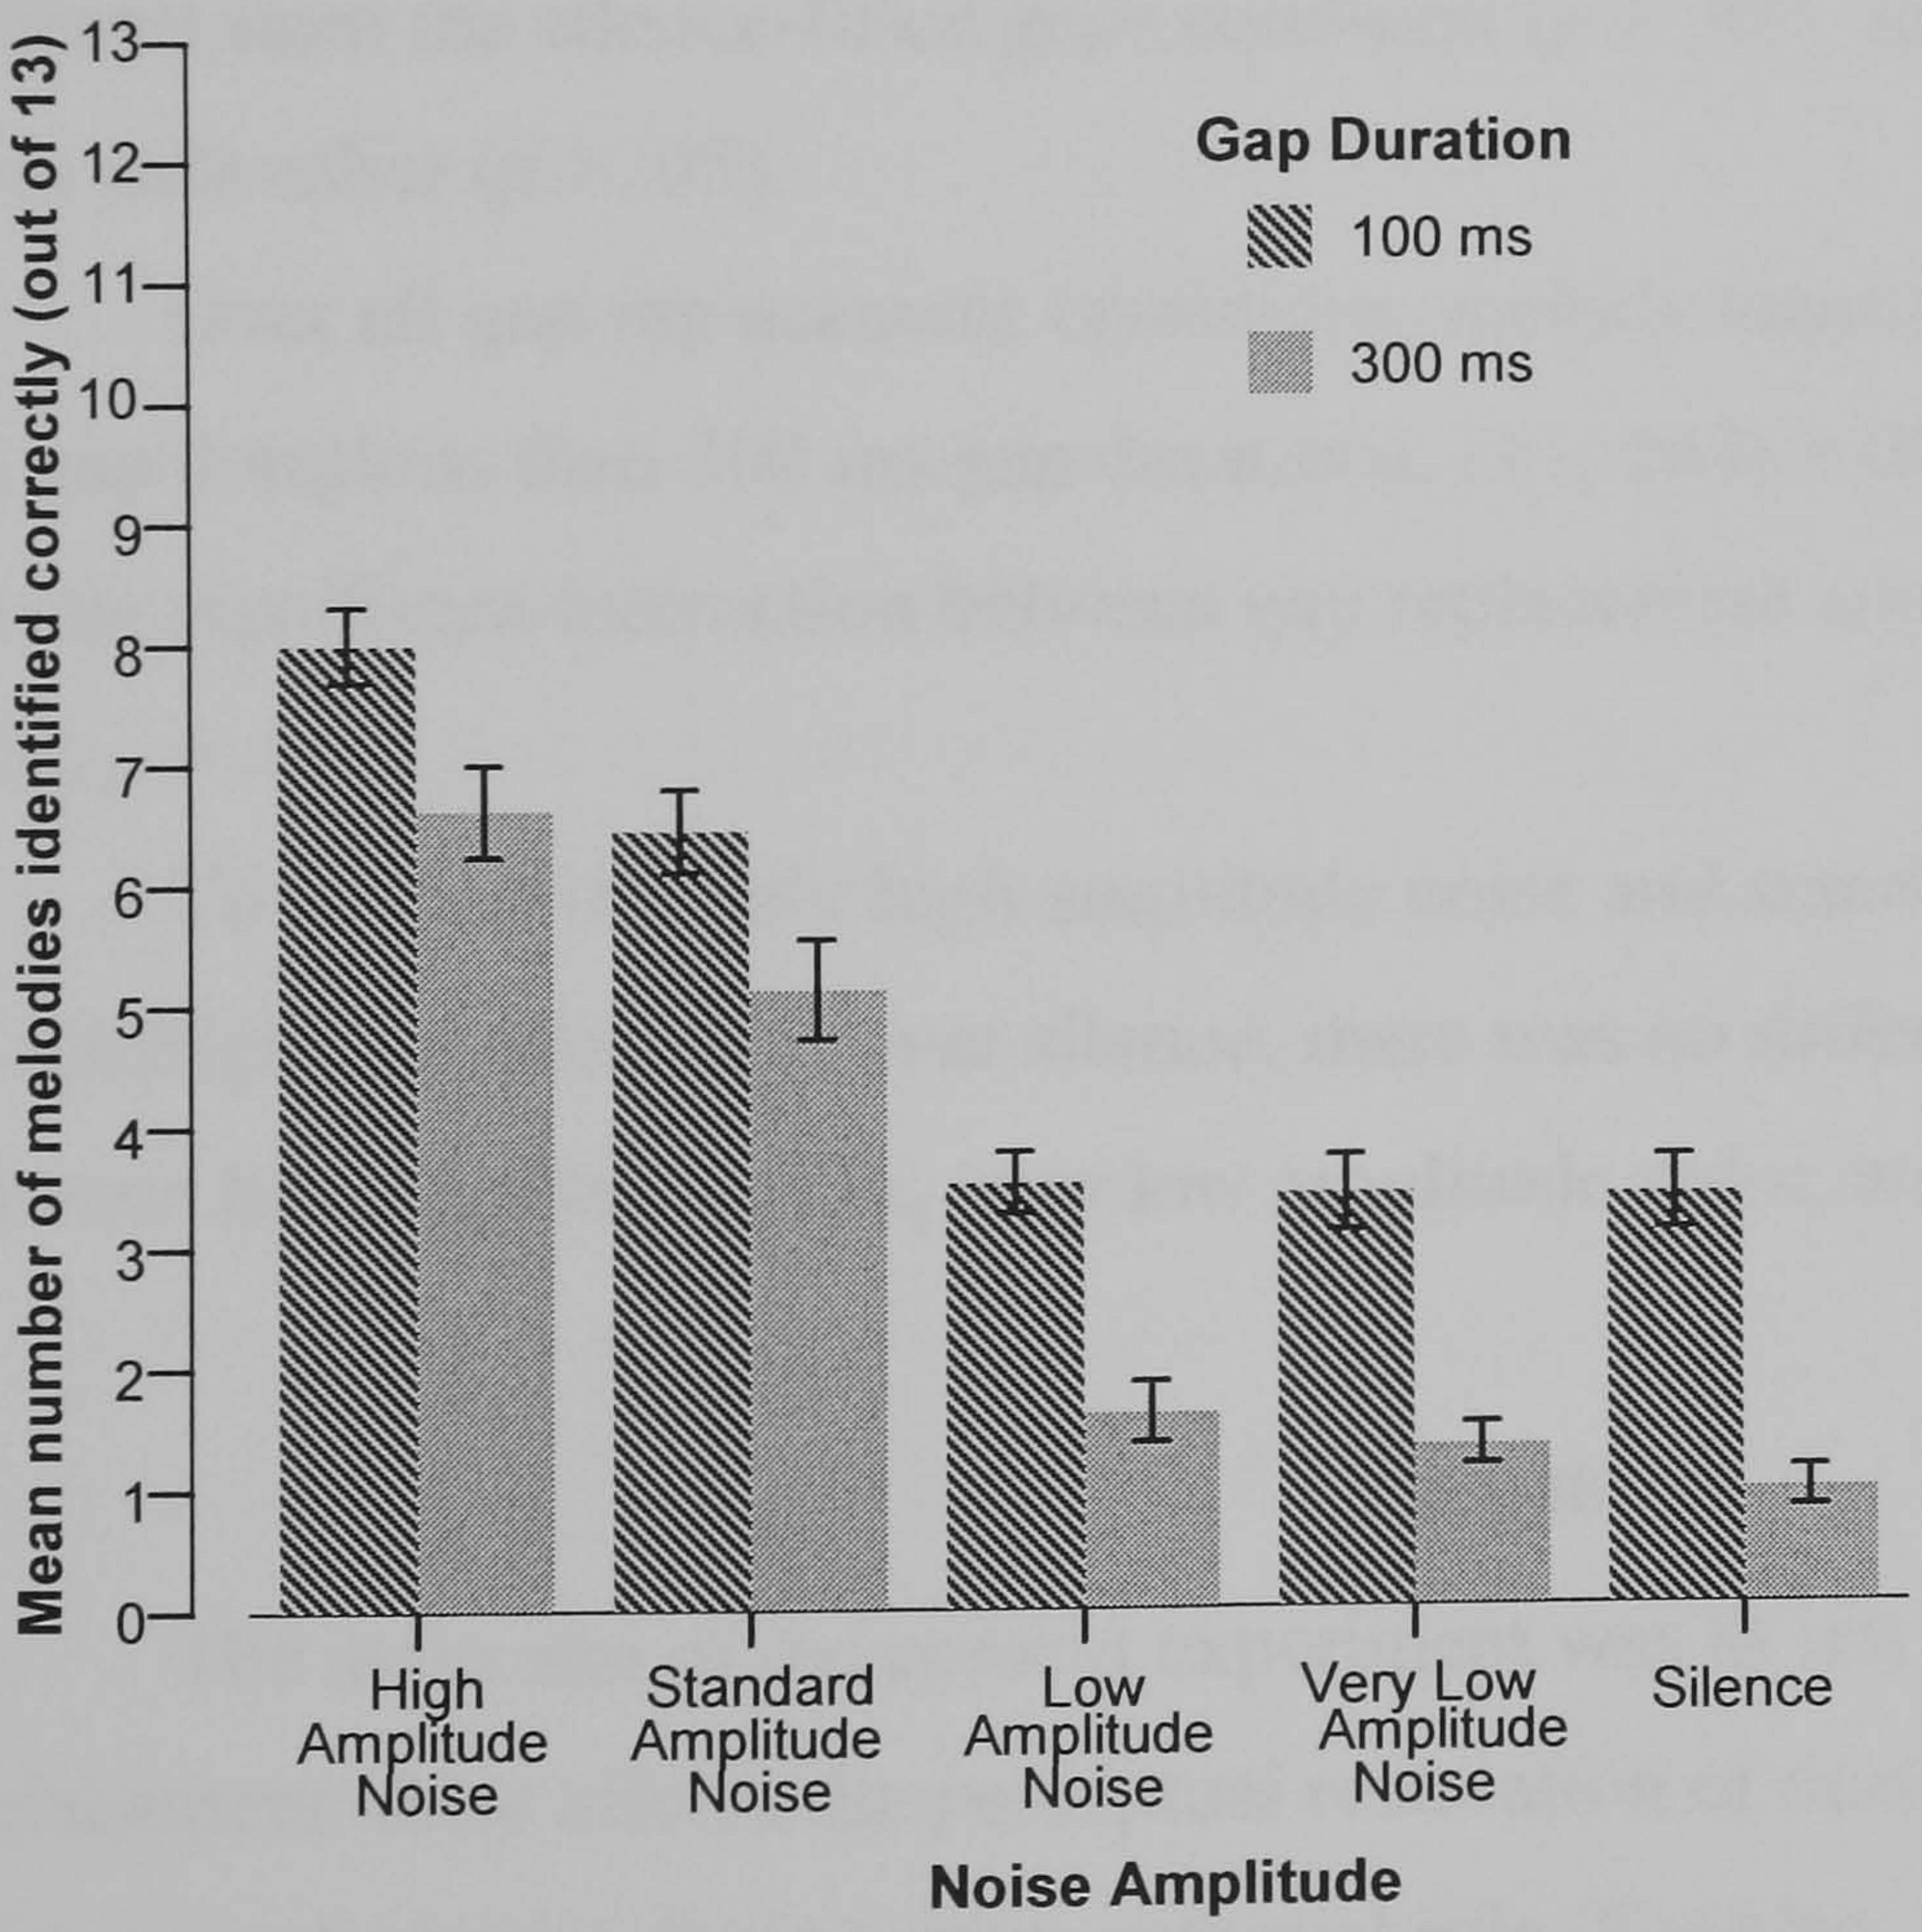

Figure 5.6. Mean number of melodies identified correctly (± *SE*) by children in each of the ten experimental conditions (high amplitude noise, standard amplitude noise, low amplitude noise, very low amplitude noise, silence; 100 ms or 300 ms gaps).

These data were analysed using a 5 (gap replacement: high amplitude noise; standard amplitude noise; low amplitude noise; very low amplitude noise; silence) × 2 (gap duration: 100 ms; 300 ms) between-subjects ANOVA. There was a significant main effect of gap replacement, $F(4, 204) = 119.22, p < .001, \eta^2 = .70$. Post hoc tests (Games-Howell) showed that melody identification in the high amplitude noise conditions was significantly superior to identification in all other gap replacement conditions except for standard amplitude noise (low amplitude noise, very low amplitude noise, silence: all *p*s $< .001$; standard amplitude noise: $p > .05$). Most importantly, high amplitude noise as a gap replacement had a highly significant advantage over silence in terms of identification ($p < .001, d = 2.96$).

Identification in the standard amplitude noise conditions was significantly superior to low amplitude noise and very low amplitude noise conditions ($p < .001$), and also showed a significant advantage over silence in identification ($p < .01, d = 1.91$). Neither low amplitude noise nor very low amplitude noise conditions were significantly different from the silence-filled gaps condition ($p > .05$), and did not significantly differ from each other ($p > .05$).

Over all gap replacement conditions, melody identification was superior for 100 ms gap durations than 300 ms gap durations, $F(1, 204) = 90.66, p < .001, \eta^2 = .31$. There was no significant interaction between gap replacement and gap duration, $F(4, 204) = 1.28, p > .05$.

To summarise, only high amplitude noise and standard amplitude noise produced an intelligibility advantage over silence; there was no difference in melody identification between low amplitude noise, very low amplitude noise, and silence conditions.

Discussion

The main aim of the present experiment was to test whether the amplitude of the replacement noise affects the perceptual restoration of musical signals in children. This was a test of whether the masking potential rule (Kashino, 2006) applies in this context, where only noise that exceeds the peak amplitude of the signal, and is thus a potential masker of the missing information, can result in perceptual restoration.

The main experimental findings with reference to melody identification were that when missing sections of the melodies were replaced with high amplitude noise (above the peak amplitude of the signal), identification was significantly superior to when the same sections were replaced by low amplitude noise and very low amplitude noise (both below the peak amplitude of the signal). Interestingly, there was no significant difference between identification of melodies in the low amplitude noise and very low amplitude noise conditions. This suggests that below the peak amplitude of the signal, there is no increase in intelligibility with increasing amplitude of replacement noise. Furthermore, the same effect of gap duration occurred for all conditions regardless of the noise amplitude. This supports the effect of gap duration found in Experiment 2, where stimuli with shorter gaps (100 ms) were easier to identify than stimuli with longer gaps (300 ms).

The significant gender difference for high amplitude noise conditions might represent a difference between boys and girls in the ability to attend to a signal containing loud bursts of noise. Whilst perceptual restoration was clearly operating for both boys and girls, it is possible that loud noise is more distracting for girls than boys, which might have had an impact on identification scores.

When data from the present experiment were compared with data from Experiment 2, it was evident that perceptual restoration, as demonstrated by a significant noise-filled gaps over silence-filled gaps advantage in melody identification, only occurred where the amplitude of the replacement noise exceeded the peak amplitude of the signal (high and standard amplitude noise conditions). When the amplitude of the replacement noise did not exceed the peak amplitude of the signal (low and very low amplitude noise conditions), there was a striking similarity in melody identification when compared to melodies with silent gaps. This suggests that in the case of children's music perception, this constraint prevents the perceptual restoration of gaps in an auditory signal in inappropriate situations, such as when the extraneous sound could not have been a potential masker of the missing information (Kashino, 2006).

As in Experiments 1 and 2, there were no differences in identification times between the conditions. Whilst the data indicate that melodies were identified more quickly with increasing amplitude of replacement noise, the variability in the data prevented this trend from reaching significance.

Examination of children's subjective reports of whether melodies were continuous or fragmented suggested there to be a relationship between the amplitude of the replacement noise and children's judgements. When missing sections were replaced with noise that was a potential masker of the missing information (high amplitude noise) the majority of children reported that they thought all of the tune was there. This suggests that they were hearing as present sections that in reality were missing from the melody. In contrast, children in the low and very low amplitude noise conditions, where the noise was not a potential masker of the missing information, most frequently reported that they thought little pieces of the tune were missing, as they were in reality. These data corroborate the correct identification scores; only where perceptual restoration occurred did the majority of children report the signals to be continuous. These findings also support those of Riecke et al. (2008), where continuity of a signal is not perceived unless the amplitude of the replacement noise exceeds that of the signal.

The perception of continuity and the restoration of intelligibility appear to be related; children who reported the melodies to be continuous scored significantly higher in terms of melody identification than those who reported the melodies to be fragmented. Whilst the continuity report data rely on a subjective measure, which may not be reliable in such young listeners, the consistency with which children responded in the different conditions does indeed suggest that children were able to make a judgement about what they thought the signal sounded like. This data is also consistent with behavioural identification data, which as Kingstone et al. (2008) suggest, allows both levels of analysis to complement one another.

Taken together, the most plausible explanation for this pattern of findings is that only when the noise is loud enough to act as a potential masker, and hide the fact that parts of the signal are actually missing, can illusory continuity of the signal can be perceived, and perceptual restoration occur. According to Bashford et al. (1992), the ability of top-down processes to contribute to perceptual restoration by restoring the content of the missing fragments is dependent upon the perception of a continuous signal. It certainly appears that perceptual restoration involves two separate, but interdependent processes as Bashford et al. (1992) and Verschuure & Brocaar (1983) suggested. First, illusory continuity of the signal is perceived so that the listener does not believe that any fragments are missing. Second, intelligibility of the signal is restored, for example using the automatic interpolation mechanism suggested by Verschuure and Brocaar (1983). Data from the present experiment suggest that the strongest perceptual restoration of intelligibility only occurs when illusory continuity of the signal is also perceived[12]. Noise that is a potential masker of the signal is more likely to produce the perception of illusory continuity than quieter replacement noise, supporting the propositions made by Remijn et al. (2007) and Shriberg (1992). In failing to act as a potential masker for the signal, low and very low amplitude noise are unable to provide sufficient bottom-up confirmation of the expected content of the missing fragments.

The operation of a noise amplitude effect on the perceptual restoration of music in children is also consistent with claims that the role of noise in perceptual restoration is to provide perceptual closure of missing gaps, mask auditory edges around the missing fragment, and to provide continued neural stimulation during the gap in the signal (e.g. Bregman, 1990; Petkov et al., 2007). To the extent that noise provides these conditions, where the amplitude of the noise exceeds that of the signal, perceptual restoration can occur. This is the first demonstration that the masking potential rule applies to the perception restoration of music, and to perceptual restoration in children.

The present findings support previous studies that have demonstrated a stronger effect of perceptual restoration with increasing intensity of the replacement noise (Powers & Wilcox, 1977; Warren & Obusek, 1971). Specifically, these data mirror those obtained by Eimas et al. (1996) with speech signals; standard white noise led to stronger perceptual restoration when compared to low amplitude noise. The similarity of the present findings to those in the speech domain strengthens propositions that the perceptual restoration process operates in parallel in different auditory domains. This is further supported by the similarity of the present findings to phonemic restoration studies in demonstrating that shorter gap durations were better restored than longer gap durations. This effect in the context of music perception is most likely due to the fact that fewer transitional cues between notes are missing when interruptions are brief, thus preserving the melodic contour of the tune.

The present data also provide some important insights into the nature of auditory perception in young children. When children who heard the melodies with noise-filled

---

[12] The direction of this relationship remains open to debate; it may not necessarily be unidirectional in the direction implied by previous theories.

gaps reported that they believed the tune to be continuous, and were then able to identify the melody, this indicates that the children were hearing parts of the signal as being present that were, in reality, absent. It is well documented that such active auditory perception operates in adults, but the present findings suggest that auditory perception at a young age also operates in this way.

The aim of this experiment was to investigate the operation of a further acoustic constraint on the perceptual restoration of music in children. The experiment has demonstrated that the amplitude of the replacement noise in perceptual restoration is crucial in determining whether restoration can occur in musical signals and with child listeners. These results support a two-process model of perceptual restoration (e.g. Bashford et al., 1992; Shahin et al., 2009; Verschuure & Brocaar, 1983) involving the restoration of both continuity and intelligibility of a fragmented signal. The restoration of both components is best achieved when the replacement noise is a potential masker of the signal it replaces. This indicates that, experimentally, perceptual restoration is not simply the product of inserting noise into missing fragments of an auditory stimulus. The noise itself must satisfy certain acoustic requirements in order for the listener to be able to make sense of what is in reality a fragmented signal. That listeners often do not notice that fragments are missing is a powerful demonstration of how the human auditory system, even in young children, has developed in order to remain efficient in noisy listening conditions.

Summary

Experiments 2 and 3 have demonstrated that whilst the perceptual restoration of music does occur in children, its operation is affected by the acoustic properties of the signal and of the noise that replaces missing fragments of the signal. Obtaining this evidence was an important aim of the present series of experiments. Such acoustic constraints on the operation of perceptual restoration appear to serve an ecological purpose in limiting the operation of perceptual restoration to the conditions where reconstruction of sensory input is adaptive. Experiment 4 aimed to test for a different kind of constraint on the operation of perceptual restoration in the music domain; a familiarity constraint. Does perceptual restoration only occur for a highly familiar, typical section of a piece of music?

Chapter 6: Experiment Four

Familiarity influences on the perceptual restoration of music in children: The perceptual restoration of initial and later phrases of familiar melodies

Introduction

Empirical work on the perceptual restoration of speech and music signals in adults has demonstrated that familiarity with the content of the intact sections of an auditory signal is an important determinant of the strength of perceptual restoration, with increased familiarity strengthening the degree to which missing input is reconstructed when replaced by noise (e.g. Samuel, 1981a; Sasaki, 1980). Theories of perceptual restoration describe the process as a result of the interaction between top-down expectations of the content of the missing fragments and bottom-up confirmation of these expectations based on what is present within the acoustic stimulus itself (e.g. Samuel, 1981a; Samuel & Ressler, 1986).

The top-down influences on perceptual restoration are very powerful; however, it is important to remember that the effect is not purely a result of the top-down flow of information. The top-down expectations based on the intact sections of a disrupted signal are equal in both the case where missing fragments are replaced by noise and the case where they are silent. Hence, the expectation-confirmation interaction between top-down and bottom-up perceptual processes is crucial to the operation of the effect. In providing confirmation of the expected content of the missing fragments by hiding their absence, noise leads to perceptual restoration whereas silence does not.

Top-down influences on the perceptual restoration of speech

The perceptual restoration of speech signals, known as phonemic restoration, has been studied extensively. Research has demonstrated that top-down familiarity influences in the form of lexical and contextual factors strongly affect perceptual restoration of missing fragments. Samuel (1981a) reported a series of experiments demonstrating that missing phonemes from high frequency words with a stronger lexical representation were, in most cases, restored to a stronger degree than missing phonemes from low frequency words, which are less familiar. According to Samuel, high frequency words are more easily retrieved from the lexicon than low frequency words, and lexical information is important in generating expectations about later acoustic events. Thus, high frequency words generate stronger expectations with the result that perceptual restoration is stronger when compared to the perceptual restoration of low frequency words.

In addition, Samuel's (1981a) experiments demonstrated that missing phonemes from longer words that provide more contextual cues were more strongly restored than missing phonemes from shorter words. Samuel also found that missing phonemes from words with a lexical entry (e.g. *basis*) were better restored than missing phonemes from phonologically-legal pseudowords (e.g. *pafis*), regardless of differences in word frequency and length. Additional support for top-down influences on perceptual restoration comes from the finding that when isolated words are presented in lists and are thus not semantically connected, strong perceptual restoration does not occur (e.g. Bashford et al., 1992; Bashford & Warren, 1987b).

The manipulation of lexical factors provides an ideal and systematic means to investigate familiarity effects on the perceptual restoration of speech. However, these methods cannot be directly translated to investigation of the perceptual restoration of musical signals, for the reason that there are no direct correlates of word frequency, sentence context, and words and nonwords in the domain of music. DeWitt & Samuel (1990) reported a series of experiments investigating top-down influences on the perceptual restoration of music, and found effects indicating that top-down influences play a similar role in musical perceptual restoration to phonemic restoration.

DeWitt & Samuel (1990) employed a discrimination method with signal detection analysis, in which listeners were asked to report whether a version of a stimulus where a note had been removed and replaced with noise, or a version where the note had not been removed but noise was superimposed upon it, were intact. To the extent that the missing note was perceptually restored, the two versions would sound identical and hence a measure of discriminability between the items would be low. Their results appeared to be inconsistent with the literature on phonemic restoration; the measure of discriminability between added and replaced versions of a stimulus was higher for familiar than unfamiliar melodies, which suggests that perceptual restoration was weaker for the familiar melodies. However, the argument is made that when using such a discriminability measure as the indicator of perceptual restoration, effects of familiarity can be difficult to interpret. When a melody is highly familiar, fewer processing resources will be allocated to processing the melody, thus greater resources can be used to determine whether a stimulus is an added or replaced item. In later experiments, where the perceptual restoration of missing notes from chords and scales were tested, DeWitt and Samuel did find evidence that increased expectations resulted in increased perceptual restoration, thus providing some, if rather limited, evidence for the operation of top-down influences on the perceptual restoration of music.

There are difficulties interpreting familiarity effects with a discrimination measure as used by DeWitt and Samuel (1990), where the task requires listeners to make a judgement about whether noise replaces or is added to a musical note, that is, whether the item is intact. In addition, such a complicated signal analysis would not be possible with child listeners, as it would be difficult to phrase such task instructions in a way that children would fully understand. Therefore, a different measure and a different familiarity manipulation are required in order to investigate familiarity effects on the perceptual restoration of music in children. Experiments 1 to 3 investigating musical perceptual restoration in 4- to 6-year-old children have used an identification measure. This method minimises the task demands for young children and provides them with an enjoyable task that does not require complex perceptual decisions to be made. However, since an identification measure requires that all stimuli are familiar and hence capable of being identified, familiarity in this context needs to be manipulated in an alternative way.

In speech perception, word recognition is affected by "front anchoring" (Anderson, 1983). This refers to the facilitation in recognition provided by presentation of the beginning section of a word as opposed to a later section. Anderson (1983) illustrates this with reference to a letter sequence: the sequence *JVLVB* will be easier to recognise through presentation of *JVLV* than *VLVB*. Similarly, in music perception, it is generally assumed that the beginning section of a piece of music exemplifies the most typical part. Peretz, Radeau & Arguin (2004) state that, "beginnings are much easier to recognise than other song parts" (p.151). They go further to explain that the first section of a song can help trigger a memory representation, and that these beginning sections are stored differently in memory, as anchor points for the rest of the song. In terms of purely melodic music, the beginning section can also be the anchor point in memory. It is highly likely that the representation in semantic memory for a particular piece of music will contain this initial phrase. A later section of music stored in the musical lexicon is usually accessed through the representation of the beginning section. Therefore, it is possible to speculate that a later phrase of a familiar melody will be less familiar than the beginning phrase, at least in terms of the salience of the stored representation. In many ways, the musical lexicon, containing stored representations of familiar music, is very similar to the speech lexicon (Peretz, 1993).

There is evidence in the literature that stored memory representations for music are organised temporally. According to Schulkind (2004), the initial notes of a melody (the "anchor") access the stored representation of the melody in the musical lexicon. Once activated, this representation is involved in the on-line generation of expectations about subsequent sections of the melody; such expectations are important for perceptual restoration. These expectations are veridical expectations, which are related to a specific stored representation of something familiar, rather than schematic expectations for musical regularities (Bharucha, 1994).

Schulkind (2004) found that melodies that did not start with the initial notes (i.e. a "non-anchor" section) were harder to identify than those that started with the initial notes. This suggests that the global organisation of a melody is established through the initial notes, and where this organisation is not activated, expectations will be weaker. This might have important implications for the perceptual restoration of melodies. It is possible to draw a parallel between the beginning (highly typical) and later (less typical) musical sections, and the high and low lexical frequency words used in the study reported by Samuel (1981a). Thus, is it possible to speculate that in a similar way, beginning phrases will generate stronger expectations about later musical events and as a result perceptual restoration will be stronger than for a later phrase from the same piece of music.

The present experiment aimed to investigate familiarity effects on the perceptual restoration of music in children, by comparing the perceptual restoration of the beginning phrase of nursery rhyme melodies, to the perceptual restoration of a later phrase. Whilst beginning phrases can be restored (as demonstrated in Experiment 1), and are seen as the most typical nature of a stored representation for a piece of music, in order for perception to be truly flexible, perceptual restoration should also occur for other sections of the piece. It was expected that whilst perceptual restoration for the later, less typical phrases would occur, it would not be as strong as for the more typical initial phrases, due to weaker expectations of the content of missing fragments being generated. Thus, this experiment also serves as a test of the role of expectations in children's music perception, in the case of both intact and disrupted signals.

Method

Participants

Sixty-six children (38 girls) between the ages of 4 years 3 months and 6 years 5 months with a mean age of 5 years 3 months (63.24 months; *SD* = 7.05) participated in this study. None of these children had participated in previous experiments. These children were taken from the Reception and Year One classes of a primary school in Surrey, UK, in a predominantly white, middle-class area. Consent for participation was gained from both the school Headteacher and the child's parent or guardian. All children were native English speakers and had normal sensory development, as reported by the school. Further details about the age and gender of the participants in each condition are shown in Table 6.1.

Table 6.1. The gender split and mean age of the children in each of the three experimental conditions.

Experimental Condition[a]	**Gender**	**Mean age in months (*SD*)**
Intact	M=11; F=11	63.41 (7.60)
Noise-filled gaps	M=12; F=10	63.27 (6.81)
Silence-filled gaps	M=5; F=17	63.05 (7.05)

[a] *n* = 22 for each condition

Design

Children were randomly allocated to one of three independent conditions. Each child heard the same 15 melodic excerpts from nursery rhymes, but these were either heard intact, with alternating 100 ms fragments removed and replaced with white noise, or with alternating 100 ms fragments removed and replaced with silence, depending on the experimental condition. In this way, the stimuli were identical to those presented in Experiment 1 except that in the present experiment the melodies were a later phrase within the nursery rhyme, rather than the beginning phrase as presented in Experiment 1 (see Appendix 4 for details of the stimuli used in the present experiment).

Materials

In the same way as for Experiment 1, the stimuli were produced by playing the monophonic melody line of 15 nursery rhymes on a MIDI keyboard. The excerpt from each nursery rhyme that would form the "later phrase" stimulus was chosen according to the following criterion. If the phrase immediately following the initial phrase of the nursery rhyme (that had been used in previous experiments) involved direct repetition of the preceding phrase, the next distinctive phrase in the nursery rhyme was chosen. The later excerpts were matched in terms of tempo to the initial phrases used in previous experiments, and the phrase length of the excerpt in terms of the number of bars was equated to that of previous experiments. The average stimulus length was 5.24 seconds.

The stimuli were recorded and edited using the same equipment and in the same way as for Experiment 1. The duration of missing fragments was 100 ms, and replacement noise amplitude was the same as in Experiment 1; 12 dB below the maximum bandwidth of the file. The same Matlab programme was used to present the stimuli in a random order and record the child's response and identification time. In addition, the participating school was sent the standard song familiarity questionnaire (Appendix 1) prior to the school visits.

Procedure

The procedure was identical to that of Experiment 1, with the addition of the subjective continuity report introduced in Experiment 3. After all 13 experimental trials had been presented, the child was asked whether they thought that "little pieces of the tune were missing or all of the tune was there" as a further test of perceptual restoration. The child's response was noted, and they were praised and thanked for their help.

Results

Correct identification scores

Each child was given a score out of 13 which represented the number of melodies they identified correctly in the experimental trials. Data were examined for gender differences (see Appendix 2) and none emerged so data were collapsed across gender. The mean number of melodies identified correctly by children in each of the three experimental conditions is shown in Figure 6.1.

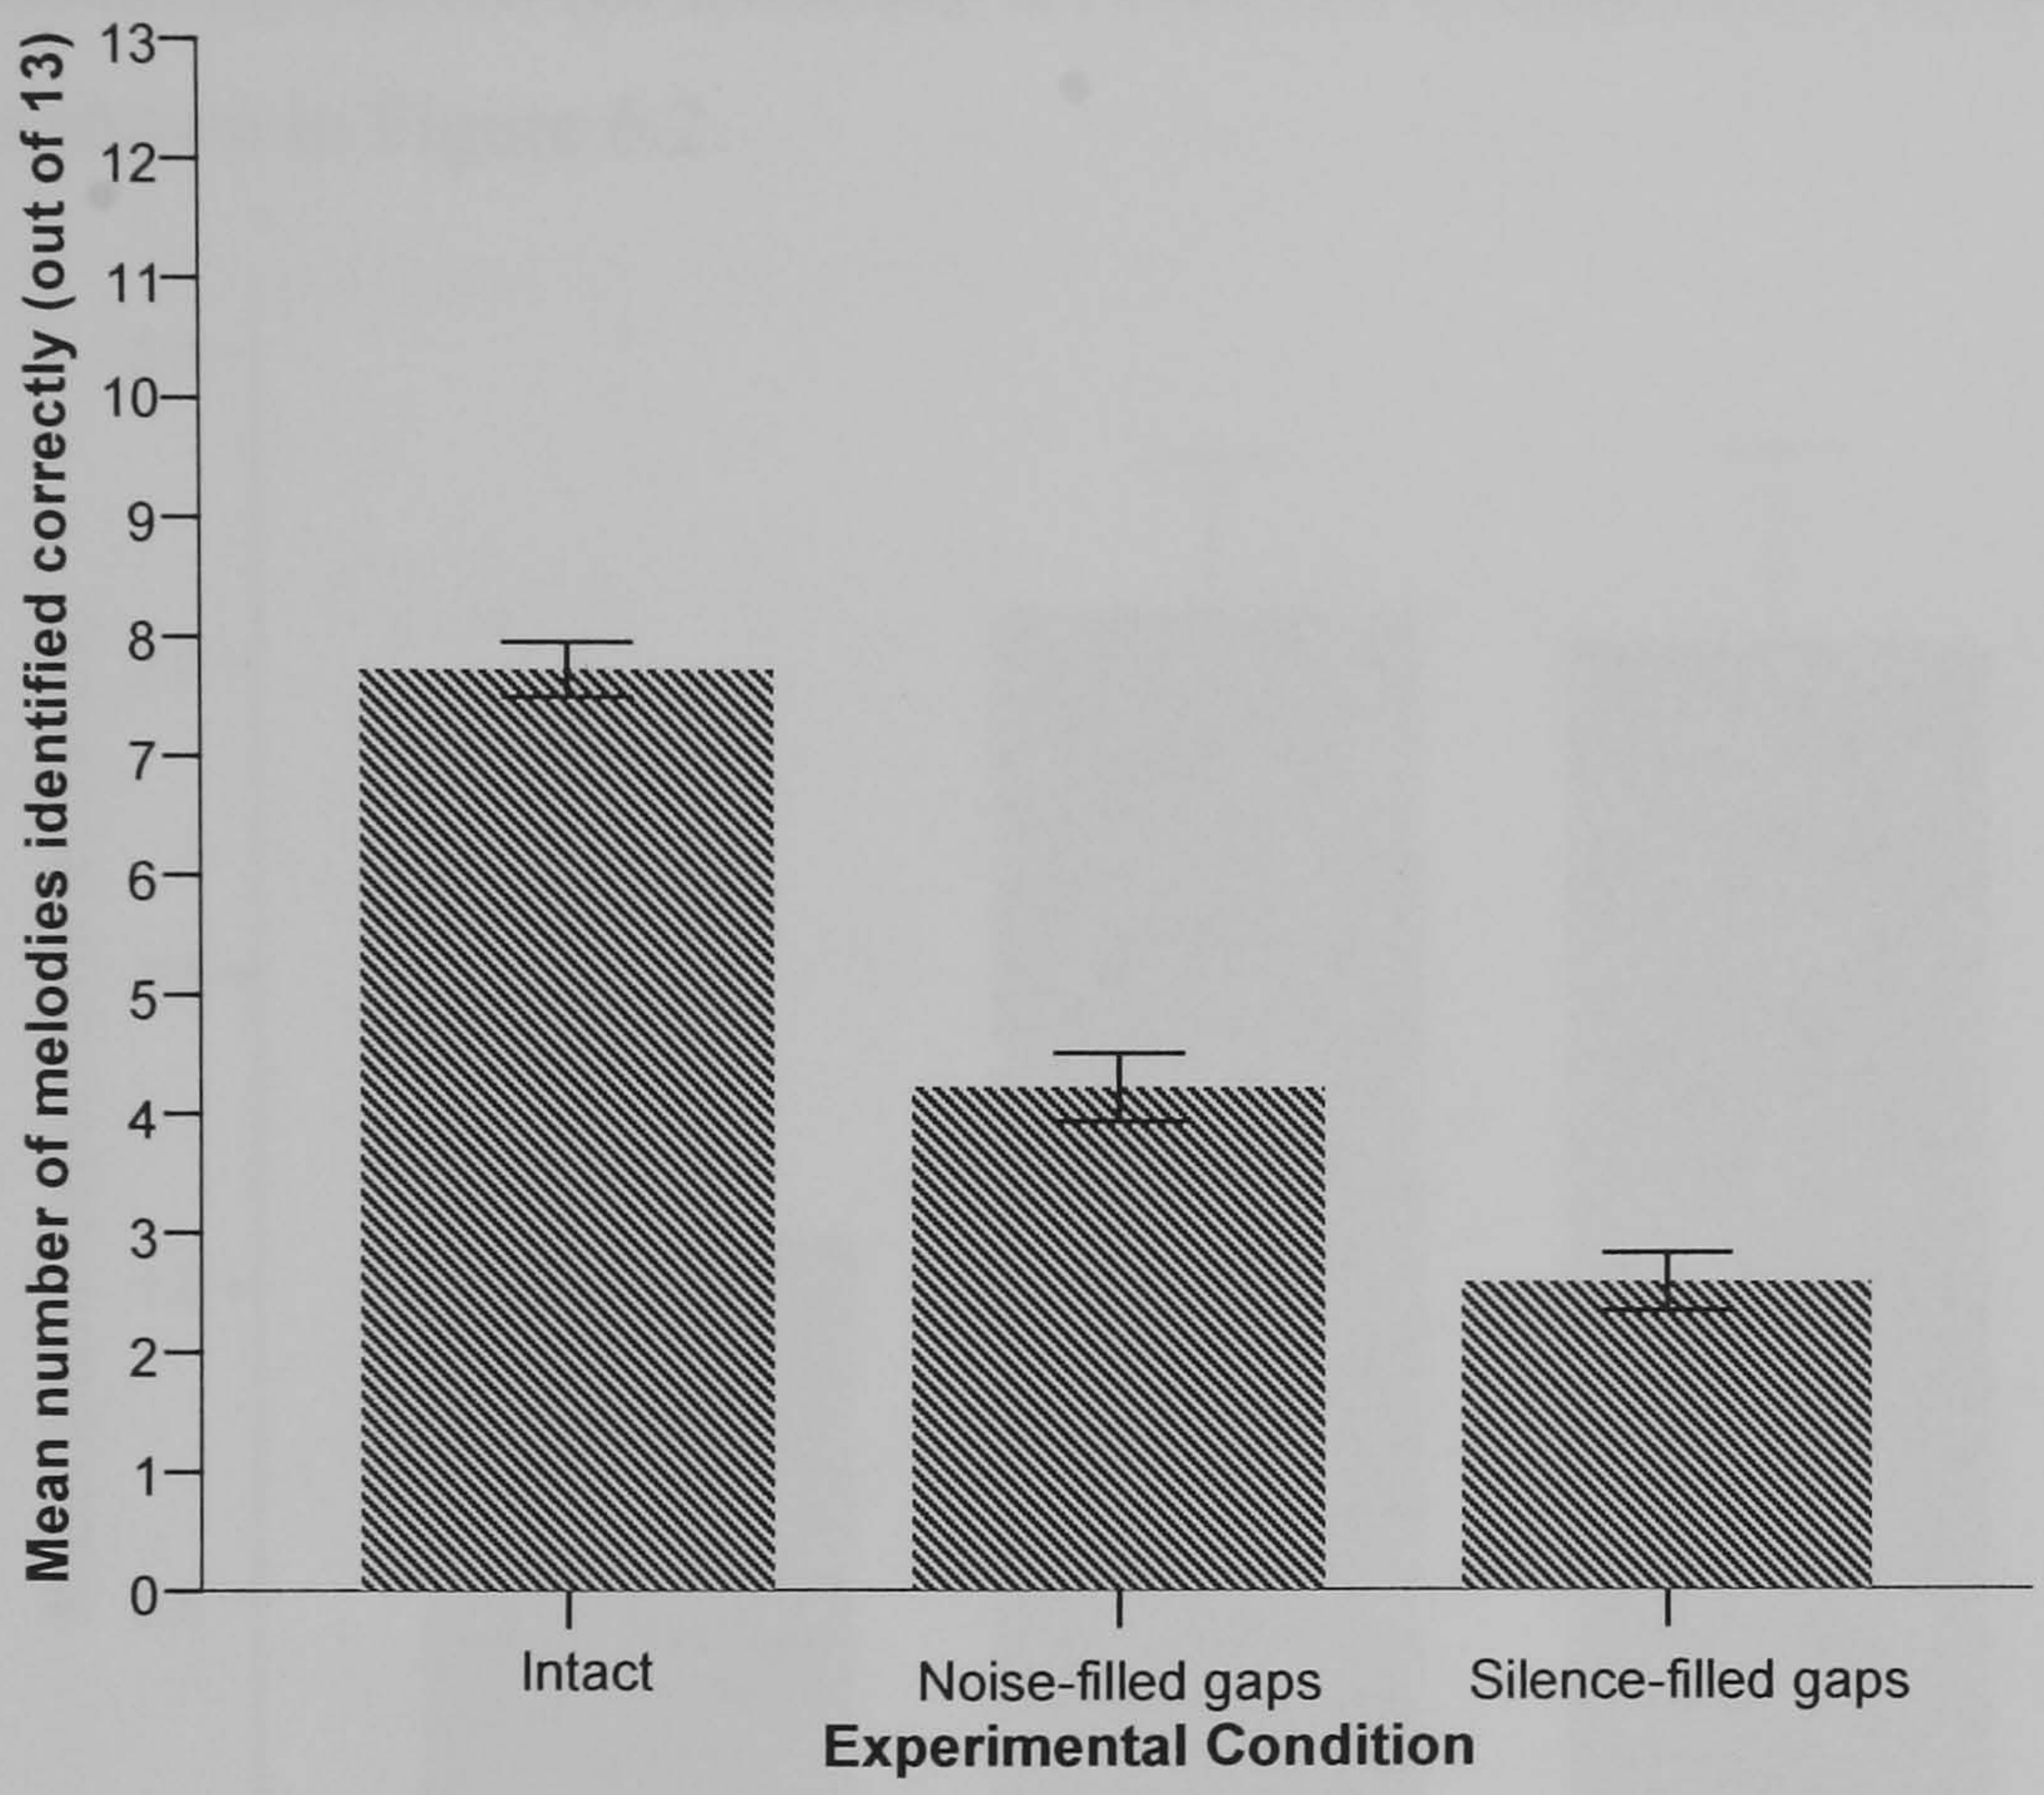

Figure 6.1. Mean number of melodies identified correctly (± *SE*) by children in the three experimental conditions (intact, noise-filled gaps and silence-filled gaps).

These data were analysed using a one-way between-subjects ANOVA, with the three-level factor of experimental condition (intact, noise-filled gaps, or silence-filled gaps). There was a significant effect of experimental condition on identification scores, $F(2, 63) = 106.57, p < .001, \eta^2 = .77$. Post hoc tests (Tukey HSD) revealed significant differences between all pairs of conditions: intact versus noise-filled gaps, $p < .001$, $d = 2.61$; intact versus silence-filled gaps, $p < .001$, $d = 4.50$; noise-filled gaps versus silence-filled gaps, $p < .05$, $d = 1.22$. Identification of melodies was best in the intact condition, but children in the noise-filled gaps condition identified significantly more melodies than children in the silence-filled gaps condition.

Identification time

In keeping with previous experiments, the measure of identification time was taken to be the difference between the stimulus duration and the time taken to identify the stimulus. No gender differences in identification time were found (see Appendix 2) so data were collapsed across gender. The mean identification time for correctly identified stimuli (in seconds), for children in each of the three experimental conditions, is shown in Figure 6.2.

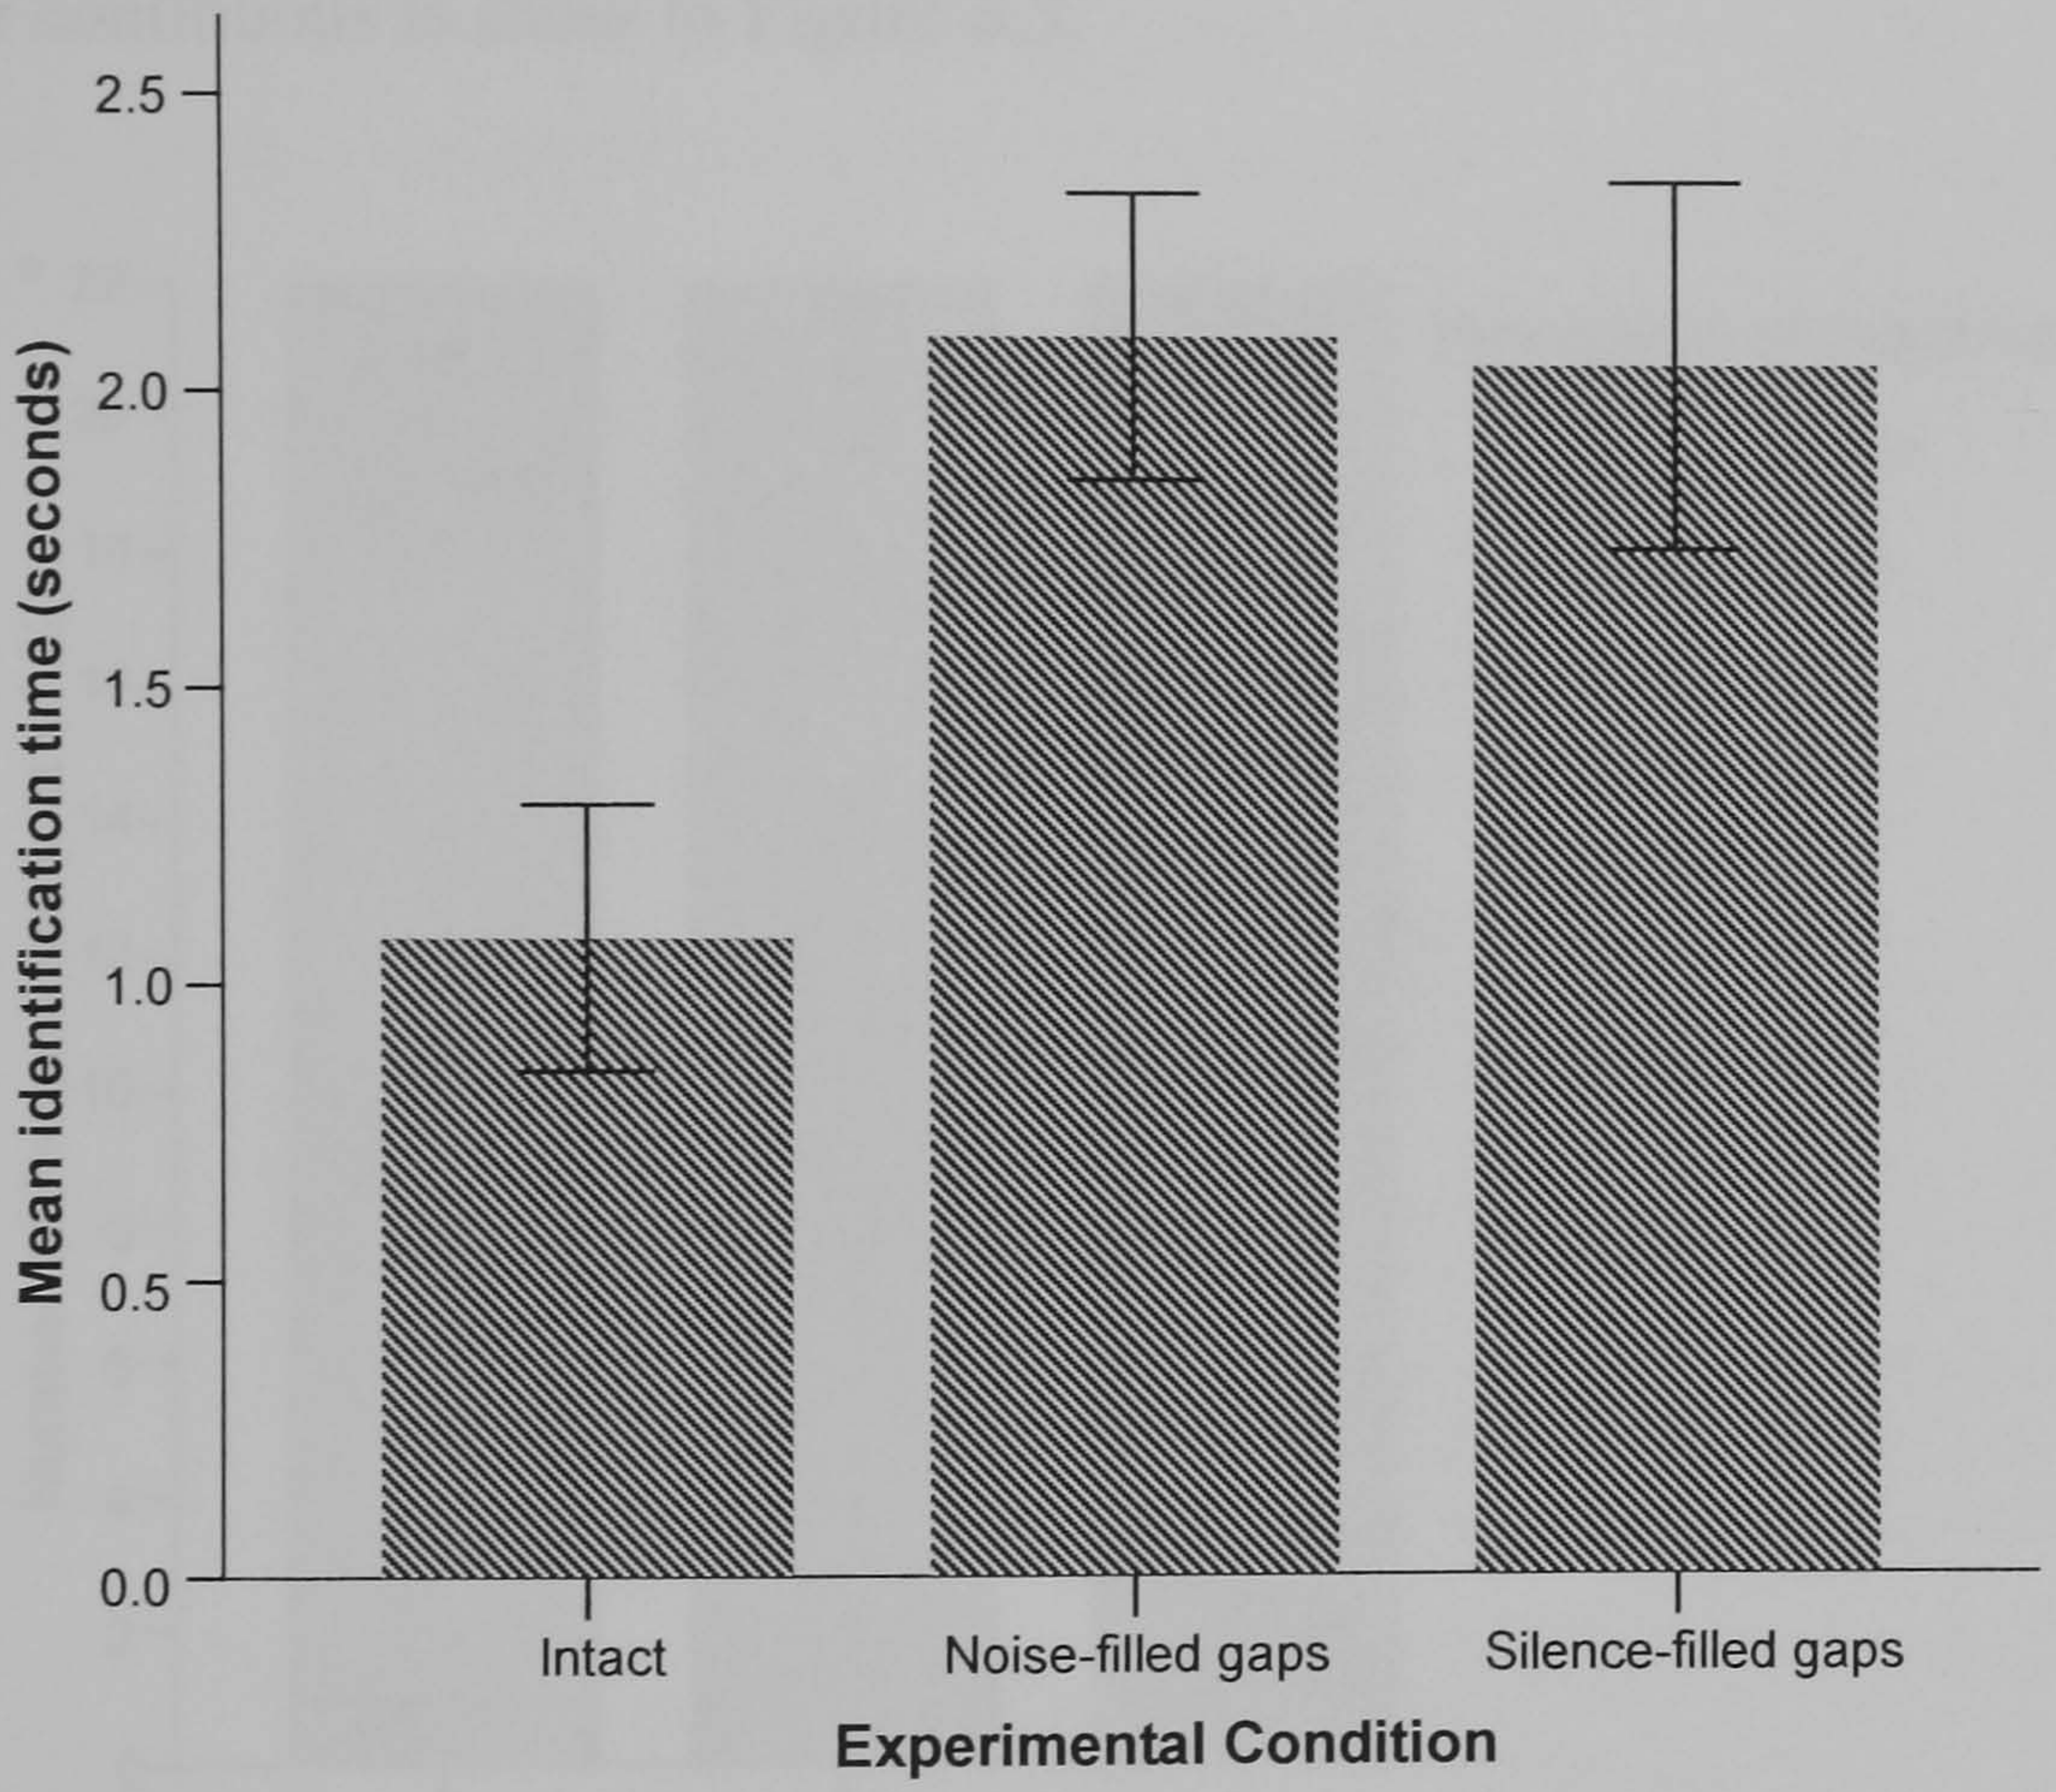

Figure 6.2. Mean identification time (identification time − stimulus duration; ± *SE*) in seconds for children in the three experimental conditions (intact, noise-filled gaps or silence-filled gaps).

A one-way between-subjects ANOVA with the three-level factor of experimental condition (intact, noise-filled gaps, or silence-filled gaps) showed a significant effect of experimental condition on identification time, *F*(2, 63) = 4.93, $p < .05$, $\eta^2 = .14$. Post hoc tests (Tukey HSD) revealed that stimuli in the intact condition were identified significantly faster than those in both the noise-filled gaps and silence-filled gaps conditions ($p < .05$, $d = 0.67$ and $p < .05$, $d = 0.90$, respectively). Identification times in noise-filled gaps and silence-filled gaps conditions did not significantly differ from each other ($p > .05$).

Perception of continuity

Children were asked whether they thought that little pieces of the tune were missing (fragmented signal) or all of the tune was there (continuous signal). The total number of children in each condition reporting that the melodies were either fragmented or continuous is show in Figure 6.3.

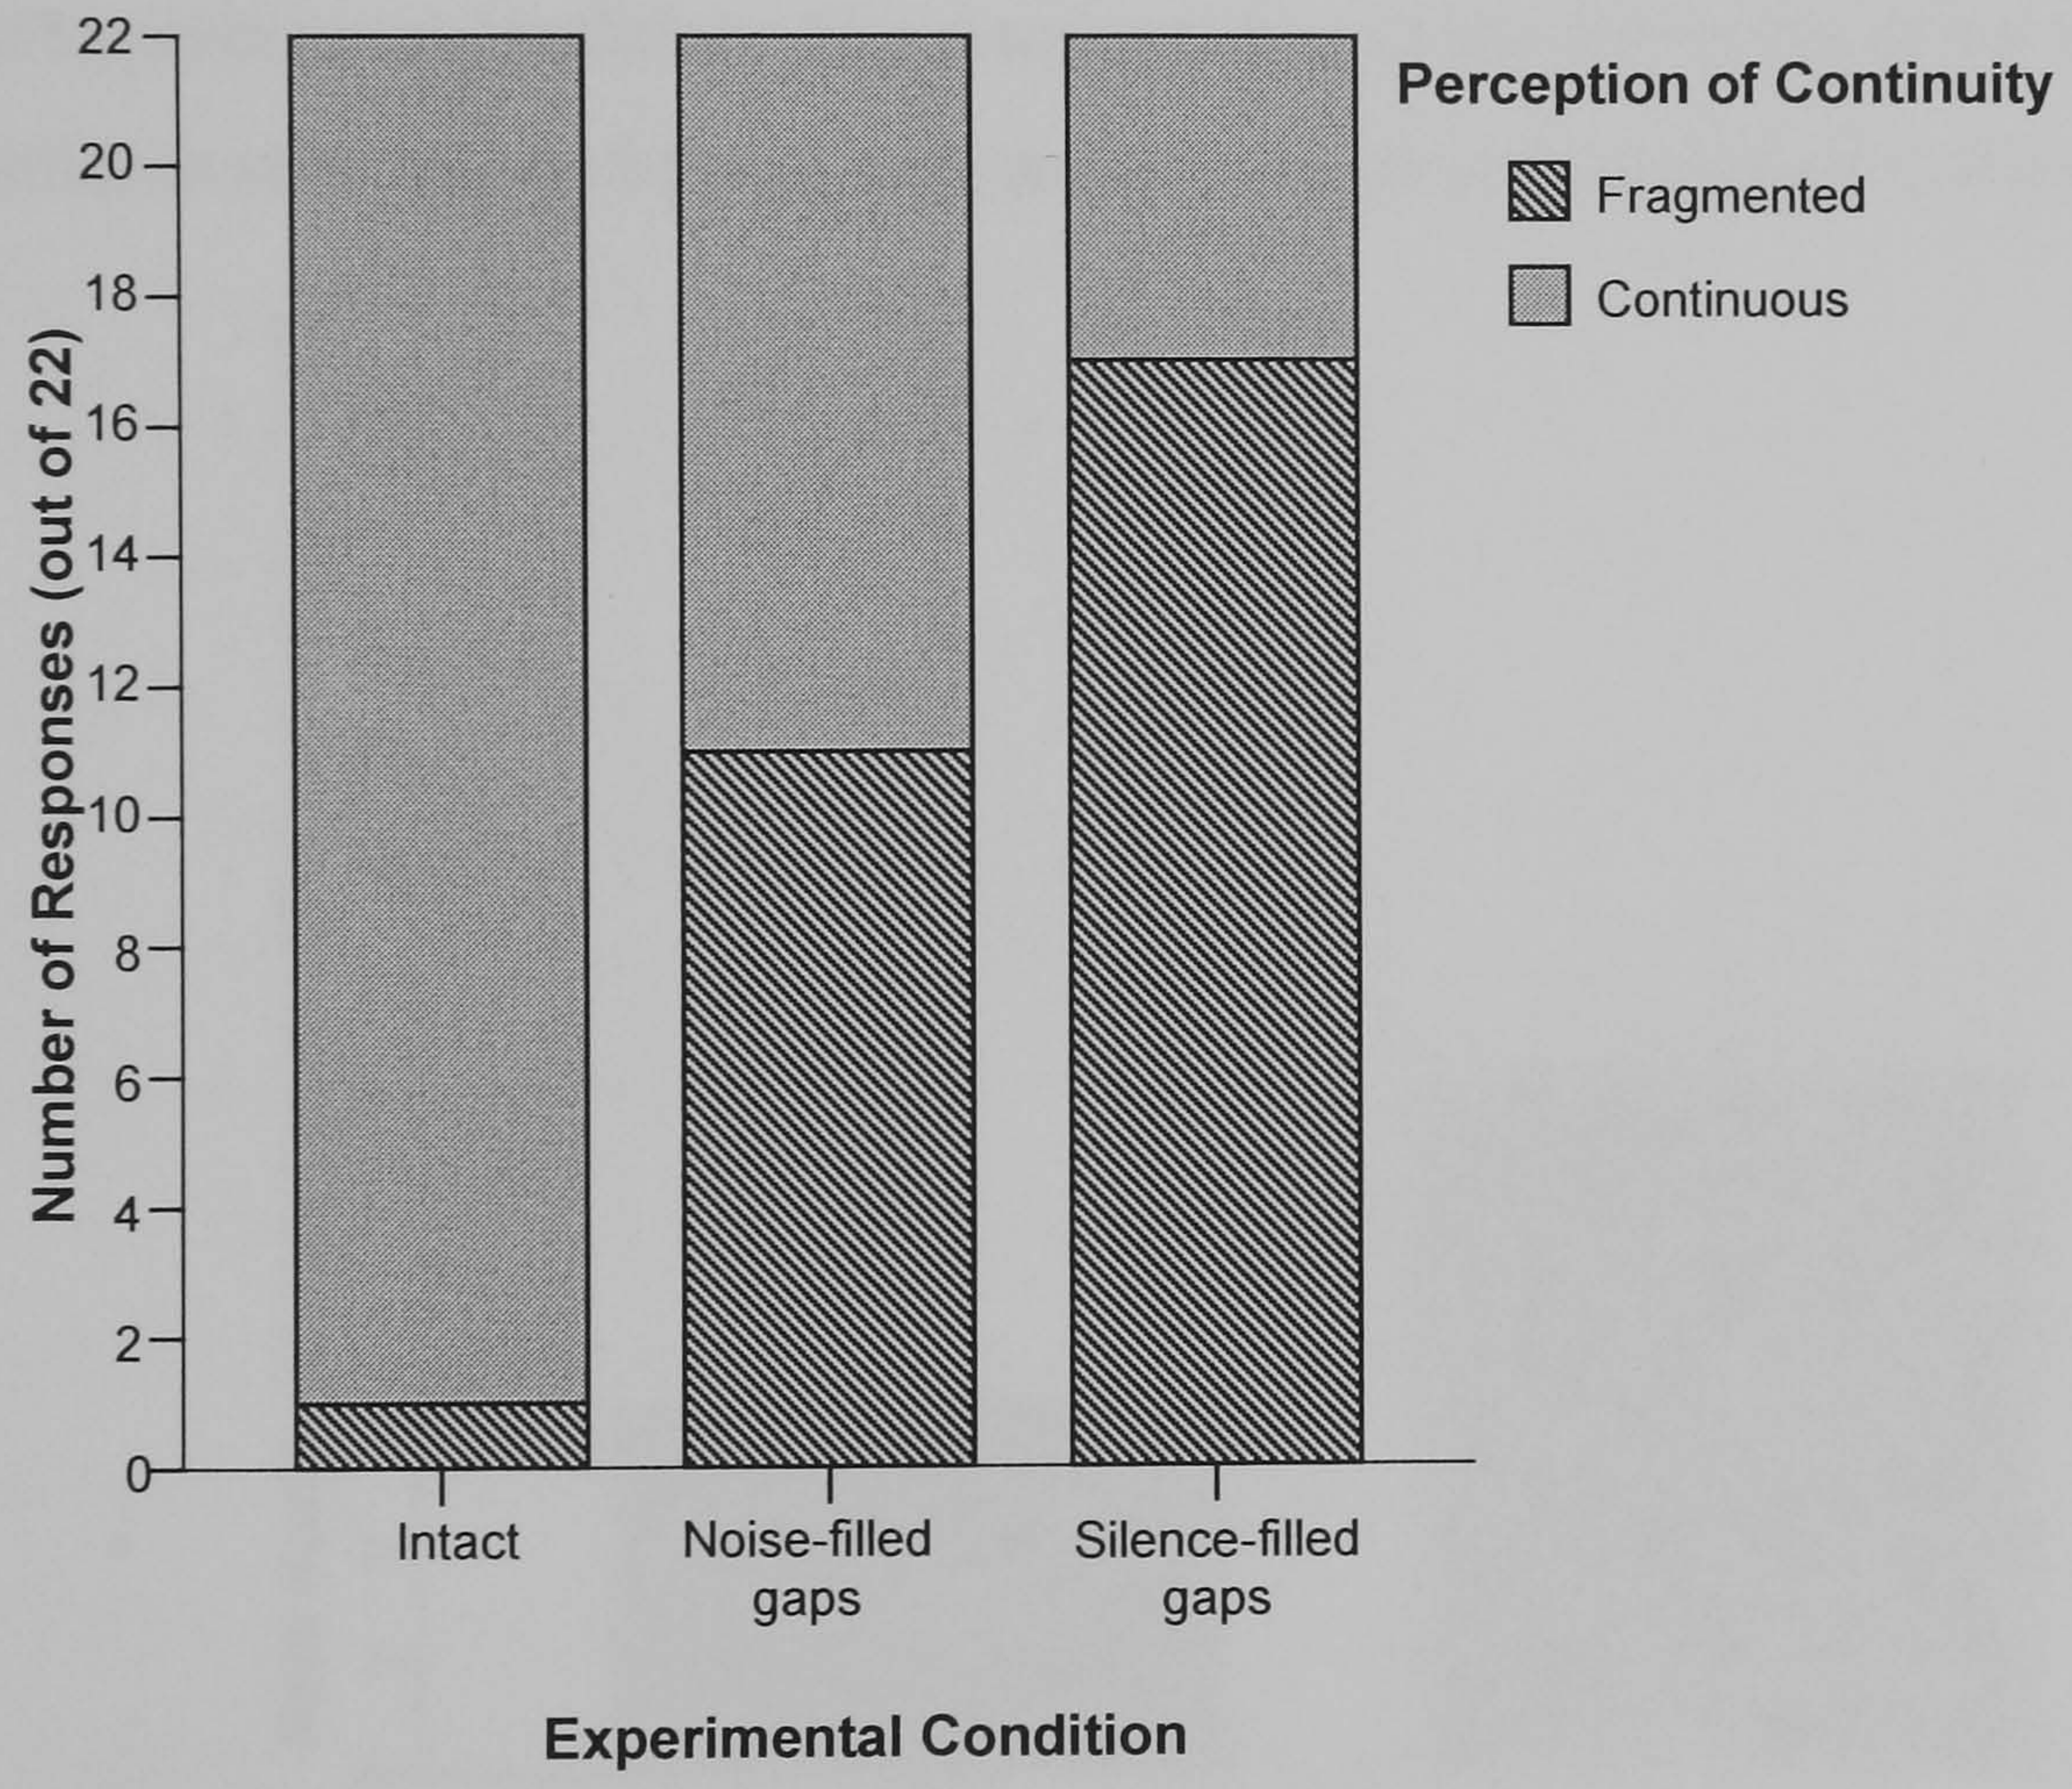

Figure 6.3. Total number of children in each condition (intact, noise-filled gaps, silence-filled gaps) reporting the melodies as fragmented and reporting the melodies as continuous.

These data were analysed using a 3 (experimental condition: intact; noise-filled gaps; silence-filled gaps) × 2 (perception of continuity: continuous; fragmented) chi-square. A significant association between experimental condition and perception of continuity was found, χ^2 (2, $N = 66$) = 24.11, $p < .001$, Cramer's $V = .60$. As Figure 6.3 shows, the majority of children in the intact condition reported the signals to be continuous, and the majority of children in the silence-filled gaps condition reported the signals to be fragmented. Exactly half of the children in the noise-filled gaps condition reported the signals to be continuous, that is, they did not think that anything was missing, whilst the other half reported the signals to be fragmented.

Based upon this previous analysis, children were categorised as either reporting the signals to be fragmented or continuous, in order to test whether there is a difference in children's identification scores according to their perception of continuity, regardless of experimental condition. The mean number of melodies identified correctly by children who reported the signals as continuous or fragmented is shown in Figure 6.4.

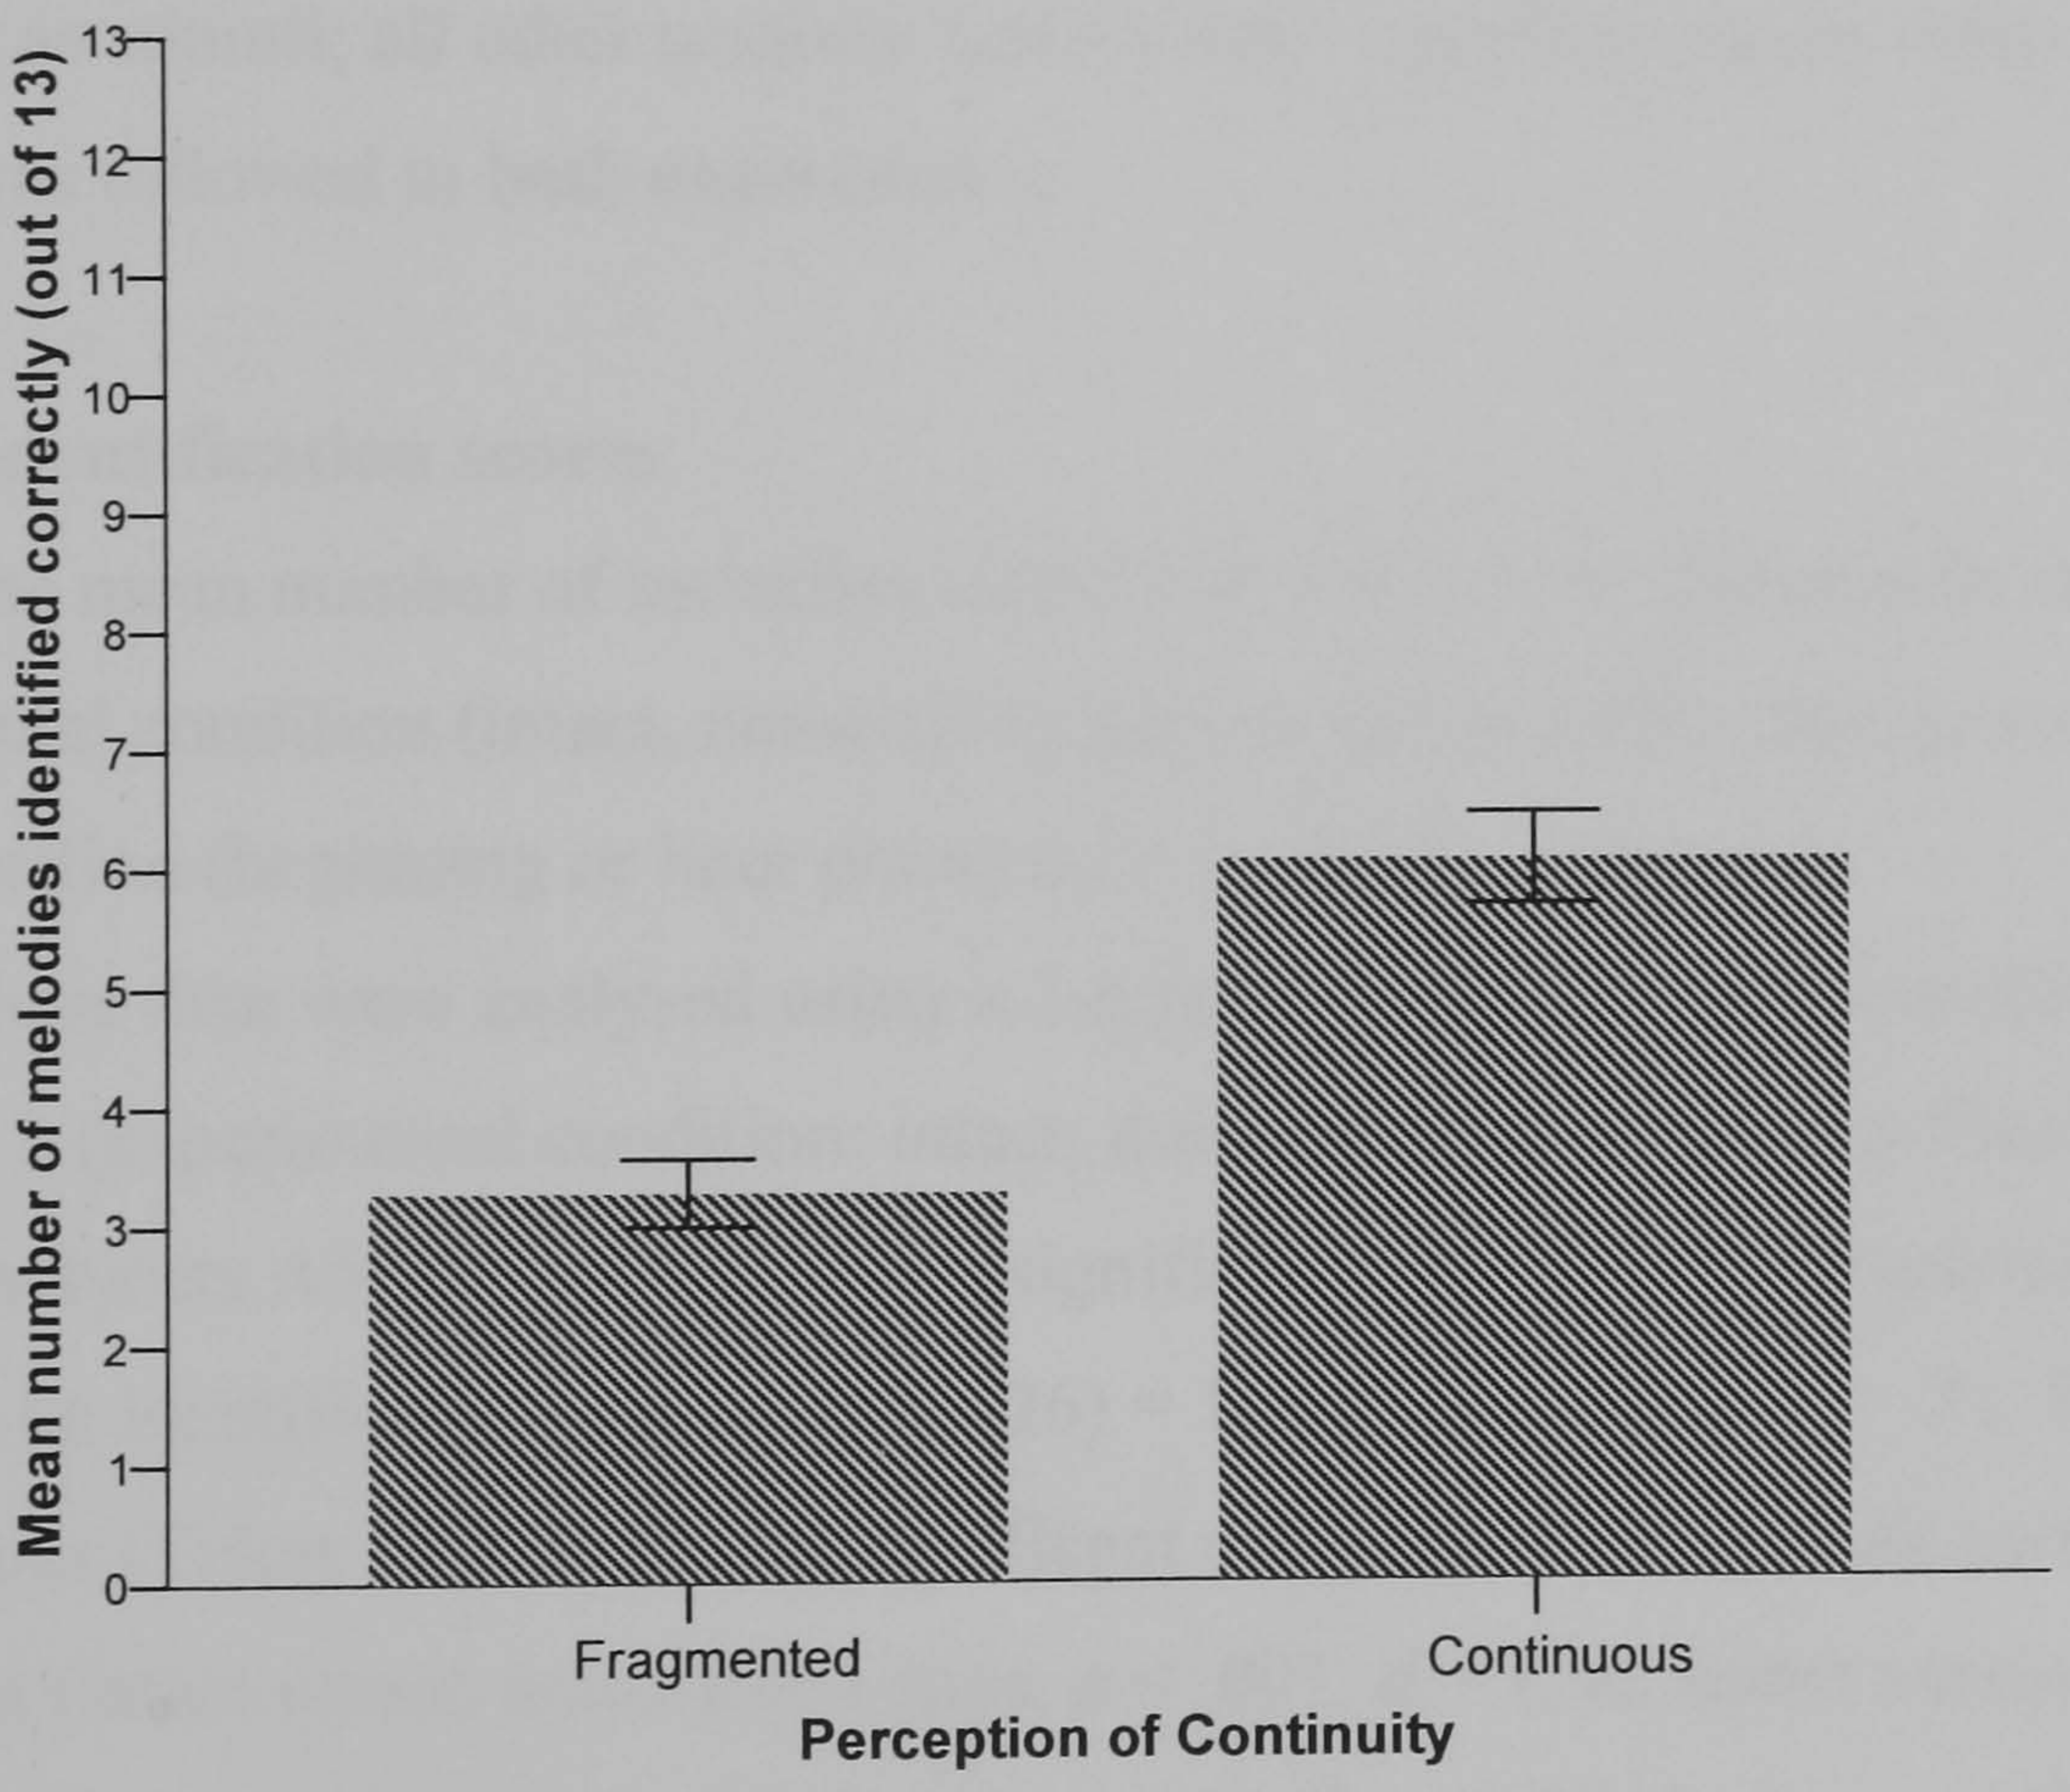

Figure 6.4. Mean number of melodies identified correctly (± *SE*) by children who reported the melodies to be fragmented and those who reported the melodies to be continuous.

An independent-samples *t* test showed that those children who reported the signals as continuous identified significantly more melodies than children who reported the signals as fragmented, $t(62) = 5.85, p < .001, d = 1.19$.

Taken together, these analyses have shown that perceptual restoration does occur for later phrases of a musical representation. It is therefore important to compare the perceptual restoration of these later phrases with perceptual restoration of the initial phrases of the nursery rhymes, as tested in Experiment 1.

Comparison between Experiment 1 and Experiment 4: The perceptual restoration of initial and later phrases of familiar melodies

This section reports a comparison between the present experiment and Experiment 1 where the beginning phrases of the same nursery rhyme melodies were presented as stimuli; all other acoustic factors were identical and the same experimental method was followed in both experiments.

Correct identification scores

The mean number of melodies identified correctly by children in each experimental condition (intact, noise-filled gaps or silence-filled gaps), and for each phrase position (beginning or later phrases), is shown in Figure 6.5.

These data were analysed using a 2 (phrase position: beginning phrase; later phrase) × 3 (experimental condition: intact; noise-filled gaps; silence-filled gaps) between-subjects ANOVA. There was a significant main effect of experimental condition on identification scores, $F(2, 126) = 150.65, p < .001, \eta^2 = .71$. Post hoc comparisons (Tukey HSD) showed significant differences between all pairs of conditions (intact versus noise-filled gaps, $p < .001, d = 1.42$; intact versus silence-filled gaps, $p < .001, d = 3.70$; noise-filled gaps versus silence-filled gaps, $p < .001, d = 1.26$).

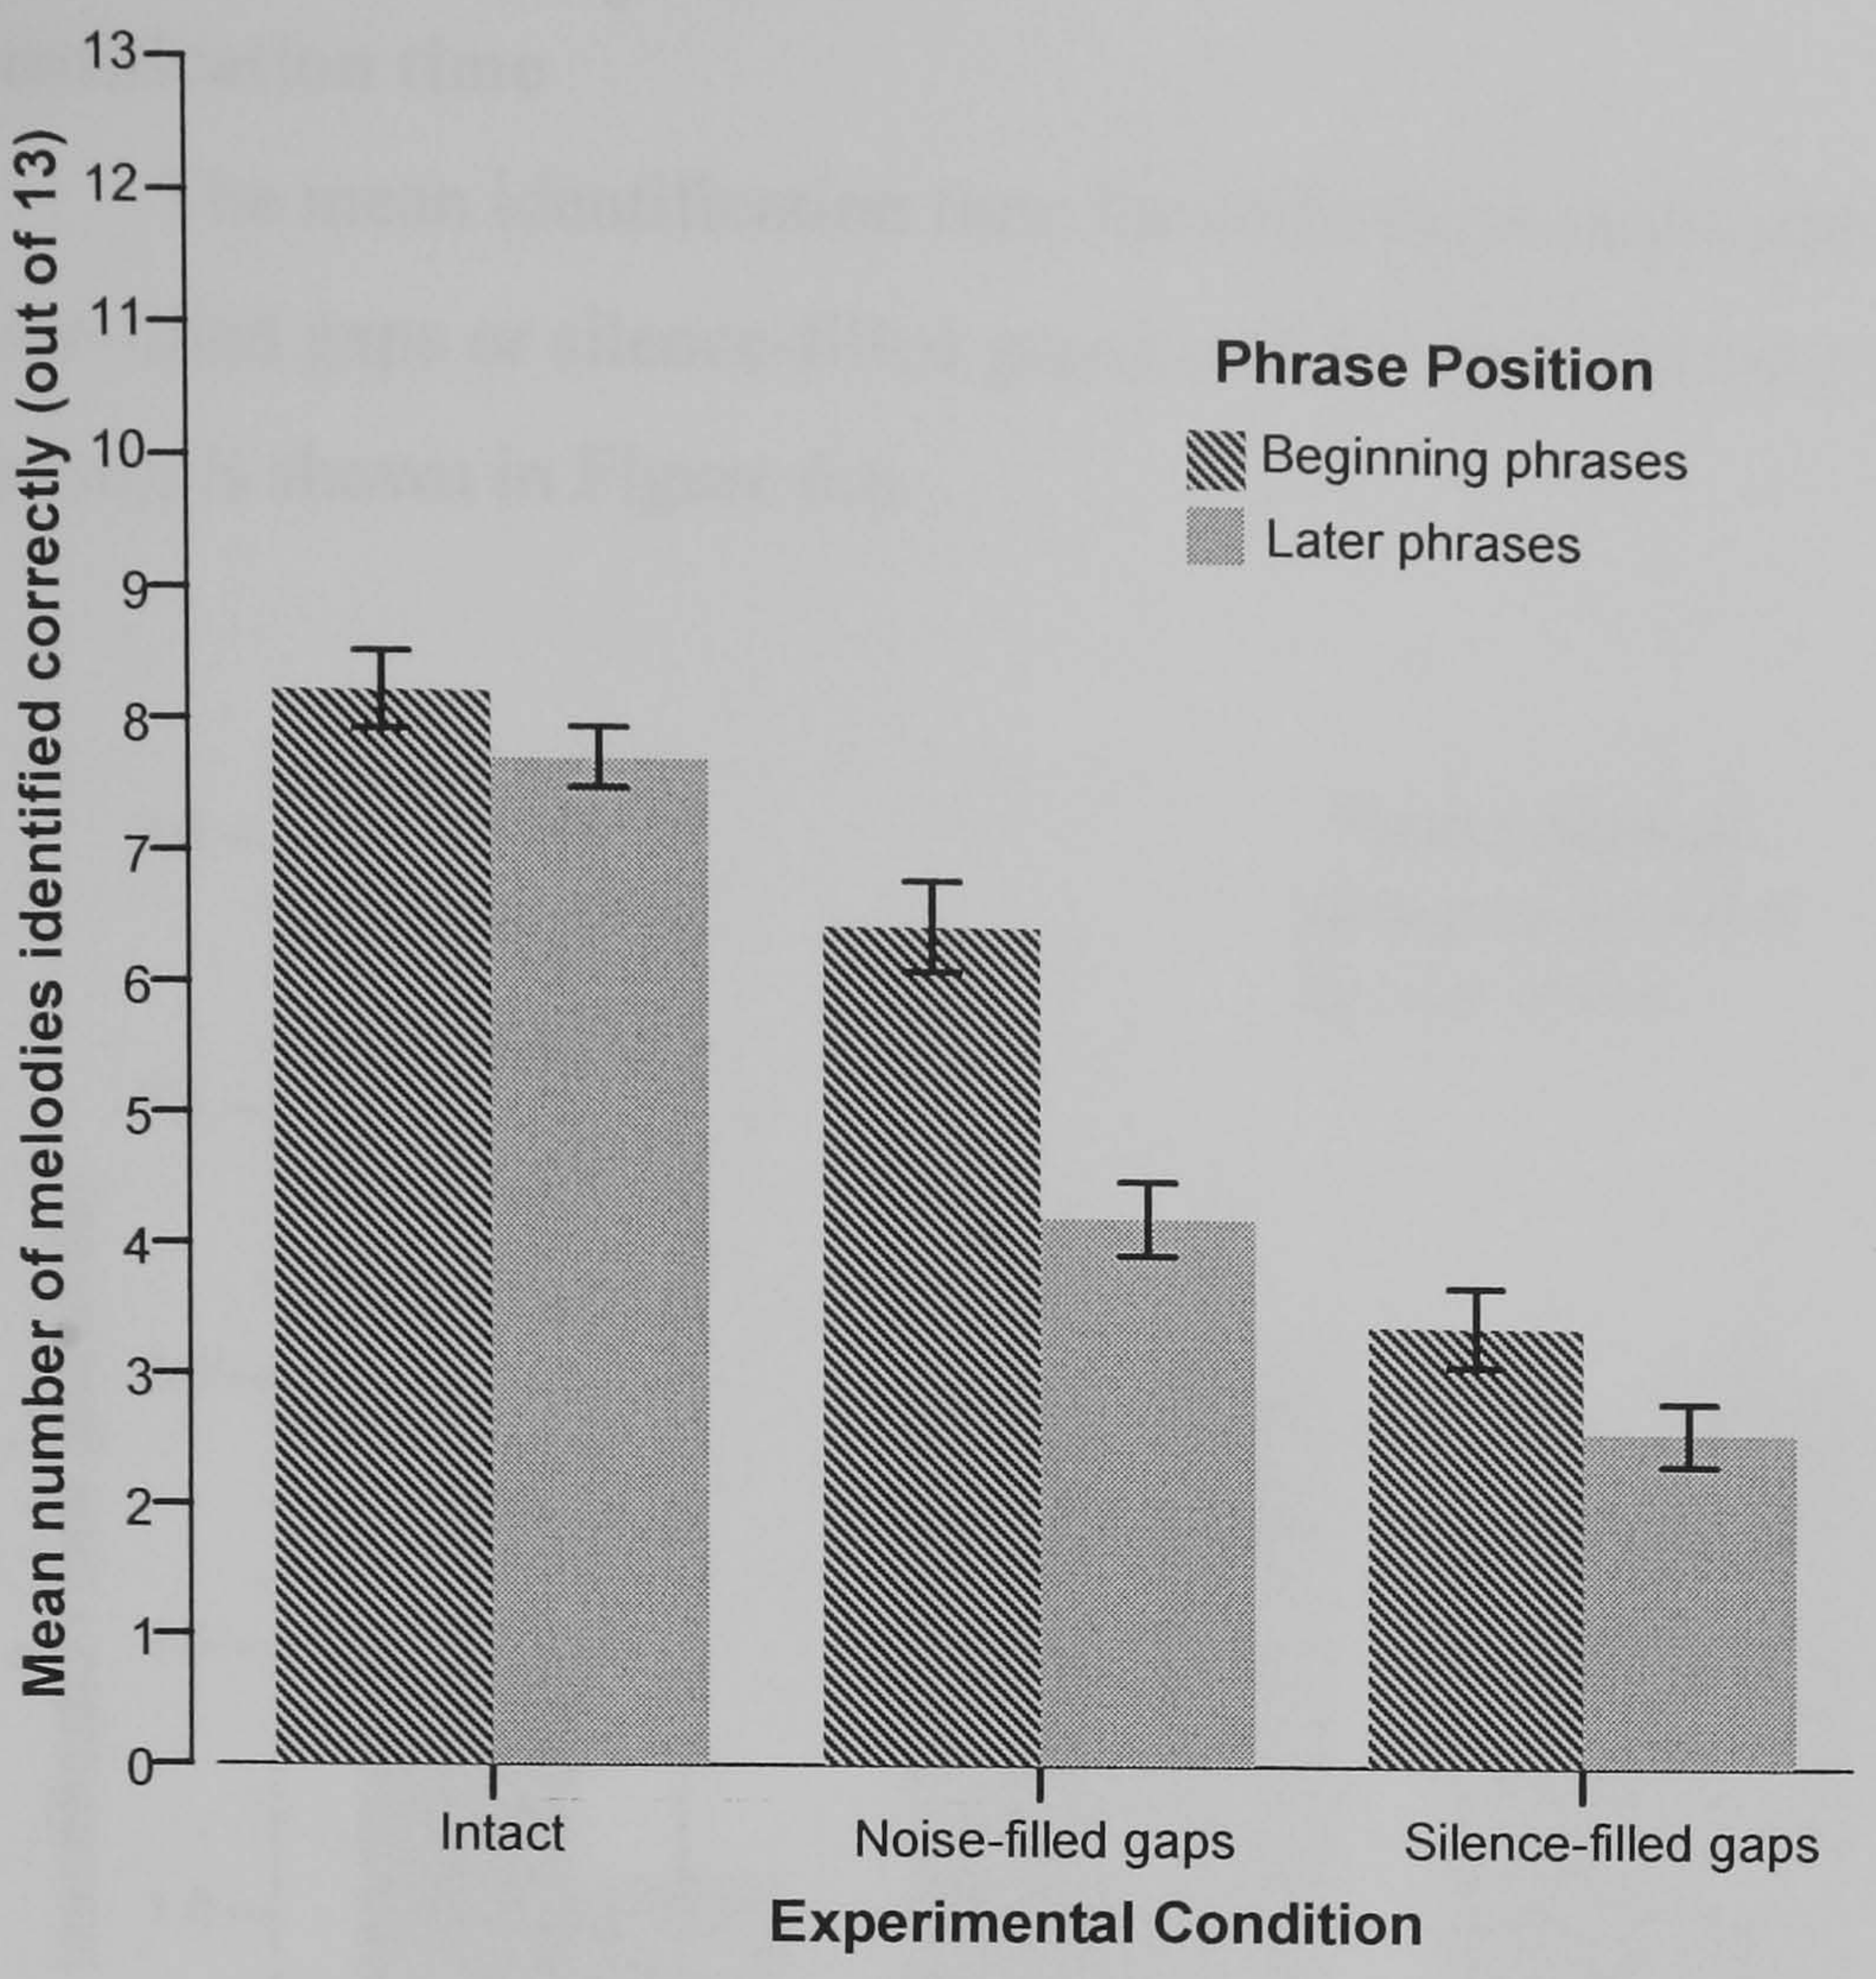

Figure 6.5. Mean number of melodies identified correctly (± *SE*) by children in the three experimental conditions (intact, noise-filled gaps and silence-filled gaps), and in the two phrase position conditions (beginning phrases and later phrases).

There was also a significant main effect of phrase position, *F*(1, 126) = 25.41, $p < .001$, $\eta^2 = .17$. As Figure 6.5 shows, overall identification scores were higher for beginning phrases than later phrases of the music. These main effects were qualified by a significant Phrase position × Experimental condition interaction: *F*(2, 126) = 5.13, $p < .01$, $\eta^2 = .08$.

Further analysis of this interaction using one-way ANOVAs revealed that there was only a significant effect of phrase position for the noise-filled gaps conditions, *F*(1,126) = 9.03, $p < .05$, $\eta^2 = .07$. There was no significant effect of phrase position for the intact or silence-filled gaps conditions, *F*(1, 126) = .45, $p > .05$ and *F*(1, 126) = 1.22, $p > .05$, respectively. This suggests that the position of the phrase within the nursery rhyme only significantly affected identification in the conditions where perceptual restoration can occur.

Identification time

The mean identification time for children in each experimental condition (intact, noise-filled gaps or silence-filled gaps), and for each phrase position (beginning or later phrase), is shown in Figure 6.6.

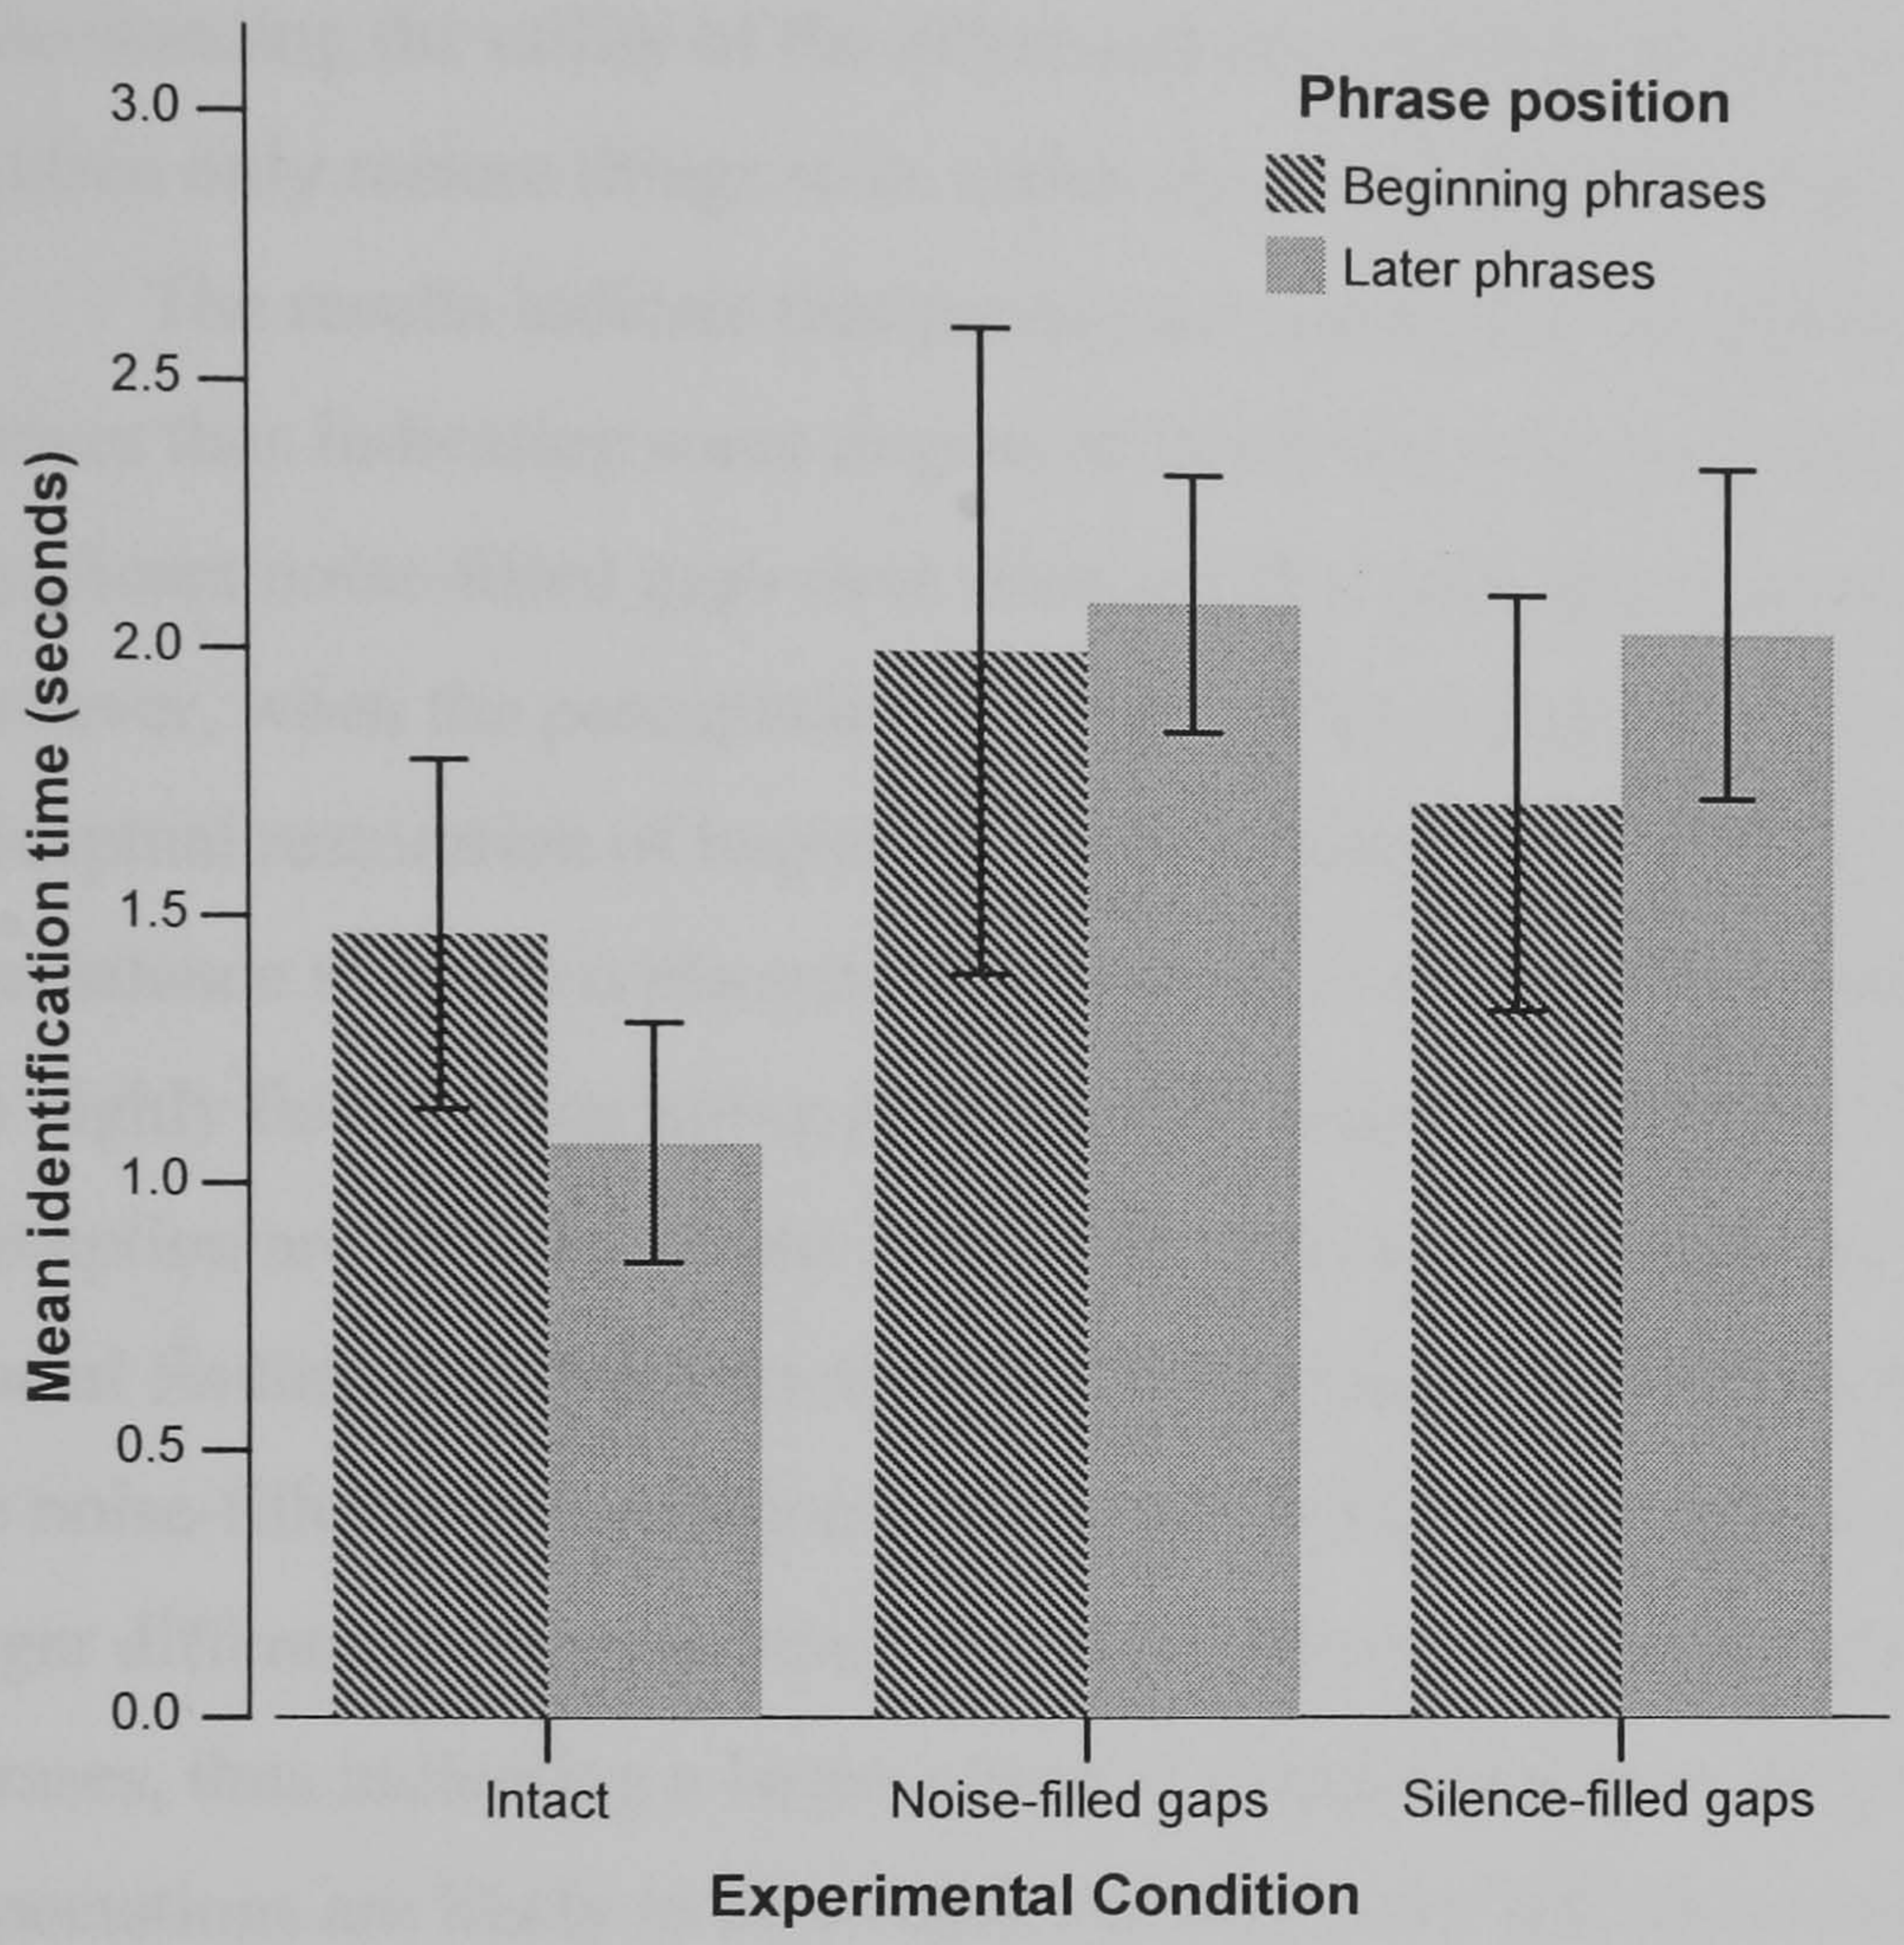

Figure 6.6. Mean identification time (identification time – stimulus duration; ± *SE*) in seconds for children in the three experimental conditions (intact, noise-filled gaps and silence-filled gaps), and in the two phrase position conditions (beginning and later phrases).

A 3 (experimental condition: intact; noise-filled gaps; silence-filled gaps) × 2 (phrase position: beginning phrase; later phrase) between-subjects ANOVA showed no significant main effect of experimental condition on identification time, $F(2, 126) = 2.78, p > .05$; the effect of phrase position on identification time was also non-significant, $F(1, 126) = .49, p > .05$. There was also no significant interaction between the factors, $F(2, 126) = 1.10, p > .05$.

Discussion

The main aim of this experiment was to investigate familiarity effects on the perceptual restoration of music in children, by comparing the perceptual restoration of highly familiar beginning phrases with phrases occurring later in the nursery rhyme that are less typical and thus are likely to be less familiar. Familiarity effects are important in understanding the utility of the effect and any limits to its operation; for example, can children only restore things with which they are highly familiar?

The results indicate that perceptual restoration did operate for later melodic phrases thus indicating some degree of flexibility in perception, as shown by a significant noise-filled gaps over silence-filled gaps advantage in identification. However, when the perceptual restoration of later phrases was compared to the perceptual restoration of beginning phrases from Experiment 1, the advantage of noise over silence as a gap replacement was small in comparison to perceptual restoration of the highly familiar beginning phrases. This suggests that top-down influences on perception are reduced where a less typical section of a piece of music is presented. The crucial finding was that phrase position only had a significant effect on identification in the noise-filled gaps conditions, where perceptual restoration occurs. There was also a larger difference between intact and noise-filled gaps conditions for later than beginning phrases, thus indicating a larger effect of signal disruptions in general where expectations are likely to be weaker and top-down influences reduced.

These findings have important implications for theories of perceptual restoration (e.g. Samuel, 1981a; Samuel & Ressler, 1986). Such theories commonly include the operation of both top-down and bottom-up processes. The data from the present experiment support such theories; because the effect of familiarity here only affected the noise-filled gaps condition where perceptual restoration occurs, this suggests that the interaction between expectation and confirmation is critical in the operation of familiarity effects. That the additional top-down influences operating for typical rather than non-typical melodic phrases had no significant effect on identification in the silence-filled gaps conditions indicates that the top-down effects were operating in conjunction with the effects of bottom-up confirmation of the increased expectations. Furthermore, the present findings cannot simply be explained as a result of the beginning sections being easier to identify. If this were the only influence operating, the effect of phrase position should have a strong effect on melody identification in the intact conditions and it does not.

To the extent that these findings can be taken as evidence of familiarity effects on the perceptual restoration of music in children, these findings are also consistent with previous research into the perceptual restoration of both music and speech signals, as well as theories of perceptual restoration in general. Research in the speech domain (e.g. Bashford et al., 1992; Bashford & Warren, 1987b; Samuel, 1981a, 1987) together with work on the perceptual restoration of music in adults (DeWitt & Samuel, 1990) found evidence of top-down effects on perceptual restoration. In the case of the present experiment, the effect of phrase position can be best explained by the proposition that the beginning section of a piece of music is important in the generation of expectations about what is likely to occur next in the piece of music (Schulkind, 2004). Having strong expectations about what is likely to occur in an acoustic event, such as listening to a melody, is crucial in making sense of disrupted or ambiguous input and is an important part of the perceptual restoration process (Samuel, 1981a). The strength of expectations appears to play a less important role in identification of undisrupted signals, at least for the nursery rhyme melodies presented here.

The time taken to identify the later melodic phrases was significantly faster in the intact than the noise-filled gaps and silence-filled gaps conditions. This result suggests that identification of disrupted stimuli is much slower for the less typical representations. However, there was no significant difference in the time taken to identify beginning and later melodic phrases. The data indicate that for noise-filled gaps and silence-filled gaps conditions, it takes longer to identify later melodic phrases than beginning melodic phrases. However, there are clear individual differences within each experimental condition in terms of identification time.

The present experiment has demonstrated that perceptual restoration operates for later, less typical melodic phrases when missing sections are replaced by noise; this serves to enhance intelligibility of melodies when compared to silence as a gap replacement. However, children's subjective reports of whether they thought the signals sounded continuous or fragmented indicated that only half of the children in the noise-filled gaps condition had the experience of the signals being continuous with nothing missing. It is clear that, over all experimental conditions, having this perception can be associated with superior melody identification, as shown by the significant identification advantage for children who perceived the signals as continuous when compared to children who perceived the signals as fragmented.

Whilst, on first inspection, these data might appear to show little support for propositions that the restoration of intelligibility depends on the perception of a continuous signal (e.g. Bashford et al., 1992), it is crucial to remember that the continuity data here represent responses from children who were presented with the later, less typical melodic phrases only. There is no comparable data for children who were presented with the beginning, more typical melodic phrases. However, perceptual restoration, as shown by identification scores, was reduced for later phrases when compared to beginning phrases. Therefore, this might indicate that the reduced expectations for later phrases are reflected in both the intelligibility and continuity components of the effect. Alternatively, the relationship between continuity and intelligibility, for signals such as speech and music, might be more complex than the model proposed by Bashford et al. (1992) suggests. There is evidence in the literature to suggest that noise can produce an intelligibility advantage relative to its absence, without the accompanying perception that nothing is missing (Powers & Wilcox, 1977). Thus, whilst the data from this experiment are still consistent with the operation of two separate processes in perceptual restoration, it does question the relationship between them.

The aim of this experiment was to investigate the influence of familiarity on the perceptual restoration of music in children, and the empirical evidence presented supports the operation of such an influence. This suggests that children's music perception, like that of adults, is both active and interactive. Not only can children reconstruct missing fragments of musical stimuli, in doing so, they are influenced by both their expectations and the acoustic information contained within the stimulus.

It is clear from the present experiment that familiarity does have an effect on perceptual restoration. Therefore, it might also be possible to increase the strength of perceptual restoration through other familiarity manipulations, such as providing two sources of familiarity rather than just one. The next experiment investigates this possibility, by adding lyrics to the melodic line of the nursery rhymes.

Chapter 7: Experiment Five

The perceptual restoration of songs in children

Introduction

Experiments 1 to 4 have demonstrated the perceptual restoration of musical melody in children. In these experiments, children between the ages of 4 and 6 years were better able to identify nursery rhyme melodies when missing fragments were replaced with masking noise, when compared to nursery rhyme melodies with the same fragments replaced by silence. In demonstrating that perceptual restoration does operate in children's music perception, these results can be taken as evidence of active, rather than passive music perception in children. Furthermore, the results suggest that the perceptual system attempts to maintain global coherence of disrupted input where possible.

However, not all music is purely melodic. Song represents one of the most popular musical forms (Bonnel, Faita, Peretz & Besson, 2001), and songs play a key role in children's learning and development, particularly in literacy development (Butzlaff, 2000; Hansen & Bernstorf, 2002; Smith, 2000). Therefore, it would be highly beneficial if a perceptual restoration mechanism operated to restore coherence of disrupted song-based input in young children.

When a representation of a familiar song is stored, this is comprised of two separate components: one melodic (the pitch and rhythm) and one semantic (the nature of the lyrics). The aim of the present experiment was to compare perceptual restoration of songs, which possess these two components together, with perceptual restoration of melody, where only the melodic component is present. Song recognition is described by Dalla Bella, Peretz & Aronoff (2003) as requiring a "mental confirmation" of what is currently being perceived, with stored representations in long-term-memory. For perceptual restoration to operate, therefore, the current perception must be restored to the degree where it can successfully be matched with the stored representation.

The integration of lyrics and melody in song memory

There is reason to expect greater perceptual restoration for songs than melodic music. Perceptual restoration is influenced by the familiarity of the stimulus; because a more familiar stimulus generates greater expectations about what is likely to be the content of the disrupted information, perceptual restoration is stronger. It is plausible to suppose that by providing the listener with the lyrics of the songs as well as the melody, there are two sources of familiarity that can generate expectations about the content of the missing fragments. This would plausibly affect the noise-filled gaps over silence-filled gaps advantage (i.e. perceptual restoration) in some specific way, rather than just improving stimulus identification across conditions. Put another way, does the addition of lyrics have the same effect on identification across conditions?

There is also evidence from the music perception literature that leads to an expectation that perceptual restoration will be stronger for songs rather than for an isolated melodic component. Research has indicated that lyrics and melody are integrated in adults' memory for songs, and that it is much easier to recall a song from memory when both of these components are presented together (Crowder, Serafine & Repp, 1990; Serafine, Crowder & Repp, 1984; Serafine, Davidson, Crowder & Repp, 1986). This occurs on the basis of association-by-contiguity (Crowder, Serafine & Repp, 1990). As songs are presented repeatedly, the melody and lyrics are heard together. This leads to tight associations between them in memory, with the result that during recall one component can activate the other (Bartlett & Snelus, 1980; Rainey & Larsen, 2002; Steinke, Cuddy & Jakobson, 2001). Therefore, in the case of disrupted input, it is possible that one component could help reduce the ambiguity of the other and hence assist compensatory processes such as perceptual restoration.

A further finding from the work by Serafine et al. (1984) is that generally, lyrics are easier to remember than melodies. They suggest that this is because many melodies in songs such as nursery rhymes and folksongs are extremely similar in terms of musical characteristics (melody, key, harmony, rhythm, etc.). In contrast, lyrics have more semantic variability. As a result, a single salient word can allow identification by the lyrics, but a single note is not sufficient to allow identification by the melody. Identification based on partial input has implications for perceptual restoration.

A priming study carried out by Peretz et al. (2004) demonstrated that in adults, lyrics and melody are highly related in a bi-directional relationship in song memory. They also found that lyrics are easier to remember than melody, which may be a result of the increased proficiency in language as compared to music in most individuals. They suggest that for these reasons, lyrics are able to activate the stored representation of a familiar song much faster than melody can activate this representation. Whether this is also the case in children remains to be investigated.

Whilst there is a substantial body of evidence to suggest that in song memory, lyrics and melody are tightly integrated, there is also evidence that in on-line processing of music, they might be independent components (Besson, Faita, Peretz, Bonnel & Requin, 1998; Bonnel et al., 2001). These findings suggest that whilst lyrics and melody are processed separately, they are later combined to form an integrated representation in memory.

Research has also investigated the integration of melody and lyrics in song memory in children. Morrongiello & Roes (1990) found that lyrics were more salient than melody in memory for songs, for both adults and 5- to 6-year-old children. This might serve an adaptive purpose; due to the similarity between many of the melodies in children's songs, primarily using lyrics as a way of distinguishing between songs in memory would prevent confusion between songs.

Similar evidence of the importance of lyrics in children's song memory is reported by Feierabend, Clark Saunders, Holahan, & Getnick (1998). In their study, 3- to 5-year-old children regularly listened to tapes of eight songs over a familiarisation period of four weeks. In one condition, the songs were always sung to lyrics. In the second condition, the songs were first sung to lyrics and then to a nonsense-syllable. In the third condition, the songs were always sung to the nonsense syllable. Visual pictures accompanied the songs, to give an association between the aural and visual song representations. The identification test required the child to pick up the picture that corresponded to the melody being heard. Children who had been familiarised to the songs with lyrics performed better at test. They conclude that song text helps children's song recognition ability.

Taken together, these findings indicate that lyrics play an important role in children's memory for songs. In particular, the finding that familiar melodies are better recognised if presented with their lyrics suggest that there will be stronger familiarity cues operating when children are required to identify nursery rhymes where both melody and lyrics are presented concurrently. Because familiarity is a key component of the perceptual restoration process, it is therefore expected that children will be better able to restore missing fragments of songs when both lyrics and melody are presented together, than when the melody is presented alone. However, the advantage of lyrics in identification of intact and undisrupted songs might not operate in the same way for identification of disrupted signals such as those used in testing perceptual restoration effects.

There were two main questions to be answered by the present experiment. First, does perceptual restoration operate for vocal as well as melodic music in children? The perceptual restoration of vocal music has not previously been investigated. Second, does stronger perceptual restoration operate where both melody and lyrics are presented together when compared to the melody alone? These issues have clear implications for children's development in terms of the importance of song-based learning.

Method

Participants

One-hundred and twenty children (60 girls) participated in this study, between the ages of 4 years 3 months and 6 years 5 months, with a mean age of 5 years, 4 months (64.93 months; *SD* = 7.18). None of the children had participated in previous experiments. The children were recruited from the Reception and Year One classes of two primary schools in Surrey, UK, both with a predominantly white, middle-class catchment area. All children were native English speakers and had normal sensory development, as reported by the school. Consent for participation was obtained from both the child's parent and the school Headteacher. Further details about the participants in each condition are shown in Table 7.1.

Table 7.1. The gender split and mean age of the children in each of the six experimental conditions.

Gap replacement	**Stimulus type**	**N**	**Gender**	**Mean age in months (*SD*)**
Intact	Lyrics & Melody	20	M=13; F=7	65.55 (4.59)
	Melody Only	20	M=11; F=9	65.60 (9.20)
Noise-filled gaps	Lyrics & Melody	20	M=8; F=12	64.85 (5.23)
	Melody Only	20	M=8; F=12	64.00 (9.07)
Silence-filled gaps	Lyrics & Melody	20	M=10; F=10	66.05 (4.48)
	Melody Only	20	M=10; F=10	63.50 (9.05)

Design

Children were randomly allocated to one of six independent conditions in a 3 × 2 between-subjects design. Each child heard the same 15 nursery rhymes (see Appendix 5), but these were modified according to the experimental condition. Two variables were manipulated; the type of gap replacement (intact, noise-filled gaps or silence-filled gaps) and the stimulus type (either the lyrics and the melody of each nursery rhyme, or the melody only, sung to the syllable /la/). Thus, both types of stimulus were played on the same "instrument", the voice.

Materials

Stimuli were sung by the author. In order to match tempo and rhythm to the stimuli used in previous experiments, the piano melody stimuli were listened to over headphones as the new stimuli were sung. The mean length of the stimuli was 5.31 seconds, with the melody and lyrics stimuli containing an average of 11.7 words. The stimuli were recorded in a soundproof booth, using a Macintosh G4 powerbook operating Garageband V3.0.4. The stimuli were then edited in Soundstudio V2.2.4, to equate the peak amplitude of each stimulus at 20 dB below the maximum bandwidth of the soundfile. Soundstudio was also used to insert noise or silence into alternating 100 ms fragments of the stimuli. The noise inserted was 12 dB below the maximum bandwidth of the file. A programme written in Matlab V5.2.1 presented the stimuli to the child in a random order, and recorded their response and identification time. The standard song familiarity questionnaire (Appendix 1) was sent to participating schools.

Procedure

The procedure was identical to that of previous experiments and included the continuity question introduced in Experiment 3; however, stimuli were referred to as children's "songs" rather than children's "tunes".

Results

Correct identification scores

Each child was given a score out of 13 representing the number of songs they identified correctly in the experimental trials. Data were examined for gender differences (see Appendix 2) but no significant differences were found hence data were collapsed across gender. The mean number of correctly identified songs by children in each condition is shown in Figure 7.1.

These data were analysed using a 3 (experimental condition: intact; noise-filled gaps; silence-filled gaps) × 2 (stimulus type: lyrics and melody; melody only) between-subjects ANOVA. There was a significant main effect of experimental condition, $F(2, 114) = 224.53, p < .001, \eta^2 = .80$. Post hoc analyses (Games-Howell) revealed significant differences between all pairs of conditions: intact versus noise-filled gaps, $p < .01$, $d = 0.79$; intact versus silence-filled gaps, $p < .001$, $d = 2.67$; noise-filled gaps versus silence-filled gaps, $p < .001$, $d = 2.08$. There was also a significant main effect of stimulus type, $F(1, 114) = 157.93, p < .001, \eta^2 = .58$. As Figure 7.1 shows, stimuli where both lyrics and melody were present were easier to identify than stimuli where the melody only was sung to /la/.

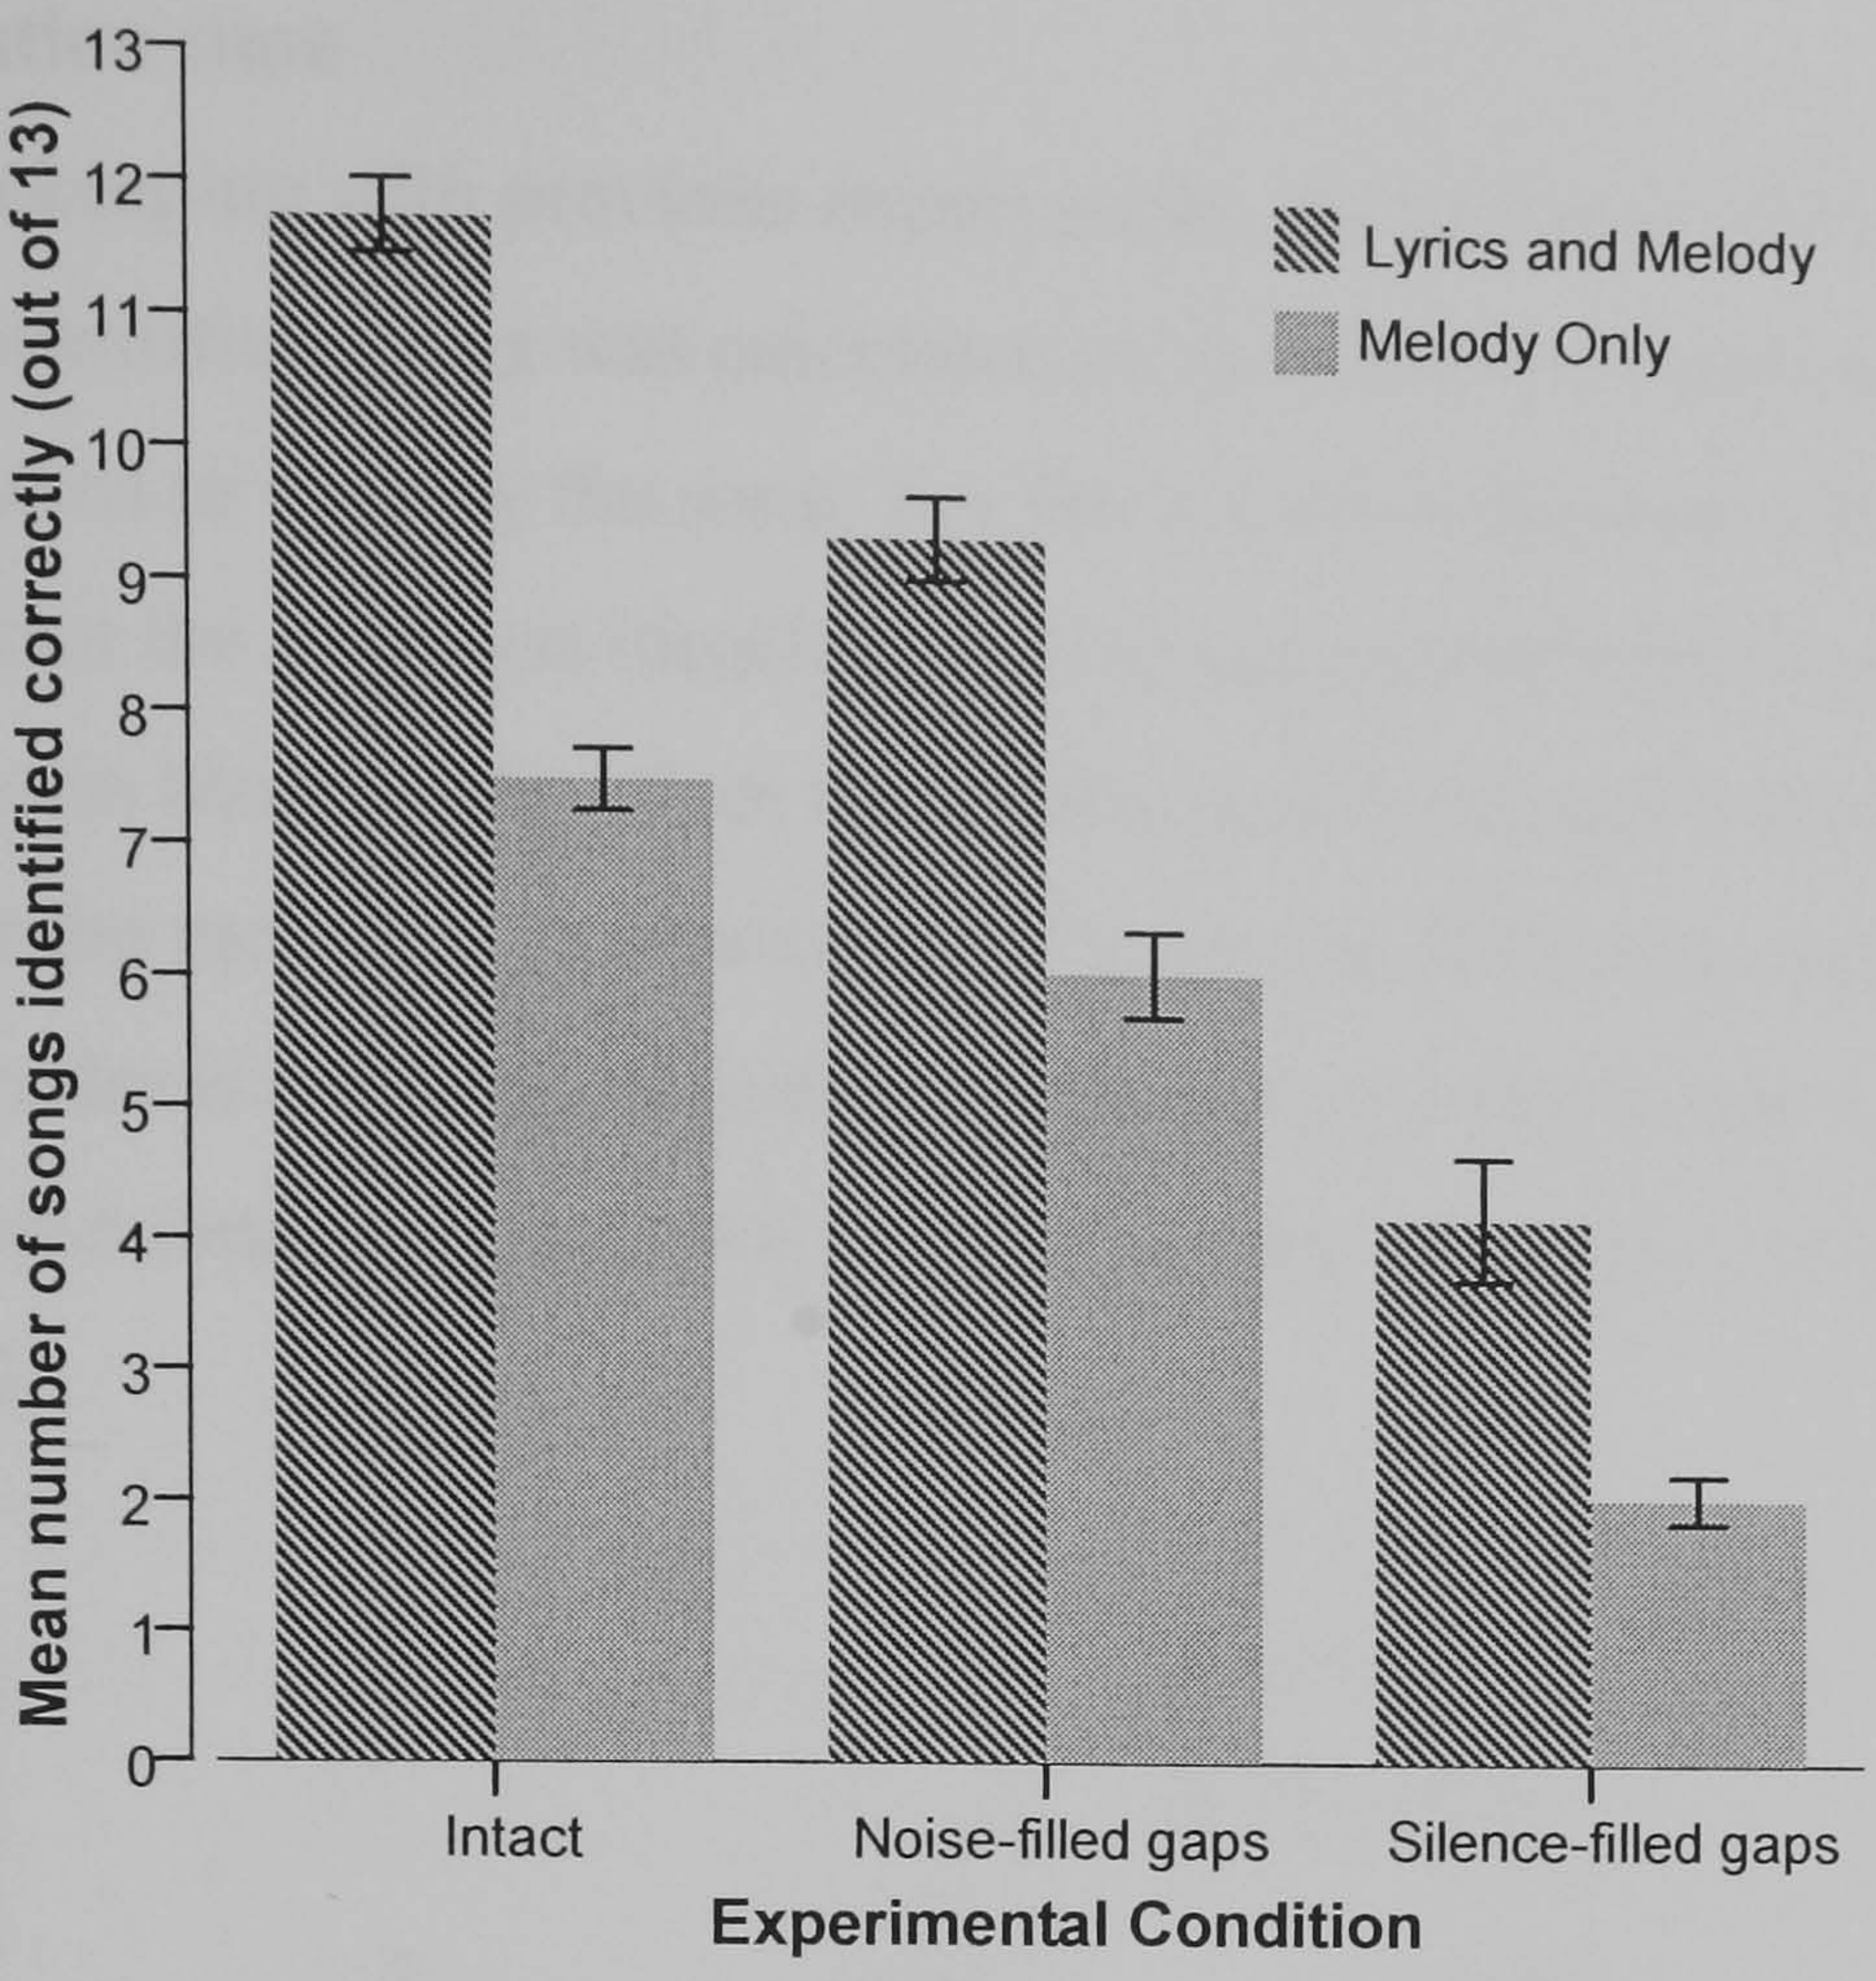

Figure 7.1. Mean number of songs identified correctly (± *SE*) by children in each of the six experimental conditions (intact, noise-filled gaps and silence-filled gaps; either lyrics and melody or melody only).

These main effects were qualified by a significant Experimental condition × Stimulus type interaction, $F(2, 114) = 5.57, p < .01, \eta^2 = .09$. Simple effects analyses (2 one-way ANOVAs) indicated that there was a significant effect of experimental condition for both lyrics and melody stimuli, $F(2, 57) = 111.98, p < .001, \eta^2 = .79$, and for melody only stimuli $F(2, 57) = 121.27, p < .001, \eta^2 = .81$. For both stimulus types, there were significant differences between all pairs of conditions (Games-Howell: all *p*s < .005). Whilst noise has a larger advantage over silence in terms of the difference in mean identification scores for lyrics and melody stimuli (a difference between means of 5.15 as compared to a difference of 4.00 for melody only stimuli), the size of the effect is larger for melody only stimuli (Games-Howell: lyrics and melody: $p < .001, d = 2.46$; melody only: $p < .001, d = 2.72$).

Identification time

In keeping with previous experiments, the measure of identification time for correctly identified songs was calculated as the difference between the song duration and the time taken to identify the song. For this reason a negative value of identification time indicates that the song was identified before the excerpt had finished. Analysis of gender differences in identification time revealed a significant difference between males and females in the melody only silence-filled gaps condition (see Appendix 2). For this reason, gender is included as a factor in the analysis that follows. The mean identification time in seconds for children in each of the six conditions is shown in Figure 7.2.

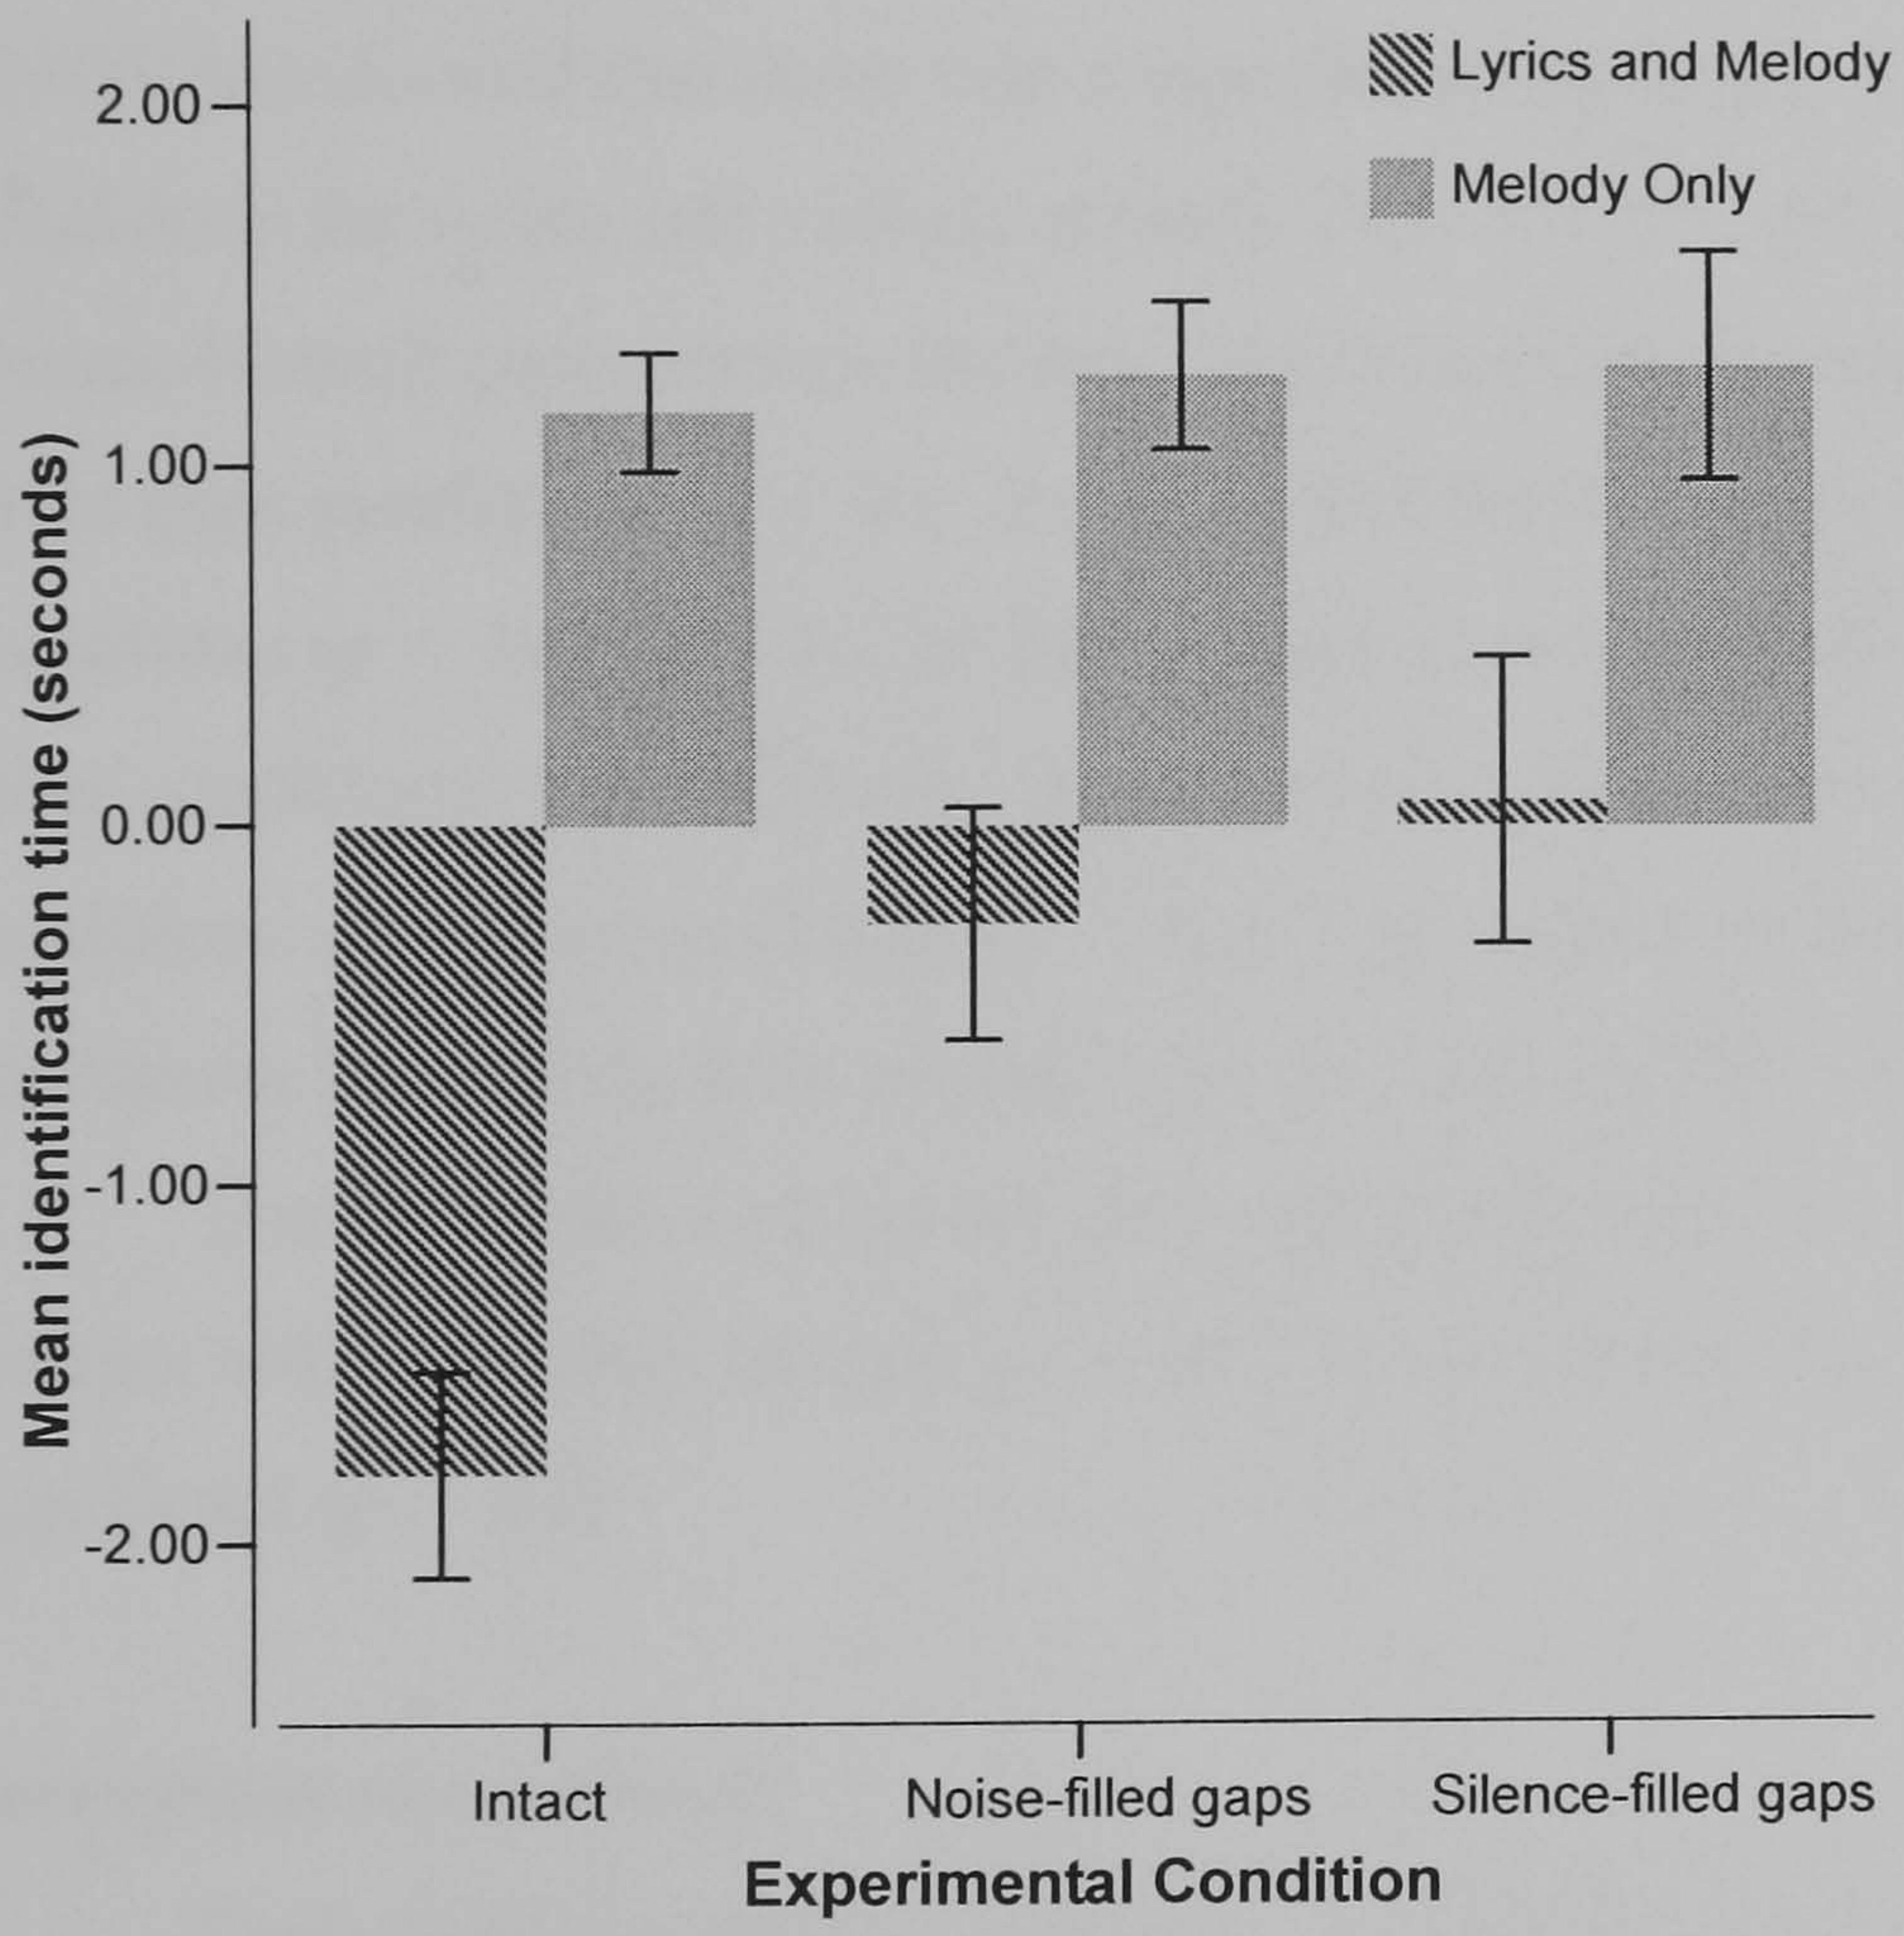

Figure 7.2. Mean identification time (identification time – stimulus duration; ± *SE*) in seconds for children in each of the six conditions (intact, noise-filled gaps and silence-filled gaps; either lyrics and melody or melody only).

These data were analysed using a 3 (experimental condition: intact; noise-filled gaps; silence-filled gaps) × 2 (stimulus type: lyrics and melody; melody only) between-subjects ANOVA. The analysis revealed a significant main effect of stimulus type, $F(1, 108) = 62.07, p < .001, \eta^2 = .37$. As Figure 7.2 shows, stimuli with both lyrics and melody were identified significantly faster than stimuli where the melody only was sung to /la/. The main effect of experimental condition was also significant, $F(2, 108) = 7.71$, $p < .01, \eta^2 = .13$. Post hoc comparisons (Games-Howell) revealed a significant difference between intact and noise-filled gaps conditions ($p < .05$, $d = 0.47$); with the differences between intact and silence-filled gaps conditions and between noise-filled gaps and silence-filled gaps conditions being non-significant ($p > .05$). There was also a significant interaction between stimulus type and experimental condition, $F(2, 108) = 7.88, p < .01, \eta^2 = .13$.

Further analysis of this interaction using simple effects analysis (2 one-way ANOVAs) showed that there was a significant difference between experimental conditions for lyrics and melody stimuli, $F(2, 57) = 11.48, p < .001, \eta^2 = .29$. Post hoc Games-Howell comparisons showed significant differences between intact and noise-filled gaps conditions ($p < .01$, $d = 1.01$) and between intact and silence-filled gaps conditions ($p < .001$, $d = 1.12$). Intact stimuli were identified faster than those in the other conditions. The difference between noise-filled gaps and silence-filled gaps conditions was non-significant ($p > .05$). The difference between experimental conditions for melody only stimuli was also non-significant; $F(2, 43) = .04, p > .05$.

The main effect of gender was non-significant ($p > .05$), and gender did not interact with any other factors (*p*s > .05). The three-way interaction was also non-significant ($p > .05$).

Perception of continuity

Each child reported whether they thought that little pieces of the tune were missing (fragmented) or that all of the tune was there (continuous). The number of children in each condition reporting that the signals were either fragmented or continuous is shown in Figure 7.3.

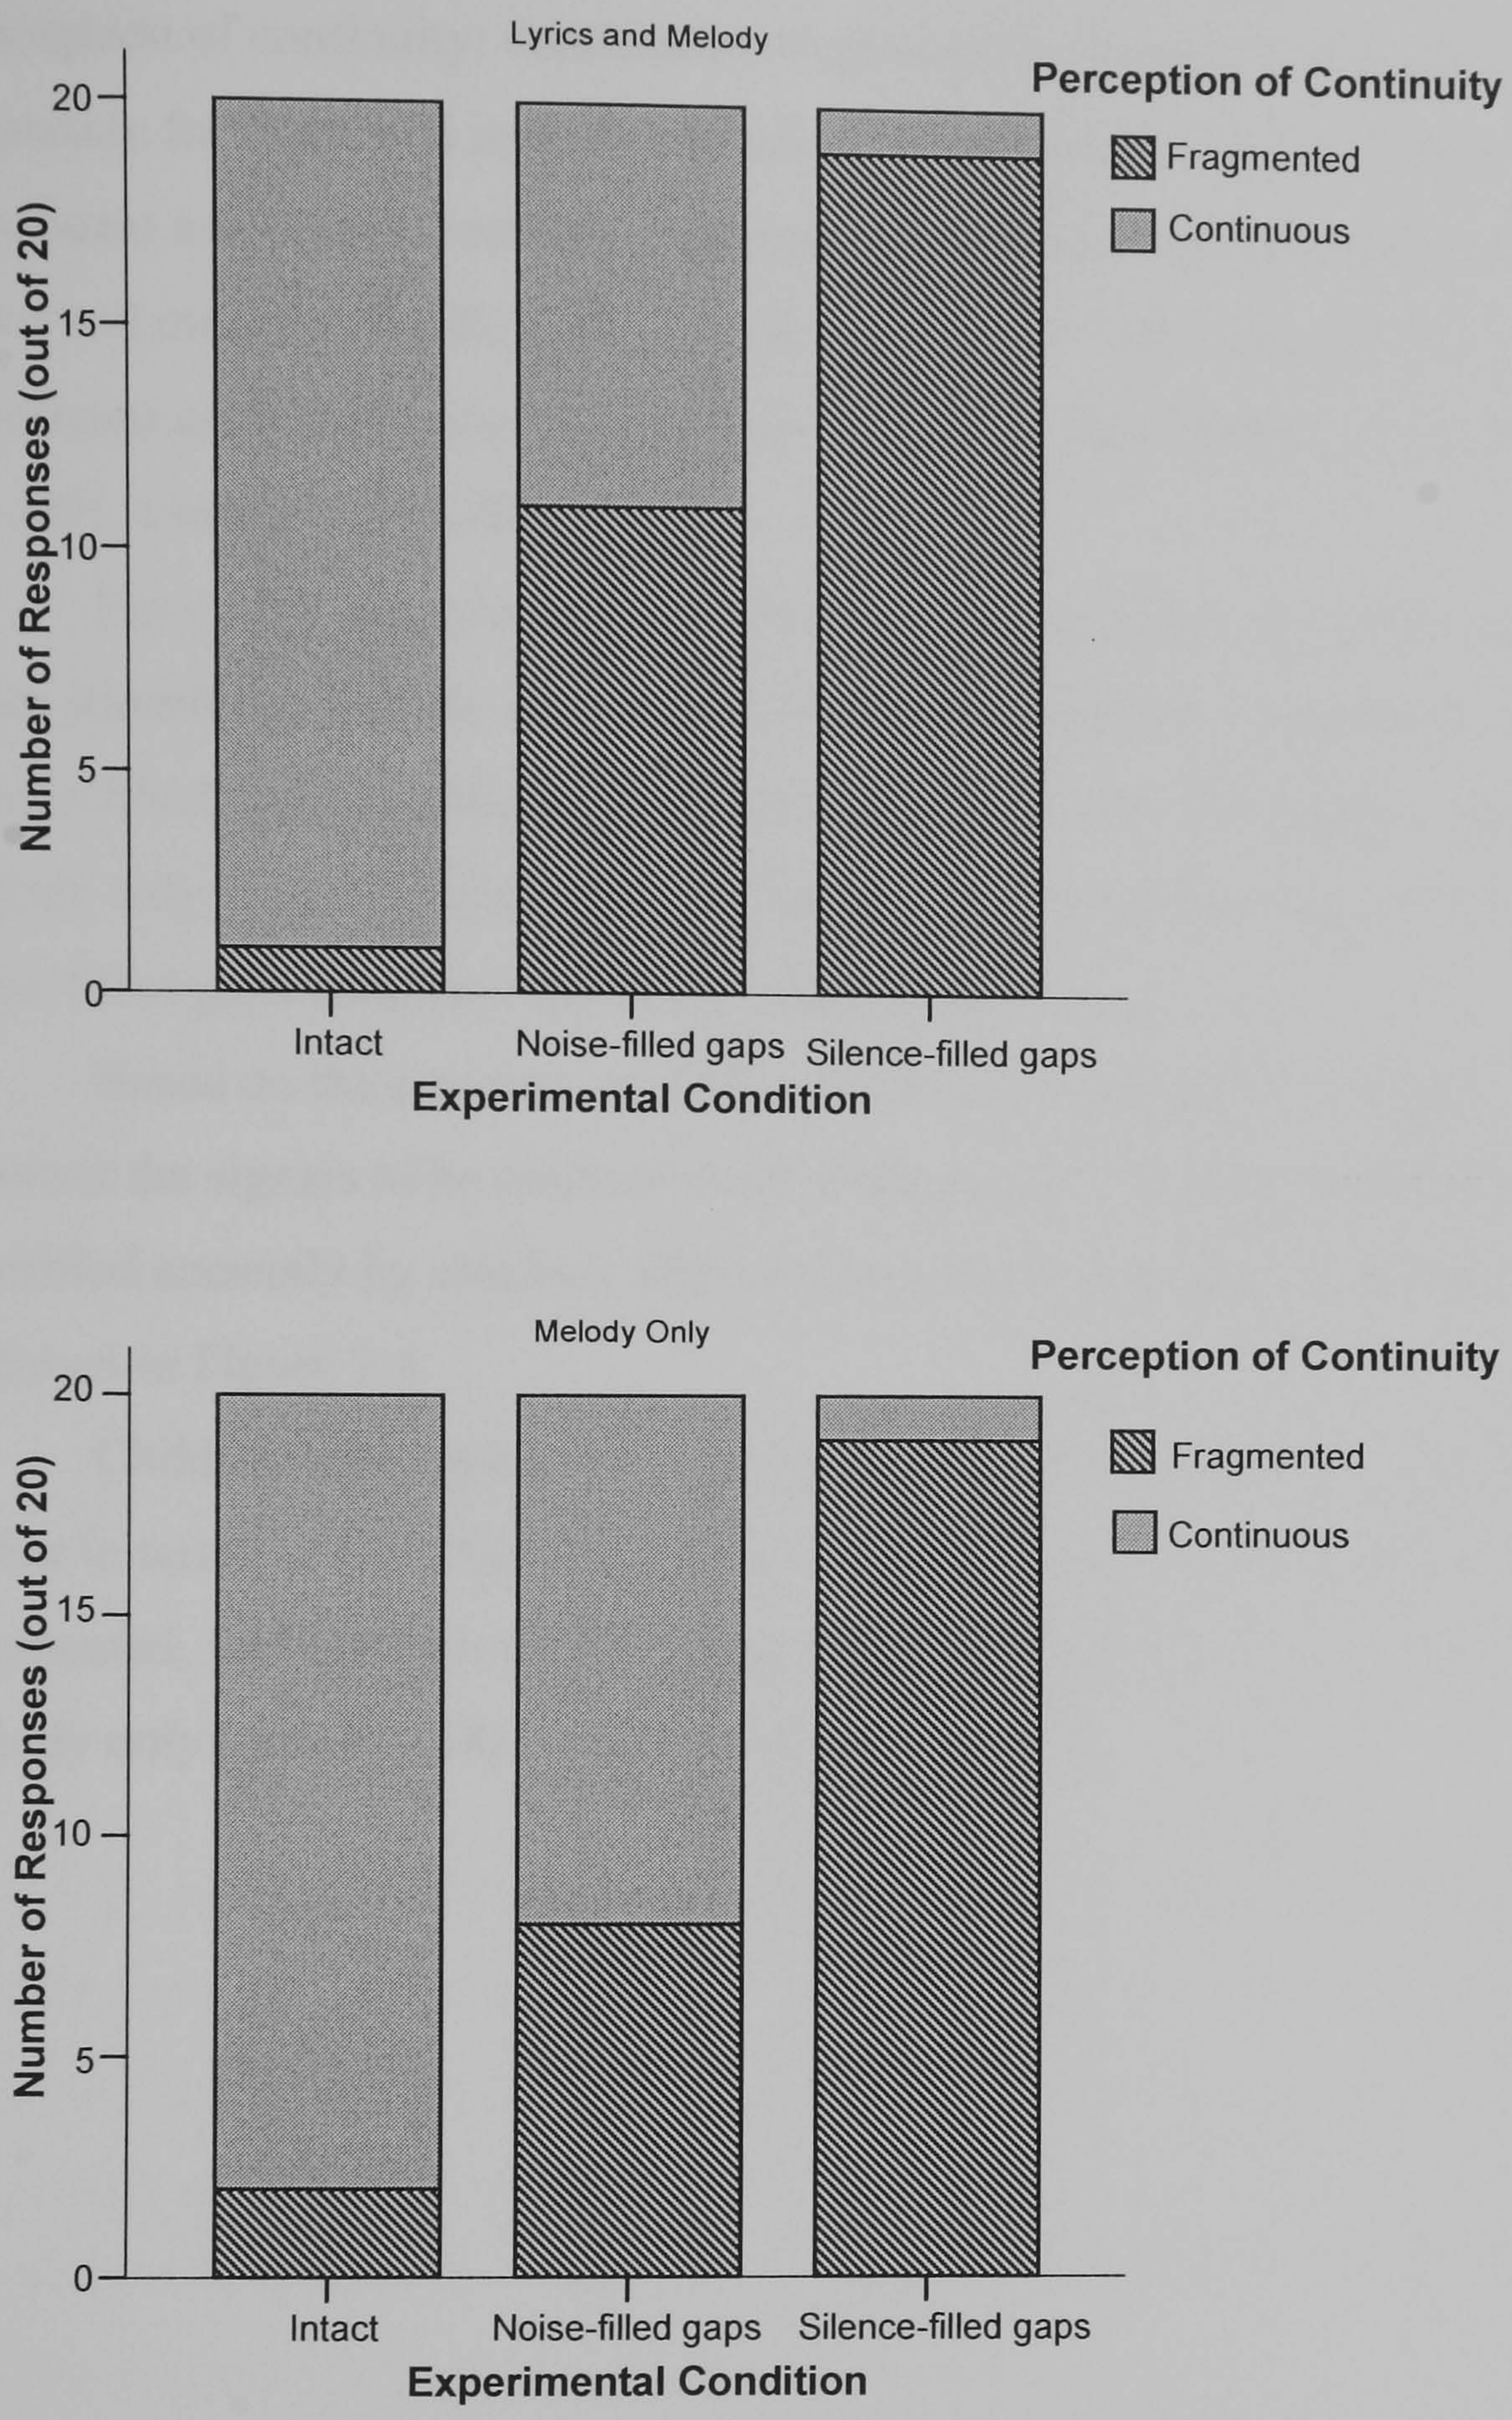

Figure 7.3. Total number of children in each condition (intact, noise-filled gaps and silence-filled gaps) reporting the songs as fragmented and reporting the songs as continuous, shown for lyrics and melody stimuli (top panel) and melody only stimuli (bottom panel).

A 3 (experimental condition: intact; noise-filled gaps; silence-filled gaps) × 2 (perception of continuity: continuous; fragmented) chi-square was run on these data, separately for lyrics and melody stimuli and melody only stimuli. There was a significant association between experimental condition and perception of continuity for lyrics and melody stimuli, χ^2 (2, $N = 60$) = 32.57, $p < .001$, Cramer's $V = .74$. A significant association was also evident for melody only stimuli, χ^2 (2, $N = 60$) = 29.77, $p < .001$, Cramer's $V = .70$.

Figure 7.3 suggests that, whilst there is little difference between melody and lyrics stimuli and melody only stimuli in terms of children's reports in intact and silence-filled gaps conditions, more children in the noise-filled gaps condition for melody only stimuli reported that the songs were continuous than did children in the noise-filled gaps condition for lyrics and melody stimuli.

Based on the previous analysis, children were grouped according to whether they reported the signals to be continuous or fragmented. The mean number of songs identified correctly by children who reported the stimuli to be continuous or fragmented is shown in Figure 7.4.

Children who reported that the stimuli were continuous scored significantly higher in terms of identification scores than children who reported the stimuli to be fragmented, for both lyrics and melody stimuli, $t(57) = 5.75$, $p < .001$, $d = 1.34$, and for melody only stimuli, $t(58) = 5.29$, $p < .001$, $d = 1.25$.

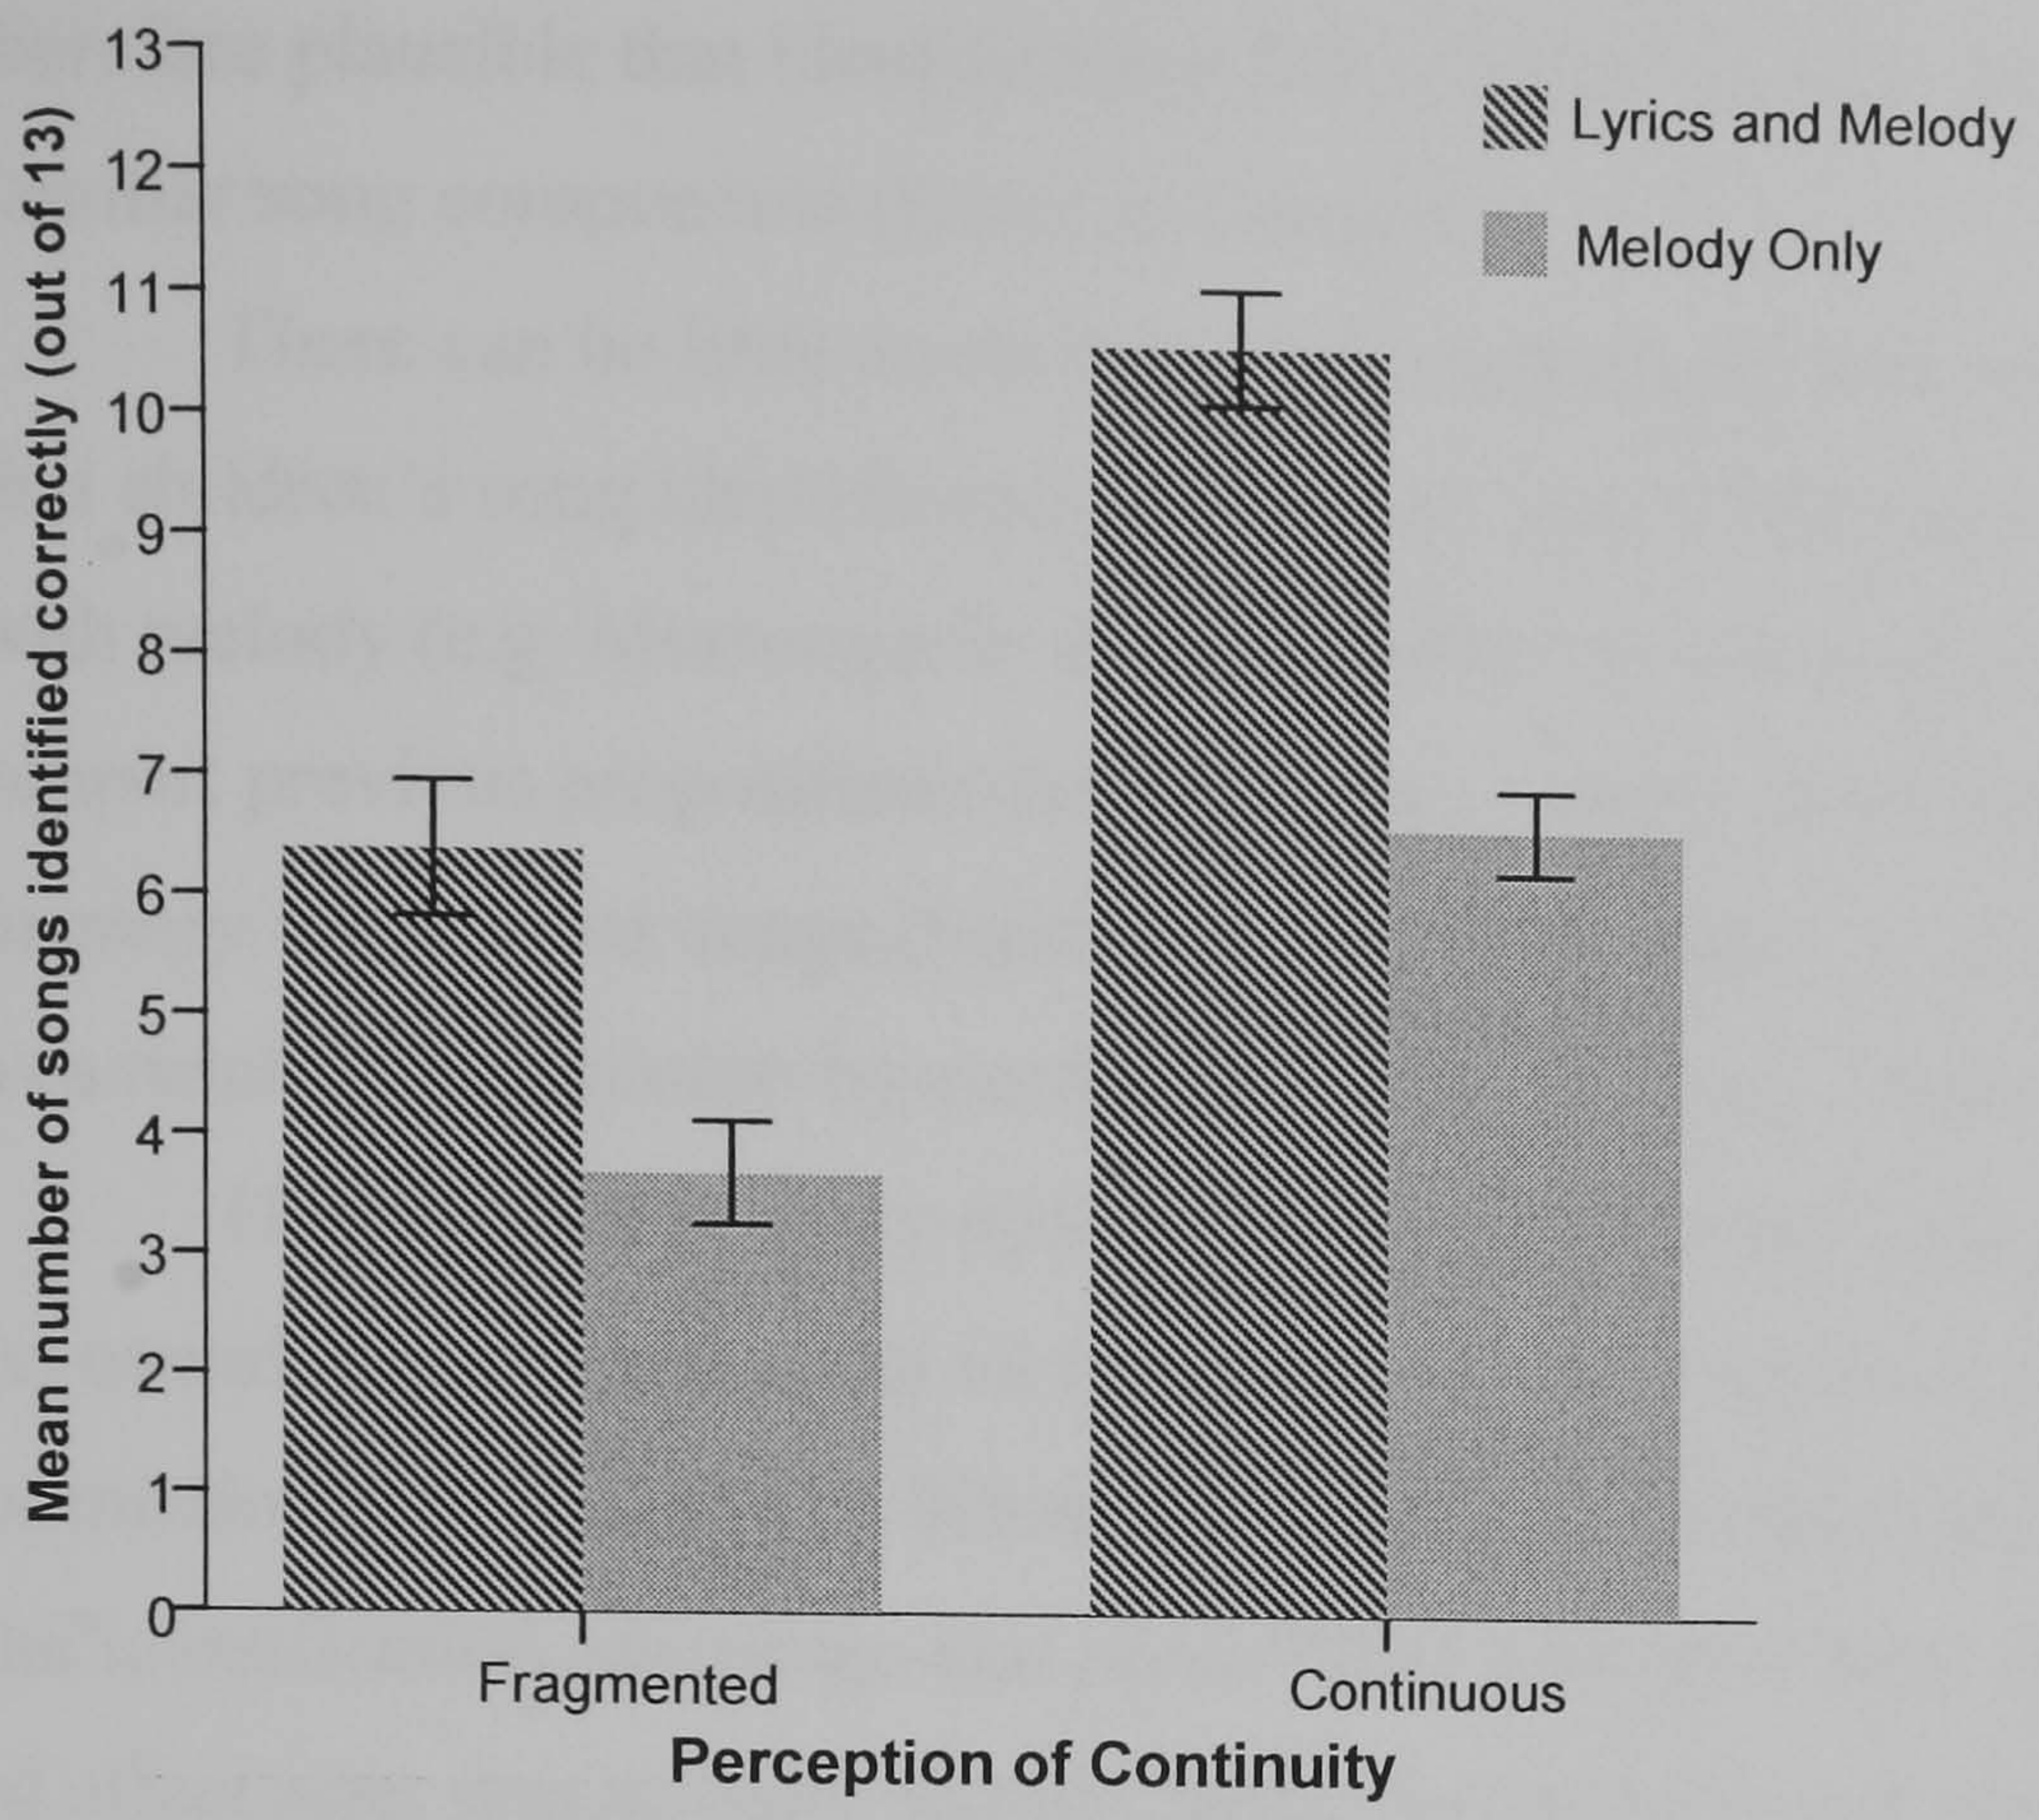

Figure 7.4. Mean number of songs identified correctly (± *SE*) by children who reported the songs to be fragmented and children who reported the songs to be continuous, for lyrics and melody and melody only conditions.

Discussion

The main aim of this experiment was to compare the perceptual restoration of children's songs, where both the lyrics and melody are presented, with the perceptual restoration of the melodic component of children's songs, where no lyrics are presented. Regardless of stimulus type (lyrics and melody or melody only), stimuli with noise-filled gaps had an identification advantage over stimuli with silence-filled gaps, indicating that perceptual restoration occurred for both songs and melodies. This extends findings from previous experiments that perceptual restoration in young children operates for familiar piano melodies, to vocal music.

The data from the present experiment show that over all conditions, song stimuli (lyrics and melody) were easier to identify than melody only stimuli. This is further strengthened by the finding that it is easier to identify lyrics and melody stimuli with noise-filled gaps than it is to identify melody only stimuli completely intact. It is therefore plausible that identification is facilitated through the presentation of two familiar song components (lyrics and melody).

There can be little doubt that these findings fit with previous research indicating that children's song identification is superior when lyrics are concurrently presented with melody (e.g. Morrongiello & Roes, 1990). It is also plausible that these findings support previous propositions that lyrics and melody are at least partially integrated in memory for familiar songs (Serafine et al., 1984; 1986); with this integration developing as a result of association-by-contiguity (Crowder et al., 1990).

However, of primary interest in the present experiment is not the effect of lyrics on overall identification, but on the ability of white noise to allow perceptual restoration of missing song fragments. When this is considered, the analysis showed that the size of the identification advantage that noise-filled gaps have over silence-filled gaps, in terms of effect size, was actually slightly larger for melody only stimuli than lyrics and melody stimuli. Crucially, however, there was no clear increase in the strength of perceptual restoration through presentation of both lyrics and melody, over presentation of the melodic component alone. Adding lyrics to the stimuli has no effect on perceptual restoration exclusively. This finding leads to two possible interpretations of the overall identification advantage for stimuli with both melody and lyrics over stimuli containing just the melody.

First, it is possible that lyrics really are the most important component of a song for identification, both in optimal and noisy listening conditions. Children presented with melody and lyrics stimuli might have completed the identification task in the present experiment by purely attending to the lyrics, with the melodic component playing a limited role in the process in all conditions. This might suggest that the larger noise-filled gaps over silence-filled gaps effect for melody only stimuli when compared to lyrics and melody stimuli occurs because the lack of lyrics makes identification of melodies with silent gaps especially poor.

Second, it is possible that the overall lyrics advantage in identification observed in the present experiment purely represents an additive effect; lyrics only provide extra familiarity cues for identification in the context of the accompanying familiar melody.

This would indicate that melody is as crucial a part of the identification process as lyrics are, in both intact and disrupted conditions.

It is not possible to determine from the present results which component (melody or lyrics) children were using as the primary basis for identification, and whether this was equivalent in all three experimental conditions (intact, noise-filled gaps and silence-filled gaps). Furthermore, if lyrics play a similar role across conditions, do lyrics help to make the melody more identifiable, or does the melody assist the identification of the lyrics? It is possible that the answer to this further question will differ depending on whether songs are heard intact, or with missing fragments replaced with noise or silence.

The significantly faster identification times for lyrics and melody stimuli when compared to melody only stimuli supports the proposition made by Peretz et al. (1994) that lyrics result in faster access to the stored representation of a familiar piece. The fact that this advantage held over all experimental conditions indicates that even when gaps are filled with silence, hearing a fragment of a word can help access the stored representation faster than a fragment of a melody can. This might be explained by the proposition of Serafine et al. (1984) that lyrics are a more distinct indication of song identity; a single word (or possibly even fragment of a word) can lead to identification whereas a musical note alone cannot lead to identification. Indeed, the melodies of children's songs are often extremely similar in terms of rhythm, pitch contour and harmonic progressions.

Whilst there is evidence of a considerable noise-filled gaps over silence-filled gaps advantage in identification scores in the present experiment, there is little evidence that perceptual restoration of intelligibility was systematically accompanied by the subjective report that the signal was continuous. Children's reports can be considered as reliable due to the high consistency with which children in intact and silence-filled gaps conditions responded, with their reports reflecting the actual status of the stimuli (continuous in the intact condition, and fragmented in the silence-filled gaps condition). In the noise-filled gaps conditions, marginally more children reported that the stimuli were continuous in the melody only condition than the lyrics and melody condition. However, those children who reported the signals to be continuous scored significantly higher than those who reported the signals to be fragmented in terms of identification scores. This suggests that noise can restore intelligibility of vocal music without the accompanying perception that nothing is missing, raising the possibility that the relationship between the two components of the perceptual restoration process is highly complex in nature.

One potential problem with the present study is that the verbal response method might have produced the identification advantage for the addition of lyrics. As the title of a song is often present in the lyrics, for the intact condition, the answer was, in most cases, present within the lyrics. This is a plausible explanation for the near-ceiling identification scores for this condition. However, the advantage that this would give to the noise-filled gaps or silence-filled gaps conditions depends on the extent to which the lyrics are transmitted through the disruptions. Furthermore, Peretz et al. (2004) raised a similar problem in their study testing priming of melody and text in familiar songs. Having found a lyrics advantage in priming later sections of songs, they questioned whether this was due to the frequent presence of the title in the lyrics of these beginning sections. However, even when they did not use the beginning section as the prime, and instead a section that did not contain the song title in the lyrics, the advantage for lyrics still held. Therefore, it is more likely that the present result indicates the increased distinctiveness of lyrics as opposed to melody.

The aim of this experiment was to test for the perceptual restoration of vocal music, and to test whether adding lyrics to a melody increases the strength of perceptual restoration. In conclusion, these findings indicate that musical perceptual restoration in children operates for vocal music. In addition, the presentation of lyrics increases identification scores when compared to melody alone, but not the strength of perceptual restoration. Whether these findings indicate a universal advantage for lyrics in identification regardless of experimental condition, or an indication that lyrics are only supported by perceptual restoration when also in the context of a familiar melody, is an important question for future research. Therefore, this result represents the starting point for further investigation of the relative influence of melody and lyrics in the process of the perceptual restoration of familiar songs.

Summary

Consideration of the results obtained in the present experiment has raised two possible interpretations of the findings. An important unresolved question that has arisen is whether perceptual restoration of songs operates purely on the basis of the lyrics, with little contribution from the melodic component. If this were the case, it is possible that children in the present experiment were identifying the melody and lyrics stimuli purely by focusing on the lyrics. Alternatively, it is possible that lyrics can only be supported by perceptual restoration when heard in the context of the accompanying melody. Finally, it is also important to investigate whether children use the same information to identify songs in undisrupted and disrupted forms. These questions form the aims of the next experiment.

Chapter 8: Experiment Six

The contribution of lyrics and melody to the perceptual restoration of songs in children

Introduction

Song identity is determined by a complex interaction between lyrics and melody (Feierabend et al., 1998). For children in particular, songs play an important role in learning and development (Racette & Peretz, 2007). It is therefore important that children can still make sense of songs when they hear them in a noisy environment, such as a bustling classroom. Research with adult and child listeners has questioned how the two primary components of a song, lyrics and melody, are related in the stored song representation. The aim of the present experiment is to investigate the relative status of each component in allowing children's identification of songs when the signal is disrupted, such as when heard in a noisy environment.

A key component of the perceptual restoration processes is the role of familiarity; the more familiar one is with the content of the signal, the stronger the expectations concerning the identity of the missing fragments. This expectation operates as part of an interaction with confirmation provided by what is present in the signal. It is unclear, however, whether such familiarity effects are holistic; in order for expectations to be generated, must the content of the signal be an exact match with a previously stored template, or in the case of songs, can familiarity effects operate separately for melody and lyrics when not presented in their usual pairing? Whilst it is well documented that familiarity effects operate in the case of perceptual restoration, consideration of this issue can further clarify the exact nature of familiarity influences needed to contribute to perceptual restoration in the music domain.

The role of melody and lyrics in song identification

There exists an accumulation of evidence that melody and lyrics are integrated in song representation (Crowder et al., 1990; Morrongiello & Roes, 1990; Peretz et al., 2004; Serafine et al., 1984; Serafine et al., 1986). The established method used to test integration effects gives listeners a brief familiarisation phase where they hear a set of previously unfamiliar songs just once, followed by a recognition task. In the recognition task, they are required to report if they have previously heard: songs in the exact format from the familiarisation phase; original lyrics paired with a new melody; an original melody paired with new lyrics; a new melody and new lyrics; and a mismatch between one set of original lyrics and another original melody.

Research has demonstrated that both adults and children are better able to recognise song components heard in their original pairing than two familiar but mismatched components where these familiar components are heard in different contexts (Morrongiello & Roes, 1990; Serafine et al., 1984). In addition, Serafine et al. (1984) reported that adult listeners found it very difficult to recognise a previously heard melody if not presented with its original lyrics; recognition of previously heard lyrics in a new melodic context was slightly better, but was still poor in comparison to recognition of originally heard complete songs. Serafine et al. also reported how participants recognised the lyrics when presented with a new melody if lyrics had been the previously heard item at familiarisation, and recognised the melody if this had been the previously heard component at familiarisation, but presented with new lyrics at test. In cases where both components at test were familiar but mismatched, participants reported that they recognised the lyrics far more often than they reported that they recognised the melody. This suggests that overall, in a new context, lyrics are more frequently recognised than melody.

Similar results were obtained by Morrongiello & Roes (1990) with both child and adult participants. Both adults and children recognised songs with a pairing of lyrics and melody as heard at familiarisation more often than the lyrics or melody of a mismatch song where familiar components were incongruously matched. This superior recognition of correctly paired components was slightly superior in adults when compared to children, suggesting that lyrics and melody are more strongly integrated in adults' song memory than children's. Children in particular seemed to show an advantage for lyrics over melodic memory; they reported that a song at test was the same as a song at familiarisation if the lyrics were the same. If the melody was the same and the lyrics were new, the song was reported as "not at all the same" as the original presentation. Finally, when both adults and children reported that a mismatch song was "somewhat the same" as the original song, they were much more likely to report that the lyrics were the same, than reporting that the melody was the same or indeed noticing that both components were in fact the same, but had been mismatched. This suggests that for both adults and children, lyrics are easier to identify than melody when presented in a different context.

Steinke et al. (2001) describe song recognition as a result of the activation of both a "melody analysis system" and a "speech analysis system". The more experience one has with a particular song, the stronger the connections between the two systems become, and so each component will cue the other more strongly. A mismatch between familiar lyrics and melodies from different songs, and between either component and its usual context, can be overcome where the melody and speech analysis systems cue the correct partner in each case.

Taken together, these results seem to indicate an advantage for lyrics in song recognition, both in their original and in a new or mismatched melodic context. This has commonly been described as a result of the increased ability of song lyrics to distinguish between different song representations than song melodies (Peretz et al., 2004). It is possible to identify a familiar melody by presentation of approximately six notes (Dalla Bella et al., 2003). Words, on the contrary, can be recognised on the basis of a few syllables (Grosjean, 1980). Morongiello & Roes (1990) suggest that this asymmetry might serve an adaptive purpose; since the melody of many children's songs is extremely similar in terms of rhythm, melodic contour and harmonic progressions, paying greater attention to the text helps children to distinguish between different songs in their repertoire. When most attention is focused on the lyrics during song presentation, this can explain poorer recognition memory for the melodic component of a song.

Whilst these results seem to point to the privileged status of lyrics as opposed to melody in song identification, it would be wrong to assume that the component primarily used for identification when signals are intact and undisrupted is necessarily the same component used when signals are disrupted. It is possible that lyrics are not the primary basis for song identification when signals are disrupted, and instead the process of identifying a familiar song relies equally on melody and lyrics. Likewise, there is a possibility that melody is the primary cue used to identify familiar songs when signals are disrupted.

A further problem with previous studies is that they have used as stimuli previously unfamiliar music which participants have heard only once in some cases. It may be that the integration effect, or the lyrics advantage, might differ in the case of highly familiar songs, and where this familiarity has accumulated over repeated exposure to the familiar song. For this reason, the present study uses songs that would be highly familiar to the children, and employs an identification task rather than simply asking participants to report whether they have previously heard the songs.

Directly putting melody and lyrics into conflict in an identification task should also help to clarify the results of Experiment 5. Songs where both lyrics and melody were presented, when compared to melodies only, led to superior identification; however, perceptual restoration was not stronger where lyrics and melody were presented together. The present experiment was therefore designed in order to answer the following three questions. Is superior identification of lyrics and melody stimuli simply due to the additive familiarity effects provided by two components rather than one? Alternatively, do lyrics really play a stronger role in perceptual restoration than melody? If they do, is this only the case when heard in the context of their expected melody?

Method

Participants

Fifty-four children (21 girls) between the ages of 4 years 8 months and 6 years 6 months participated in this study, with a mean age of 5 years 7 months (67.46 months; *SD* = 7.56); none of the children had participated in previous experiments. The children were from the Reception and Year One classes of a primary school in a predominantly white, middle-class area of Surrey, UK. The school Headteacher and the child's parent or guardian gave consent for each child to participate. All children were native English speakers with normal sensory development, as reported by the school. Further details about the participants in each condition are given in Table 8.1.

Table 8.1. The gender split and mean age of the children in each of the three experimental conditions.

Experimental Condition[a]	**Gender**	**Mean age in months (*SD*)**
Intact	M=9; F=9	68.17 (7.37)
Noise-filled gaps	M=12; F=6	65.94 (8.35)
Silence-filled gaps	M=12; F=6	68.28 (7.74)

[a] *n* = 18 for each condition

Design

Children were randomly allocated to conditions in a 3 × 3 mixed design. Each child was asked to identify 18 stimuli; six stimuli of each of three types: old (original) lyrics, old (original) melody, and mismatch; tested within-subjects (see Appendix 6). Each child did so in one of three independent conditions of the task (intact, noise-filled gaps or silence-filled gaps; between-subjects).

Materials

All stimuli were recorded and edited as for Experiment 5. The participating school was sent the song familiarity questionnaire (Appendix 1) prior to participation. The six most familiar songs from all previous experiments (i.e. those that were identified most frequently) were selected to be the stimuli for this experiment (see Appendix 7). Because the content of the songs was manipulated in this experiment, it was decided that children needed to hear the most familiar songs in order to be able to make identification judgements on the basis of a single familiar component (lyrics or melody). However, in creating the mismatch stimuli, the set of possible lyric partners for the six melodies included all stimuli from the set in order to allow appropriate matches to be made in all cases.

Once the six most familiar songs had been selected, one song from this set, *If You're Happy and You Know It*, was substituted for *London Bridge*. This was because the rhythm of the previous melody for *If You're Happy and You Know It* contained many repeated notes of the same pitch and note duration and therefore to write a new and distinctive melody resulted in a stimulus that in terms of length was out of the range of previous melodies. Therefore, this stimulus was substituted with the next most familiar song from previous experiments, *London Bridge* (these two songs barely differed in frequency of correct responses over all previous experiments).

So as to present one demonstration trial of each stimulus type, three demonstration trials were presented, including the two songs used for this purpose in previous experiments (*Baa Baa Black Sheep* and *The Grand Old Duke of York*). In addition, the next most familiar song from previous experiments not being used in the experimental trials (*Row, Row, Row Your Boat*) was also presented as a demonstration trial. Which song of the three represented each stimulus type was decided randomly. A programme written in Matlab V5.2.1 presented the stimuli in a pseudo random order ensuring that no two stimuli of the same type would follow one another (e.g. *Twinkle Twinkle Little Star* old melody followed by *Old MacDonald had a Farm* old melody), and no two examples of the same song would follow each other (e.g. *The Wheels on the Bus* old lyrics followed by *The Wheels on the Bus* mismatch).

Old Lyrics, New Melody stimuli

For these stimuli, the original lyrics from each nursery rhyme were sung to newly composed melodies, for example, the lyrics of *Twinkle Twinkle Little Star* were accompanied by a new melody. The melodies were written to be comparable to the original melodies in terms of their musical characteristics, and consisted of both a new pitch structure and rhythm. The average stimulus length was 5.28 seconds. New melodies were written in the same key as the original melodies, and had the simple rhythmic structure characteristic of nursery rhyme melodies. In addition, measures of musical characteristics (as reported by Steinke et al., 2001) were calculated for the original melodies and the new melodies, for purposes of comparison.

Measures included: the number of note onsets; the average interval size; the contour complexity (a measure of how many times the pitch direction changes, calculated by dividing the number of pitch direction changes by the total number of note onsets); the number of different note durations (a measure of rhythmic complexity); the percentage of notes accounted for by the most common note durations; and the tonal strength (a measure of how well the melody represents the key in which it is written; calculated using the algorithm reported by Krumhansl, 1990). Once calculated, it was ensured that the value of each characteristic for the new melodies fell within the range of the values for the original melodies (see Appendix 8 for values of each characteristic for original and new melodies).

Old Melody, New Lyrics stimuli

To produce these stimuli, new lyrics were written to accompany the original melodies, for example, the new lyrics *Can you see the busy bee, buzzing round the apple tree* were sung to the original melody of *Twinkle Twinkle Little Star.* In writing the new lyrics, the number of syllables was matched exactly to the original lyrics, and rhyming structure was matched to the original lyrics. The stimuli had an average length of 5.21 seconds, and contained an average of 12.33 words. The lyrics from the original songs were grouped into semantic categories according to the subject of the song (animals, people, places, transport, parts of the body, and miscellaneous). The new lyrics for each melody were written to belong to one of these semantic categories, but not the category that the original lyrics belonged to (see Appendix 7). Newly composed melodies and newly written lyrics were verified by an experienced schoolteacher for plausibility as nursery rhymes and suitability for the age range.

Mismatch stimuli

For mismatch stimuli, each original melody was paired with a different set of original lyrics, for example, the lyrics to *Humpty Dumpty* were sung to the melody of *Twinkle Twinkle Little Star.* These stimuli had an average length of 5.19 seconds. Where possible, no alterations were made to the rhythm of the melody or lyrics to accommodate the new component. However, this was not always possible and in some cases, slight rhythmic or syllabic alterations were made (see Appendix 9). Importantly, this only involved tying or splitting a syllable across two notes; no lyrics were removed and the melodic contour remained unaltered.

Procedure

Each child was tested individually in a quiet room in the school. The child was told that they would be playing a music game, where they had to work out what some children's songs were, and were given the standard instructions used across all previous experiments. Each child was presented with three demonstration trials (*Row, Row, Row Your Boat* as an old lyrics stimulus type, *Baa Baa Black Sheep* as an old melody stimulus type, and *The Grand Old Duke of York* as a mismatch stimulus type.) No further explanation was given if the child did not identify the song in the demonstration trials, in order to avoid biasing children to respond in a particular way to all trials regardless of stimulus type.

The 18 experimental trials followed (six of each stimulus type), and children were asked to name each song when it was presented, again using the standard instructions previously reported. No time limit was imposed upon children giving a response. If a child did not identify the song correctly, or gave no response, this was scored as incorrect and the next trial followed immediately. No feedback was given as to whether the child's response was correct or incorrect.

The experimenter entered into the programme whether the child had given a correct response on each trial, and whether this correct response was based on the melody or the lyrics of the stimulus (a "melody" or "lyrics" response). In order to see if children use lyrics over melody even when they are unfamiliar, it was noted separately whether, on old melody trials, children made an identification response characterised by repeating some of the new unfamiliar lyrics as their answer, to make a song title.

After all 18 experimental trials had been presented, the child was asked whether they thought that "little pieces of the song were missing" (the perception of a fragmented signal) "or all of the song was there" (the perception of a continuous signal) as a further test of perceptual restoration. This was not asked separately for each stimulus type but instead referred to all trials presented. This was to prevent children thinking they had to vary their response across trials. The child's response was noted, and they were praised and thanked for their help.

Results

Correct identification scores

Each child was given a score out of six according to the number of songs of each stimulus type that they identified correctly in the experimental trials. There were no gender differences in identification scores (see Appendix 2) and so data were collapsed across gender for further analyses. The mean number of stimuli of each type identified correctly by children in each experimental condition is shown in Figure 8.1.

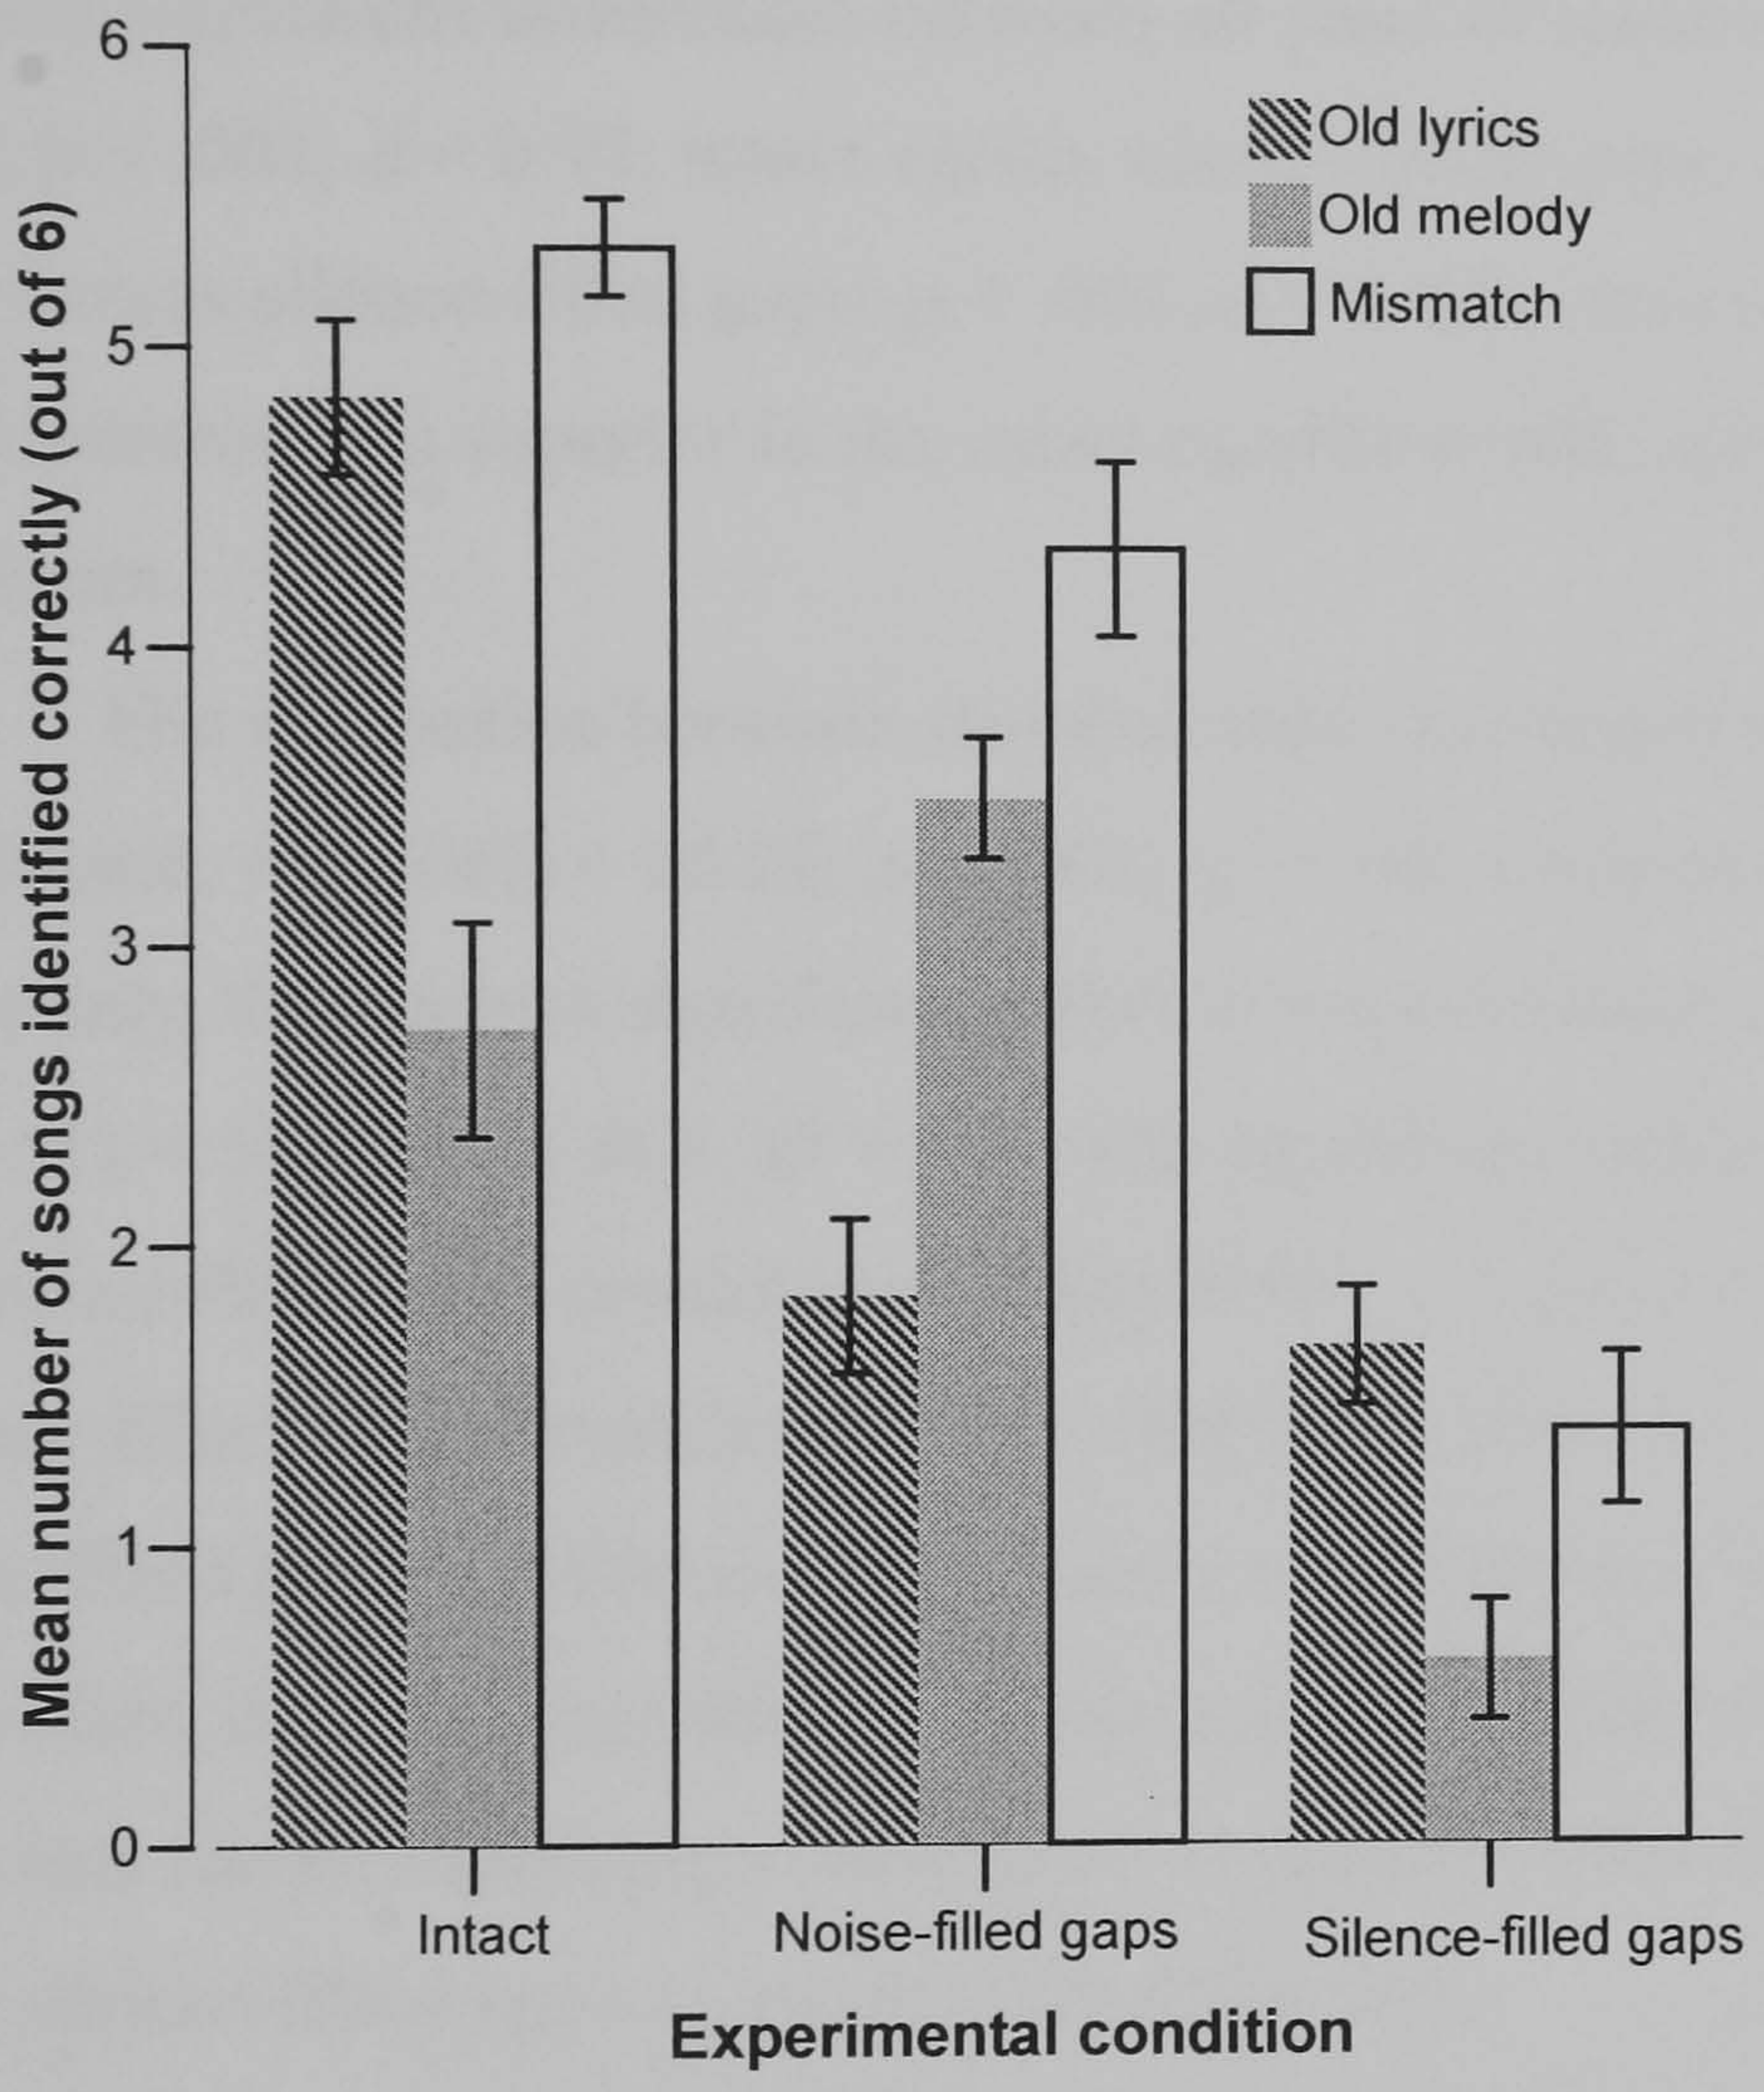

Figure 8.1. Mean number of songs identified correctly (± *SE*) by children in each of the three experimental conditions (intact, noise-filled gaps, silence-filled gaps) and for the three stimulus types (old lyrics, old melody, mismatch).

These data were analysed using a 3 (stimulus type; old lyrics, old melody, mismatch) × 3 (experimental condition; intact, noise-filled gaps, silence-filled gaps) mixed ANOVA with stimulus type as the repeated measure. There was a significant main effect of stimulus type on identification, $F(2, 102) = 26.52, p < .001, \eta^2 = .34$. Post hoc tests[13] (Sidak) revealed significant differences between old lyrics stimuli and mismatch stimuli ($p < .001, d = 0.46$), and between old melody stimuli and mismatch stimuli ($p < .001, d = 0.71$). The difference between old lyrics and old melody stimuli was non-significant ($p > .05$). Overall, identification of mismatch stimuli was the best whereas identification of old melody stimuli was the worst.

There was also a significant main effect of experimental condition on stimulus identification, $F(2, 51) = 101.68, p < .001, \eta^2 = .80$. Post hoc tests (Games-Howell) showed significant differences between all pairs of conditions (intact versus noise-filled gaps: $p < .001, d = 0.97$; intact versus silence-filled gaps: $p < .001$; $d = 2.79$; noise-filled gaps versus silence-filled gaps: $p < .001, d = 1.88$). Across all stimulus types, identification was superior in the intact condition and inferior in the silence-filled gaps condition.

The interaction between stimulus type and experimental condition was significant, $F(4, 102) = 20.08, p < .001, \eta^2 = .44$. Considering each stimulus type separately, there was a significant effect of experimental condition for old lyrics stimuli, $F(2, 51) = 55.01, p < .001, \eta^2 = .68$, with significant differences evident between intact and noise-filled gaps conditions (Tukey HSD: $p < .001, d = 2.73$) and between intact and silence-filled gaps conditions (Tukey HSD: $p < .001, d = 2.88$). The difference between noise-filled gaps and silence-filled gaps conditions was non-significant (Tukey HSD: $p > .05$.) Here, the intact condition had superior identification scores to both the noise-filled gaps and silence-filled gaps conditions. Crucially, noise-filled gaps had no advantage over silence-filled gaps in terms of identification.

For the old melody stimuli, there was also a significant effect of experimental condition, $F(2, 33) = 52.01, p < .001, \eta^2 = .56$, with significant differences between intact and silence-filled gaps conditions and between noise-filled gaps and silence-filled gaps conditions (Games-Howell: $p < .001, d = 1.38$ and $p < .001, d = 3.37$, respectively). The difference between intact and noise-filled gaps conditions was non-significant (Games-Howell: $p > .05$). Here, however, identification scores in the noise-filled gaps

---

[13] For within-subjects factors, Sidak post-hoc tests are used where the assumption of sphericity is met; Bonferroni tests are used where the assumption of sphericity is broken.

condition were the highest, and identification scores in the silence-filled gaps condition were the lowest.

Finally, for mismatch stimuli, there was also a significant effect of experimental condition, $F(2, 32) = 82.51, p < .001, \eta^2 = .73$. There were significant differences between all pairs of conditions: intact versus noise-filled gaps (Games-Howell: $p < .05$, $d = 0.81$); intact versus silence-filled gaps (Games-Howell: $p < .001$, $d = 3.61$); and noise-filled gaps versus silence-filled gaps (Games-Howell: $p < .001$, $d = 2.38$). Here, the intact condition had the highest mean identification score, and the silence-filled gaps condition the lowest. Taken together, these results show that perceptual restoration did not operate equally for all stimulus types. Further analysis of children's responses may help to clarify this finding.

Responses made on the basis of melody and lyrics

Overall identification scores do not show what information within the stimulus children were using as the basis for their identification response. For this reason, when a correct identification response was made, it was noted whether this response was based on the lyrics or the melody component of the stimulus. The mean number of correct responses made on the basis of lyrics and on the basis of melody for each stimulus type and in each experimental condition is shown in Figure 8.2. The figure also shows "identification" responses made by repeating some of the unfamiliar lyrics as an answer on old melody trials. It is clear that children were responding in a highly systematic way to old lyrics and old melody stimuli, responding on the basis of the component that was familiar (the lyrics component of old lyrics trials and the melodic component of old melody trials). For this reason, statistical analysis of response types is restricted to the mismatch condition.

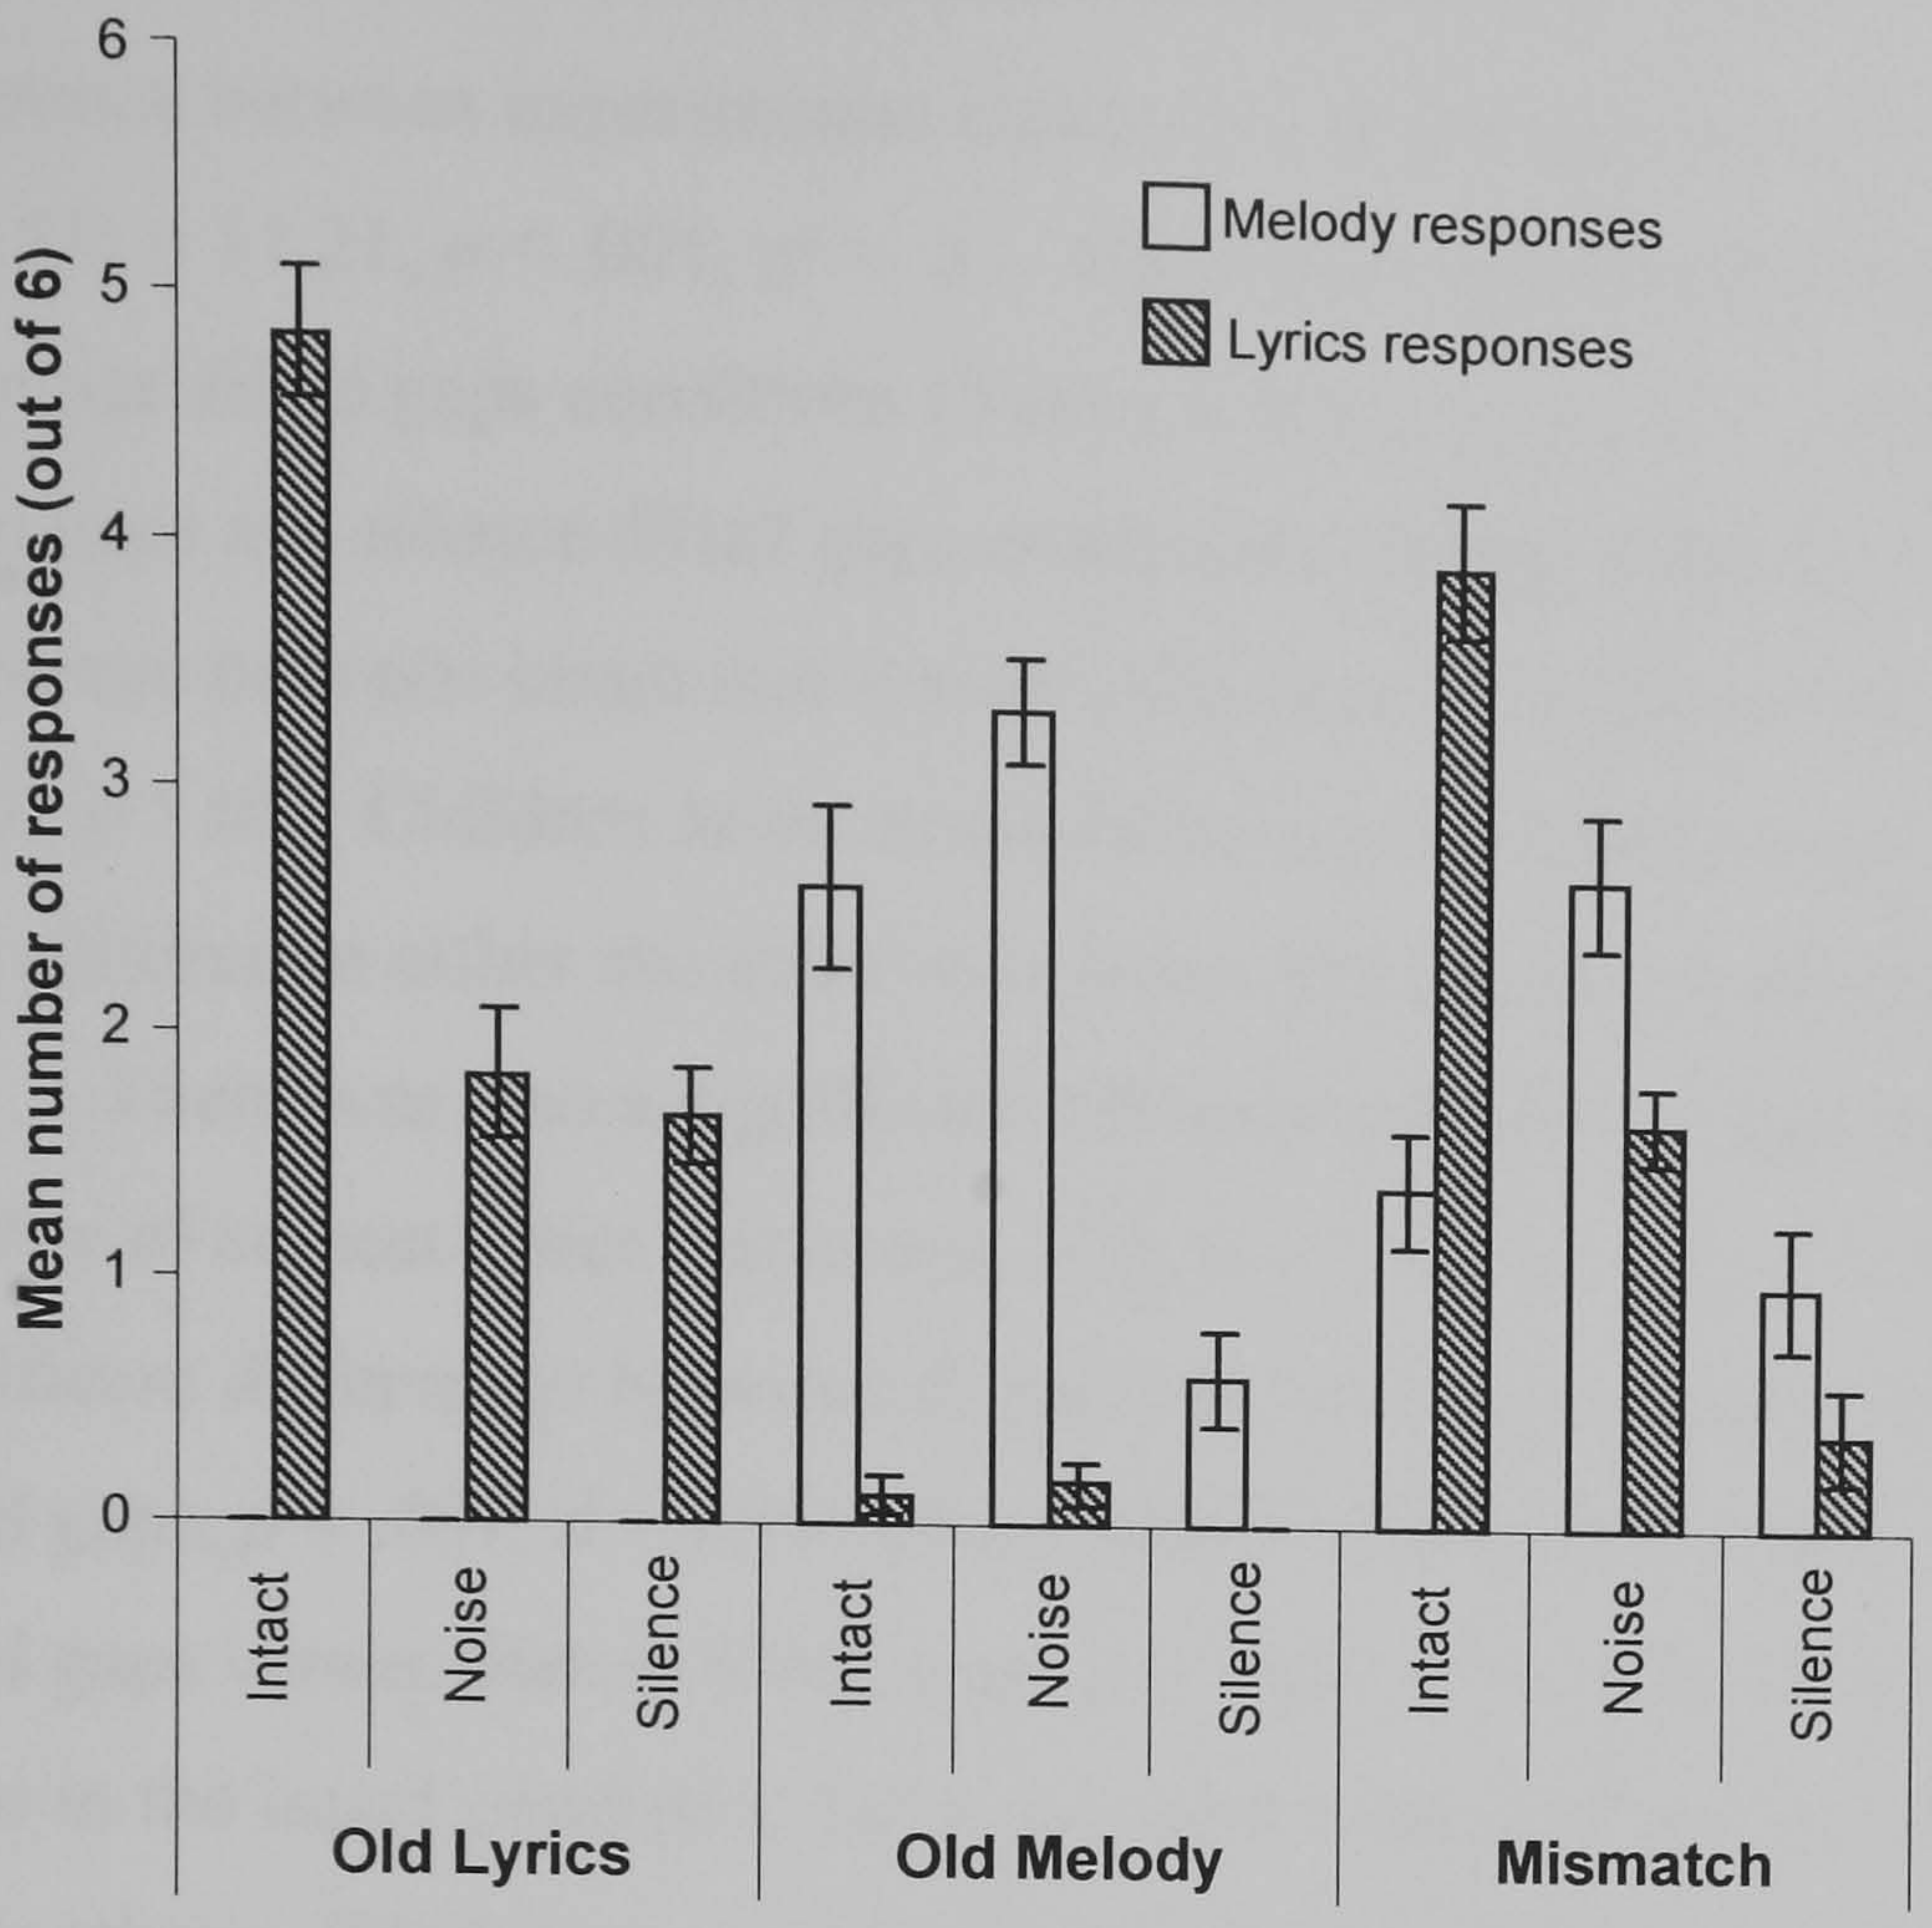

Figure 8.2. Mean number of responses (± *SE*) made on the basis of melody and lyrics by children in each experimental condition (intact, noise-filled gaps, and silence-filled gaps), and for each stimulus type (old lyrics, old melody, and mismatch).

Melody and lyrics responses to mismatch trials were analysed using a 3 (experimental condition: intact; noise-filled gaps; silence-filled gaps) × 2 (response type: melody; lyrics) mixed ANOVA, with response type as the repeated measure. There was a significant main effect of experimental condition, $F(2, 51) = 71.12, p < .001, \eta^2 = .74$. There were significant differences between all pairs of conditions (Tukey HSD: intact versus noise-filled gaps, $p < .05, d = 0.45$; intact versus silence-filled gaps, $p < .001, d = 1.79$; noise-filled gaps versus silence-filled gaps, $p < .001, d = 1.52$). The main effect of response type was non-significant, $F(1, 51) = 1.74, p > .05$, because overall a similar number of melody and lyrics responses were made. Most importantly, the interaction between experimental condition and response type was significant, $F(2, 51) = 22.28, p < .001, \eta^2 = .47$.

Examination of this interaction using one-way ANOVAs revealed a significant difference between experimental conditions in the number of correct melody responses, $F(2, 51) = 11.21, p < .001, \eta^2 = .31$. There were significant differences between intact and noise-filled gaps conditions (Tukey HSD: $p < .01, d = 1.07$) and between noise-filled gaps and silence-filled gaps conditions (Tukey HSD: $p < .001, d = 1.40$). The difference between intact and silence-filled gaps conditions was non-significant (Tukey HSD: $p > .05$). Children in the noise-filled gaps condition made more melody responses than children in either the intact or silence-filled gaps conditions.

There was also a significant difference between experimental conditions in the number of correct lyrics responses, $F(2, 51) = 68.90, p < .001, \eta^2 = .73$. There were significant differences between all pairs of conditions (Tukey HSD: intact versus noise-filled gaps, $p < .001, d = 1.96$; intact versus silence-filled gaps, $p < .001, d = 3.06$; noise-filled gaps versus silence-filled gaps, $p < .001, d = 1.50$). Most lyrics responses were made in the intact condition, but more were made in the noise-filled gaps condition than in the silence-filled gaps condition. Overall, significantly more lyrics responses were made in the intact condition when compared with the noise-filled gaps condition, whilst significantly more melody responses were made in the noise-filled gaps condition when compared with the intact condition.

Identification time

As the songs were not equal in duration, the measure of identification time for correctly identified stimuli was taken to be the difference between the song duration and the time taken to identify the song. Data were collapsed across gender, as there were no significant gender differences in identification times (see Appendix 2). The mean identification time in seconds for children in each of the three experimental conditions is shown in Figure 8.3. If these data were to be analysed using a within-subjects factor based on stimulus type, those children who did not identify any songs of one or more stimulus types correctly would be excluded from this analysis (data from 3 children in the noise-filled gaps condition and 13 children in the silence-filled gaps condition would be excluded). Because data from only five children in the silence-filled gaps condition would be included, and thus cannot be analysed using a parametric test, the first analysis considers only intact and noise-filled gaps conditions.

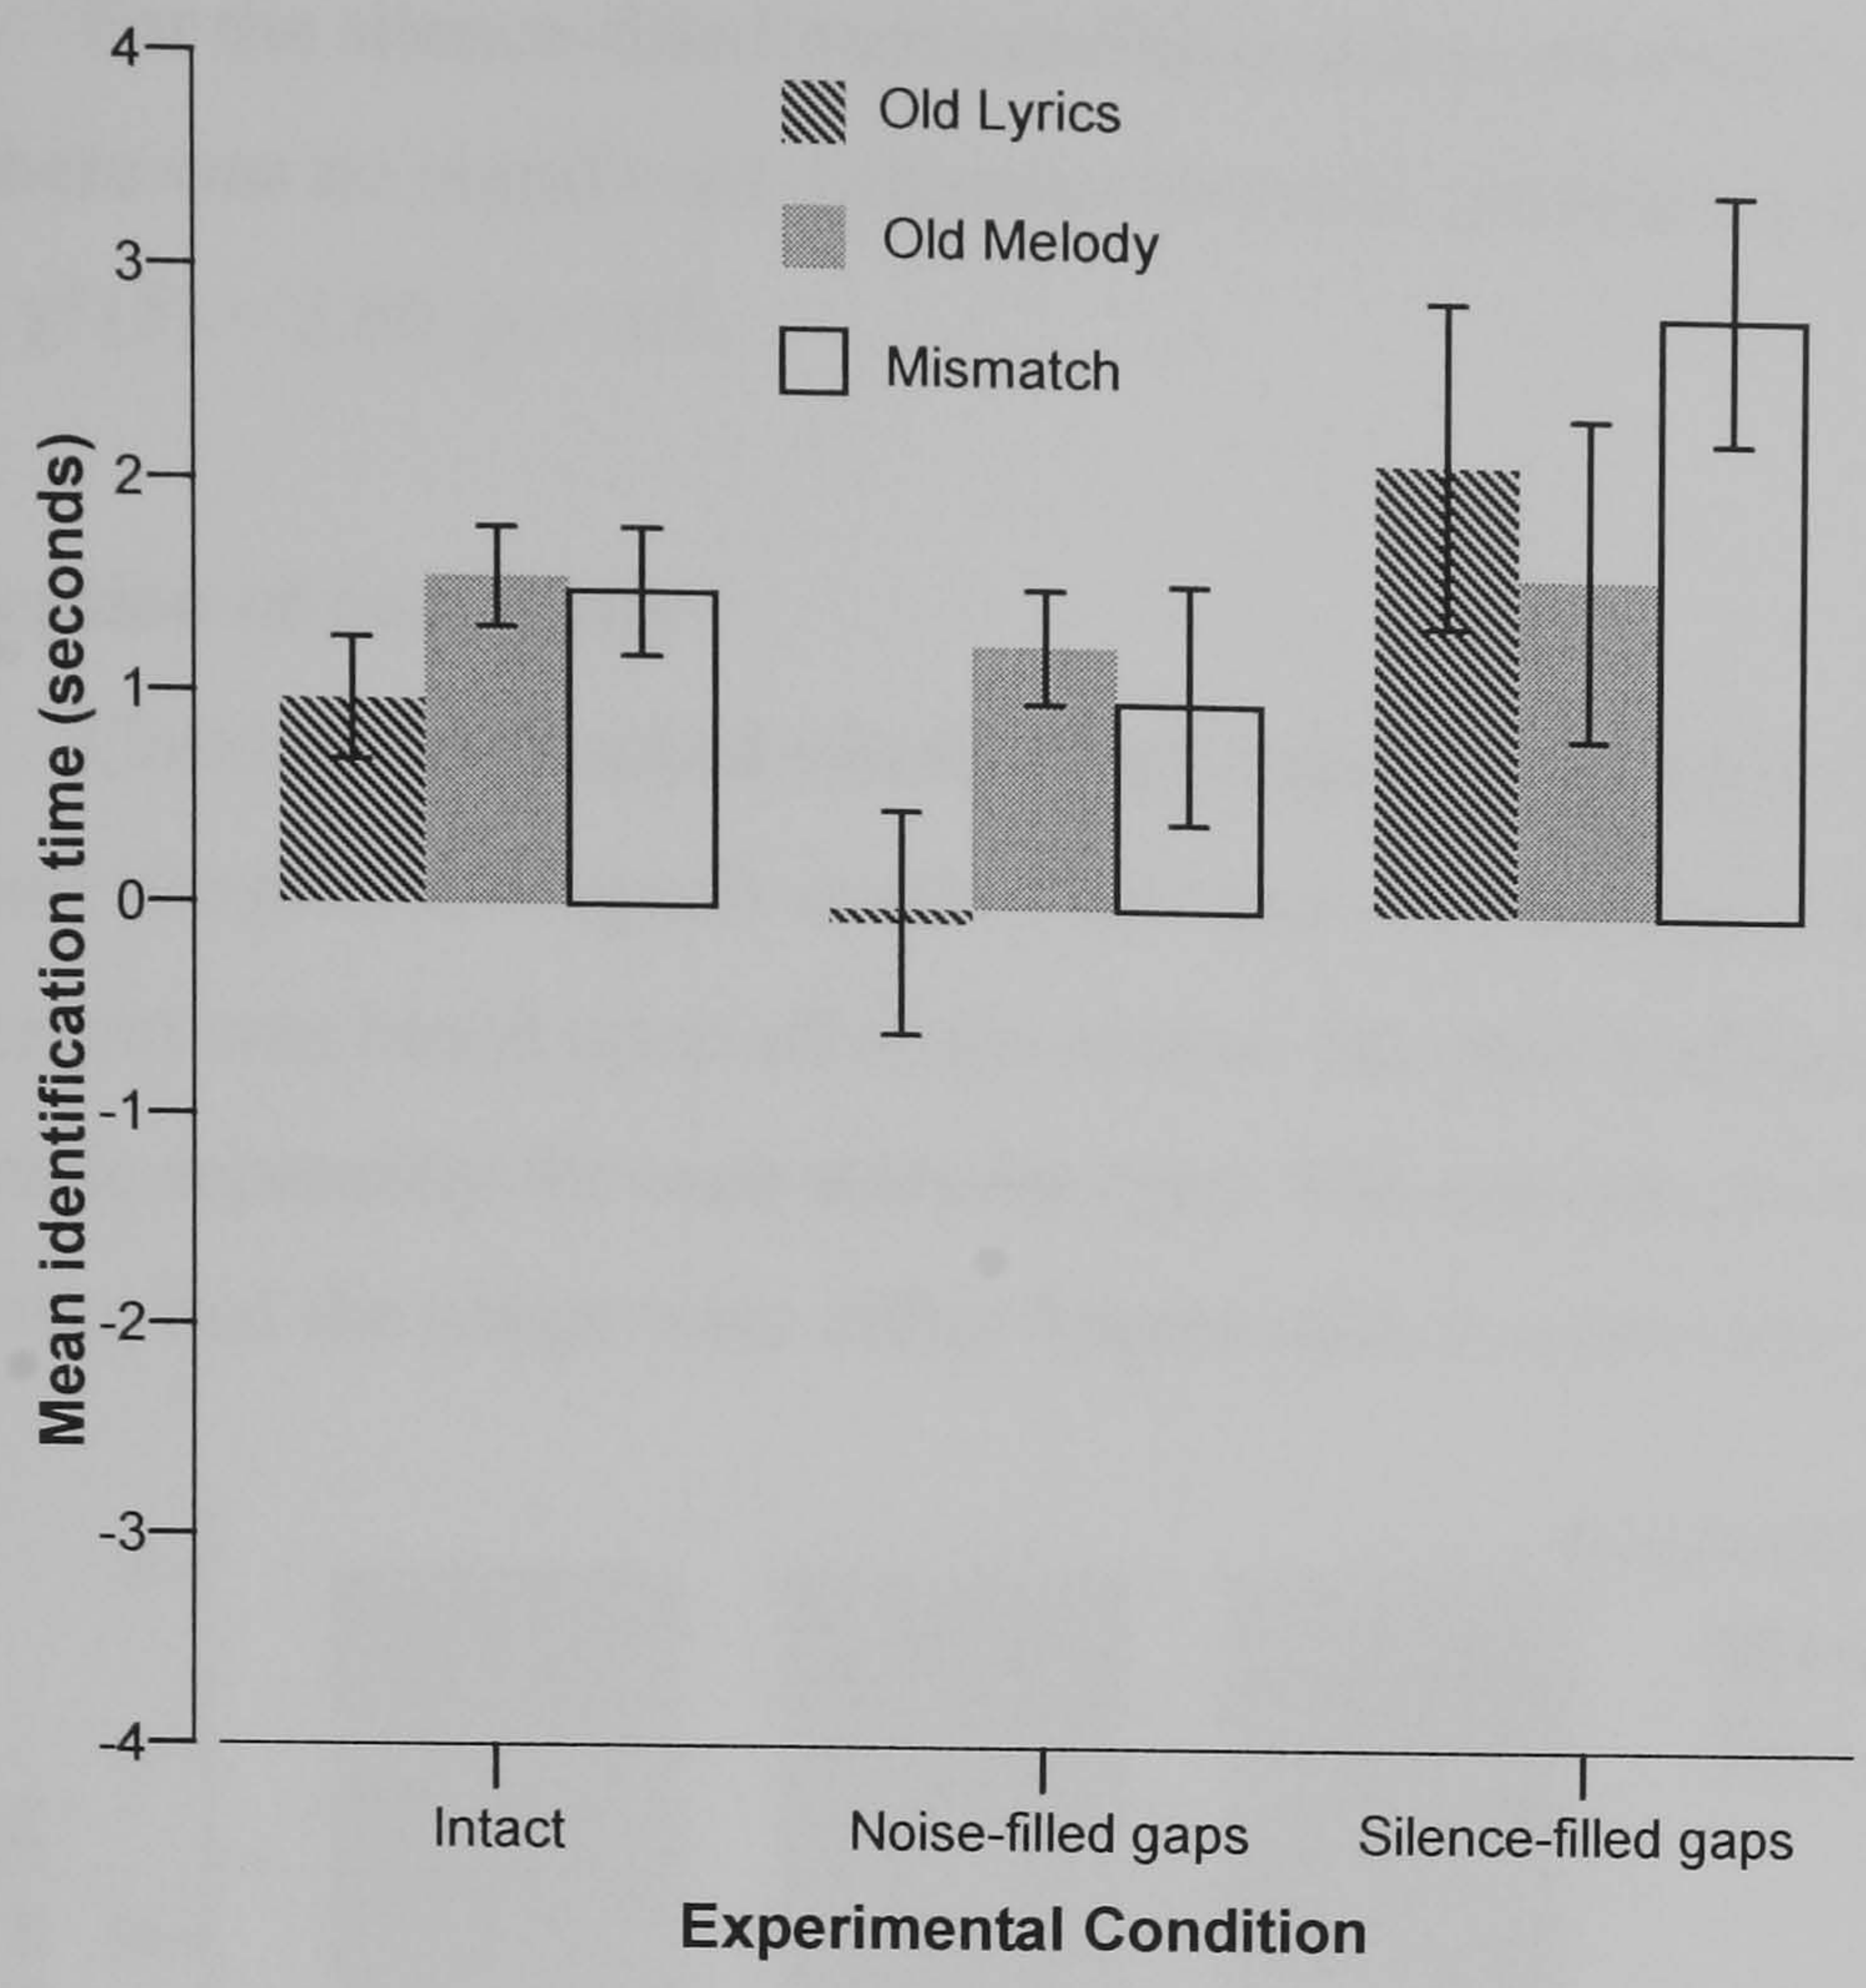

Figure 8.3. Mean identification time (identification time − stimulus duration; ± *SE*) for old lyrics, old melody and mismatch stimuli, for children in the three experimental conditions (intact, noise-filled gaps and silence-filled gaps).

These data were analysed using a 3 (stimulus type: old lyrics; old melody; mismatch) × 2 (experimental condition: intact; noise-filled gaps) mixed ANOVA with stimulus type as the repeated measure. The main effect of stimulus type was significant, $F(2, 62) = 7.68, p < .001, \eta^2 = .20$. Old lyrics stimuli were identified significantly faster than both old melody and mismatch stimuli (Sidak: $p < .001$, $d = 0.54$, and $p < .05$, $d = 0.44$, respectively). There was no significant difference between the time taken to identify old melody and mismatch stimuli (Sidak: $p > .05$).

There was no significant effect of experimental condition on identification time, $F(1, 31) = 2.00, p > .05$; the interaction between stimulus type and experimental condition was also non-significant, $F(2, 62) = 1.07, p > .05$.

For the silence-filled gaps condition, a non-parametric Friedman test revealed that there was no significant difference between stimulus types in terms of identification time, χ^2 (5) = 2.80, $p > .05$.

Perception of continuity

Children were asked whether they thought that little pieces of the tune were missing (fragmented signal) or all of the tune was there (continuous signal). This judgement was based upon all of the stimuli that they had been presented with and was not made separately for each stimulus type. The number of children in each condition reporting that the songs were either fragmented or continuous is show in Figure 8.4.

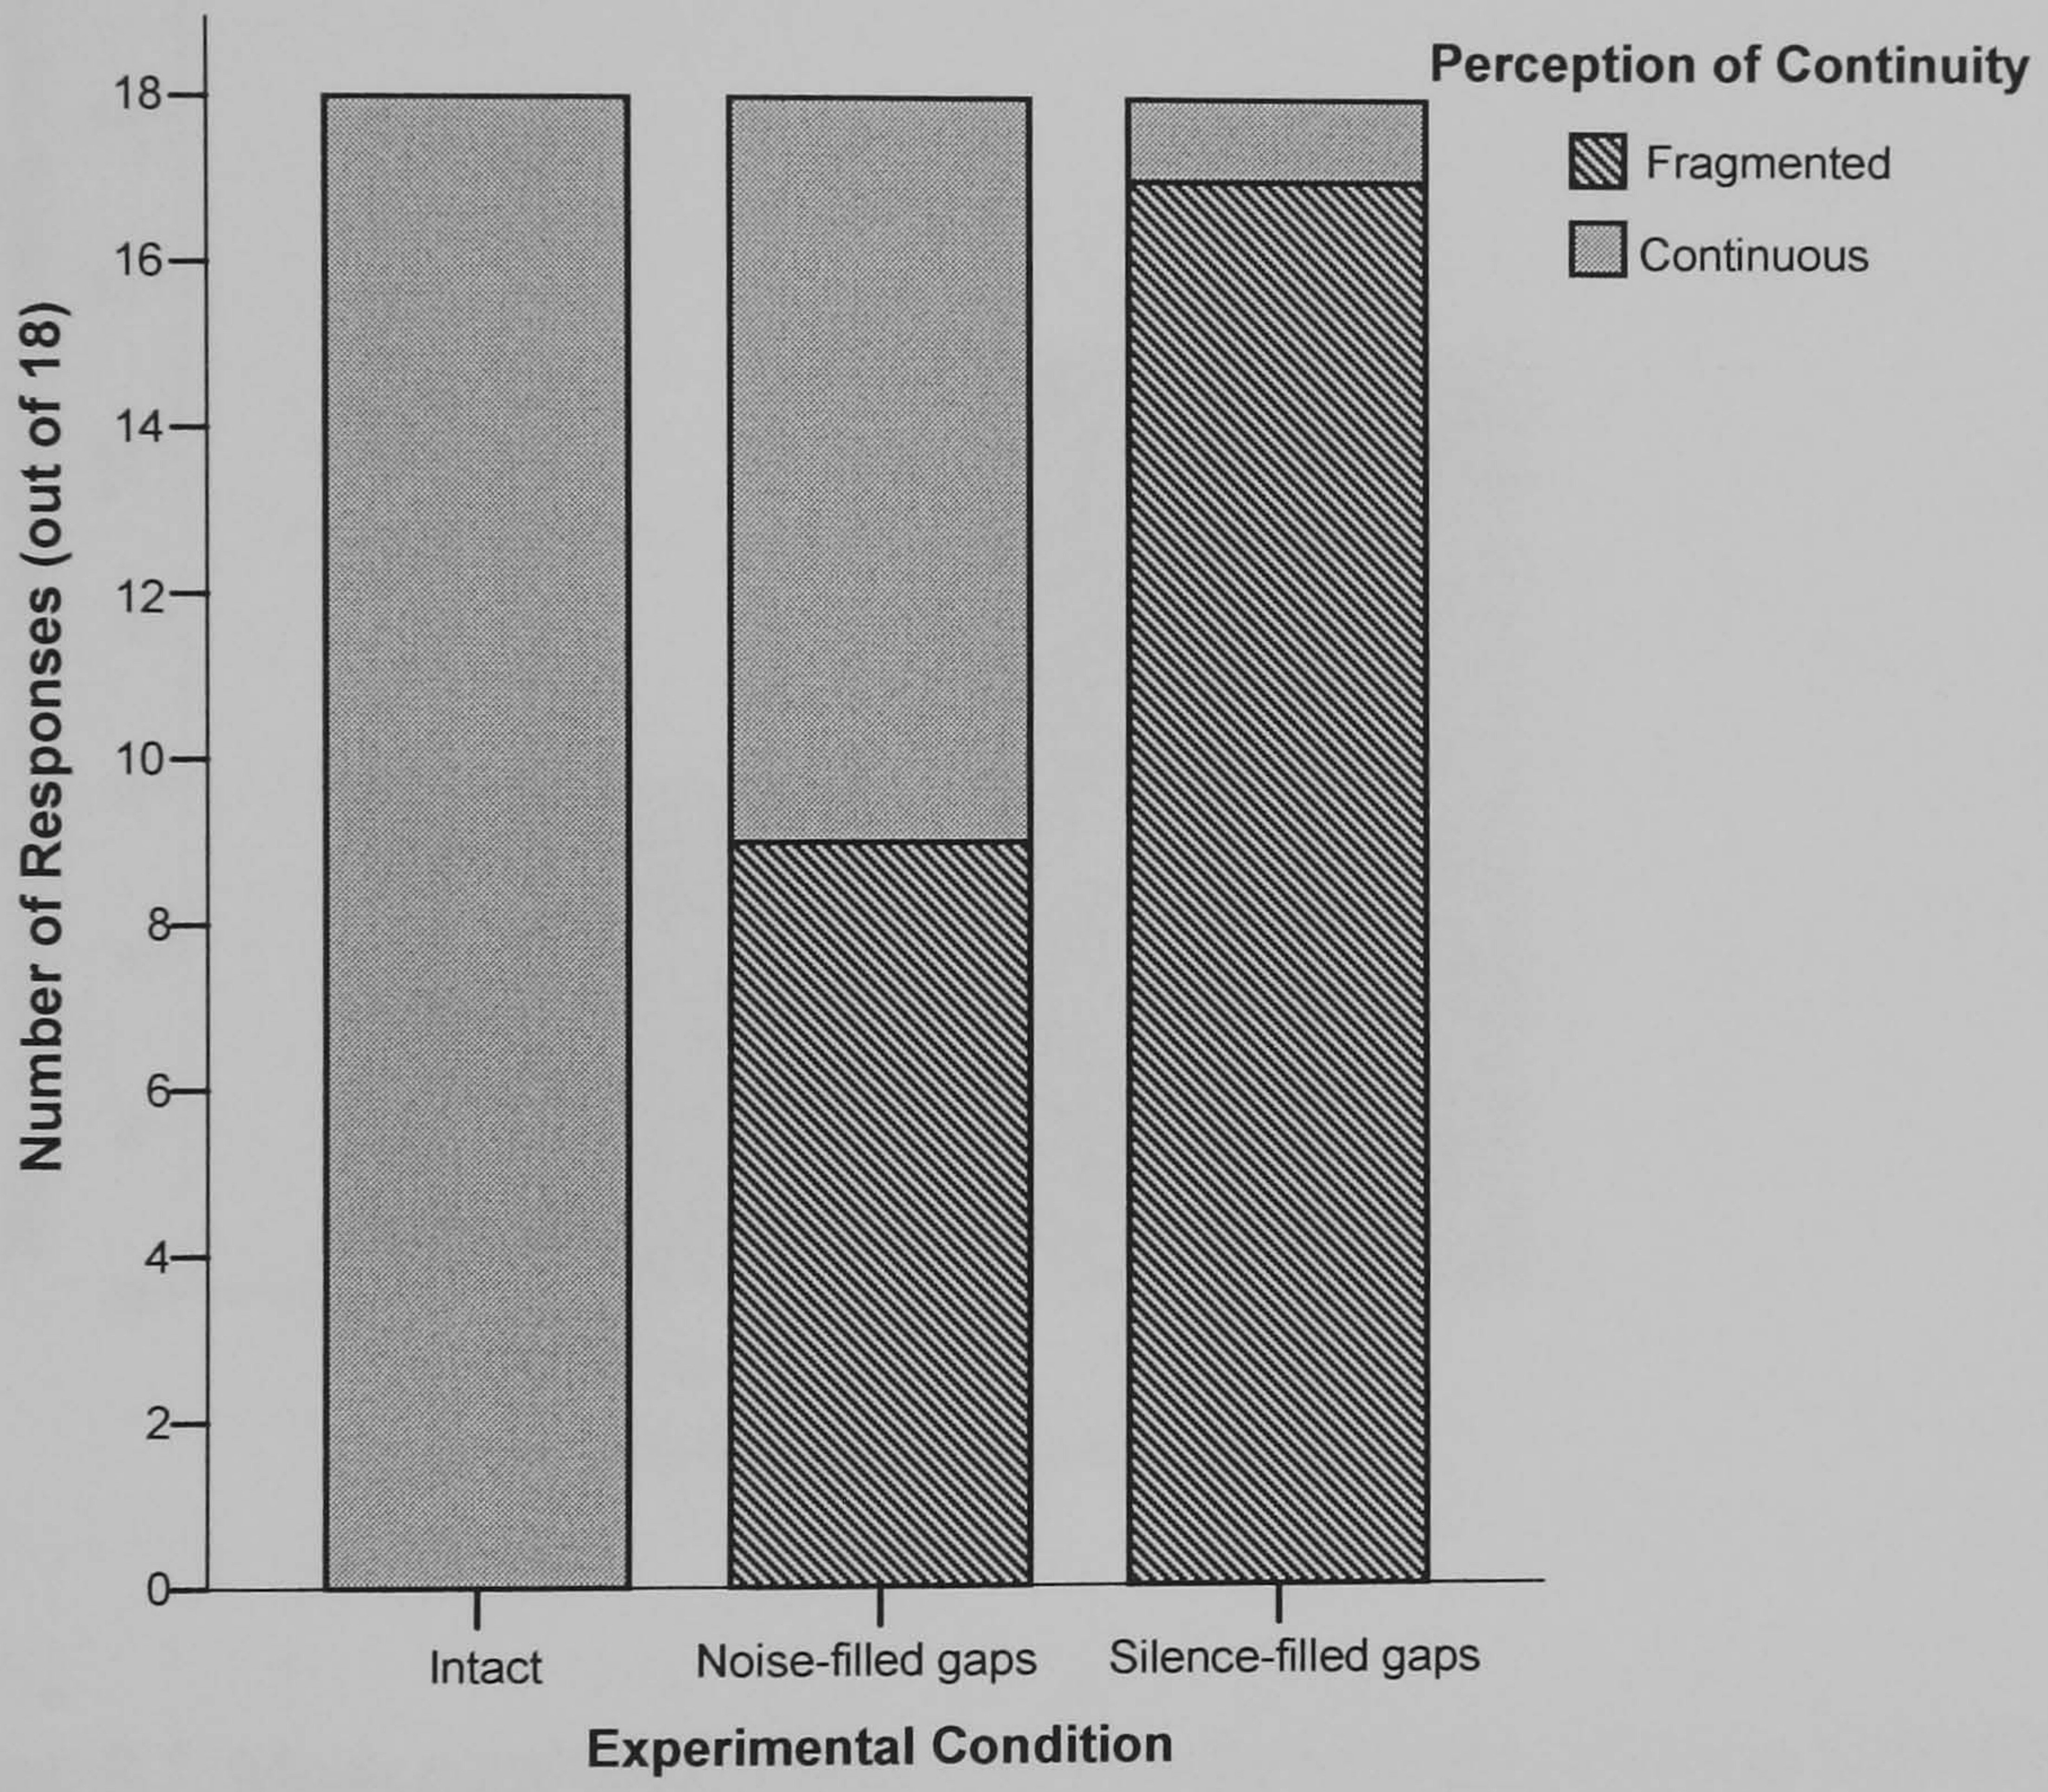

Figure 8.4. Total number of children in each condition (intact, noise-filled gaps, silence-filled gaps) reporting the songs as fragmented and reporting the songs as continuous.

These data were analysed using a 3 (experimental condition: intact; noise-filled gaps; silence-filled gaps) × 2 (perception of continuity: continuous; fragmented) chi-square. A significant association was evident, χ^2 (2, $N = 54$) = 32.19, $p < .001$, Cramer's $V = .77$. As Figure 8.4 shows, all of the children in the intact condition reported the signals to be continuous, and the majority of the children in the silence condition reported the signals to be fragmented. Half of the children in the noise condition reported the signals to be continuous, that is, they did not think that anything was missing, whilst the other half reported the signals to be fragmented.

Based on the previous analysis, children were categorised as either reporting the signals to be fragmented or continuous. The mean number of songs (across all stimulus types) identified correctly by those children who reported the signals as fragmented or continuous is shown in Figure 8.5.

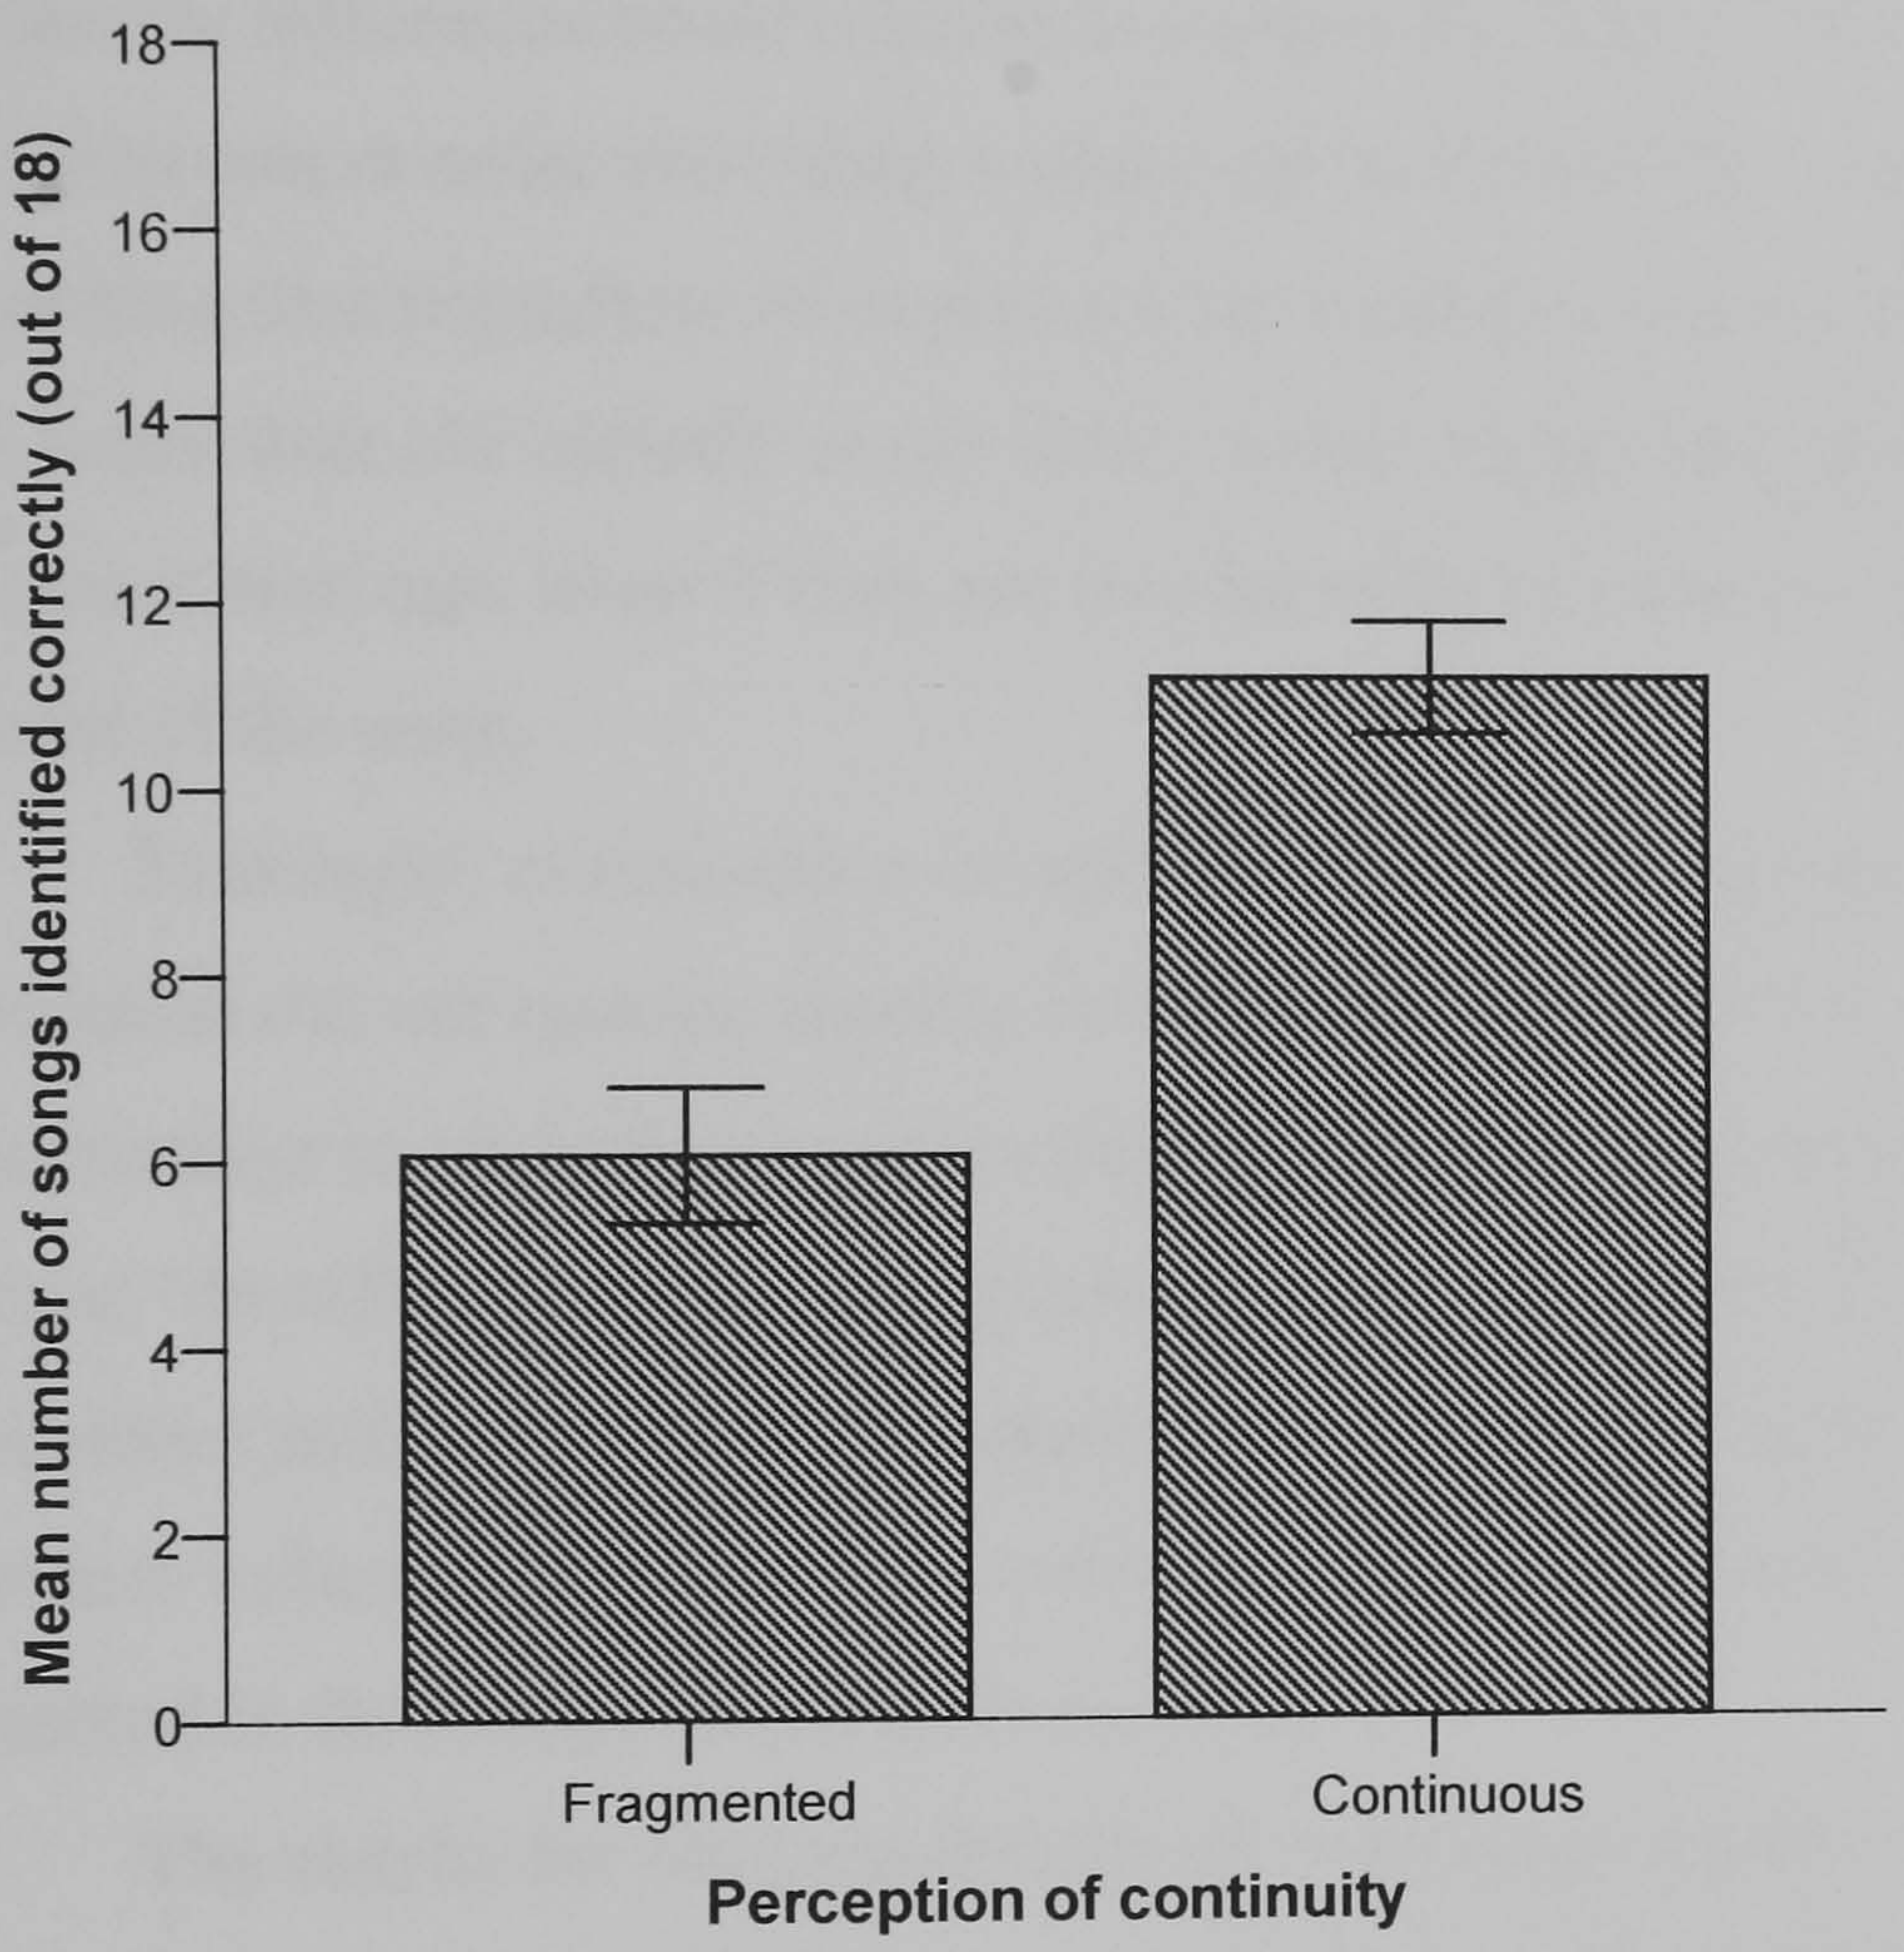

Figure 8.5. Mean number of songs (± *SE*) across all stimulus types identified correctly by children who reported the songs to be fragmented and those who reported the songs to be continuous.

An independent-samples *t* test showed that those children who reported the signals as continuous identified significantly more songs correctly than children who reported the signals as fragmented, $t(52) = 5.42, p < .001, d = 0.49$.

Discussion

The main aim of the present experiment was to investigate the effect of familiarity with the lyrics and melody of a song separately on song identification, when listening to songs intact, and when listening to disrupted versions of the songs.

First, the main effect of experimental condition where there was a significant identification advantage of songs with noise-filled gaps over the same songs with silence-filled gaps can be taken as evidence that perceptual restoration was operating here. Even when melody and lyrics were presented in new contexts, some top-down familiarity influences could operate to restore the content of the missing sections, with the replacement noise providing bottom-up confirmation of these expectations. It is also interesting that regardless of experimental condition, identification of mismatch songs was better than old melody or old lyrics songs, suggesting that two familiar components are better than one, even if they are incongruent, in generating expectations about the content of the song.

Strikingly, examination of each stimulus type separately revealed that perceptual restoration did not operate equally for all stimulus types. For old lyrics stimuli, there was no advantage in identification provided by replacing missing sections with noise over silence; identification was equally poor in both conditions. In contrast, when the signals were intact and undisrupted, children could easily identify the familiar lyrics. This therefore indicates that lyrics are not supported by perceptual restoration when not presented in their original melodic context.

The results for old melody stimuli were quite different. Here, clear evidence of perceptual restoration was demonstrated; children who heard the stimuli with gaps replaced by noise were significantly better at identifying the familiar melody than children who heard the same stimuli with the gaps replaced by silence. This indicates that top down familiarity influences can operate and lead to perceptual restoration when only the melodic component of a song is familiar. Equally interesting is the uncharacteristically poor identification of old melody songs in the intact condition. These children frequently reported that they did not know these songs, whereas children in the noise-filled gaps condition were able to identify them on the basis of the melodic component. How might this result be explained? Due to the fact, as reported in previous research (e.g. Morrongiello & Roes, 1990) that lyrics are used for song identification more than melodies, it is highly likely that in the intact condition children employed their normal strategy of focusing on the lyrics. When the child realised that these were unfamiliar, the child treated the entire song, including the melody, as a "new song", and reported that they did not know it. It appears that when signals are disrupted, children might process the signal more analytically, hence their perception of the melody is not overridden by dominant processing of the lyrics.

With reference to totally undisrupted signals, Serafine et al. (1984) found that adults had difficulty recognising a familar melody when presented with different lyrics. Furthermore, Morrongiello and Roes (1990) found that children had difficulty recognising familiar lyrics with a different melody and familiar melodies with different lyrics. In the intact condition of the present experiment, children were quite able to identify familiar lyrics when paired with a new melody. The possible reason for this difference is that the song components in the present experiment were highly familiar to the children, whereas in Morrongiello and Roes' study they were produced for the experiment and familiarised over a short period. However, the children in their study did find it easier to identify old lyrics than old melody songs, a finding consistent with the intact condition in the present experiment. This does suggest an advantage for lyrics in song identification, at least when songs are heard in clear and undisrupted form. Previous work has suggested that this lyrics advantage results in faster speed of access to stored representations based on lyrics than on melody (Peretz et al., 2004). This is partially supported by the results from intact and noise-filled gaps conditions, where old lyrics stimuli were identified faster than the other stimulus types.

Taken together, these two results indicate that top-down familiarity influences can still operate on perceptual restoration when only one component of a song is familiar, but only if the familiar component is melody. Familiarity is an important part of the perceptual restoration process, and this result shows that an exact match with a previously stored song template is not needed in order for familiarity effects to operate. The components of melody and lyrics are separable to the extent that familiarity with one is enough, but which component is familiar is critical in determining the operation of the effect. This has implications for definitions of familiarity in perceptual restoration: holistic units might not be the critical factor. This result suggests that it is not just overall familiarity that is important in perceptual restoration, but the exact nature of what one is familiar with.

Finally, in the mismatch condition, there was again evidence of perceptual restoration. When two components were familiar but mismatched, children in the noise-filled gaps condition scored significantly higher in terms of identification scores than children in the silence-filled gaps condition. Some interesting results emerged from analysis of whether children made their correct responses on the basis of identifying the melody or the lyrics. For old lyrics stimuli, children responded on the basis of lyrics regardless of experimental condition. This is consistent with earlier work by Serafine et al. (1984). Similarly, in the old melody condition, children mainly responded on the basis of the melody in all three experimental conditions. Again, this is consistent with Serafine et al.'s result. These two findings are hardly surprising; children identify the stimulus using the component that is familiar. However, there was an expectation that children in the old melody condition might simply repeat some of the new lyrics as their answer, but only a few children did so.

The interesting, and somewhat surprising, result came from analysis of children's responses to the mismatch stimuli. In the intact condition, children responded on the basis of the lyrics, that is they gave as their answer the song that the lyrics came from, rather than the song that the melody came from. This again is consistent with Serafine et al. (1984), and also reflects the responses made by children in the study by Morrongiello & Roes (1990) to mismatch songs. It is also consistent with the proposition that when signals are clear and undisrupted, children focus on processing the lyrics to aid identification, not the melody. However, in the noise-filled gaps condition (and also in the silence-filled gaps condition although the response type difference is based on a very small number of correct responses), children identified the melody of the mismatch songs more than they identified the lyrics. This is suggestive of a strategy shift when trying to identify disrupted songs when compared to intact songs. This result further indicates that there is more than one way of accessing the musical lexicon in the process of song identification. It appears that what is most useful for identifying an intact song might not be most useful when the song is disrupted, such as in the stimuli presented here with noise- and silence-filled gaps. Children use the component best supported by perceptual restoration, the melody, because their usual strategy of focusing on the lyrics appears to be less efficient for disrupted signals.

Taken together, these results imply that melody is better supported by perceptual restoration than lyrics. This suggests that strong top-down familiarity cues operate for melody, and that the result from Experiment 5 where songs with lyrics and melody were easier to identify than songs with melody only, is indicative of an additive effect of lyrics on top of melody, rather than the privileged status of lyrics in aiding perceptual restoration.

However, familiarity might not be the only explanation for this finding. On a cognitive level, it is possible that the lyrics improve overall identification not because they are used more than melody, but because lyrics support the rhythmic and pitch structure of the melody, and hence make the melody more salient (Bentley, 1966). In view of the fact that the melody appears to be particularly important for identifying disrupted input, this could explain why the combination of lyrics and melody leads to superior identification than the melody alone.

It is unclear from the present results whether melody is better transmitted though noise than lyrics, or whether in noisy conditions processing focuses on the melody because it is best supported by perceptual restoration. A possible explanation centres around the preference for maintaining global coherence that the perceptual system demonstrates. It is well known that disrupted words can be made coherent through perceptual restoration, since there is evidence that the perceptual restoration of speech occurs, even in children (e.g. Newman, 2004). However, it may be that making a melody coherent is easier, more automatic or less effortful than making words coherent in noise. This coherence could affect the interaction between expectation (a product of familiarity) and confirmation provided by the masking potential of the replacement noise.

Consideration of both children's identification scores and their reports of the continuity of the signals becomes particularly important in considering a potential problem with this experiment. Comparing identification of stimuli across conditions is a useful way of investigating the operation of perceptual restoration. However, it is possible that because the content of the nursery rhymes was manipulated here, children could be unable to identify tunes or lyrics, but restoration of continuity, in terms of the perception of a continuous signal, could still occur. Whilst subjective judgements from young children need to be interpreted with a degree of caution, it appears from the present results that this is unlikely to be the case. Only half of the children in the noise-filled gaps condition thought that the signals were continuous rather than fragmented. In addition, the lack of a clear majority for "continuous" responses in the noise-filled gaps condition most likely represents the lack of perceptual restoration operating for some stimuli presented; had this judgement been made separately for each stimulus type different results may well have emerged.

To summarise, some possible explanations for questions posed earlier in this chapter are provided. First, it was asked whether there might be separate top-down familiarity effects operating for the perceptual restoration of melody and lyrics when in different contexts, or whether familiarity influences require an exact match with the complete song representation. The present result indicates that an exact match is not needed, but there is an asymmetry in as much as familiarity with the melody is enough to lead to perceptual restoration; the same is not true for lyrics. Second, it was asked whether lyrics still have an advantage and dominance in song identification in the case where songs are disrupted. Clearly, they do not. Third, and as an extension of this previous question, it was asked which component would be used for identification with a disrupted signal. Here, it seems that melody has an advantage. How this advantage operates cannot be understood based on the present results, but represents an important topic for future research.

Finally, it was asked whether lyrics play a key role in song identification and whether such an advantage only occurs in the context of the accompanying melody. Thus, the present experiment was designed to clarify the reasons for the finding from Experiment 5 that songs with melody and lyrics were easier to identify, but did not lead to stronger perceptual restoration, than melodies alone. On the basis of the present findings, it is possible that this result represents a purely additive effect of two familiar components over one, where it is likely that lyrics serve to make melodic characteristics more salient. It is also possible that lyrics did not play the same role in identification across the conditions where they were presented in Experiment 5.

With further reference to the findings from Experiment 5, it is highly unlikely based on the findings from the present experiment that the perceptual restoration of songs operates on the basis of the lyrics alone. Superior identification of songs as opposed to melodies, at least for the noise-filled and silence-filled gaps conditions in Experiment 5, cannot simply be the result of children attending only to the lyrics and using them as the basis for identification. If that were the case, then perceptual restoration would have operated in the present experiment in any case where the lyrics were familiar. This result did not emerge. It is also likely that if children use lyrics regardless of melody, when new unfamiliar lyrics were presented with a familiar melody (old melody trials), children would simply repeat some of the lyrics as the basis of their identification response. Very few children did so in any experimental condition. Instead, the result from Experiment 5 can be best explained as the result of an additive effect, where lyrics can only be used for identification of disrupted stimuli when heard in the context of the correct accompanying melody.

These findings indicate that the rules governing music perception of undisrupted signals must not be applied to perception of disrupted signals without first considering whether strategies used might need to be adapted. The present result certainly suggests that a strategy shift occurs when listening to songs with fragments missing, where different song components are most useful for song identification when songs are intact and when they are disrupted. For this reason, the perceptual restoration of music seems to represent a special case of music perception where, in order for comprehension to remain unaffected, different strategies must be employed to those used when an undisrupted signal is perceived.

However, whilst it is clear from the results of the present study that melody is the primary identification cue used in the context of a disrupted signal, it is not clear whether it is the melodic contour (pitch) or the temporal variations (rhythm) of a melody that is the most important in assisting identification of disrupted songs. Both components were changed in order to create new melody stimuli. The following experiment attempts to discover whether rhythm or pitch is more important for perceptual restoration than the other, or whether congruent expectations on both a pitch and rhythmic level are required for the perceptual restoration of musical input in young children.

Chapter 9: Experiment Seven

The role of pitch contour and rhythm in the perceptual restoration of melodies in children

Introduction

Musical signals such as songs and tunes are composed of many different components (e.g. lyrics, melody, harmony, timbre, etc.). The results of Experiment 6 suggest that melody is used more than lyrics for identifying disrupted songs, whereas the opposite is true for intact songs. However, the melody of a song is itself composed of two components, the patterns of pitch and the patterns of rhythm (i.e. frequency and temporal information, respectively). These are the primary dimensions of music (Krumhansl, 2000), and are what make a tune distinctive. Even very young children make the distinction between these two components (Bentley, 1966). It is possible that in the same way as there is an asymmetry in the role of the melody and lyrics of a song in the process of perceptual restoration, there might be a similar asymmetry in the role of the frequency and temporal components of melody itself in the process of musical perceptual restoration in young children.

Pitch- and rhythm-based expectations: The *what* and *when* of music perception

Expectations about what would most likely be heard next in a piece of music are crucial for musical perceptual restoration. Expectations in the context of music perception are described by Schmuckler & Boltz (1994) as: "...a listener's ability to anticipate, with varying degrees of specificity, upcoming musical events on the basis of previously heard musical events" (p.313). Such expectations develop for both pitch and rhythm as a piece of music proceeds (Bigand, 1997; Lee, 1985; Remez, Rubin, Berns, Pardo & Lang, 1994).

It has been proposed that whilst pitch-based expectations are crucial to melodic identification in adults, the rhythmic context influences the pitch-based expectations that contribute to identification. This is because the rhythmic context tells the listener when to expect certain pitch sequences, and sets the pitch-based expectations in a temporal context (Fraisse, 1982; Jones & Boltz, 1989; Schmuckler & Boltz, 1994; Smith &

Cuddy, 1989). The temporal dimension of a melody is just as important as the pitch contour dimension in generating expectations about what is to follow. This is termed the "temporal expectancy hypothesis" (Kidd, Boltz & Jones, 1984). According to this hypothesis, the rhythmic structure of a melody is used to direct attention to specific points in time at which a pitch event is expected. By setting individual sounds within a temporal context, this temporal structure and organisation means that expectations influence not only what is expected, but also when sounds are expected to occur (Warren & Warren, 1971).

It is proposed by Kidd et al. (1984) that the concept of expectation binds pitch and rhythm together, and that an incongruous rhythmic context seriously disrupts the process of pitch expectation. Boltz (1993) argues that the temporal structure of a melody guides the processing of the pitch, and makes pitch relationships more salient, and as such the temporal dimension of a melody might play a stronger role in melody identification than pitch. This view was strengthened by findings that when pitch and temporal expectations conflicted, temporal expectations were the most important for adults, because they guided the process of pitch expectation. According to Longuet-Higgins and Lee (1982), the continual generation and updating of expectations is the "core" of rhythm perception, with similar claims for pitch perception being made by Krumhansl (1995).

There exist in the literature two theories about the relationship of pitch and rhythm (Bigand, 1997). According to the single-component model, pitch and rhythm are integrated in terms of processing, whereas the two-component model represents pitch and rhythm as separable components with separate processing resources. If the single-component model is correct, it is possible that expectations will contain an integration of rhythm and pitch. If, however, the two-component model is correct, it is possible that expectations for pitch and rhythm can operate independently of each other. This distinction could have important implications for the perceptual restoration of melodies.

Evidence for the single-component model comes from studies where rhythm and pitch processing appear to interact on some level and at some stage of processing, even if they are separate at an earlier stage of processing and later integrated (e.g. Foxton, Nandy & Griffiths, 2006; Peretz & Morais, 1989). Evidence for the two-component model mainly comes from neuropsychological studies indicating separable processing of the two components (e.g. Parsons, 2003), and dissociable effects of brain damage (e.g. Peretz & Kolinsky, 1993; Piccirilli, Sciarma & Luzzi, 2000). Peretz (2006) argues that there exists a functional separability between pitch and rhythm in music perception. This has even been described in terms of separate sub-modules for rhythm and pitch within a general music module (Peretz & Coltheart, 2003).

According to Narmour (1977), music is “multiply implicative”, because expectations develop for separate components of a musical signal (e.g. pitch, harmony, rhythm). However, these individual levels of expectation interact to produce holistic expectations that cannot be accounted for by the sum of the individual levels of expectation. This suggests that whilst components such as rhythm and pitch might generate their own expectations, these interact to generate expectations for upcoming musical events in a familiar tune.

If rhythm and pitch processing are separable, the question arises as to which is most important in melody identification, and which provides the best access code to the stored representation of a familiar piece of music (Hébert & Peretz, 1997). According to Dowling (1999), pitch contour is the most salient feature of melody from infancy through to childhood. Further evidence for the primary role of pitch in melody identification comes from research with both children and adults (e.g. Dalla Bella et al., 2003; Morrongiello et al., 1985; Trehub, 2001), but the primacy of pitch-based information might be specific to Western music (Peretz, 1993).

Conversely, according to Hargreaves (1986), rhythm is the most fundamental aspect of a melody, with rhythmic responses to music developing before any other, in the form of kinaesthetic responses to a steady pulse. This relationship between rhythm and motor functioning is further emphasised by Clarke (1999) and Thackray (1972). Peretz (2006) describes rhythm as the basic “essence of music” (p.14), with Hannon & Trainor (2007) arguing that temporal structure is more fundamental to music than pitch. It is likely, however, that conflicting views of the role of pitch and rhythm in melody identification occur as a result of conclusions being based on the results of a variety of different paradigms and tasks used to assess the function of rhythm and pitch.

When considering which component is more useful for melody identification, the most conclusive evidence comes from previous studies where the two have been put into conflict to see which is most useful. Hébert & Peretz (1997) manipulated the rhythmic and pitch components of familiar music to see which adults found the most useful for melody recognition. The pitch contour of each familiar melody was played on isochronous notes (of equal duration), and the rhythm of each melody was played at a constant pitch. Mismatch stimuli were also presented, where the familiar pitch contour from one melody was played to the familiar rhythm of another melody. Hébert and Peretz found that identification of the stimuli was superior where the pitch contour was the familiar component than where the rhythm was the familiar component. Rhythmic information alone was found to be poor at accessing the stored representation of the music.

In the case of the mismatch stimuli, listeners could not override the pitch information to focus on the rhythm even when explicitly instructed to do so. Despite this, the results indicated that even though pitch contour seems to predominate over rhythm, identification was the best when pitch and rhythm were presented in their usual combination. Therefore, it appears that rhythm does play a role in melody identification, even if it is an insufficient basis for identification in isolation. Hébert and Peretz argue that pitch information is the minimum amount of information required for activation of a stored melody representation. Rhythm alone is not enough. They suggest that this could be due to the similarity of rhythms when compared to pitch patterns, leading to less distinctiveness of different rhythmic patterns at the encoding stage.

These results confirm earlier findings by White (1960), who manipulated pitch and rhythm in a similar way by presenting familiar tunes with all notes being of equal duration, and with all notes played at the same pitch. In a similar way to Hébert and Peretz (1997), White found that eliminating the rhythmic information had least impact on adults' identification. This was also described as a result of the similarity between many rhythmic patterns. This rhythmic similarity is particularly true of nursery rhyme melodies; hence, it is reasonable to expect that similar results might emerge if young children are presented with a similar task.

Whilst the research reviewed here leads to several possibilities of the role of pitch and rhythm in musical perceptual restoration, Experiment 6 revealed that what is important for identification of intact and disrupted information can be completely different. It is also possible that the information used by children for melody identification will be different to that used by the adult participants in the previous research reported here. Furthermore, of primary interest in the present experiment is not which component is most useful for melody identification, but which is best supported by perceptual restoration.

If pitch and rhythm really are separable components of a familiar melody, each component would be expected to generate its own expectations from its own representation. Thus, perceptual restoration could possibly operate on the basis of familiarity with the pitch or rhythm alone, if the expectations are strong enough to contribute to the restoration process. Previous research would indicate that the pitch contour will be most useful for identification, but might not necessarily be best supported by perceptual restoration. However, if rhythm and pitch interact then it might be possible that whilst expectations for each component are important, the correct combination of the two is required for identification, and for perceptual restoration to occur. The present experiment was designed as a test of these two possibilities in the context of the perceptual restoration of music in young children.

Method

Participants

Sixty children (30 girls) between the ages of 4 years 3 months and 6 years 2 months participated in this experiment, with a mean age of 5 years 3 months (63.00 months; *SD* = 6.74). None of these children had participated in previous experiments. The children were recruited from the Reception and Year One classes of a primary school in Surrey, UK, and of a primary school in Dorset, UK, both in predominantly white, middle-class areas. Consent for each child's participation was obtained from both the child's parent and the school Headteacher. All participants were native English speakers with normal sensory development, as reported by the school. Further details about the participants in each experimental condition are given in Table 9.1.

Table 9.1. The gender split and mean age of the children in each of the three experimental conditions.

Experimental Condition[a]	**Gender**	**Mean age in months (*SD*)**
Intact	M=10; F=10	62.70 (7.09)
Noise-filled gaps	M=10; F=10	62.80 (6.07)
Silence-filled gaps	M=10; F=10	63.50 (7.32)

[a] *n* = 20 for each condition

Design

Children were randomly allocated to conditions in a 3 × 3 mixed design. Each child was asked to identify 18 stimuli; six stimuli of each of three types: old (original) pitch, old (original) rhythm, and mismatch; tested within-subjects (see Appendix 10 for examples of the different stimulus types). Each child did so in one of three independent conditions of the task (intact, noise-filled gaps or silence-filled gaps; between-subjects). Every attempt was made to keep the design of the present experiment as similar as possible to that of Experiment 6. For this reason, familiar pitch patterns were accompanied by new rhythms and familiar rhythms accompanied by new pitch patterns for old pitch and old rhythm stimuli, respectively. This was in contrast to setting familiar pitch patterns to an isochronous rhythm (notes of equal duration) and familiar rhythms to notes of constant pitch, as was the case in previous research (e.g. Hébert & Peretz, 1997; White, 1960).

Materials

The stimuli were produced on an Oxygen8 MIDI keyboard and edited in accordance with the procedures used in previous experiments. The same six nursery rhymes presented in Experiment 6 were used as stimuli in the present experiment, with a single exception. One stimulus, *The Hokey Cokey*, was substituted for *Head, Shoulders, Knees and Toes*; the rhythmic component of *The Hokey Cokey* alone was not distinctive as the rhythm contained many isochronous notes. Furthermore, for the same reasons as described for Experiment 6, all nursery rhymes in the original set were considered as possible mismatch partners for the mismatch stimuli.

Three demonstration trials were created, one of each stimulus type. As for Experiment 6, *Baa Baa Black Sheep*, *The Grand Old Duke of York*, and *Row, Row, Row Your Boat* were used for the purpose of demonstration trials. Which nursery rhyme of the three represented each stimulus type was decided randomly. All stimuli were presented in a pseudo-random order using a programme written in Matlab V5.2.1, to ensure that no two stimuli of the same type would follow one another (e.g. *Twinkle Twinkle Little Star* old pitch followed by *Old MacDonald had a Farm* old pitch), and no two examples of the same song would follow each other (e.g. *The Wheels on the Bus* old pitch followed by *The Wheels on the Bus* mismatch). The two participating schools were sent the song familiarity questionnaire (Appendix 1) prior to participation.

Old Pitch, New Rhythm stimuli

For these stimuli, a new rhythm was written to accompany the pitch contour of the original nursery rhyme. Thus, the pitch contour (the pattern of ups and downs in the melody) was familiar, but the rhythm (the temporal variation in the notes) was unfamiliar. The tempo of the new rhythm was matched to that of the original rhythm; the stimuli had an average length of 5.01 seconds. New rhythms consisted of the simple rhythmic structure characteristic of nursery rhyme melodies. Furthermore, it was ensured that no three notes in succession consisted of the same time value as in the original melody (e.g. if the original rhythm contained a pattern of two quavers followed by a crotchet at a particular point within the rhythm, this pattern was not repeated within the new rhythm). New, distinctive rhythms were produced by varying both the note durations, and in some cases the time signature, in order to change the rhythmic structure.

Old Rhythm, New Pitch stimuli

To create these stimuli, a new pitch pattern was written to accompany the original rhythm. The pitch patterns used were those that were composed as part of the old lyrics stimuli in Experiment 6. In using the pitch component of these stimuli, it had previously been checked that they were comparable to the melodies from the original nursery rhymes in terms of musical characteristics (see Appendix 8). New stimuli were matched in terms of tempo to the original nursery rhyme stimuli, and the length of the excerpt was matched to within 500 ms of the length of the original excerpt. The average length of the stimuli was 4.91 seconds.

Mismatch stimuli

Each original pitch pattern was paired with the rhythm from a different nursery rhyme; the stimuli had an average length of 5.24 seconds. Where possible, no alterations were made to the number of notes within the pitch pattern to accommodate the new rhythm. However, this was not always possible and in some cases, slight changes were made (see Appendix 11); this involved splitting a note value in half to create an extra note. The contour of the pitch component, however, remained unaltered and the rhythmic component was unchanged.

Procedure

Each child was tested individually in a quiet room in the school. The child was told that they were going to play a music game, where their task was to work out what some children's music was. Each child was first presented with three demonstration trials (*Row, Row, Row your Boat* as an old rhythm stimulus type, *Baa Baa Black Sheep* as an old pitch stimulus type, and *The Grand Old Duke of York* as a mismatch stimulus type). In order to avoid biasing a child to respond in a consistent way across all stimulus types, no further explanation was given if the child did not identify the nursery rhyme from the demonstration trial. Instead, the next trial followed immediately.

Following presentation of the demonstration trials, the child was presented with the 18 experimental trials (six of each stimulus type). The standard instructions presented across all previous experiments were used to ask the child to name each stimulus when it was presented. No time limit was imposed upon the child giving an answer, but if the child did not identify the stimulus, or gave no response, this was scored as incorrect and the next trial followed. No feedback was given as to whether the child's response was correct or incorrect.

The experimenter entered into the programme whether the child had given a correct response on each trial, and whether this response was made on the basis of the pitch or rhythmic element of the stimulus. Once all 18 experimental trials had been presented, the child was asked the continuity question, that is, whether they thought that "little pieces of the tune were missing" or "all of the tune was there". This judgement was based on all trials presented and was not made separately for each stimulus type, for the reasons reported for Experiment 6. The child's response was noted, and they were praised and thanked for their participation.

Results

Correct identification scores

Each child was given a score out of six according to the number of songs of each stimulus type that they identified correctly. There were no gender differences in identification scores (see Appendix 2) and so data were collapsed across gender for further analyses. Figure 9.1 shows the mean number of stimuli of each type identified correctly by children in each experimental condition.

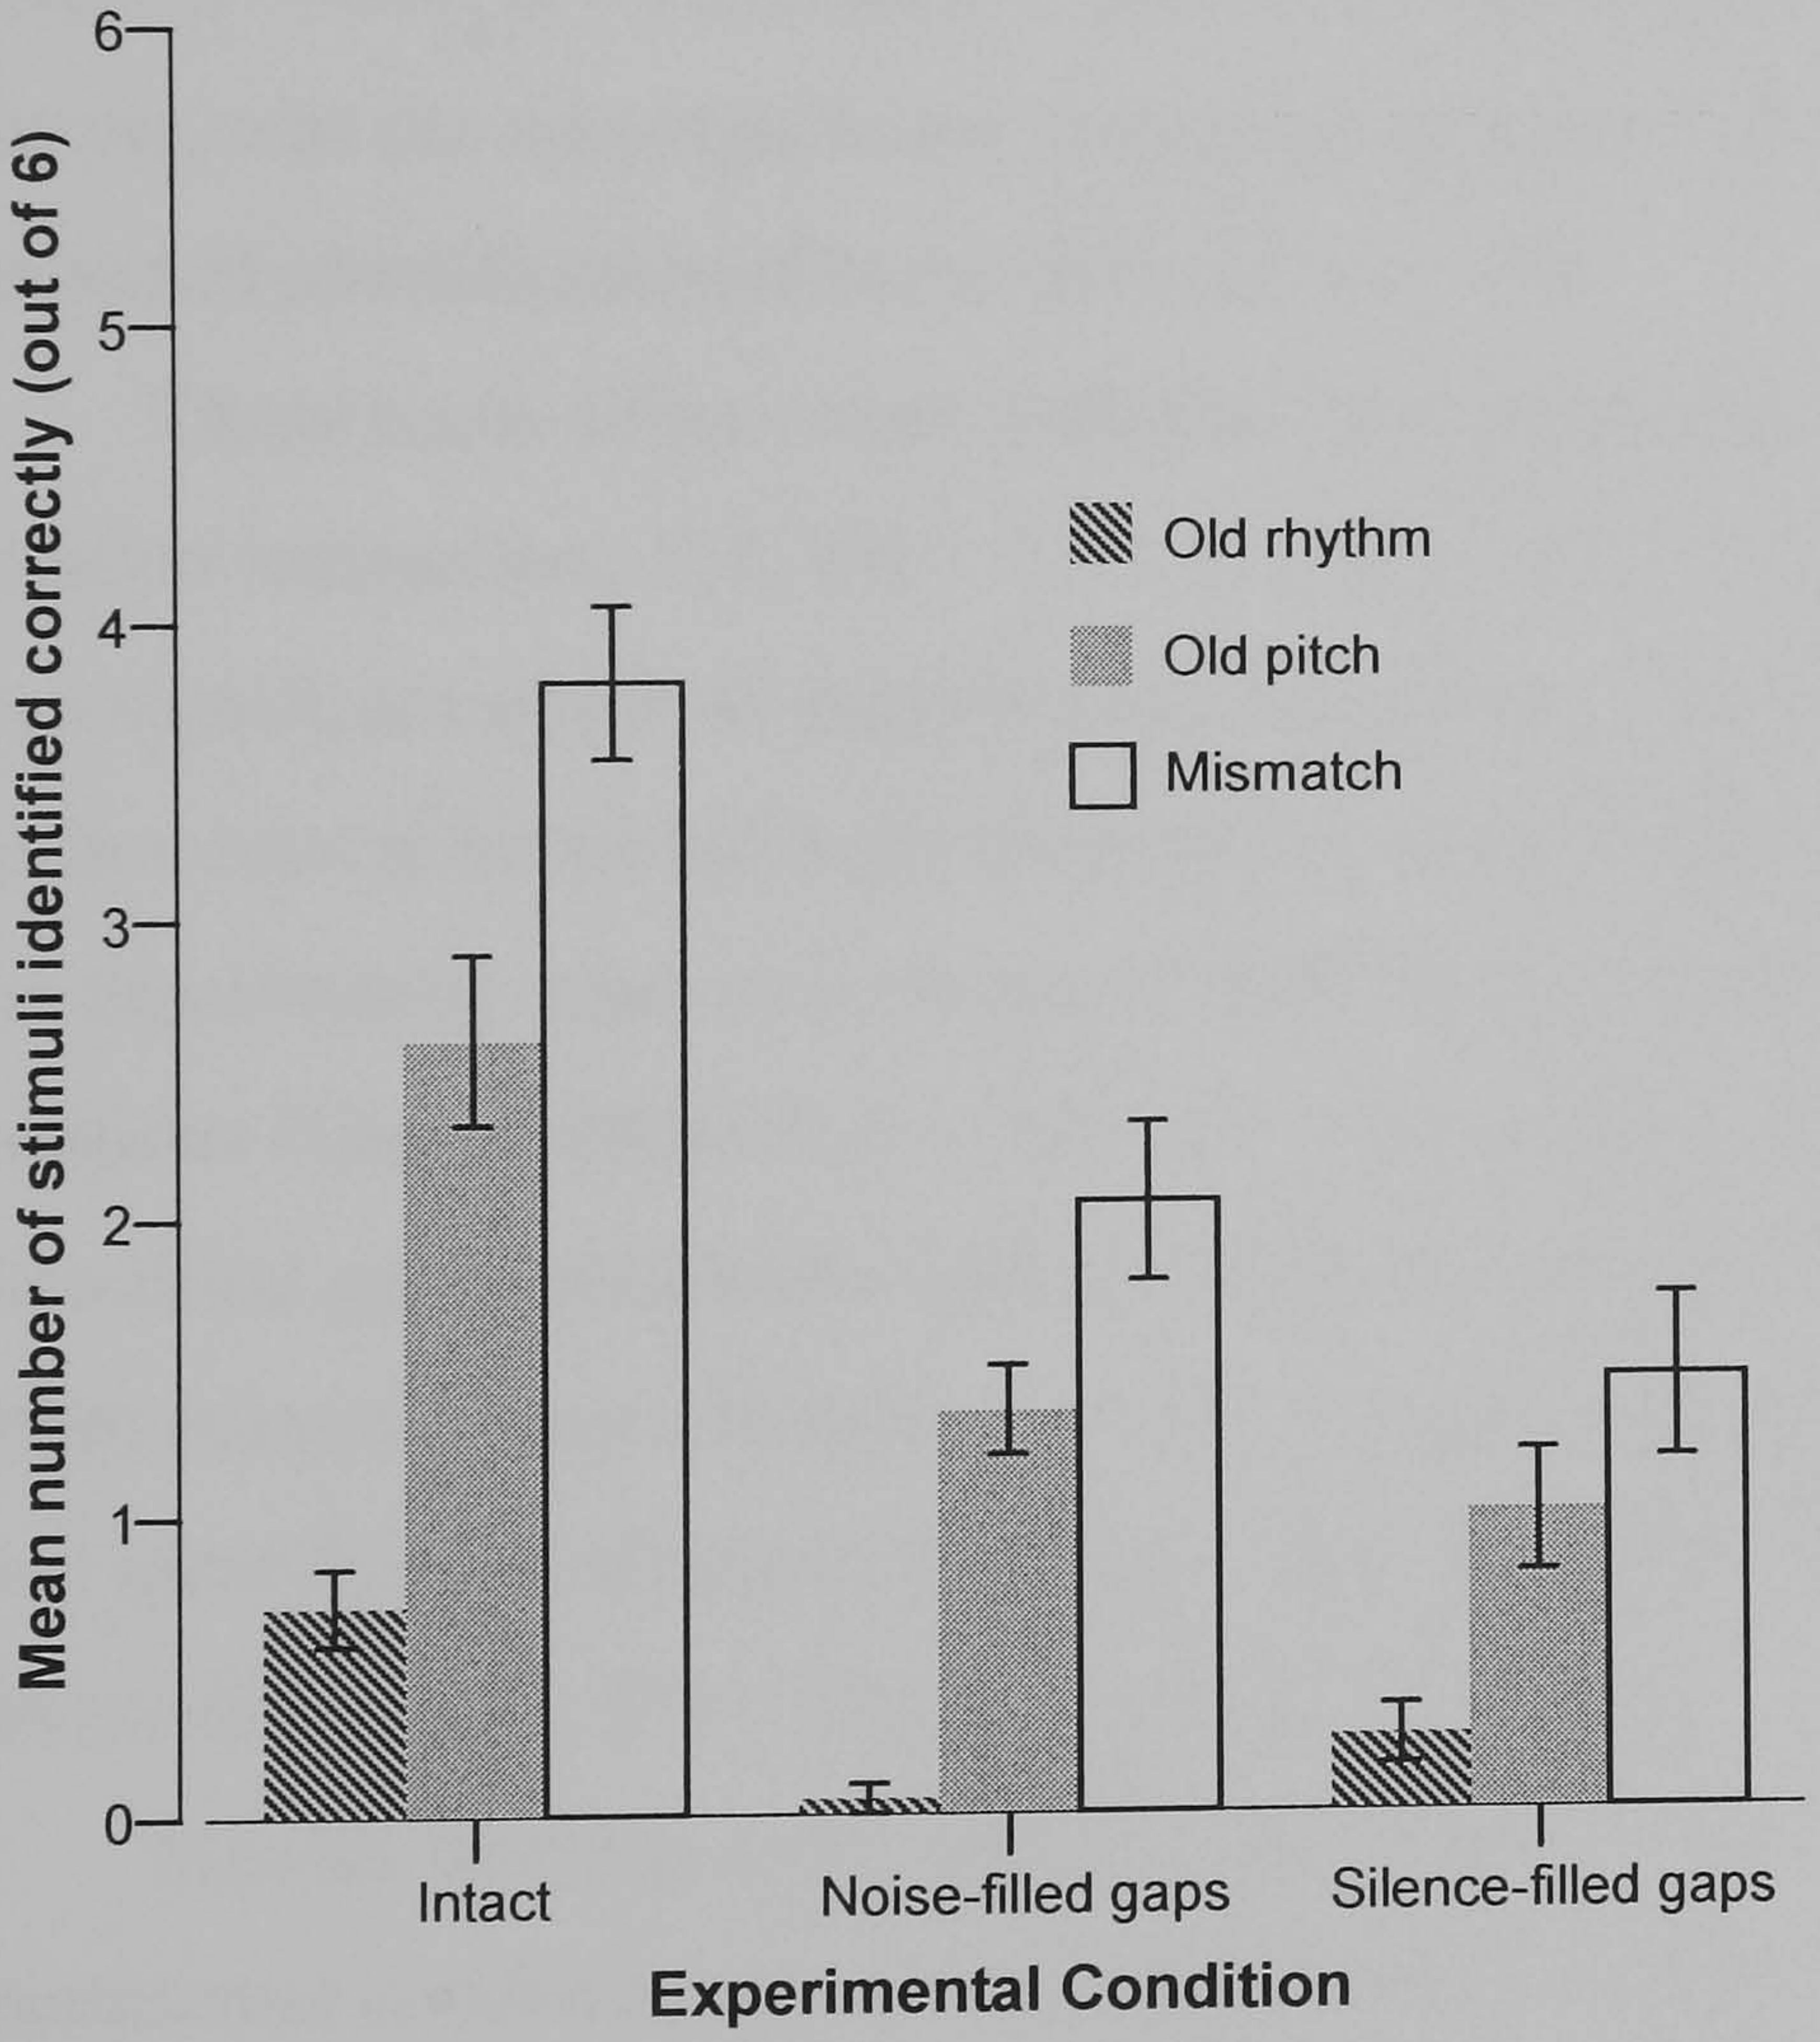

Figure 9.1. Mean number of stimuli identified correctly (± *SE*) by children in each of the three experimental conditions (intact, noise-filled gaps, silence-filled gaps) and for the three stimulus types (old rhythm, old pitch, mismatch).

These data were analysed using a 3 (stimulus type: old rhythm; old pitch; mismatch) × 3 (experimental condition: intact; noise-filled gaps; silence-filled gaps) mixed ANOVA with stimulus type as the repeated measure. There was a significant main effect of stimulus type on identification, $F(2, 89) = 68.55, p < .001, \eta^2 = .55$. Post hoc tests (Bonferroni) revealed significant differences between all stimulus types: old rhythm versus old pitch ($p < .01, d = 1.11$); old rhythm versus mismatch ($p < .001, d = 1.36$); old pitch versus mismatch ($p < .005, d = 0.51$). Identification of mismatch stimuli was significantly superior, and identification of old rhythm stimuli significantly inferior, to the other stimulus types.

The effect of experimental condition on identification was also significant, $F(2, 57) = 60.19, p < .001, \eta^2 = .68$. Identification in the intact condition was significantly superior to that of both the noise-filled gaps and silence-filled gaps conditions (Games-Howell: $p < .001, d = 1.22$ and $p < .001, d = 1.47$, respectively). Most interesting, however, was the non-significant advantage of noise-filled gaps over silence-filled gaps in terms of identification (Games-Howell: $p > .05$).

These main effects were qualified by a significant Stimulus type × Experimental condition interaction, $F(4, 89) = 4.69, p < .005, \eta^2 = .14$. For old rhythm stimuli, there was a significant effect of experimental condition, $F(2, 42) = 11.59, p < .001, \eta^2 = .29$. Post hoc tests (Games-Howell) showed that identification scores in the intact condition were significantly superior to those of both the noise-filled gaps and silence-filled gaps conditions (Games-Howell: $p < .001, d = 1.14$ and $p < .01, d = 0.79$, respectively). Noise-filled gaps provided no identification advantage over silence-filled gaps for old rhythm stimuli (Games-Howell: $p > .05$). In fact, a significant quadratic trend supports lower identification scores in the noise-filled gaps condition than in the silence-filled gaps condition, $F(1, 57) = 12.59, p < .001, \eta^2 = .16$.

Similar findings were obtained for old pitch stimuli, where a significant effect of experimental condition was also found, $F(2, 57) = 14.57, p < .001, \eta^2 = .34$. The intact condition was significantly superior to both noise-filled gaps and silence-filled gaps conditions (Tukey HSD: $p < .005, d = 0.98$ and $p < .001, d = 1.26$, respectively). Again, there was a non-significant difference between noise-filled gaps and silence-filled gaps conditions (Tukey HSD: $p > .05$).

For mismatch stimuli, there was also a significant effect of experimental condition, $F(2, 57) = 20.96$, $\eta^2 = .42$. Identification scores in the intact condition were significantly better than in the noise-filled gaps or silence-filled gaps conditions (Tukey HSD: $p < .001$, $d = 1.47$ and $p < .001$, $d = 1.91$, respectively). The noise-filled gaps condition had no identification advantage over the silence-filled gaps condition for mismatch stimuli (Tukey HSD: $p > .05$).

Thus, the crucial finding from this analysis is that there was no identification advantage provided by noise-filled gaps over silence-filled gaps for any of the three stimulus types.

Responses made on the basis of pitch and rhythm

Each time a stimulus was identified correctly, this response was recorded as either being made on the basis of the pitch pattern or on the basis of the rhythm. Figure 9.2 shows the mean number of pitch and rhythm responses for each stimulus type and by children in each condition.

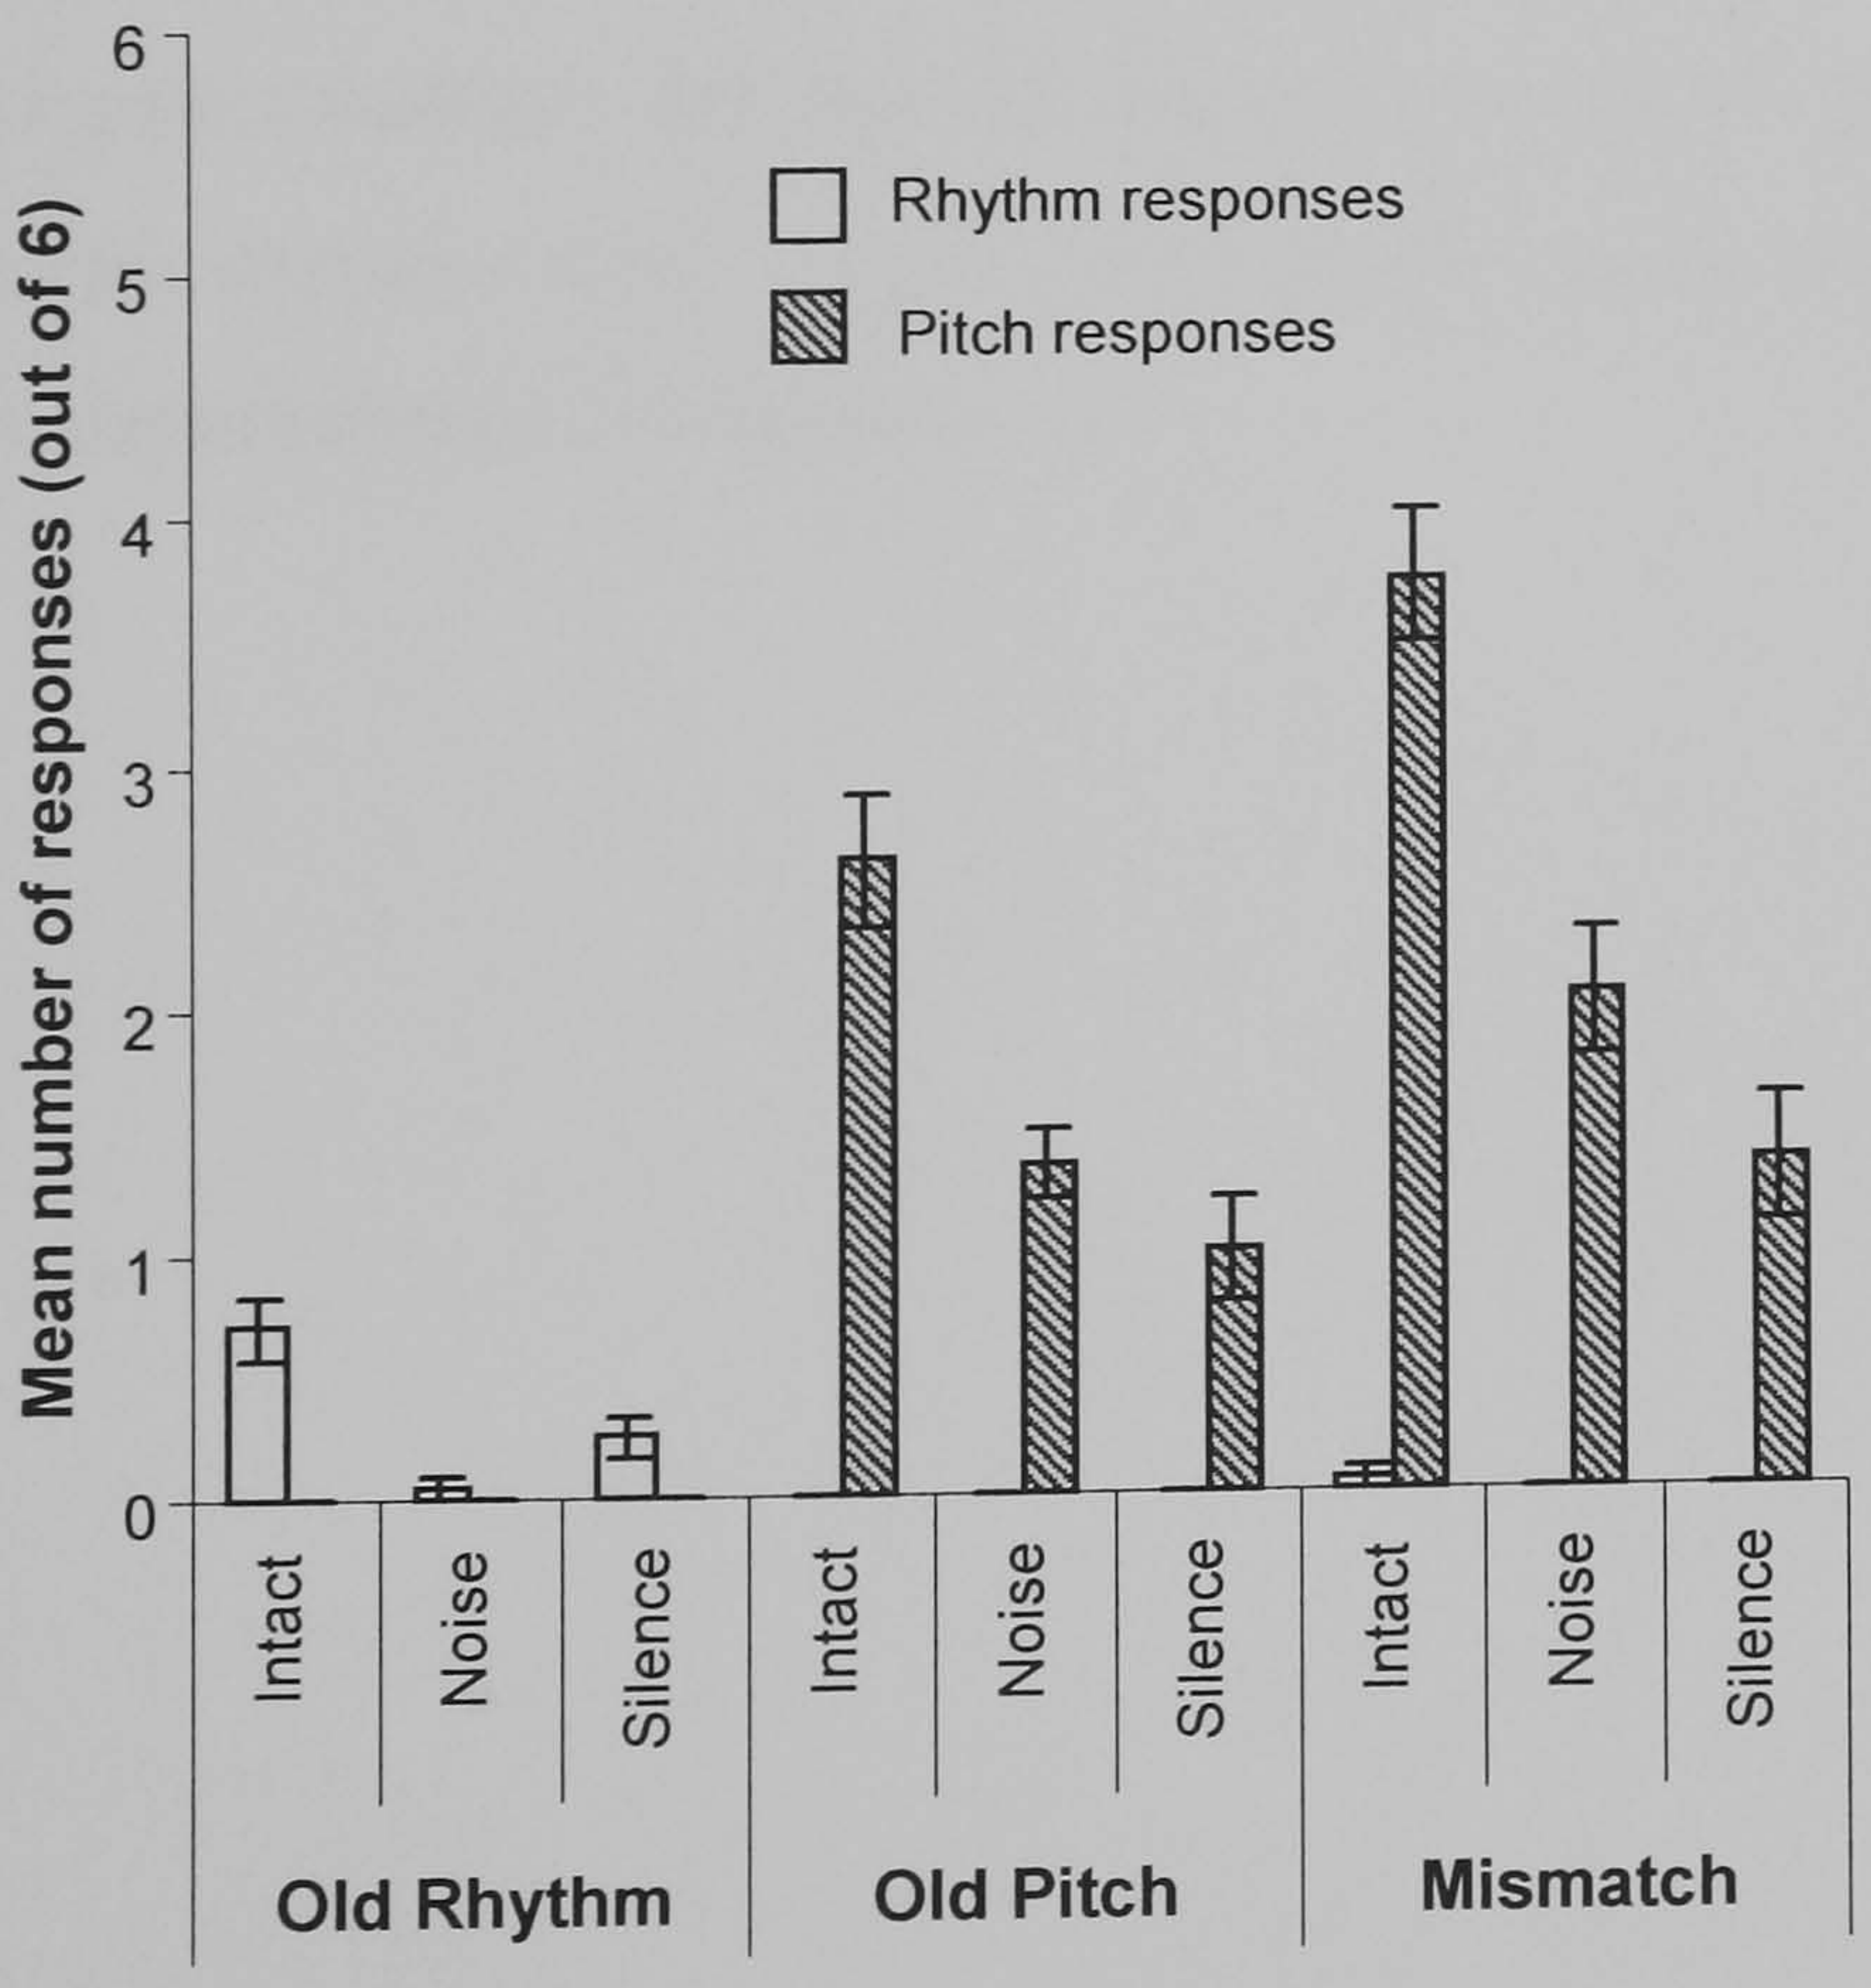

Figure 9.2. Mean number of responses (± *SE*) made on the basis of rhythm and pitch by children in each experimental condition (intact, noise-filled gaps, silence-filled gaps), and for each stimulus type (old rhythm, old pitch, mismatch).

Overall, very few rhythm responses were made. Since only one type of response was made for both old rhythm and old pitch trials, analysing responses in these conditions would only repeat the correct response analysis reported earlier. Even on mismatch trials, the vast majority of responses were based on the pitch, and not the rhythm component of the stimulus. For this reason, responses on mismatch trials were analysed using a Wilcoxon test. Of the participants that identified one or more mismatch trials correctly, all made more pitch than rhythm responses, $Z = 6.37$, $r = 0.58$, $p < .001$.

Identification time

In keeping with previous experiments, the identification time for correctly identified stimuli was calculated as the difference between the stimulus duration and the time taken to identify the stimulus. There were no gender differences in identification time (see Appendix 2), so data were collapsed across gender. This analysis excludes one child from the silence condition who failed to identify any stimuli correctly.

This analysis of identification time data does not consider each stimulus type separately, as only one correct response was made for old rhythm trials in the noise-filled gaps condition therefore it was not possible to construct a within-subjects factor based on stimulus type. Figure 9.3 shows the mean identification time for each of the three experimental conditions.

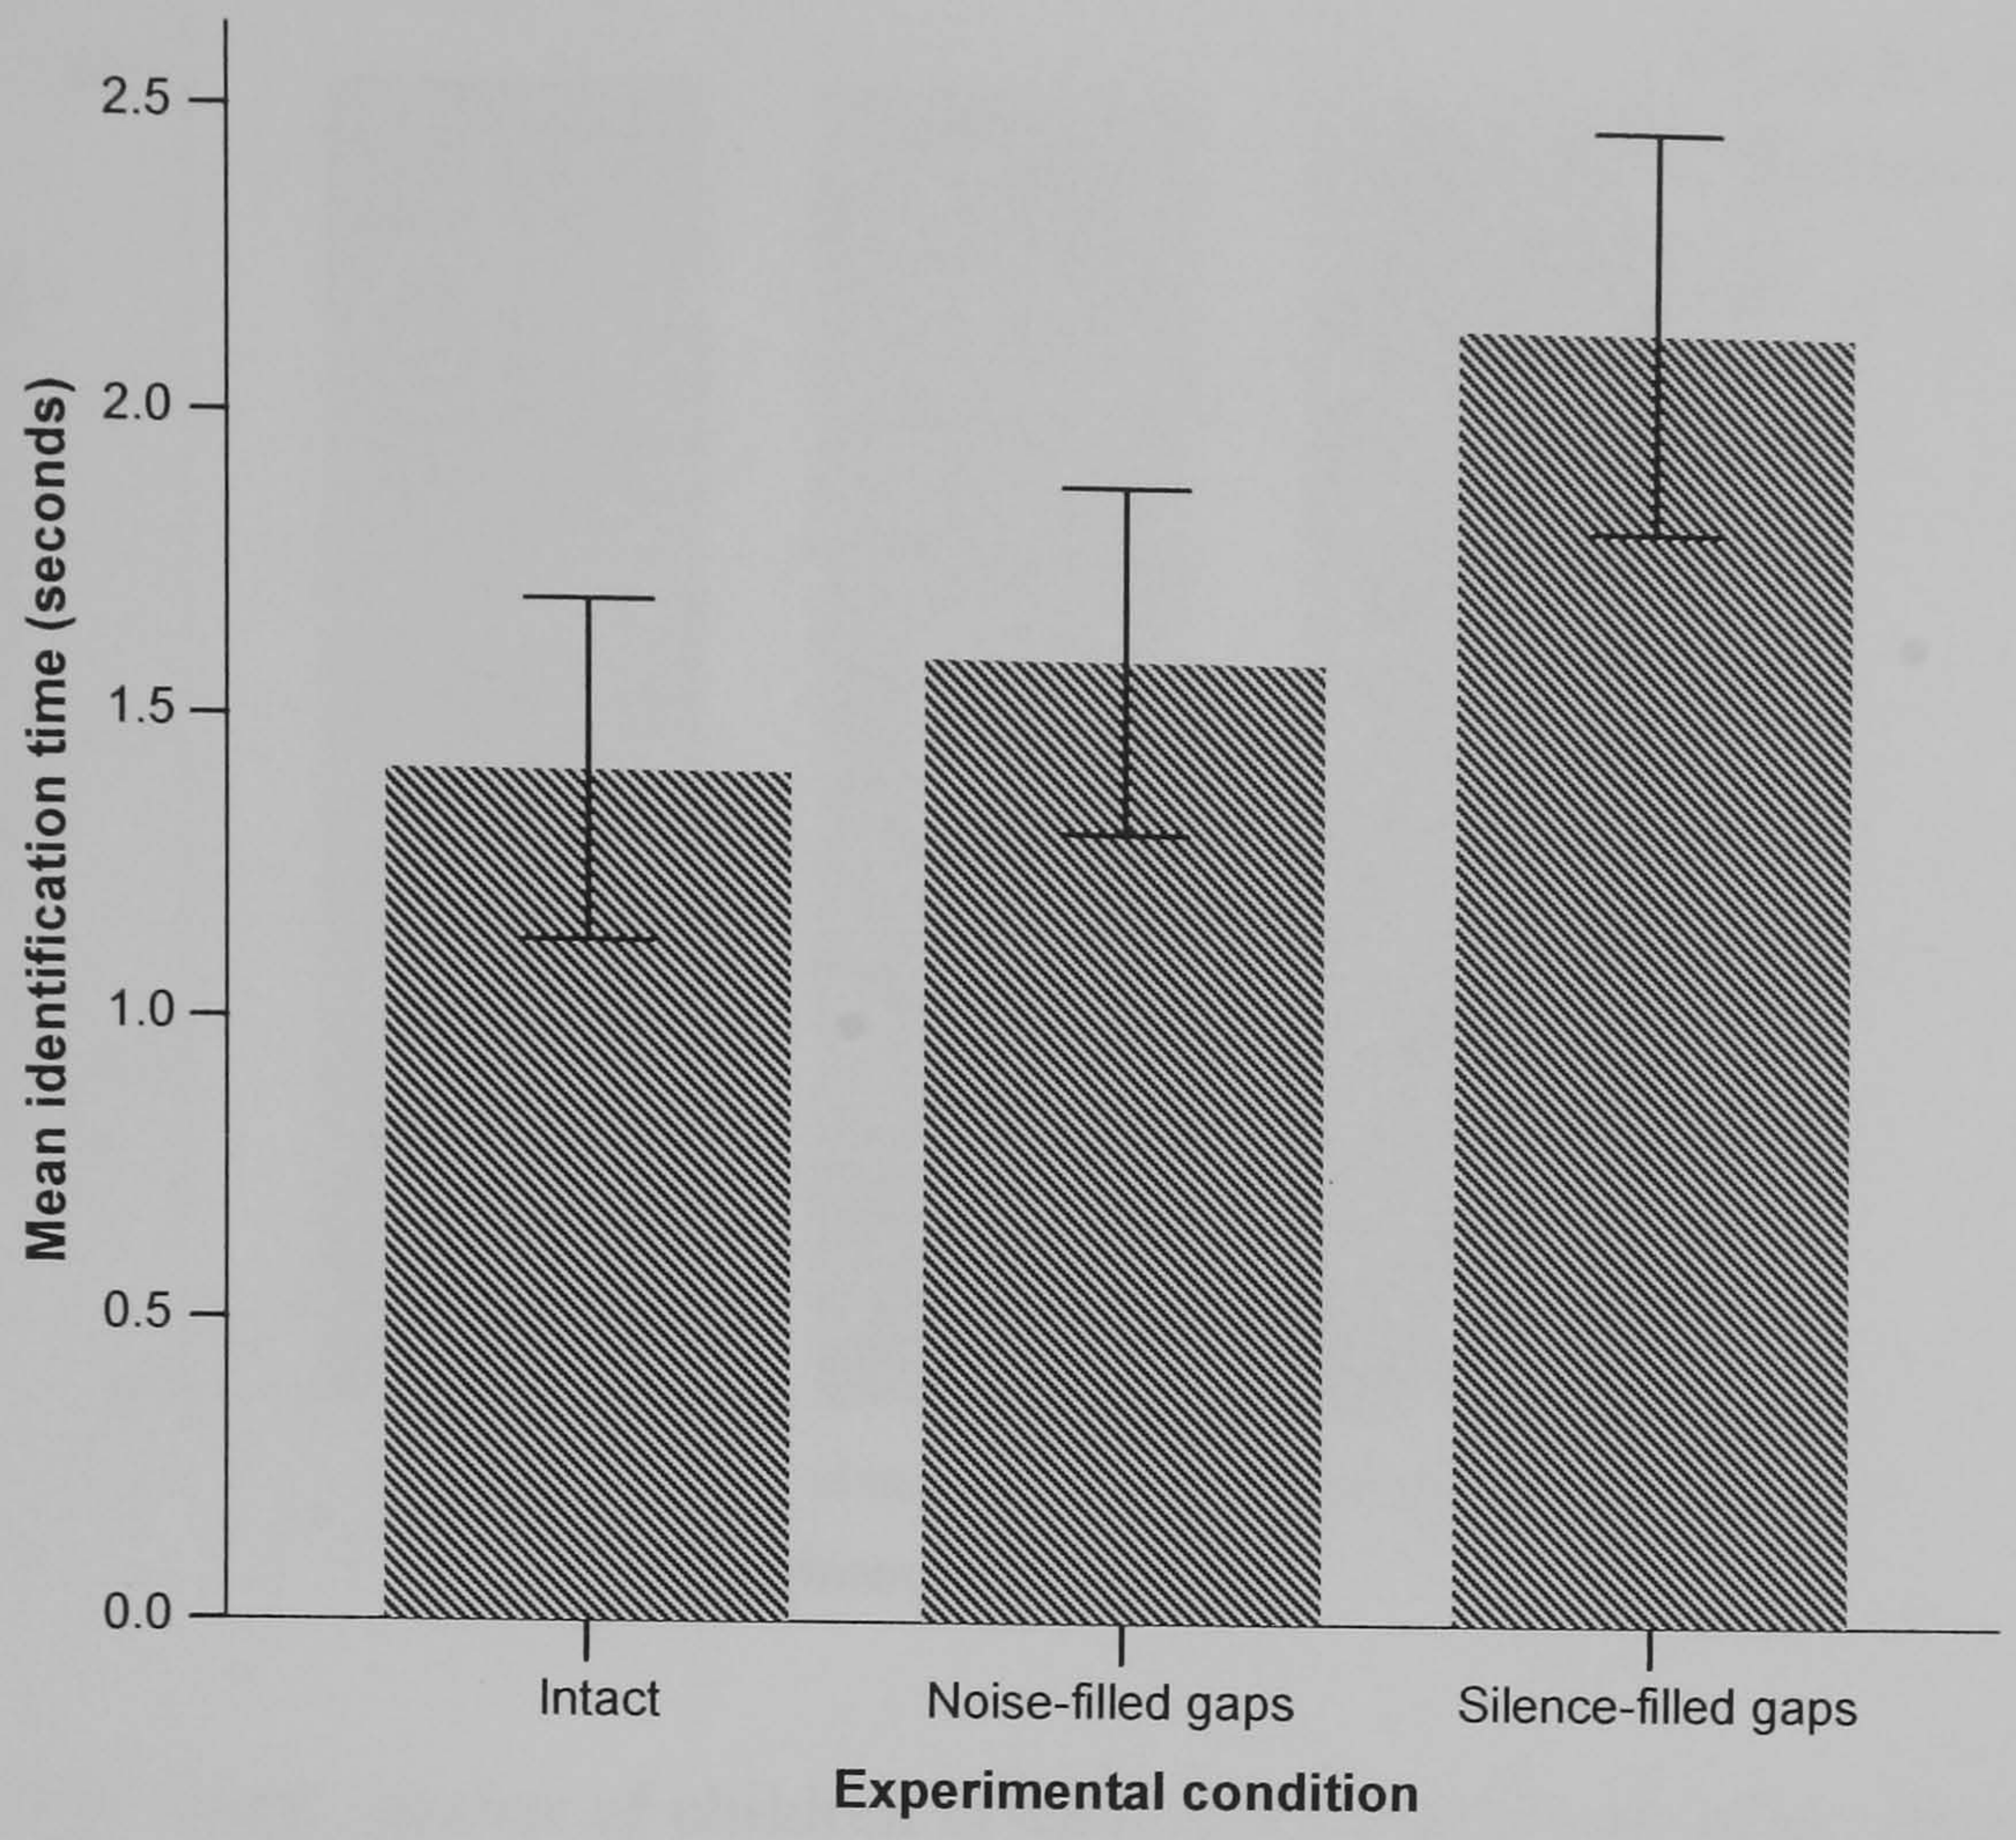

Figure 9.3. Mean identification time (identification time − stimulus duration; ± *SE*) for children in the three experimental conditions (intact, noise-filled gaps and silence-filled gaps).

Although Figure 9.3 suggests that stimuli were identified faster in the intact and noise-filled gaps conditions when compared to the silence-filled gaps condition, there was no significant difference between the average identification times of the three conditions, $F(2, 56) = 1.66, p > .05$.

Perception of continuity

Children were asked whether they thought that little pieces of the tune were missing (fragmented signal) or all of the tune was there (continuous signal). The response to this question was not made separately for each stimulus type, but across all stimuli presented. The number of children in each experimental condition reporting that the stimuli were either fragmented or continuous is shown in Figure 9.4.

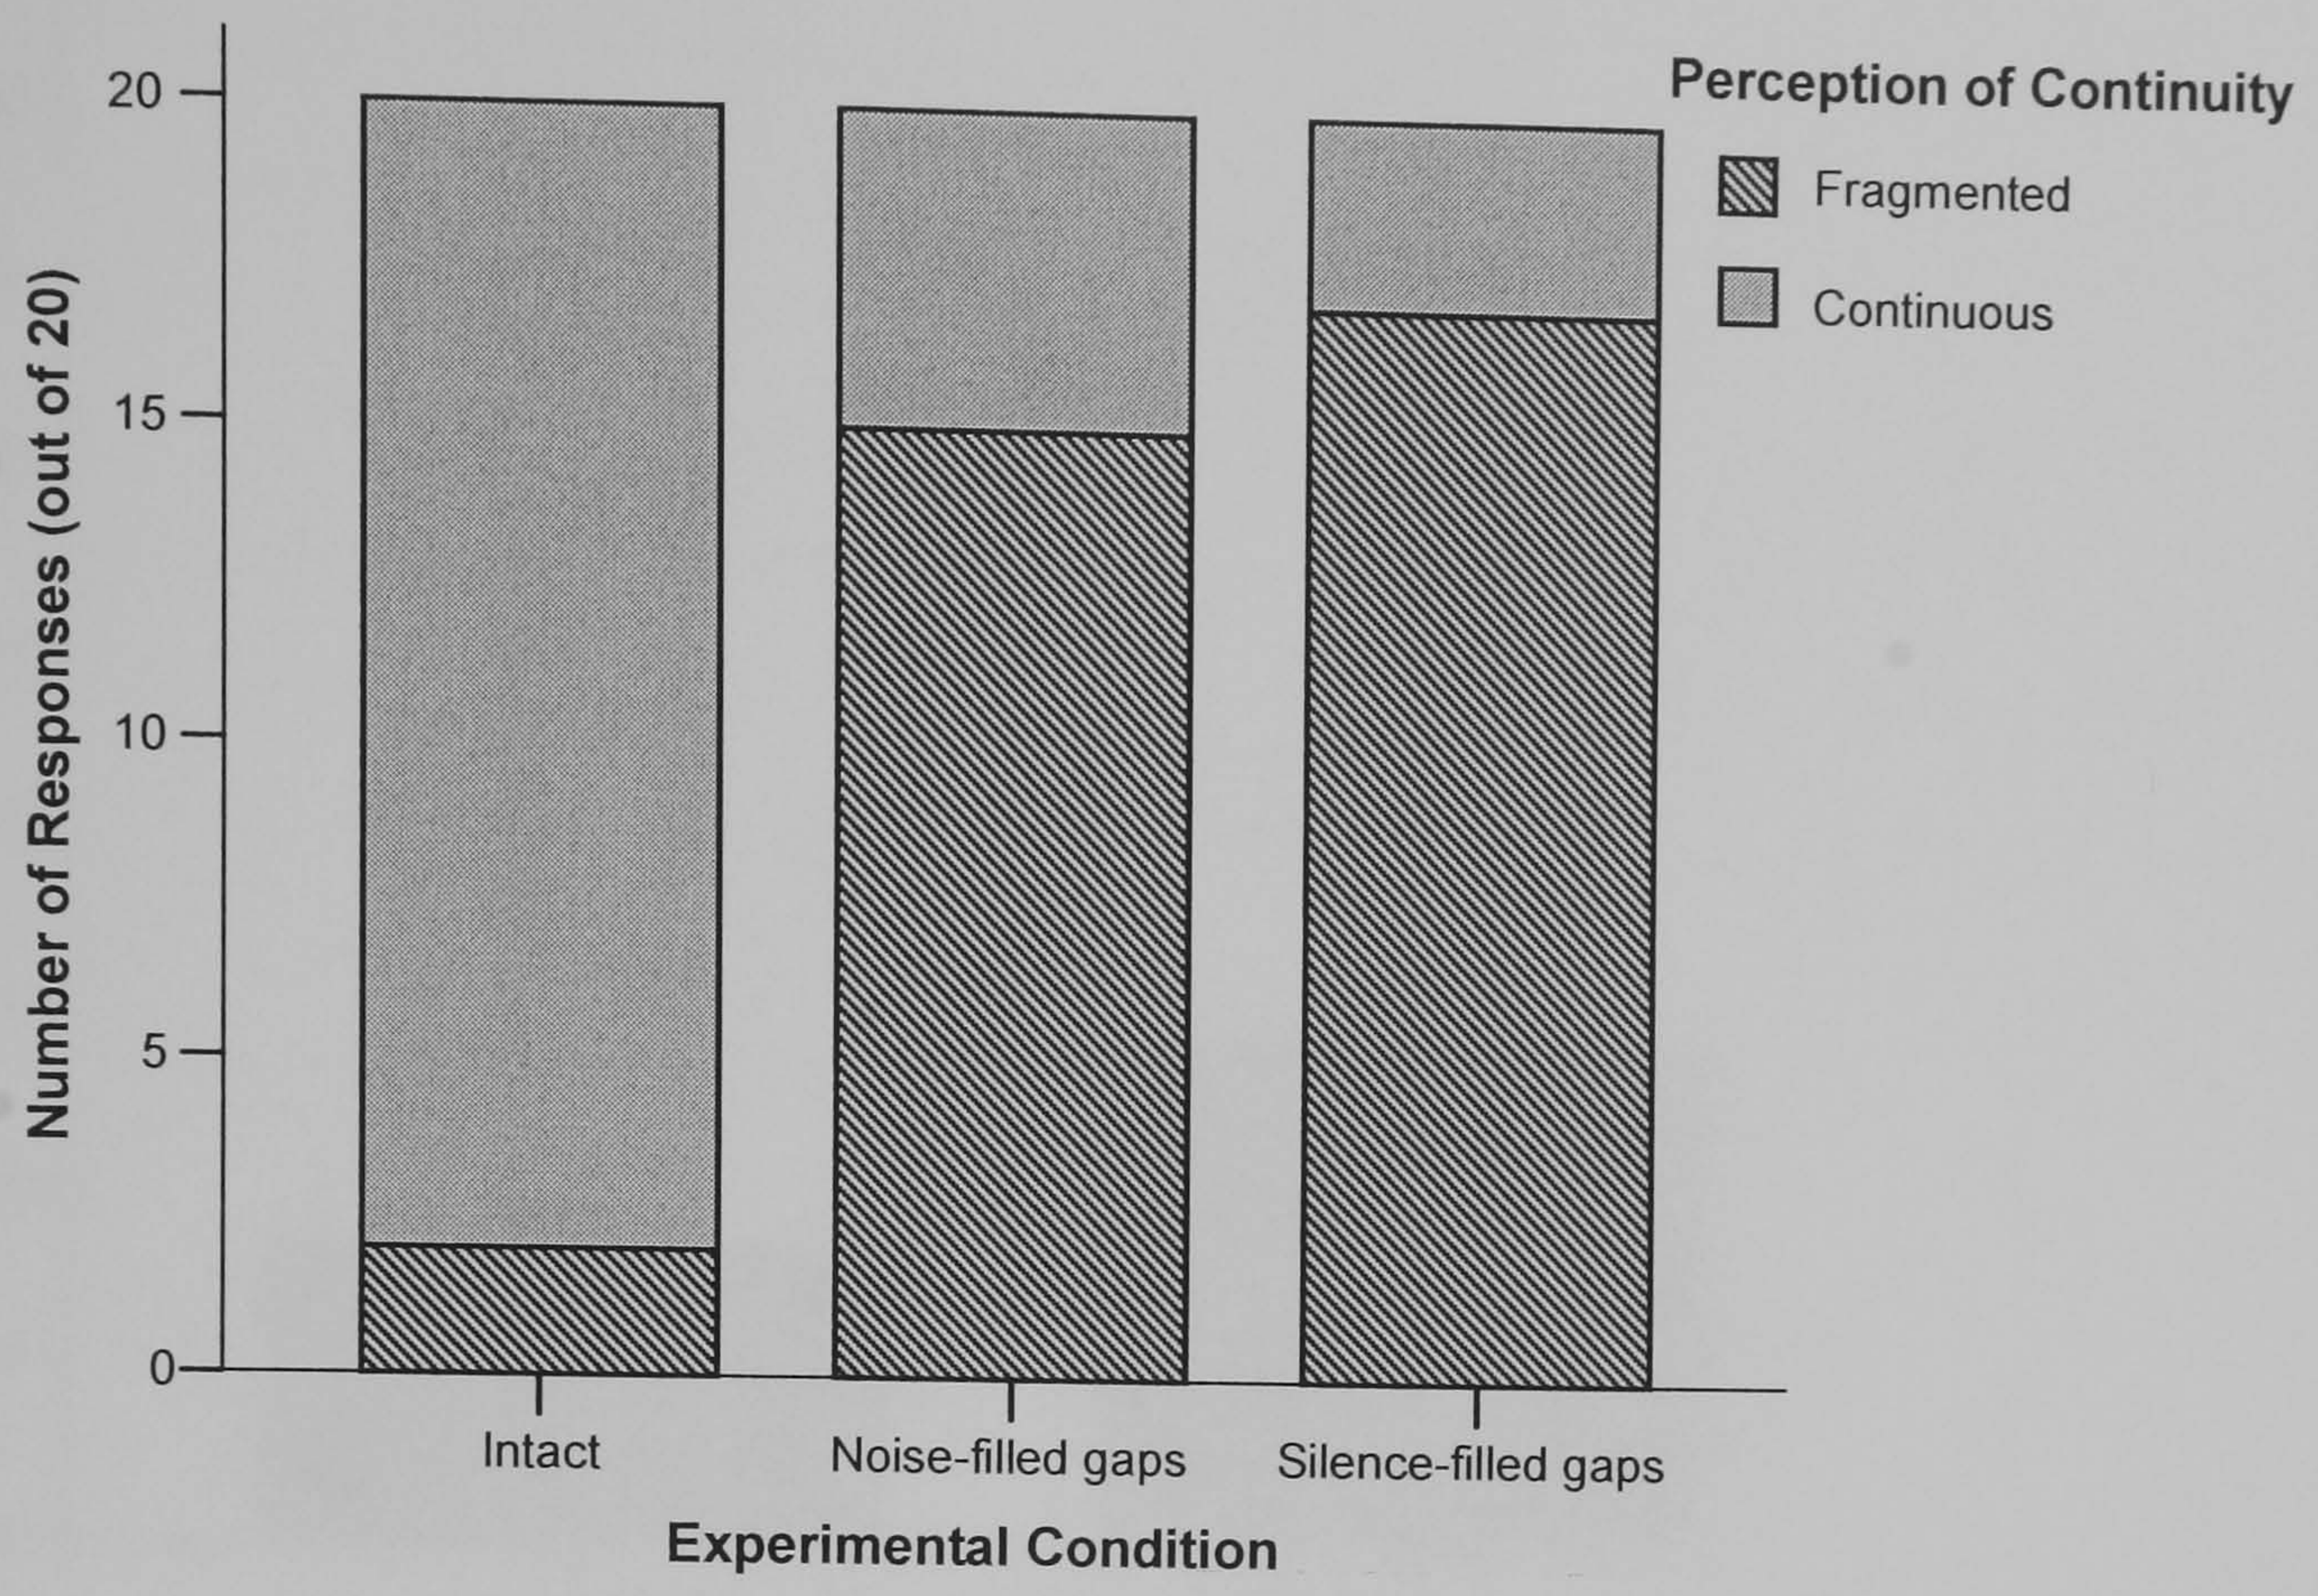

Figure 9.4. Total number of children in each condition (intact, noise-filled gaps and silence-filled gaps) reporting the stimuli as fragmented and reporting the stimuli as continuous.

These data were analysed using a 3 (experimental condition: intact; noise-filled gaps; silence-filled gaps) × 2 (perception of continuity: fragmented; continuous) chi-square, and a significant association between experimental condition and the perception of continuity was found, χ^2 (2, $N = 60$) = 27.01, $p < .001$, Cramer's $V = .67$. Figure 9.4 shows that only in the intact condition did the majority of children report the signals to be continuous; in both noise-filled gaps and silence-filled gaps conditions the majority of children reported the signals to be fragmented.

Based on the previous analysis, children were put into two groups according to whether they reported the signals to be fragmented or continuous. The mean number of stimuli (across all stimulus types) identified correctly by children in the "fragmented" and "continuous" groups is shown in Figure 9.5.

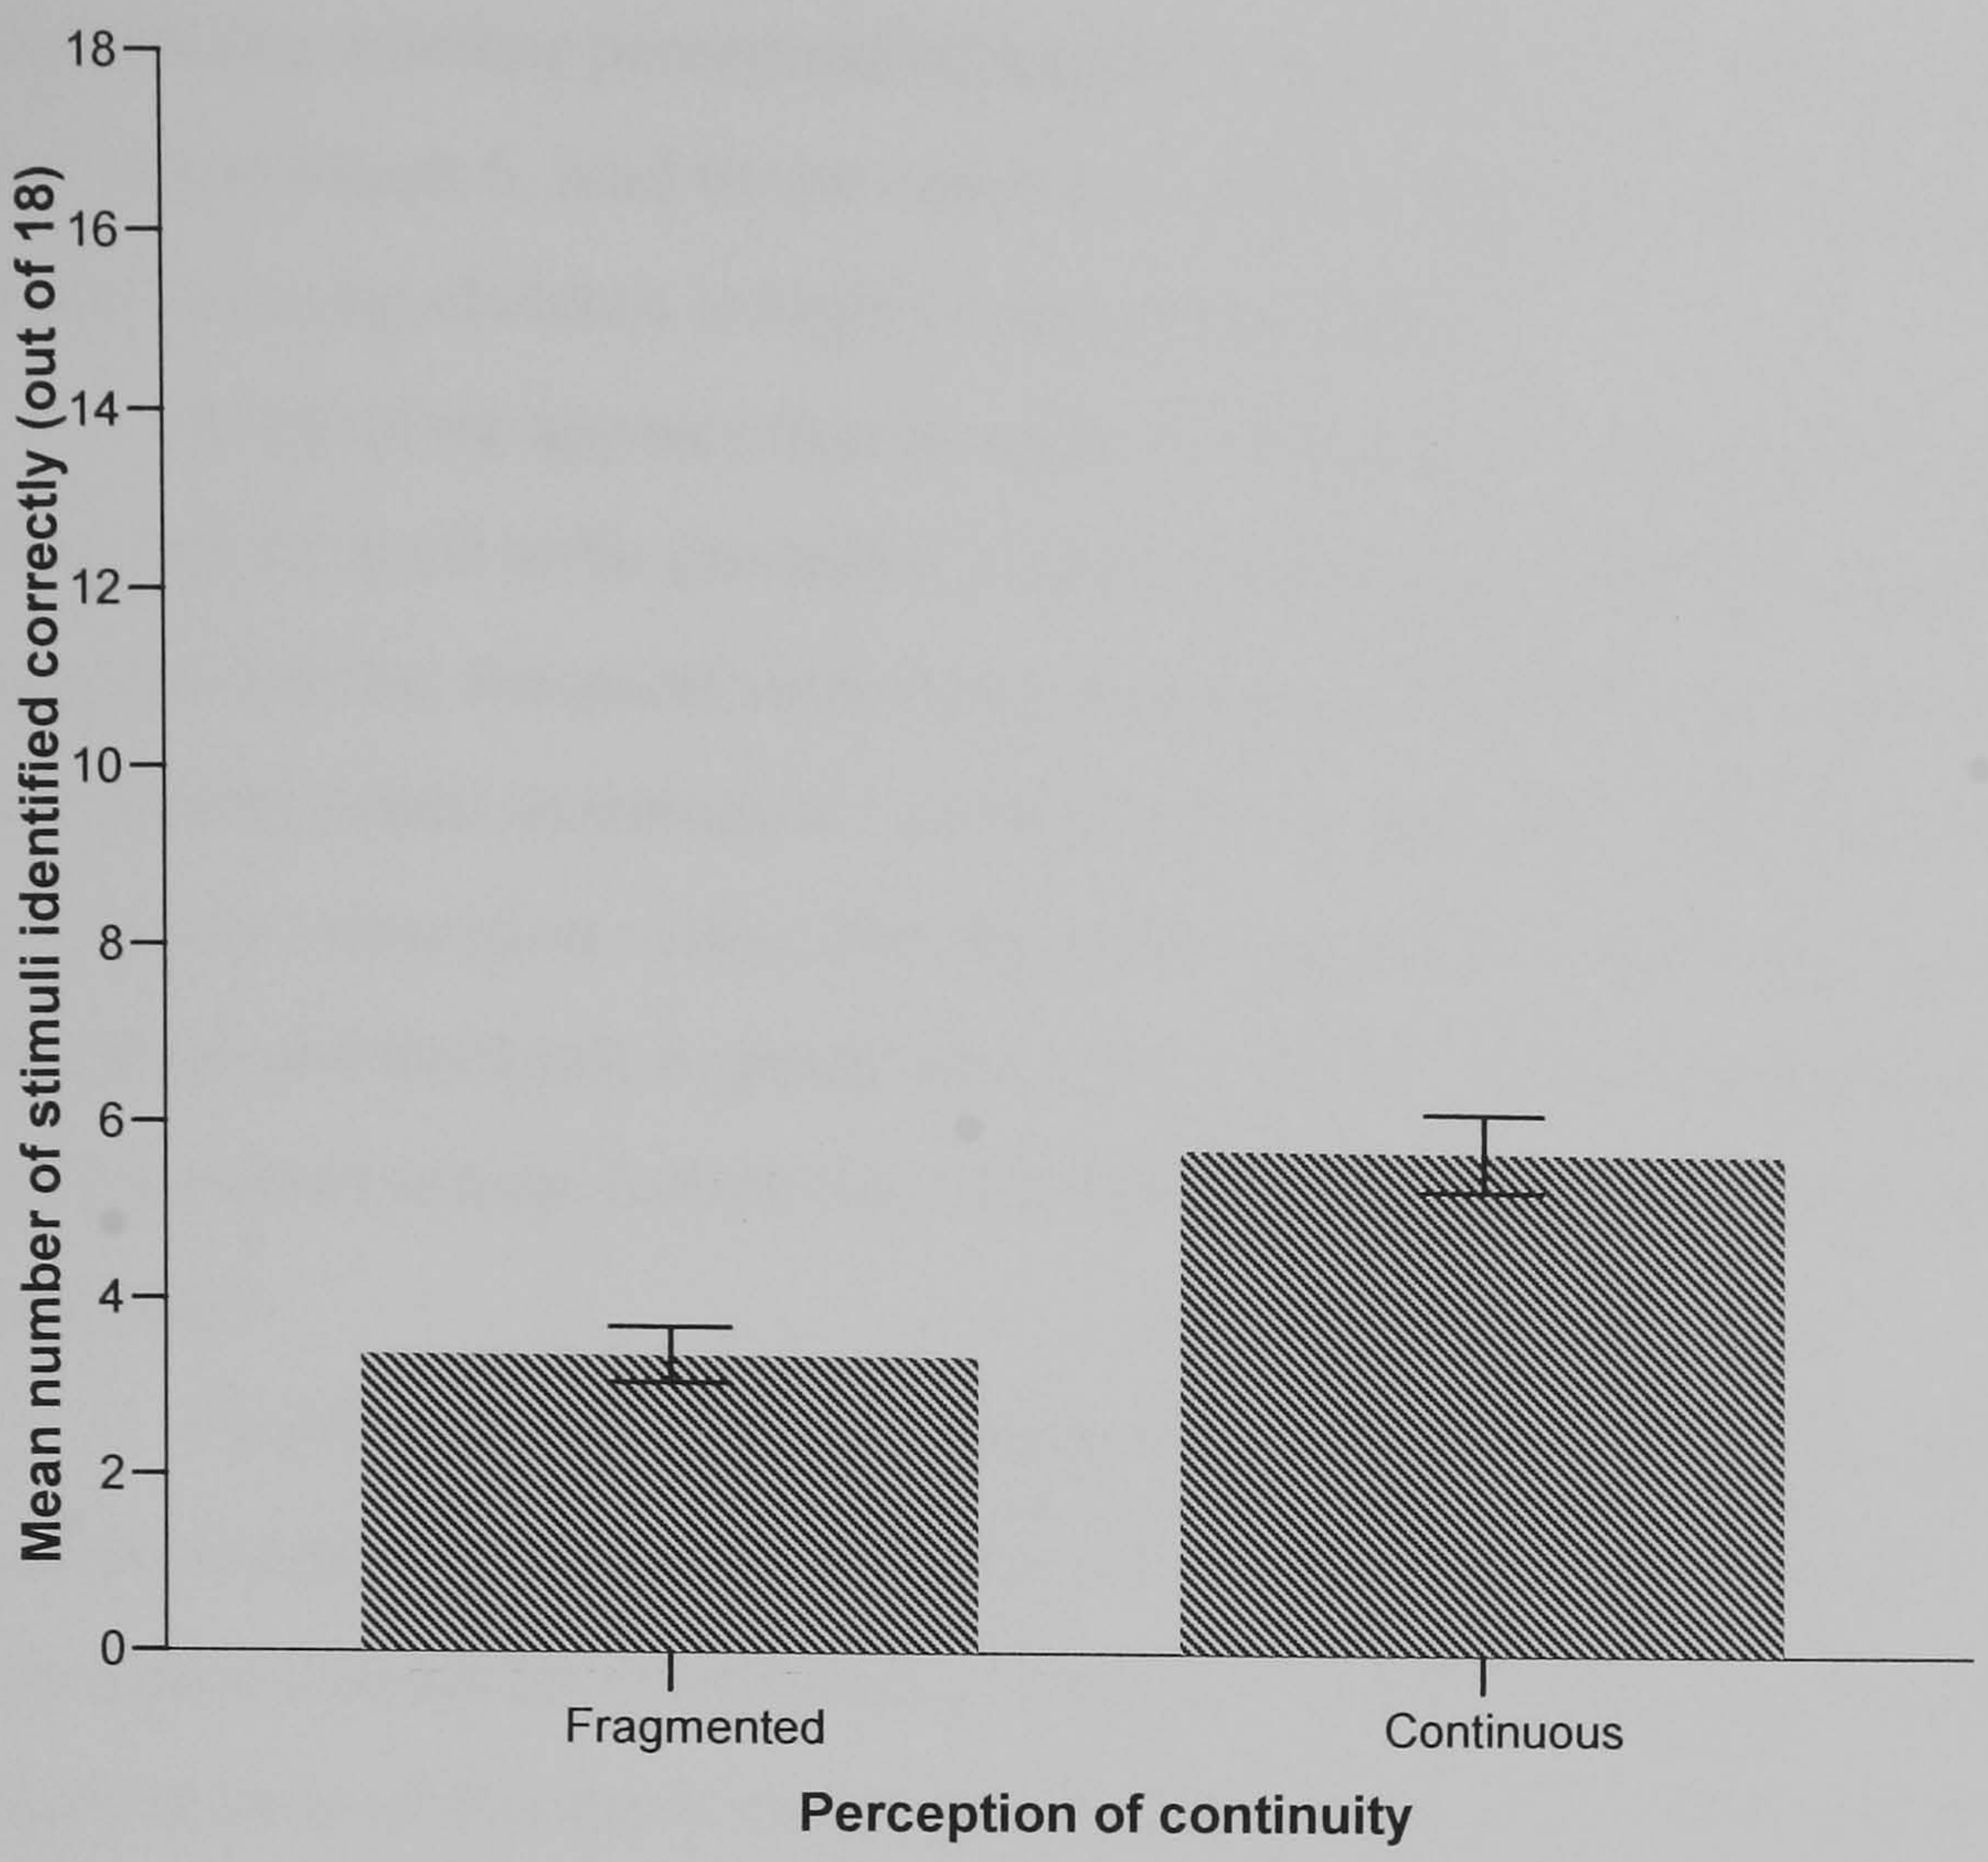

Figure 9.5. Mean number of stimuli (± *SE*) across all stimulus types identified correctly by children who reported the signals to be fragmented and by children who reported the signals to be continuous.

As Figure 9.5 suggests, children who reported that the signals were continuous identified significantly more stimuli correctly than children who reported that the signals were fragmented, $t(58) = 4.48$, $p < .001$, $d = 1.05$.

Discussion

The primary aim of the present experiment was to investigate the role of the frequency (pitch) and temporal (rhythm) components of a melody in musical perceptual restoration. The most striking result from this experiment was the failure to demonstrate the operation of perceptual restoration (as shown by no significant noise-filled gaps over silence-filled gaps advantage in stimulus identification) for any stimulus type. This further demonstrates that exactly what is familiar in a musical signal is critical in determining whether perceptual restoration can occur, and together with the findings from Experiment 6, lead to the conclusion that the operation of perceptual restoration of music in young children is highly context-dependent.

It therefore appears that in order for perceptual restoration to occur, pitch-based expectations need to be grounded in the correct temporal context, and vice versa. This suggests that the temporal expectancy hypothesis (Kidd et al., 1984), where both pitch- and rhythm-based expectations should be congruent, also applies to the operation of perceptual restoration. These findings also support the claim made by Narmour (1977) that pitch and rhythmic expectations interact in the process of melody identification. The findings demonstrate further that this interaction also applies to the process of perceptual restoration.

Children in the intact condition were able to identify some stimuli on the basis of just the rhythm or the pitch contour alone, thus demonstrating that the lack of perceptual restoration cannot be interpreted as the result of presenting an impossible task. However, identification of rhythmic and pitch components in new contexts, although superior to other conditions, was still quite poor in the intact condition. Nevertheless, children were more proficient when the familiar component was pitch contour and not rhythm. This supports propositions that pitch contour is more useful than rhythm for melody identification (e.g. Dalla Bella et al., 2003; Hébert & Peretz, 1997; Morrongiello et al., 1985; Trehub, 2001).

So why might rhythm be less useful than pitch for melody identification? The rhythmic element of a familiar melody has been described as less distinctive than the pitch pattern, thus making it a less efficient basis for melody identification (Hébert & Peretz, 1997). Indeed, nursery rhymes, as a genre of music, are extremely similar in terms of rhythmic characteristics. This might limit the utility of the rhythmic component as a means of identifying a familiar melody. Stimuli other than nursery rhymes with greater rhythmic variability would be required in order to test this explanation for the inferiority of rhythm in identification when compared to pitch contour. Furthermore, it is important to note that the dominance of pitch over rhythm in melody identification might be specific to Western music. In some cultures, rhythm is the predominant feature of music (Hébert & Peretz, 1997).

As was found for Experiment 6, mismatch stimuli where two components were familiar rather than just one, were easier to identify than either old pitch or old rhythm stimuli. Nearly all responses to mismatch stimuli were made on the basis of pitch, not rhythm. However, the finding that mismatch stimuli were identified better than old pitch stimuli therefore seems somewhat surprising. In both cases, it was the pitch component that allowed identification to occur, and in both cases, this pitch component was equally familiar. It is therefore possible that the accompanying rhythms in the mismatch trials, whilst incongruent, were nevertheless treated by the children as familiar. These rhythms might have provided a better context for pitch-based expectations to operate than the unfamiliar accompanying rhythms in the old pitch stimuli.

Overall, the superior identification when pitch contour was the familiar component, and the finding that rhythm alone was a poor basis for identification, supports findings from Hébert & Peretz (1997) and White (1960). It is also important to note that no differences in the pattern of responses (pitch/rhythm) between intact and disrupted signals (noise-filled gaps and silence-filled gaps conditions) were evident; such an asymmetry in the basis of identification responses was one of the most striking findings in Experiment 6. The lack of perceptual restoration operating here would predict that a similar strategy shift would not occur, but nevertheless correct responses in the present experiment were largely based on the same component across conditions, a finding in contrast to that of Experiment 6.

When the findings from the present experiment are compared with Experiments 1 to 6 where pitch and rhythmic components have not been mismatched, it is evident that whilst pitch-based information appears to predominate over rhythmic information (at least for intact signals), the correct combination of pitch and rhythm lead to the best melody identification. This partially supports single-component models of the relationship between pitch and rhythm (e.g. Peretz & Morais, 1989), where pitch and rhythm interact at some stage of processing. Taken further, it is also evident that perceptual restoration in children can only operate where expectations on both pitch and rhythmic levels are congruent, as was the case for the melodies presented in Experiment 1. Here, the melodies contained the correct pairing of rhythm and melody and strong evidence of the operation of perceptual restoration was obtained; melodies were significantly easier to identify where missing fragments were replaced by noise than where the same fragments were replaced with silence.

It is also likely that the advantage of melody over lyrics in perceptual restoration demonstrated in Experiment 6 was due to the correct combination of rhythm and pitch in the melody, and not the predominant role of one over the other. It is plausible that where expectations are such an important part of the perceptual restoration process, expectations generated on pitch and rhythmic levels serve to make each other more salient. At least in terms of the expectations that contribute to perceptual restoration in young children, the concept of melody seems to be complete and not further decomposable into its constituent parts. Therefore expectations about later acoustic events not only focus on *what* children expect to hear in a familiar piece of music, but also *when* children expect to hear it.

These findings, together with those from Experiment 6, have potentially important implications for how the role of familiarity in perceptual restoration is interpreted. In Experiment 6, perceptual restoration operated where the melodic component of the song was familiar, but could not operate when only the lyrics were familiar. When the melodic component was broken down into rhythmic and pitch components, perceptual restoration could not operate on the basis of one component alone. Therefore, the concept of familiarity needs to be considered carefully. Research findings on the perceptual restoration of speech have viewed familiarity as a single concept, mainly because such separation of a signal into the components that constitute it is not possible in the same way with a speech signal as with a musical signal. However, research with musical signals might indicate that it is wrong to view the role of familiarity in perceptual restoration in this way, at least with reference to perceptual restoration in children.

What the present series of experiments appears to demonstrate is that familiarity with a global representation (e.g. a nursery rhyme) and familiarity with the components of the representation (e.g. lyrics and melody; rhythm and pitch contour) all play different roles in perceptual restoration. Therefore, when considering familiarity influences on perceptual restoration in children, the crucial factor is not just familiarity with the stimulus, but exactly which parts of the stimulus children are familiar with.

There was no significant difference in the time taken to identify stimuli (across all types) between intact, noise-filled gaps and silence-filled gaps conditions. This might appear surprising given the large difference in identification scores between the intact condition and noise-filled gaps and silence-filled gaps conditions. However, as has been the case in previous experiments, there appears to be large variation in identification times within each condition.

Analysis of children's responses to the continuity question showed that only in the intact condition did the majority of children report that they thought the signals were continuous. In the noise-filled gaps condition, children made similar responses to children in the silence-filled gaps condition, reporting that they thought little pieces of the tune were missing. This is interesting, as where perceptual restoration does not occur based on analysis of identification scores, children also fail to think that the signals sound continuous when missing sections are replaced by noise.

It is possible that although children in the noise-filled gaps condition were not able to identify stimuli based upon the rhythmic or pitch component alone, perceptual restoration could still operate to restore the continuity of the disrupted signal. However, this is unlikely given that very few children in the noise-filled gaps condition reported that they thought the signals were continuous. This highlights the importance of the two measures of perceptual restoration (identification scores and continuity reports), and also strengthens the proposition that these two components of perceptual restoration, intelligibility and continuity, are closely linked. It appears to be the case that where intelligibility is not restored through perceptual restoration, neither is the continuity of a disrupted signal. This is further supported by the finding that regardless of experimental condition, those children that reported the signals to be continuous scored significantly higher in terms of identification scores than those children who reported the signals to be fragmented.

The aim of Experiments 6 and 7 was to determine more explicitly the role of familiarity in the operation of perceptual restoration in children's music perception, and to consider which components of the signal are important in generating expectations about the content of missing fragments. The emerging picture of the perceptual restoration of music in young children from Experiments 1-7 is that whilst perceptual restoration of musical input does occur in young children, its operation is context-dependent. Moreover, whilst the concept of familiarity is important in perceptual restoration, it is not just familiarity that is important, but the exact nature of what one is familiar with. The experiment that follows was designed to address the possibility of age developments in the ability to compensate for disruptions to musical signals through perceptual restoration.

Chapter 10: Experiment Eight

The effect of age on the perceptual restoration of music

Introduction

Experiments 1 to 6 successfully demonstrated that the perceptual restoration of missing musical input occurs in children aged between 4 and 6 years, indicating that music perception in children is active, not passive. This feature of music perception in children is shared with adults (DeWitt & Samuel, 1990). However, the human perceptual system in both children and adults is highly influenced by experience (Aslin & Smith, 1988). Where perceptual restoration is believed to be influenced by expectations based upon prior experience and passive exposure to a particular auditory domain (Kashino, 2003), there exists a plausible hypothesis that as such perceptual restoration will be influenced by the degree of such prior experience, and hence age. Furthermore, there are also believed to be developmental changes in perceptual processing strategies and the use of context in the process of perception (Warren & Warren, 1971) that could have implications for the operation of perceptual restoration.

Age differences in the perceptual restoration of speech are usually explained in terms of an increase in the ability to exploit the redundancy in speech and specifically in terms of speech perception, there is evidence that the size of the lexicon influences the ability to interpret ambiguous speech. The size of the lexicon in 5- to 6-year-old children is considerably smaller than that of adults (20,300 words in comparison with 80,300 words, as reported by Walley, 1988), with the implication that previous knowledge and expectations may not contribute to perception in children to the same degree as in adults. Knowledge-based expectations are also a critical part of the perceptual restoration of musical signals (DeWitt & Samuel, 1990), thus the question of possible age developments in the perception restoration of music needs to be addressed.

This experiment was designed in order to investigate whether the ability to restore missing musical input when replaced by noise develops with age. Nursery rhyme melodies are universally familiar[14], and should therefore be no less familiar to young children than adults. Thus, differences in the size of the musical lexicon are unlikely to account for any age differences in the current context. Of primary interest in the present experiment is whether children are more affected by disruptions to musical signals than adults, and, regardless of the overall familiarity of the stimuli, whether replacing missing fragments of musical signals with noise provides the same intelligibility advantage over silence-filled gaps at different ages. Thus, whilst in Experiments 1 to 7 the primary interest has been the size of the noise-filled gaps/silence-filled gaps difference, here the intact/noise-filled gaps difference for each age group is also of interest, as a measure of the general effect of signal disruptions on music perception.

Age developments in the perceptual restoration of speech

Whilst the issue of developmental change has not been addressed with reference to musical perceptual restoration, there is some evidence of age developments in the perceptual restoration of speech. In a study reported by Ackroff (1981), the test of perceptual restoration employed required children to locate the position of a missing phoneme replaced by noise within a word, as in the study reported by Warren (1970). According to this measure, children of six years of age showed evidence of perceptual restoration for 50% of test stimuli, with this rate increasing to 80% of stimuli for 8-year-old children.

Some experiments have demonstrated that children show weaker perceptual restoration abilities than adults in the context of the speech domain. Walley (1988) presented 5- to 6-year-old children and adults with words where a phoneme had either been replaced by noise, or had noise added to it. Word recognition was assessed, and children were less able to use perceptual restoration to recognise disrupted words than adults, indicating that children might be less able to use prior lexical knowledge and expectations in the process of perception than adults. However, the method of this study failed to address the possible advantage that noise-filled gaps might have over silence-filled gaps in children's perception.

---

[14] Nursery rhymes are seen as appropriate experimental stimuli in tests of music perception in both adults and children, and for this reason formed part of the stimulus set presented by DeWitt & Samuel (1990) in their test of musical perceptual restoration in adults.

In contrast to Walley's findings, Koroleva et al. (1991) reported that perceptual restoration in children might be stronger than in adults. In this experiment, 5- to 6-year-old Russian children and adults were presented with words where a phoneme was replaced by noise, and were required to report whether the word was intact or incomplete. Koroleva et al. claimed that children showed greater evidence of perceptual restoration than adults, as shown by their greater tendency to report words as intact rather than fragmented. However, this conclusion cannot be accepted for two reasons pertaining to the methods used. First, it is possible that the words presented were child-orientated words and thus were more familiar to children than adults; English translations are not provided by the authors. Second, this result might reflect a simple response bias where children were more likely than adults to report words as intact even if they did not perceive them to be so.

Gender differences in the perceptual restoration of speech were reported by Koroleva et al. (1996), supporting claims of gender differences in children's verbal abilities. Girls of 5 to 6 years of age, in contrast to boys of the same age, showed evidence of stronger perceptual restoration when the phoneme to be restored (e.g. consonant) was more similar to the replacement noise (thus demonstrating better use of bottom-up acoustic cues). Furthermore, girls showed stronger perceptual restoration for missing phonemes in the middle of a word, thus demonstrating superior use of top-down contextual cues for perceptual restoration in comparison to boys. However, these differences, reputedly due to better verbal abilities in young girls than boys, are therefore specific to the perceptual restoration of speech. Nevertheless, the authors argue that these factors suggest perceptual restoration in girls, when compared to boys, to be more similar to that in adults; yet the first in the present series of experiments found no gender differences in children's musical perceptual restoration.

In a more recent study, Newman (2004) compared the perceptual restoration abilities of adults and 5-year-old children. Multiple alternating fragments of spoken sentences were removed and replaced with either noise or silence, and listeners' repetition accuracy of the sentences was assessed. Children were better able to repeat the sentence when missing fragments were replaced by noise than silence, thus demonstrating the operation of perceptual restoration. Furthermore, the relative size of the advantage of noise-filled gaps over silence-filled gaps in terms of repetition accuracy was comparable for adults and children, even though overall children were more affected by disruptions to the sentences than adults were, as shown by their inferior repetition accuracy of the disrupted sentences. Whilst children can combine bottom-up and top-down information in the process of perception, these abilities might be less developed in comparison to the same abilities in adults, with crucial implications for children's learning in noisy environments. Further evidence that the lexical knowledge that contributes to perceptual restoration develops over time comes from a study by Newman (2006), which failed to find evidence of perceptual restoration in 2- to 3-year-old children. Newman concluded that the use of previous knowledge in the process of perception in these very young children differs to that of adults, and that there might be a period of reorganisation in development where children come to place more emphasis on top-down knowledge in the perceptual process.

The research reviewed above does appear to suggest that, for the perceptual restoration of speech at least, there might be developments in the ability with age. However, there is no evidence that the same might be true for perceptual restoration of music. There is evidence that the perceptual restoration of music occurs in adults (DeWitt & Samuel, 1990; Kaminska & Mayer, 1993; Sasaki, 1980), and Experiments 1 to 6 found evidence of a similar ability in 4- to 6-year-old children. However, adults in previous experiments and children in the present series of experiments have not been tested using the same stimulus type or the same measure of perceptual restoration, with the result that adequate age comparisons cannot be made.

There are reports in the literature that adults and children differ in the size of their musical knowledge bases that would influence musical expectations generated during listening to music (Krumhansl & Keil, 1982), and that with increased experience comes more stabilised representations of musical regularities (Bartlett & Dowling, 1980). In addition, the ability to perceive and remember melodies improves with increasing exposure to music (Schellenberg, Bigand, Poulin-Charronnat, Garnier & Stevens, 2005). However, Davidson (1994) argues that perception of the basic features of music is universal and not affected by formal training, suggesting that perception of these features operates in the same way in adults and children. Furthermore, both adults and children process auditory signals with a preference for global features, and Gestalt principles of perceptual organisation (e.g. closure, which is seen as an explanation for why perceptual restoration occurs) apply to perception in both children and adults (Bregman, 1990).

It is important to consider that adults and older children could have superior perceptual restoration abilities not as a result of age per se, but because of increased musical training that occurs as a result of age. DeWitt and Samuel (1990), when investigating the perceptual restoration of music in adults, included a musical experience and training questionnaire (METQ) to assess the extent to which formal musical training and informal musical experience influenced performance on their perceptual restoration task. This task required listeners to report whether a version of a stimulus where a musical note was replaced by noise, or a version where noise was added to the note, was intact.

A beneficial effect of musical training on perceptual restoration was a reasonable assumption for DeWitt and Samuel (1990) to make, since musical training has been found to facilitate music processing (Schön et al., 2004). However, DeWitt & Samuel found no relationship between musical background and musical perceptual restoration ability. This might appear to suggest that the perceptual restoration of music, in requiring no specialised musical ability, is a universal ability. However, this result might have been specific to the task presented. It is possible that differences in perceptual restoration with increasing musical training might emerge through the use of a different measure of perceptual restoration, such as the identification task employed in the present series of experiments. For this reason, the present experiment will include an analysis of the effect of musical training on perceptual restoration.

To conclude, previous research has indicated that the ability to fill in missing sensory input in the speech domain is, for the most part, superior in adults when compared to children. No such analysis has been carried out for the perceptual restoration of music. This has implications for understanding the role of auditory perceptual processes in perceptual restoration, and how they might change through development. In order to address these issues, the standard perceptual restoration task as reported in Experiment 1 was administered to a group of older children (9 to 11 years of age) and a group of adults, in order make direct comparisons with the performance of the 4- to 6-year-old children tested in Experiment 1. In addition, a measure of musical training was taken from the older children and adults in order to investigate whether this factor has an influence on musical perceptual restoration ability as measured using an identification task.

Method

Participants

Sixty children (29 girls) between the ages of 9 years 4 months and 11 years 1 month, with a mean age of 10 years 1 month (121.92 months; *SD* = 6.76) participated in this experiment as the "older children" group. These children were recruited from the Year 5 and Year 6 classes of a primary school in Surrey, UK, and of a primary school in Dorset, UK. Both schools were situated in predominantly white, middle-class areas, and all children that participated had normal sensory development, as reported by the school. In addition, 30 undergraduate students (27 females) from the University of Surrey participated in this experiment as the adult group. The adult group ranged in age from 18 years 3 months to 20 years 11 months, and had a mean age of 19 years 6 months (234.47 months; *SD* = 9.40). They were offered course credit for their participation. All participants were native English speakers who reported normal hearing.

Consent for the participation of the older children was obtained from both the child's parent and the school Headteacher. Participants in the adult group were given an information sheet about the experiment (Appendix 12) before being asked to sign a consent form (Appendix 13), and they were given a debrief sheet to read at the end of the experiment (Appendix 14). Further details about the gender and age of participants in each condition is given in Table 10.1.

Table 10.1. The gender split and mean age of participants in each of the two age groups and three experimental conditions.

Age group	**Experimental condition**	N	**Gender**	**Mean age in months (*SD*)**
Older Children	Intact	20	M=12; F=8	121.60 (6.27)
	Noise-filled gaps	20	M=9; F=11	121.85 (6.81)
	Silence-filled gaps	20	M=10; F=10	122.30 (7.48)
Adults	Intact	10	M=0; F=10	234.40 (8.98)
	Noise-filled gaps	10	M=1; F=9	233.60 (6.50)
	Silence-filled gaps	10	M=2; F=8	235.40 (12.65)

Design

The design of the experiment consisted of two between-subjects independent variables: the age group of the participant (older children or adults), and the experimental condition (intact, noise-filled gaps or silence-filled gaps). All participants heard the same nursery rhyme melodies that had been presented in Experiment 1 (see Table 3.2), but they were modified according to the experimental condition, being heard intact, or with alternating 100 ms fragments replaced either with white noise or silence.

Materials

The stimuli and presentation programme from Experiment 1 were used for the present experiment in order for the data from the present experiment to be compared directly to the data from young children reported for Experiment 1. The schools from which the older children were recruited were asked to fill out the song familiarity questionnaires used in previous experiments (Appendix 1), whilst the adult participants were asked to fill out a similar questionnaire at the end of their testing session (see Appendix 15).

Procedure

Each participant was tested individually, in a quiet room. Task instructions were modified from those used in Experiments 1 to 7 so that they would be appropriate for the age groups being tested. Participants in both older children and adult groups were told that they would be taking part in a music identification game:

> "You will hear a series of children's tunes and you will be asked to identify each tune as soon as you know what it is."

Each participant was presented with 2 demonstration trials (*Baa Baa Black Sheep* and *The Grand Old Duke of York*) followed by 13 experimental trials. As in previous experiments, no further explanation was given if participants did not identify the demonstration trials. On experimental trials, no time limit was imposed on responding, and no feedback was given as to whether the answer was correct or incorrect. Failure to give a response was coded as incorrect.

The order of the testing session differed slightly for the two age groups. The older children group were first asked whether they took, or had taken in the past for longer than one year, lessons in a musical instrument or singing lessons (as a measure of musical training), and then they carried out the identification task. Finally, they were asked the standard continuity question used in Experiments 3 to 7, using the same instructions as presented in previous experiments.

The adult group first read the information sheet and signed the consent form, and then took part in the identification task. Next, they answered the continuity question, phrased using the standard instructions, and provided information regarding any musical training. Finally, they completed the song familiarity questionnaire[15] before reading the debrief sheet.

---

[15] All stimuli were reported as familiar by the adult participants.

Results

The data from the present experiment were combined with data from Experiment 1 in order for data from three different age groups, all tested on the same task, to be compared.

Correct identification scores

Each participant was given a score out of 13 that corresponded to the number of melodies they identified correctly in the experimental trials. There were no gender differences in identification scores for any age group (see Appendix 2) so data were collapsed across gender. Figure 10.1 shows the mean number of melodies identified correctly by participants in each age group and experimental condition.

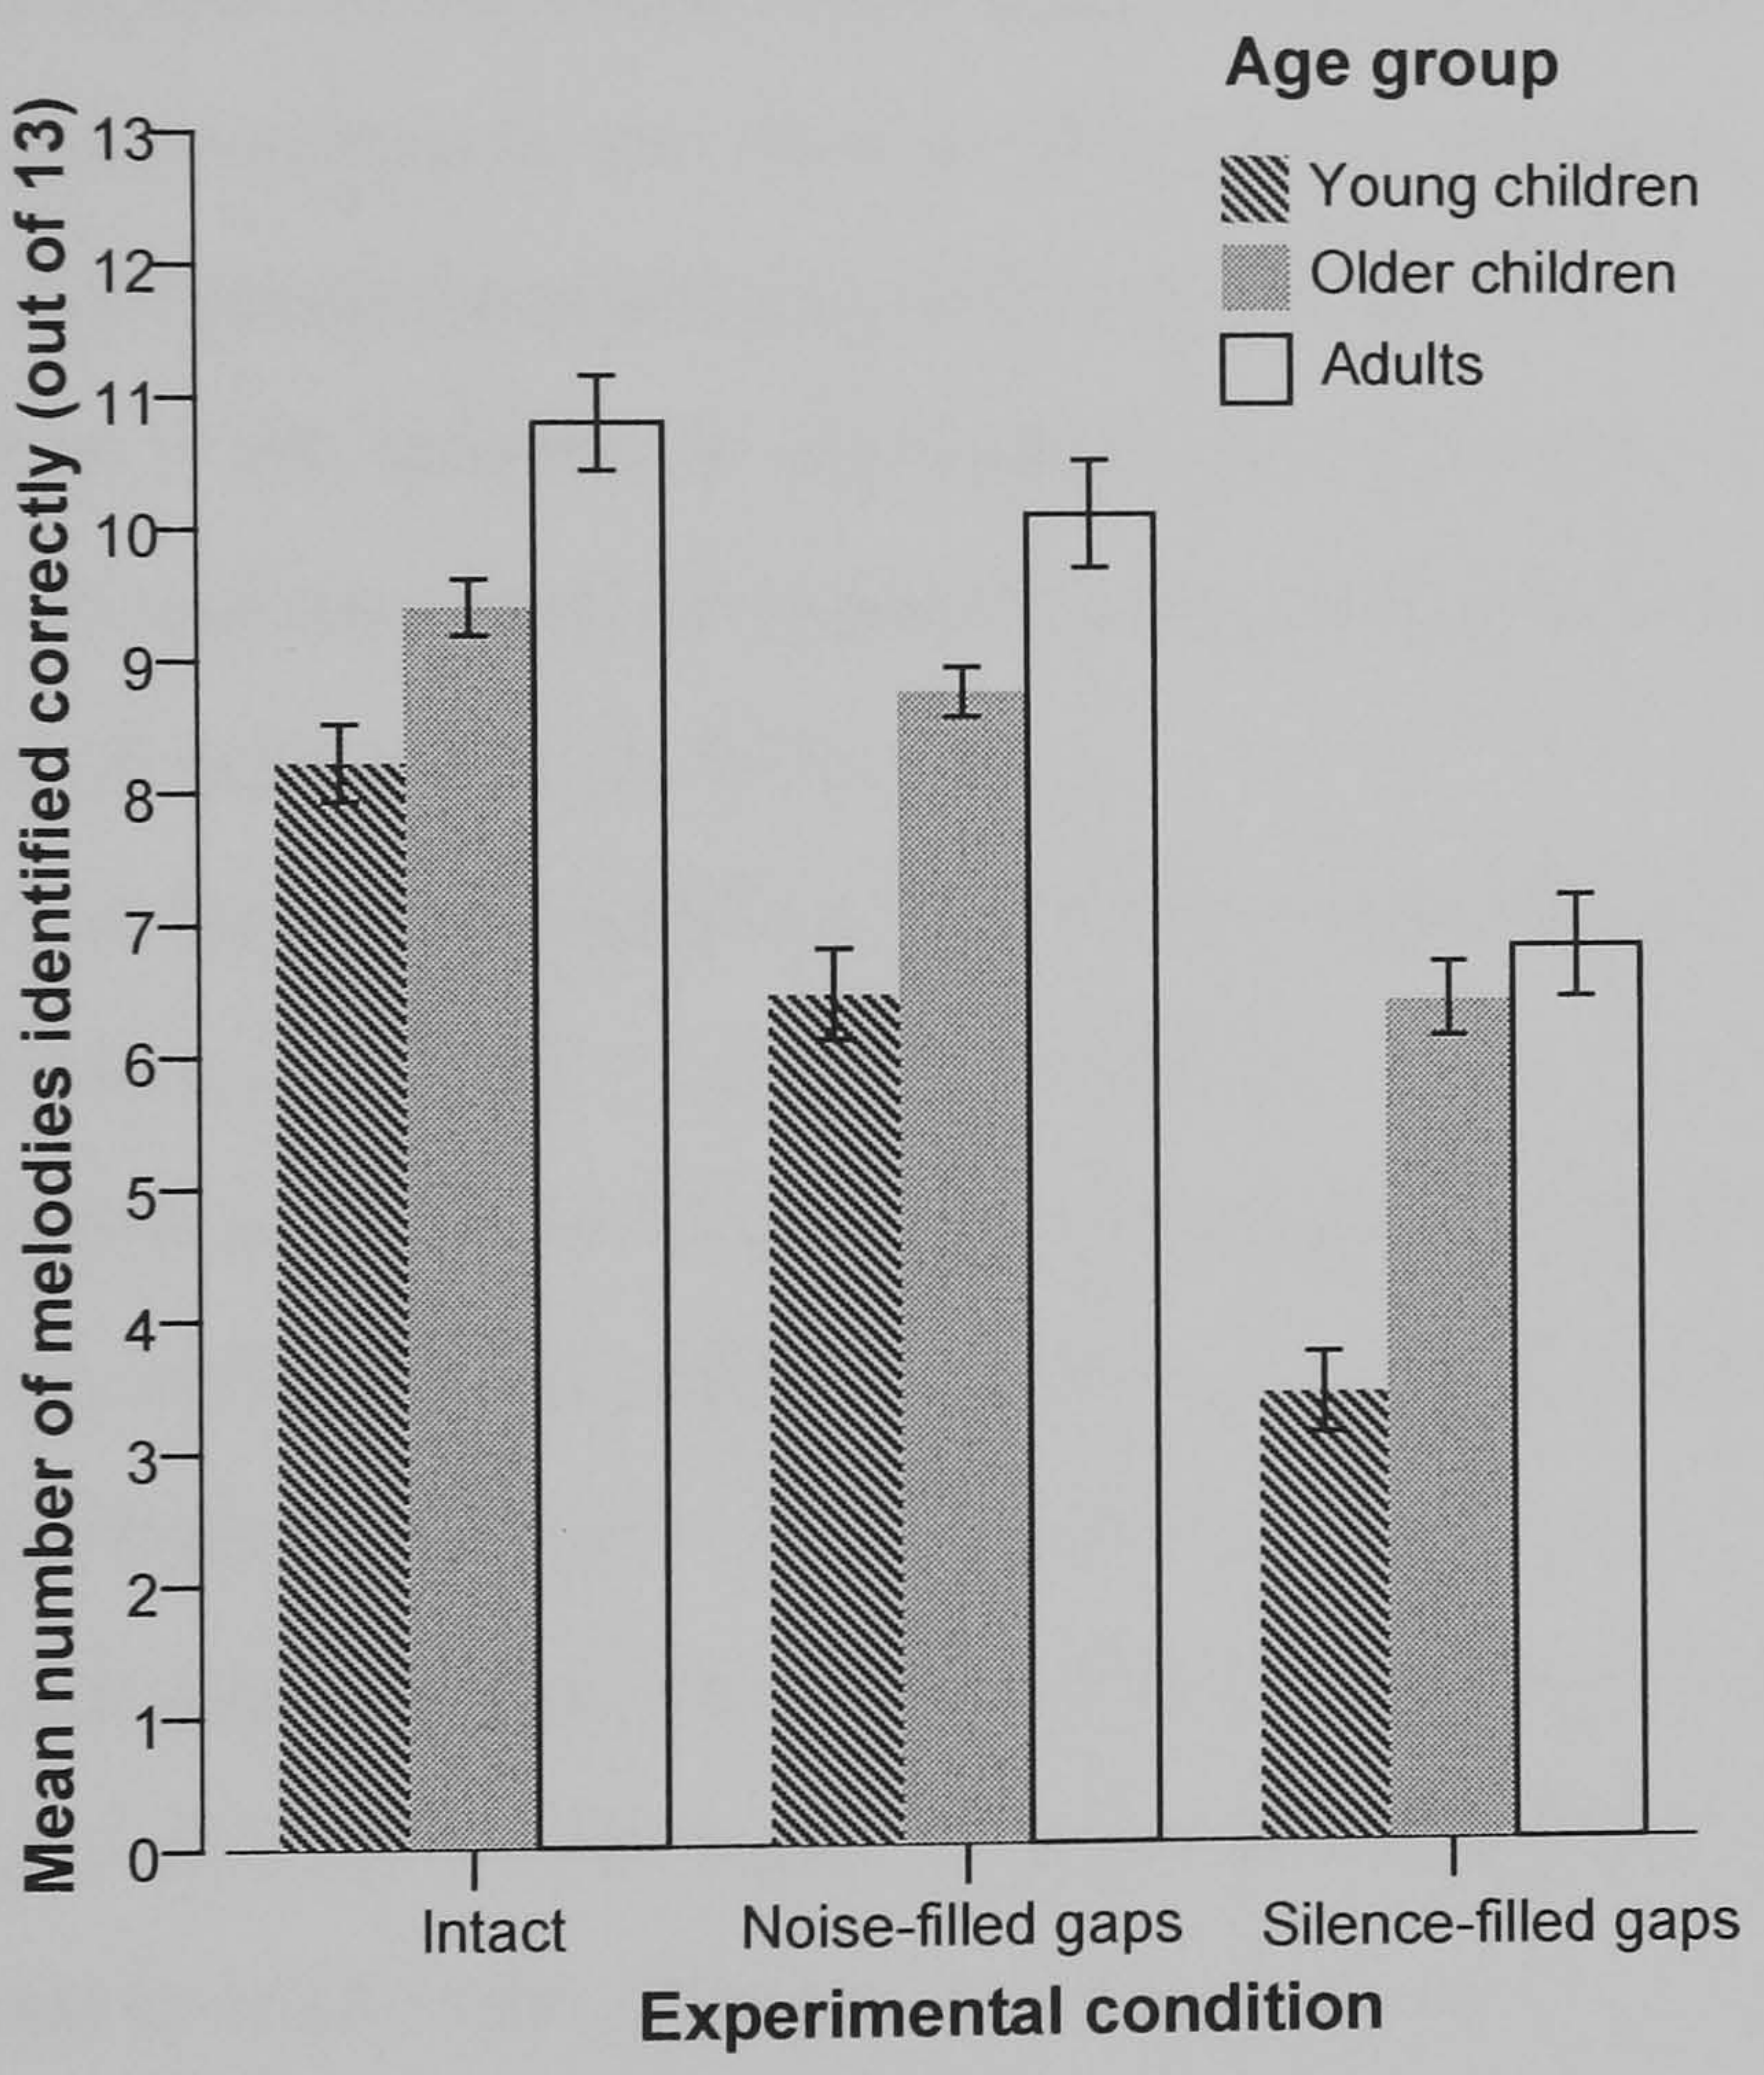

Figure 10.1. Mean number of melodies identified correctly (± *SE*) by young children, older children, and adults, in each of the three experimental conditions (intact, noise-filled gaps, silence-filled gaps).

These data were analysed using a 3 (age group: young children; older children; adults) × 3 (experimental condition: intact; noise-filled gaps; silence-filled gaps) between-subjects ANOVA. There was a significant effect of age group on identification, $F(2, 147) = 79.78, p < .001, \eta^2 = .52$. The identification scores of young children were significantly lower than those of older children and adults (Tukey HSD: $p < .001, d = 0.87$ and $p < .001, d = 1.29$, respectively). In addition, there was also a significant difference between the identification scores of older children and adults (Tukey HSD: $p < .01, d = 0.49$), with adults identifying more melodies correctly than older children.

There was also a significant effect of experimental condition on identification, $F(2, 147) = 117.49, p < .001, \eta^2 = .62$. Post hoc tests (Tukey HSD) revealed significant differences between all pairs of conditions: intact versus noise-filled gaps ($p < .001, d = 0.58$); intact versus silence-filled gaps ($p < .001, d = 1.94$); noise-filled gaps versus silence-filled gaps ($p < .001, d = 1.38$). Overall, identification in the intact condition was superior to that of the noise-filled gaps condition, which in turn had superior identification scores to the silence-filled gaps condition.

Age group interacted significantly with experimental condition, $F(4, 147) = 2.96, p < .05, \eta^2 = .08$. In order to investigate both the effect of age group within experimental condition, and the effect of experimental condition within age group, this interaction is explored in both ways, in this order.

For the intact condition, there was a significant effect of age group on identification, $F(2, 49) = 16.93, p < .001, \eta^2 = .41$. The scores of the young children were significantly lower than those of the older children (Tukey HSD: $p < .01, d = 0.85$) and of the adults (Tukey HSD: $p <.001, d = 1.87$). Adults also outperformed the older children (Tukey HSD: $p < .05, d = 1.23$).

An effect of age was also evident in the noise-filled gaps condition, $F(2, 33) = 34.19, p < .001, \eta^2 = .56$. Again, young children had lower identification scores when compared to both older children (Games-Howell: $p < .001, d = 1.41$) and adults (Games-Howell: $p < .001, d = 2.24$). The adult group had significantly higher identification scores than the older children (Games-Howell: $p < .05, d = 1.05$).

In the silence-filled gaps condition, there was also an effect of age group, $F(2, 49) = 34.90, p < .001, \eta^2 = .59$. As was the case for the other conditions, the identification scores of young children were significantly lower than those of the older children (Tukey HSD: $p < .001, d = 2.08$) and the adults (Tukey HSD: $p < .001, d = 2.36$). However, in contrast to intact and noise-filled gaps conditions, for the silence-filled gaps condition there was no significant difference between the older children and adults in their identification scores (Tukey HSD: $p > .05$).

When the interaction is examined across experimental condition, the young children showed an effect of experimental condition, $F(2, 63) = 59.37, p < .001, \eta^2 = .65$. There were significant differences between all pairs of conditions (Tukey HSD: all *p*s < .001). There was also a significant effect of experimental condition for both older children and adults, $F(2, 57) = 46.28, p < .001, \eta^2 = .62$, and $F(2, 27) = 30.73, p < .001, \eta^2 = .70$, respectively. However, for both the older children and adults, the difference between identification scores in intact and noise-filled gaps conditions was non-significant (Tukey HSD: $p > .05$). All age groups demonstrated a noise-filled gaps over silence-filled gaps advantage in identification (Tukey HSD: all *p*s < .01), but the size of the effect differed by age group: young children, $d = 1.87$; older children, $d = 1.85$; adults, $d = 2.56$.

To summarise, an effect of age group was evident for all experimental conditions, and an effect of experimental condition was evident for all age groups. However, in the silence-filled gaps condition the older children did not differ to the adults; and for the older children and adult groups, identification of melodies with missing fragments replaced by noise was not significantly inferior to identification of intact melodies.

Identification time

In keeping with Experiments 1 to 7, mean identification times as presented here represent the difference between stimulus length and the time taken to identify the stimulus, in order to control for differences in stimulus length. Data are collapsed across gender as no significant effects were found (see Appendix 2). The mean identification times by age group and experimental condition are shown in Figure 10.2.

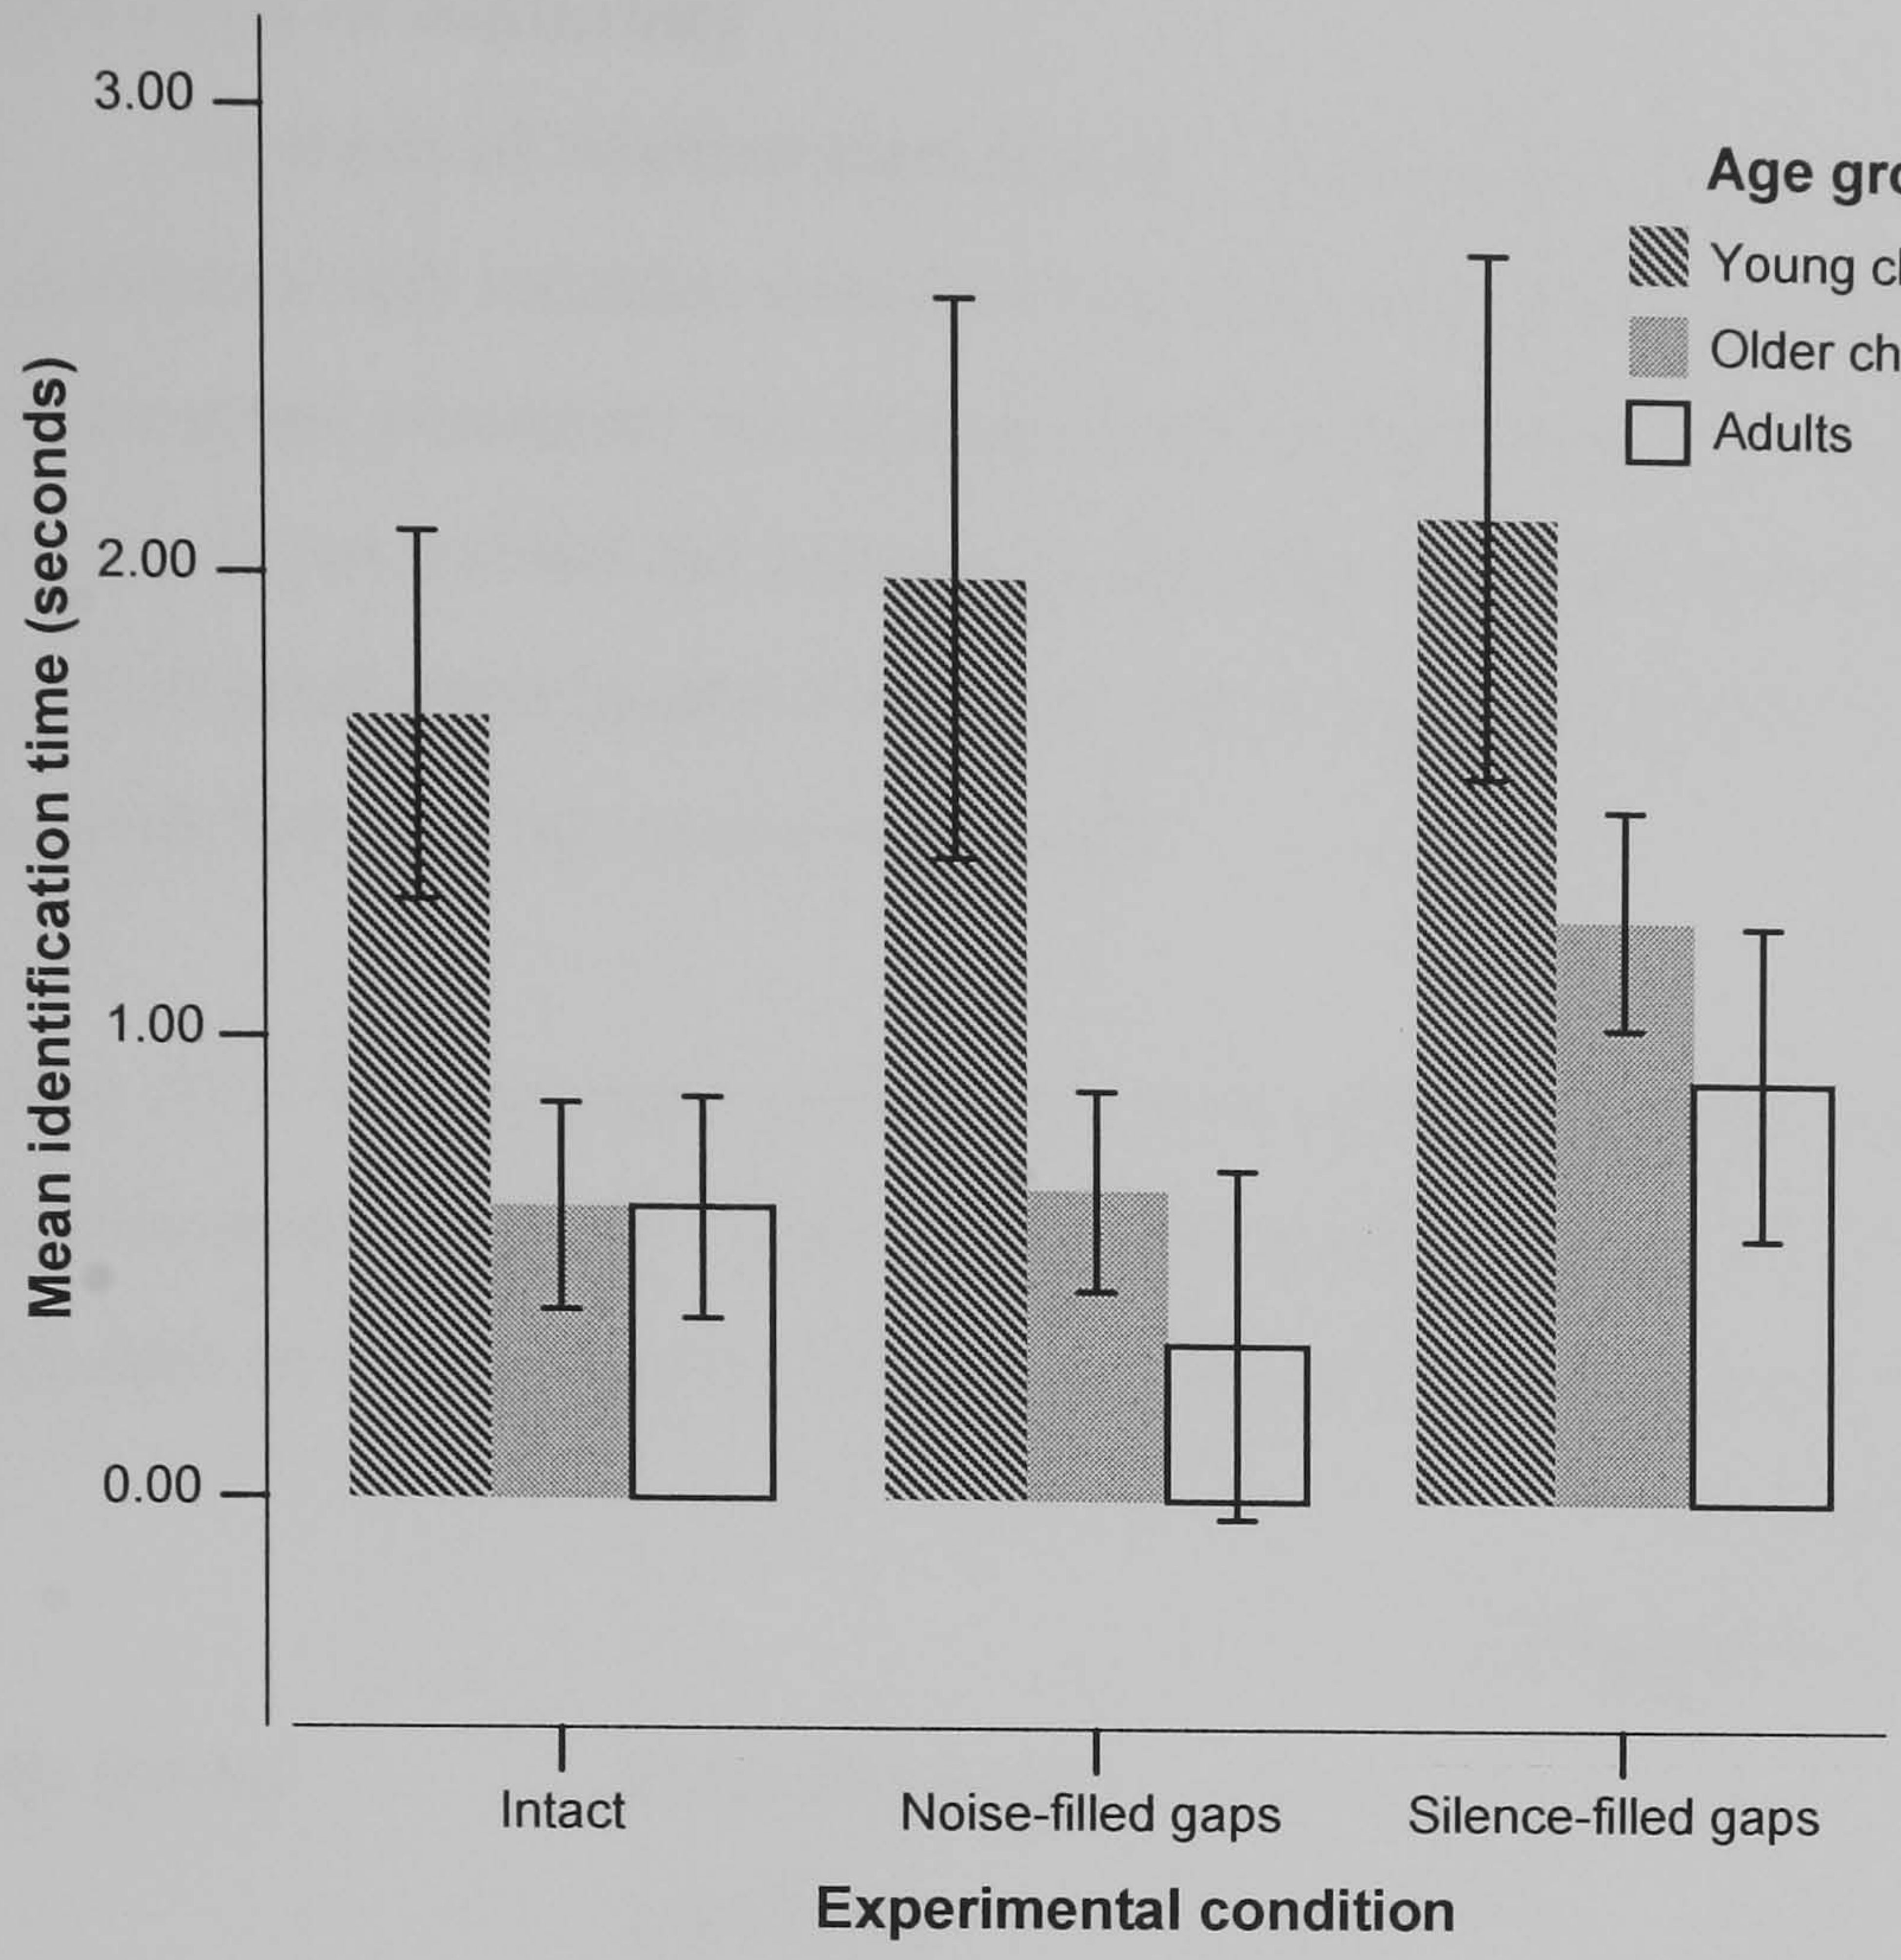

Figure 10.2. Mean identification time (identification time − stimulus duration; ± *SE*) for young children, older children, and adults, in each of the three experimental conditions (intact, noise-filled gaps, silence-filled gaps).

These data were analysed using a 3 (age group: young children; older children; adults) × 3 (experimental condition: intact; noise-filled gaps; silence-filled gaps) between-subjects ANOVA. There was a significant effect of age group on identification time, $F(2, 147) = 8.08, p < .001, \eta^2 = .10$. The young children were significantly slower at identifying the melodies than both the older children (Games-Howell: $p < .01, d = 0.44$) and the adults (Games-Howell: $p < .01, d = 0.54$), but there was no difference between the identification times of the older children and adults (Games-Howell: $p > .05$).

There was no effect of experimental condition on identification time, $F(2, 147) = .98, p > .05$; nor did experimental condition interact with age group, $F(4, 147) = .15, p > .05$.

Perception of continuity

Analysis of whether participants reported the signals as continuous or fragmented only includes data from the older children and adult group, as this part of the experimental procedure was not involved in the design of Experiment 1. In addition, this analysis is not carried out separately for older children and adults due to the small cell sizes for adult participants. However, Table 10.2 does report the frequency of each response for each age group separately.

Table 10.2. Total number of participants in each age group (older children, adults) and experimental condition (intact, noise-filled gaps, silence-filled gaps) reporting the melodies as fragmented and reporting the melodies as continuous.

		Perception of Continuity	
		Fragmented	**Continuous**
Age group	**Experimental condition**		
Older children[a]	Intact	0 (0%)	20 (100%)
	Noise-filled gaps	5 (25%)	15 (75%)
	Silence-filled gaps	20 (100%)	0 (0%)
Adults[b]	Intact	2 (20%)	8 (80%)
	Noise-filled gaps	2 (20%)	8 (80%)
	Silence-filled gaps	7 (70%)	3 (30%)

[a]*n* = 20 in each condition; [b]*n* = 10 in each condition

Interestingly, whilst responses made by the older children differ between intact and noise-filled gaps conditions, the adults' pattern of responses is the same in intact and noise-filled gaps conditions. However, the overall direction of responses is the same between age groups. Over both age groups, the number of participants in each condition who reported the signals as fragmented and the number who reported the signals as fragmented is shown in Figure 10.3.

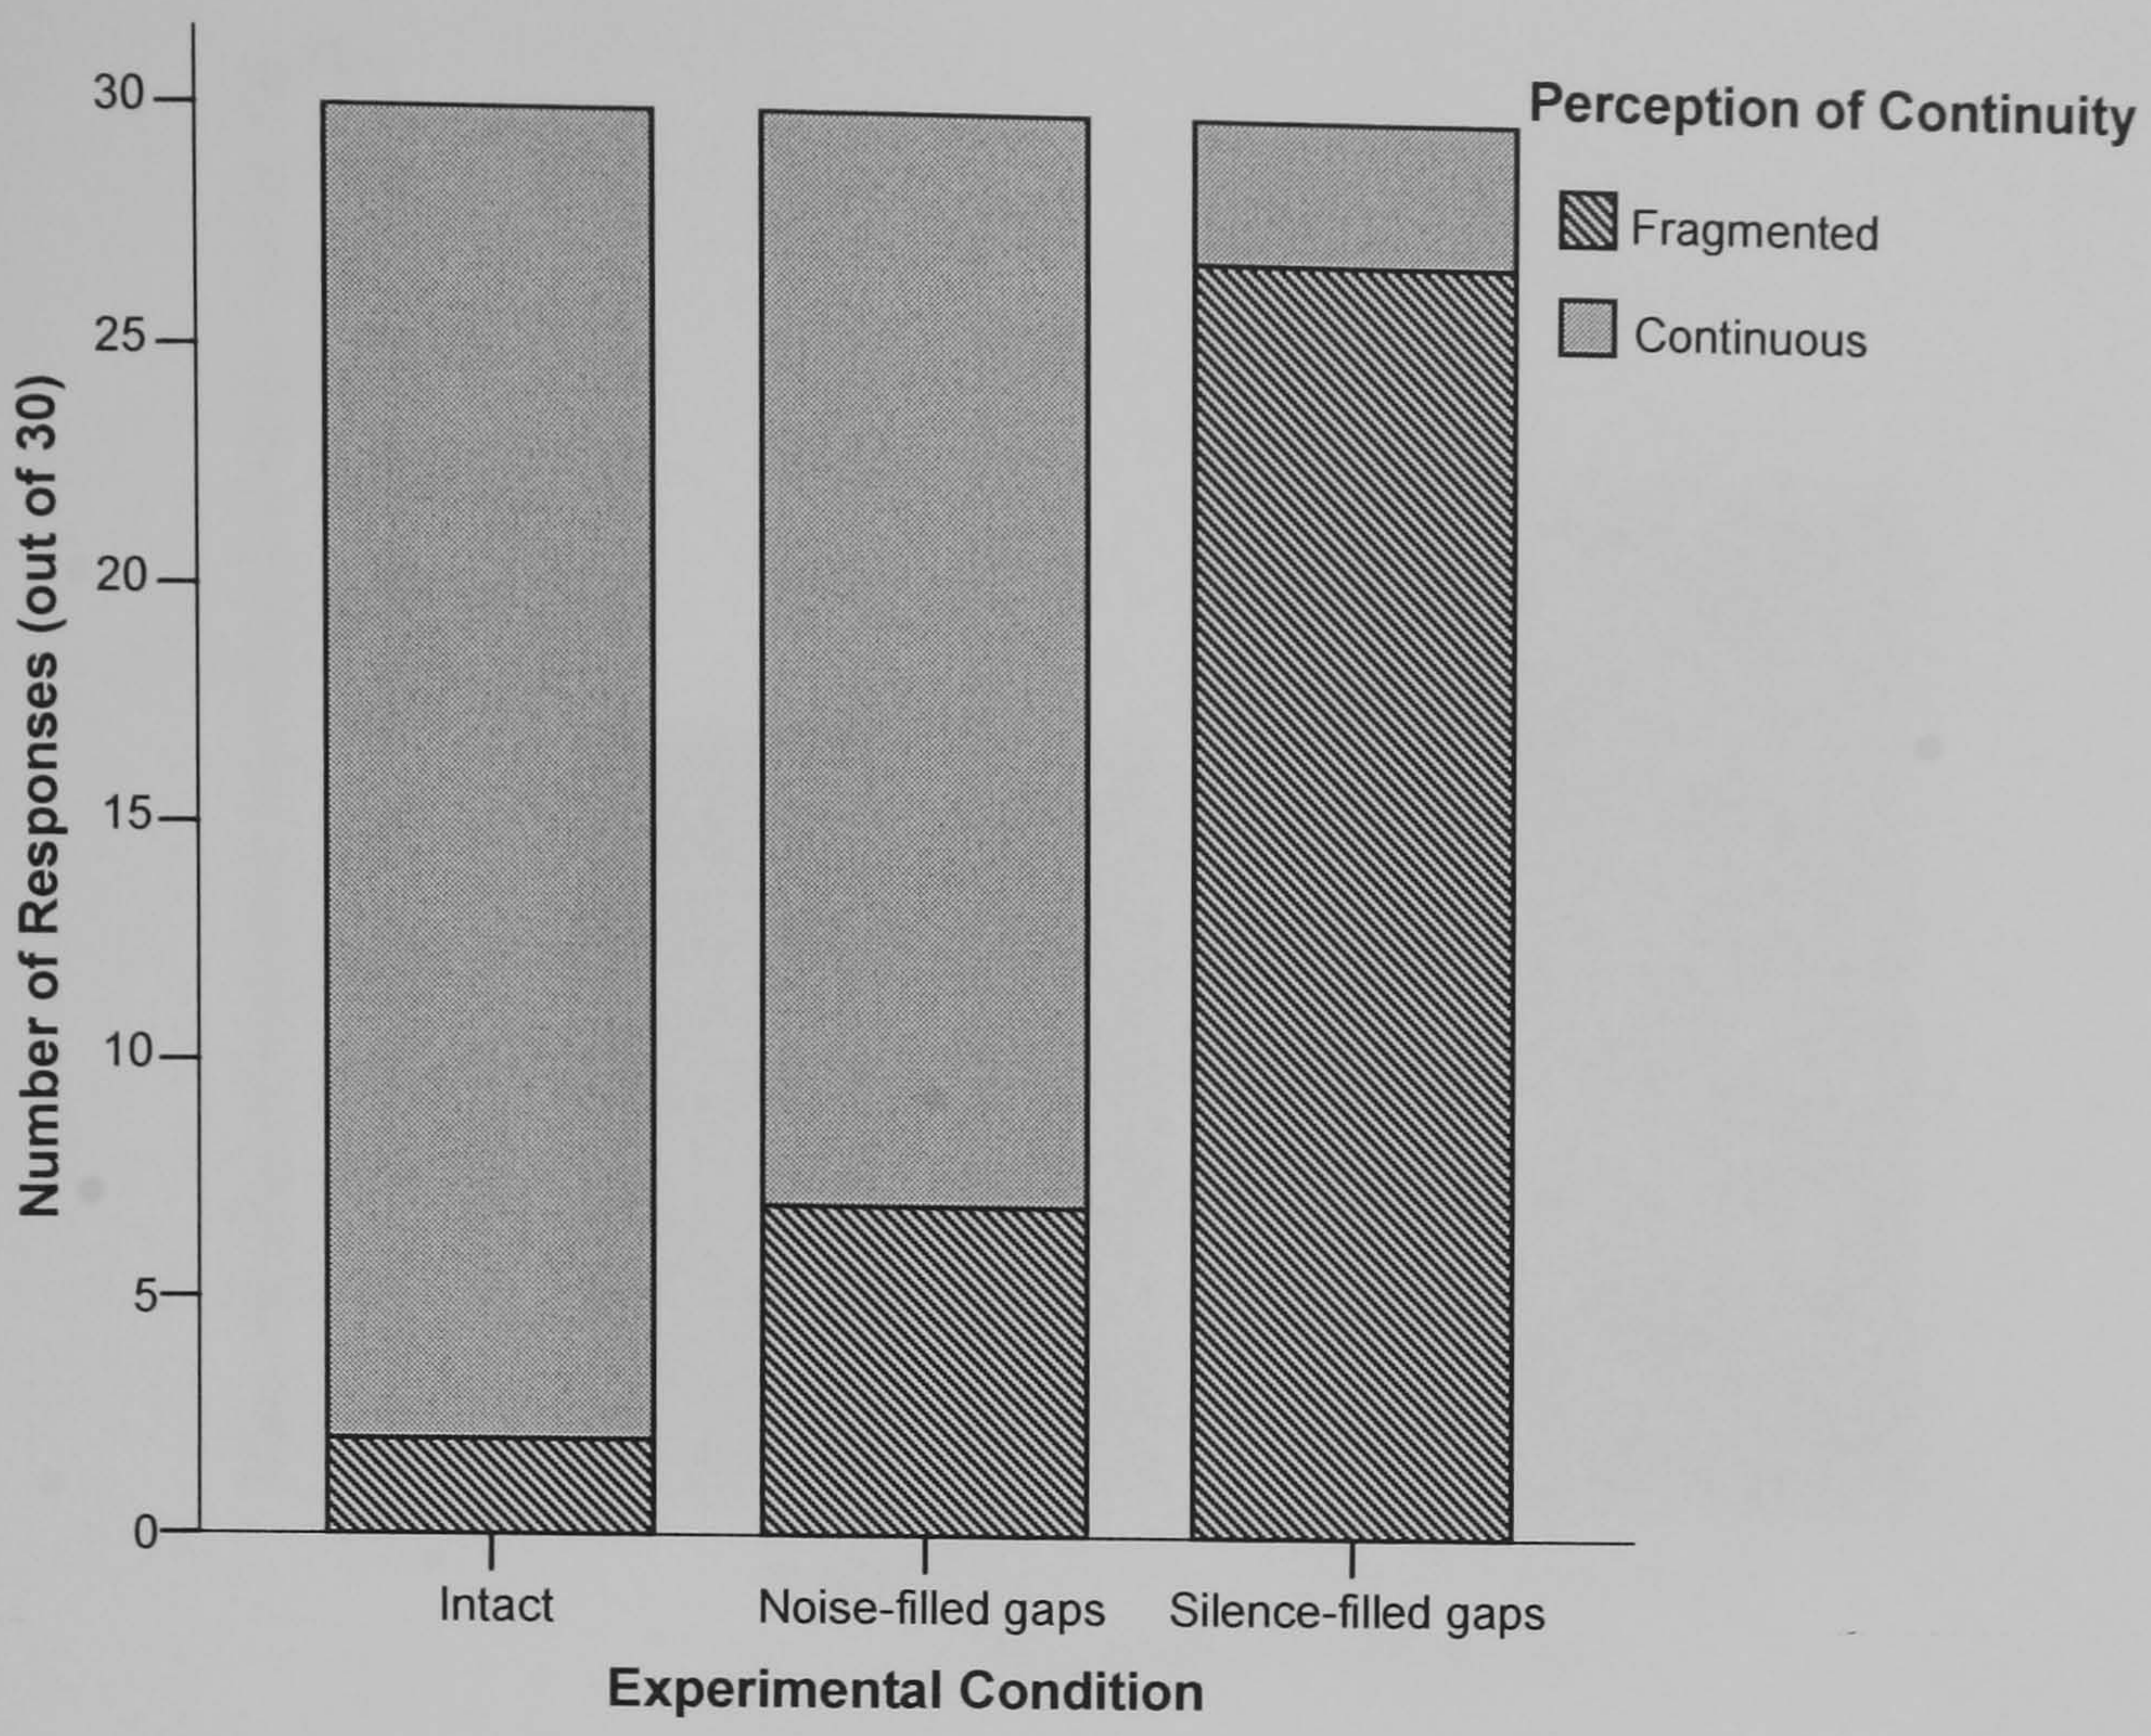

Figure 10.3. Total number of participants in each experimental condition (intact, noise-filled gaps, silence-filled gaps) reporting the melodies as fragmented and reporting the melodies as continuous.

These data were analysed using a 3 (experimental condition: intact; noise-filled gaps; silence-filled gaps) × 2 (perception of continuity: fragmented; continuous) chi-square. There was a significant association between experimental condition and the perception of continuity, χ^2 (2, $N = 90$) = 48.61, $p < .001$, Cramer's $V = .74$. In both intact and noise-filled gaps conditions, the majority of participants reported the signals to be continuous, whereas in the silence-filled gaps condition, the majority of participants reported the signals to be fragmented.

Following this analysis, participants were grouped according to whether they reported the signals as fragmented or continuous. The mean number of melodies identified correctly by participants in the "fragmented" and "continuous" groups is shown in Figure 10.4.

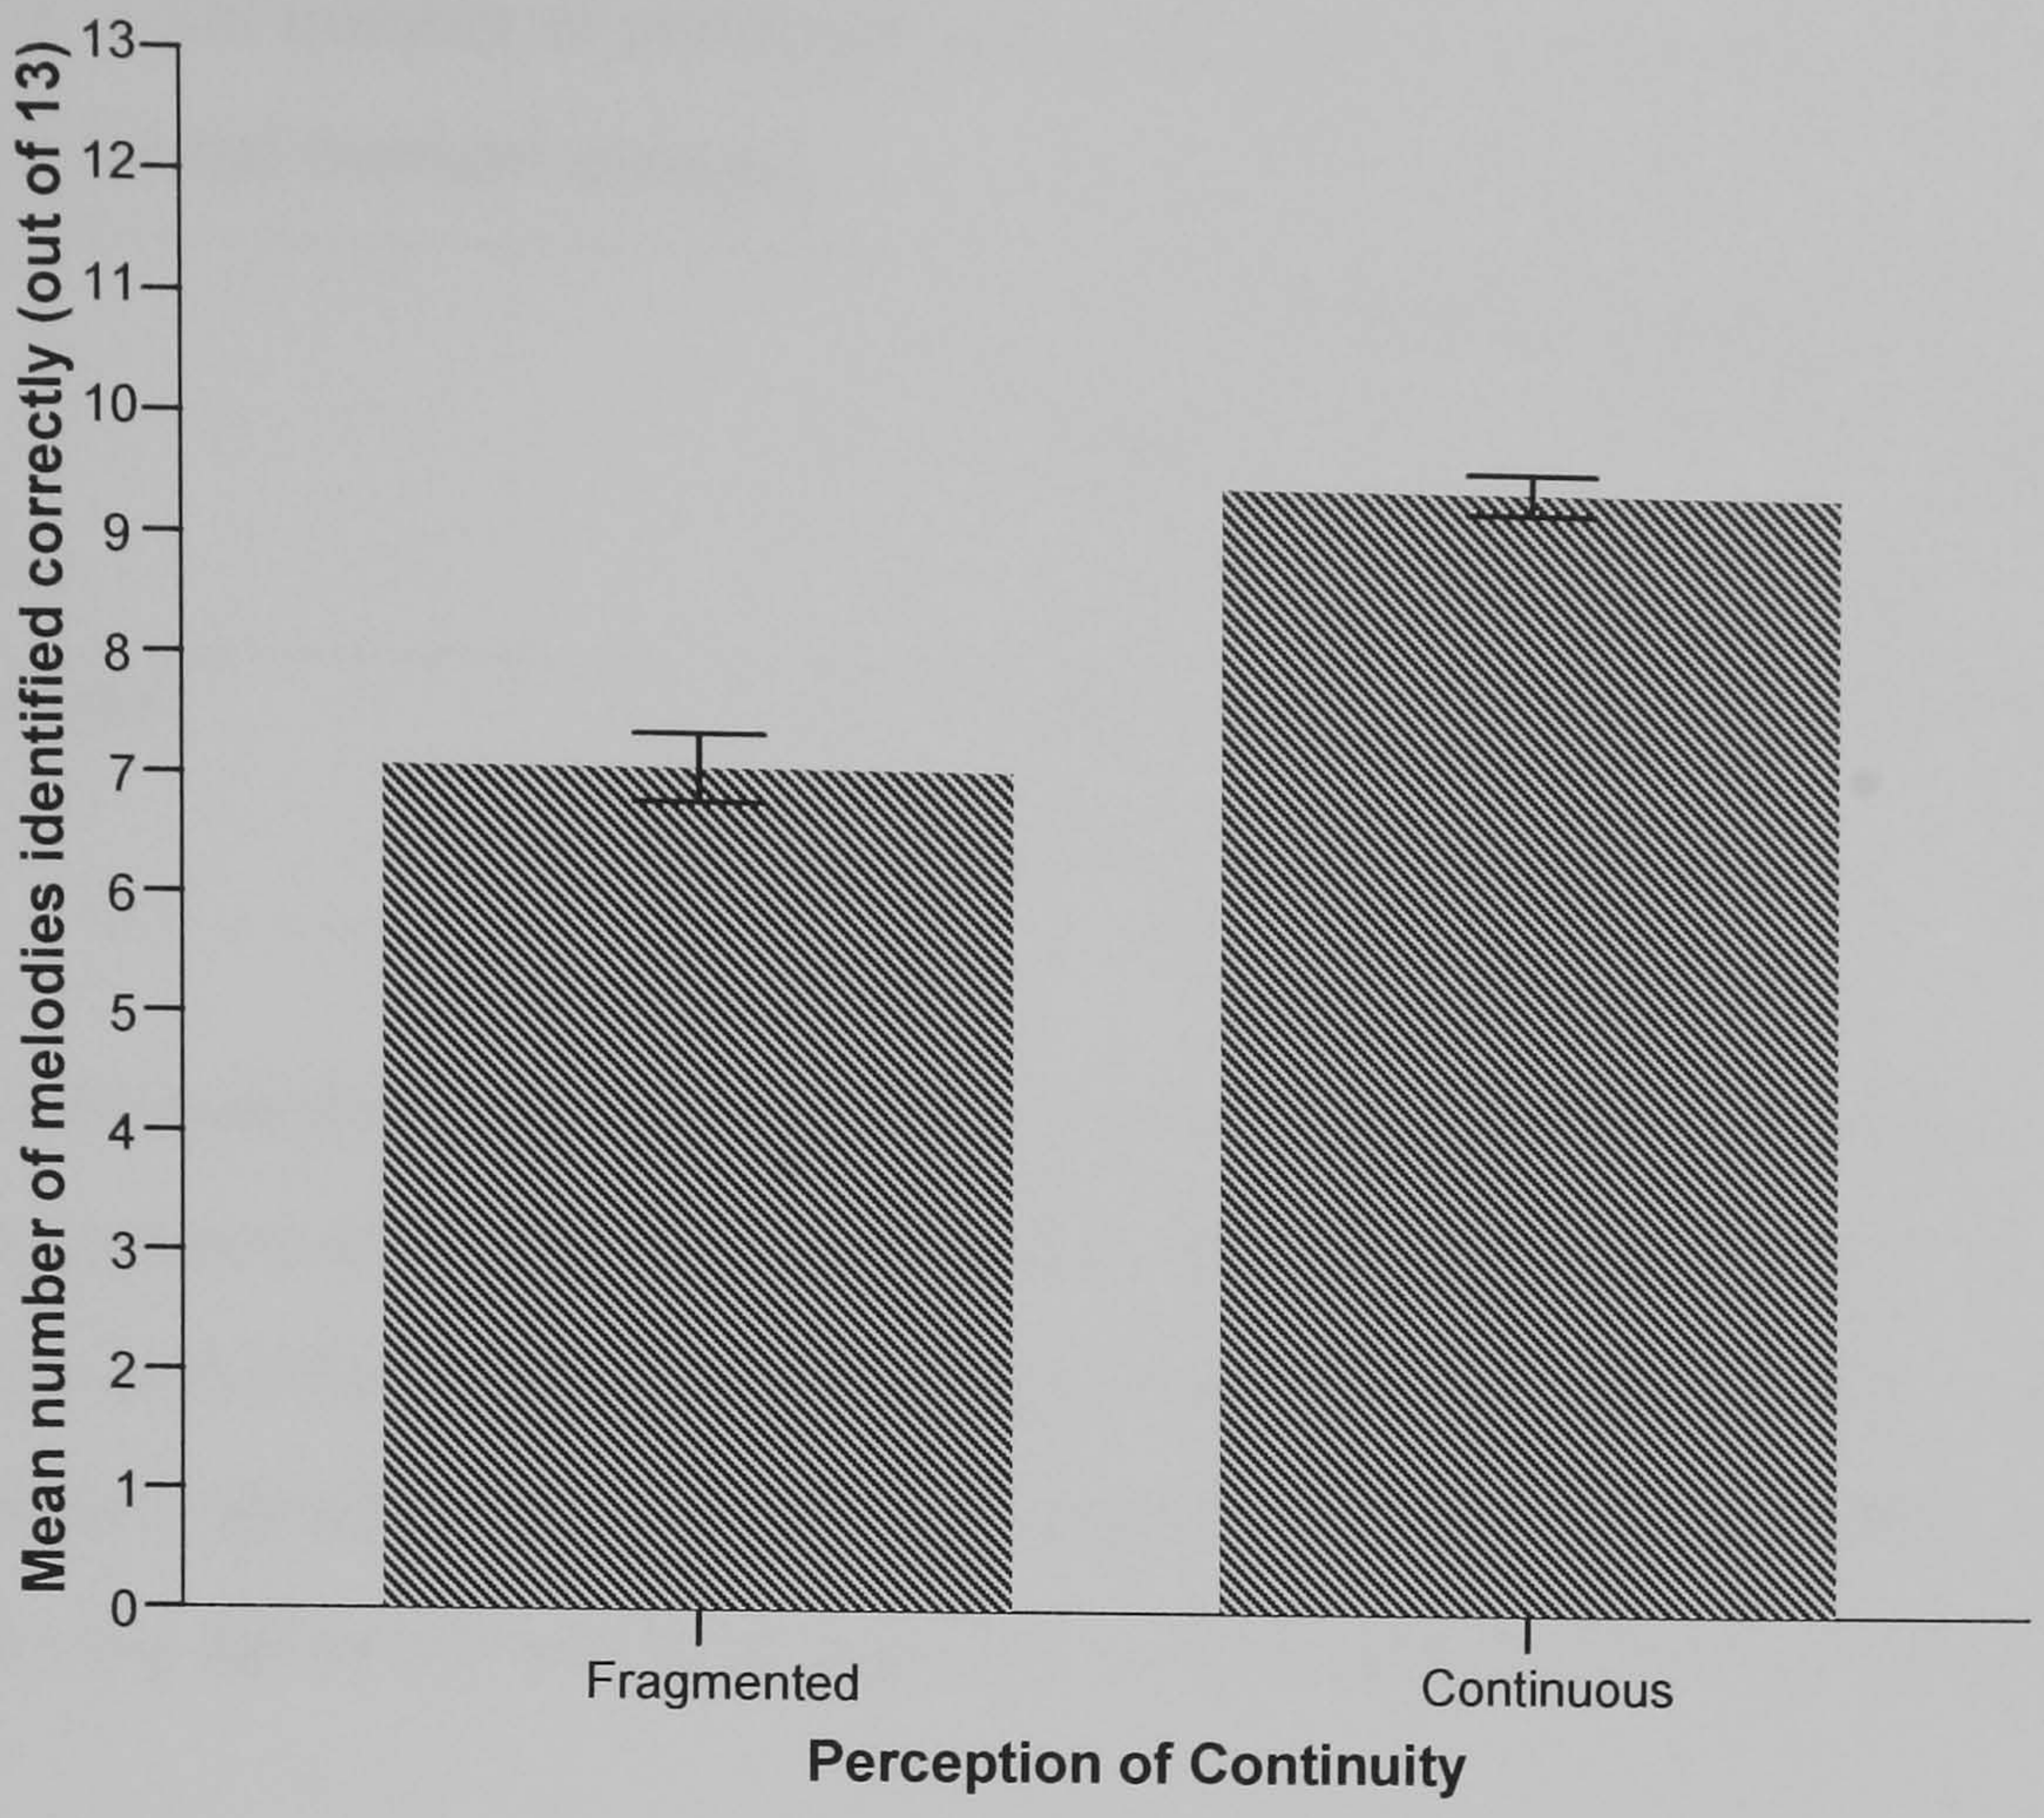

Figure 10.4. Mean number of melodies identified correctly (± *SE*) by participants who reported the melodies to be fragmented and participants who reported the melodies to be continuous.

An independent-samples *t* test showed that the participants who reported the signals as continuous had significantly higher identification scores than the participants who reported the signals as fragmented, $t(59) = 7.21, p < .001, d = 1.40$.

The effect of musical training

Both older children and adults were asked if they took, or had taken in the past for longer than one year, instrumental music lessons as a measure of formal musical experience. The frequency of responses is shown in Table 10.3.

Table 10.3. Total number of participants in each age group reporting presence or absence of formal musical training.

	Musical Training	
	Yes	**No**
Age group		
Older children	27	33
Adults	17	13

Total melody identification scores were compared by experimental condition, and by experience of formal musical training. This analysis was carried out with data from adults and older children pooled due to the small sample size of the adult group. The mean number of melodies identified correctly by participants in each experimental condition who did or did not have experience of musical training is shown in Figure 10.5.

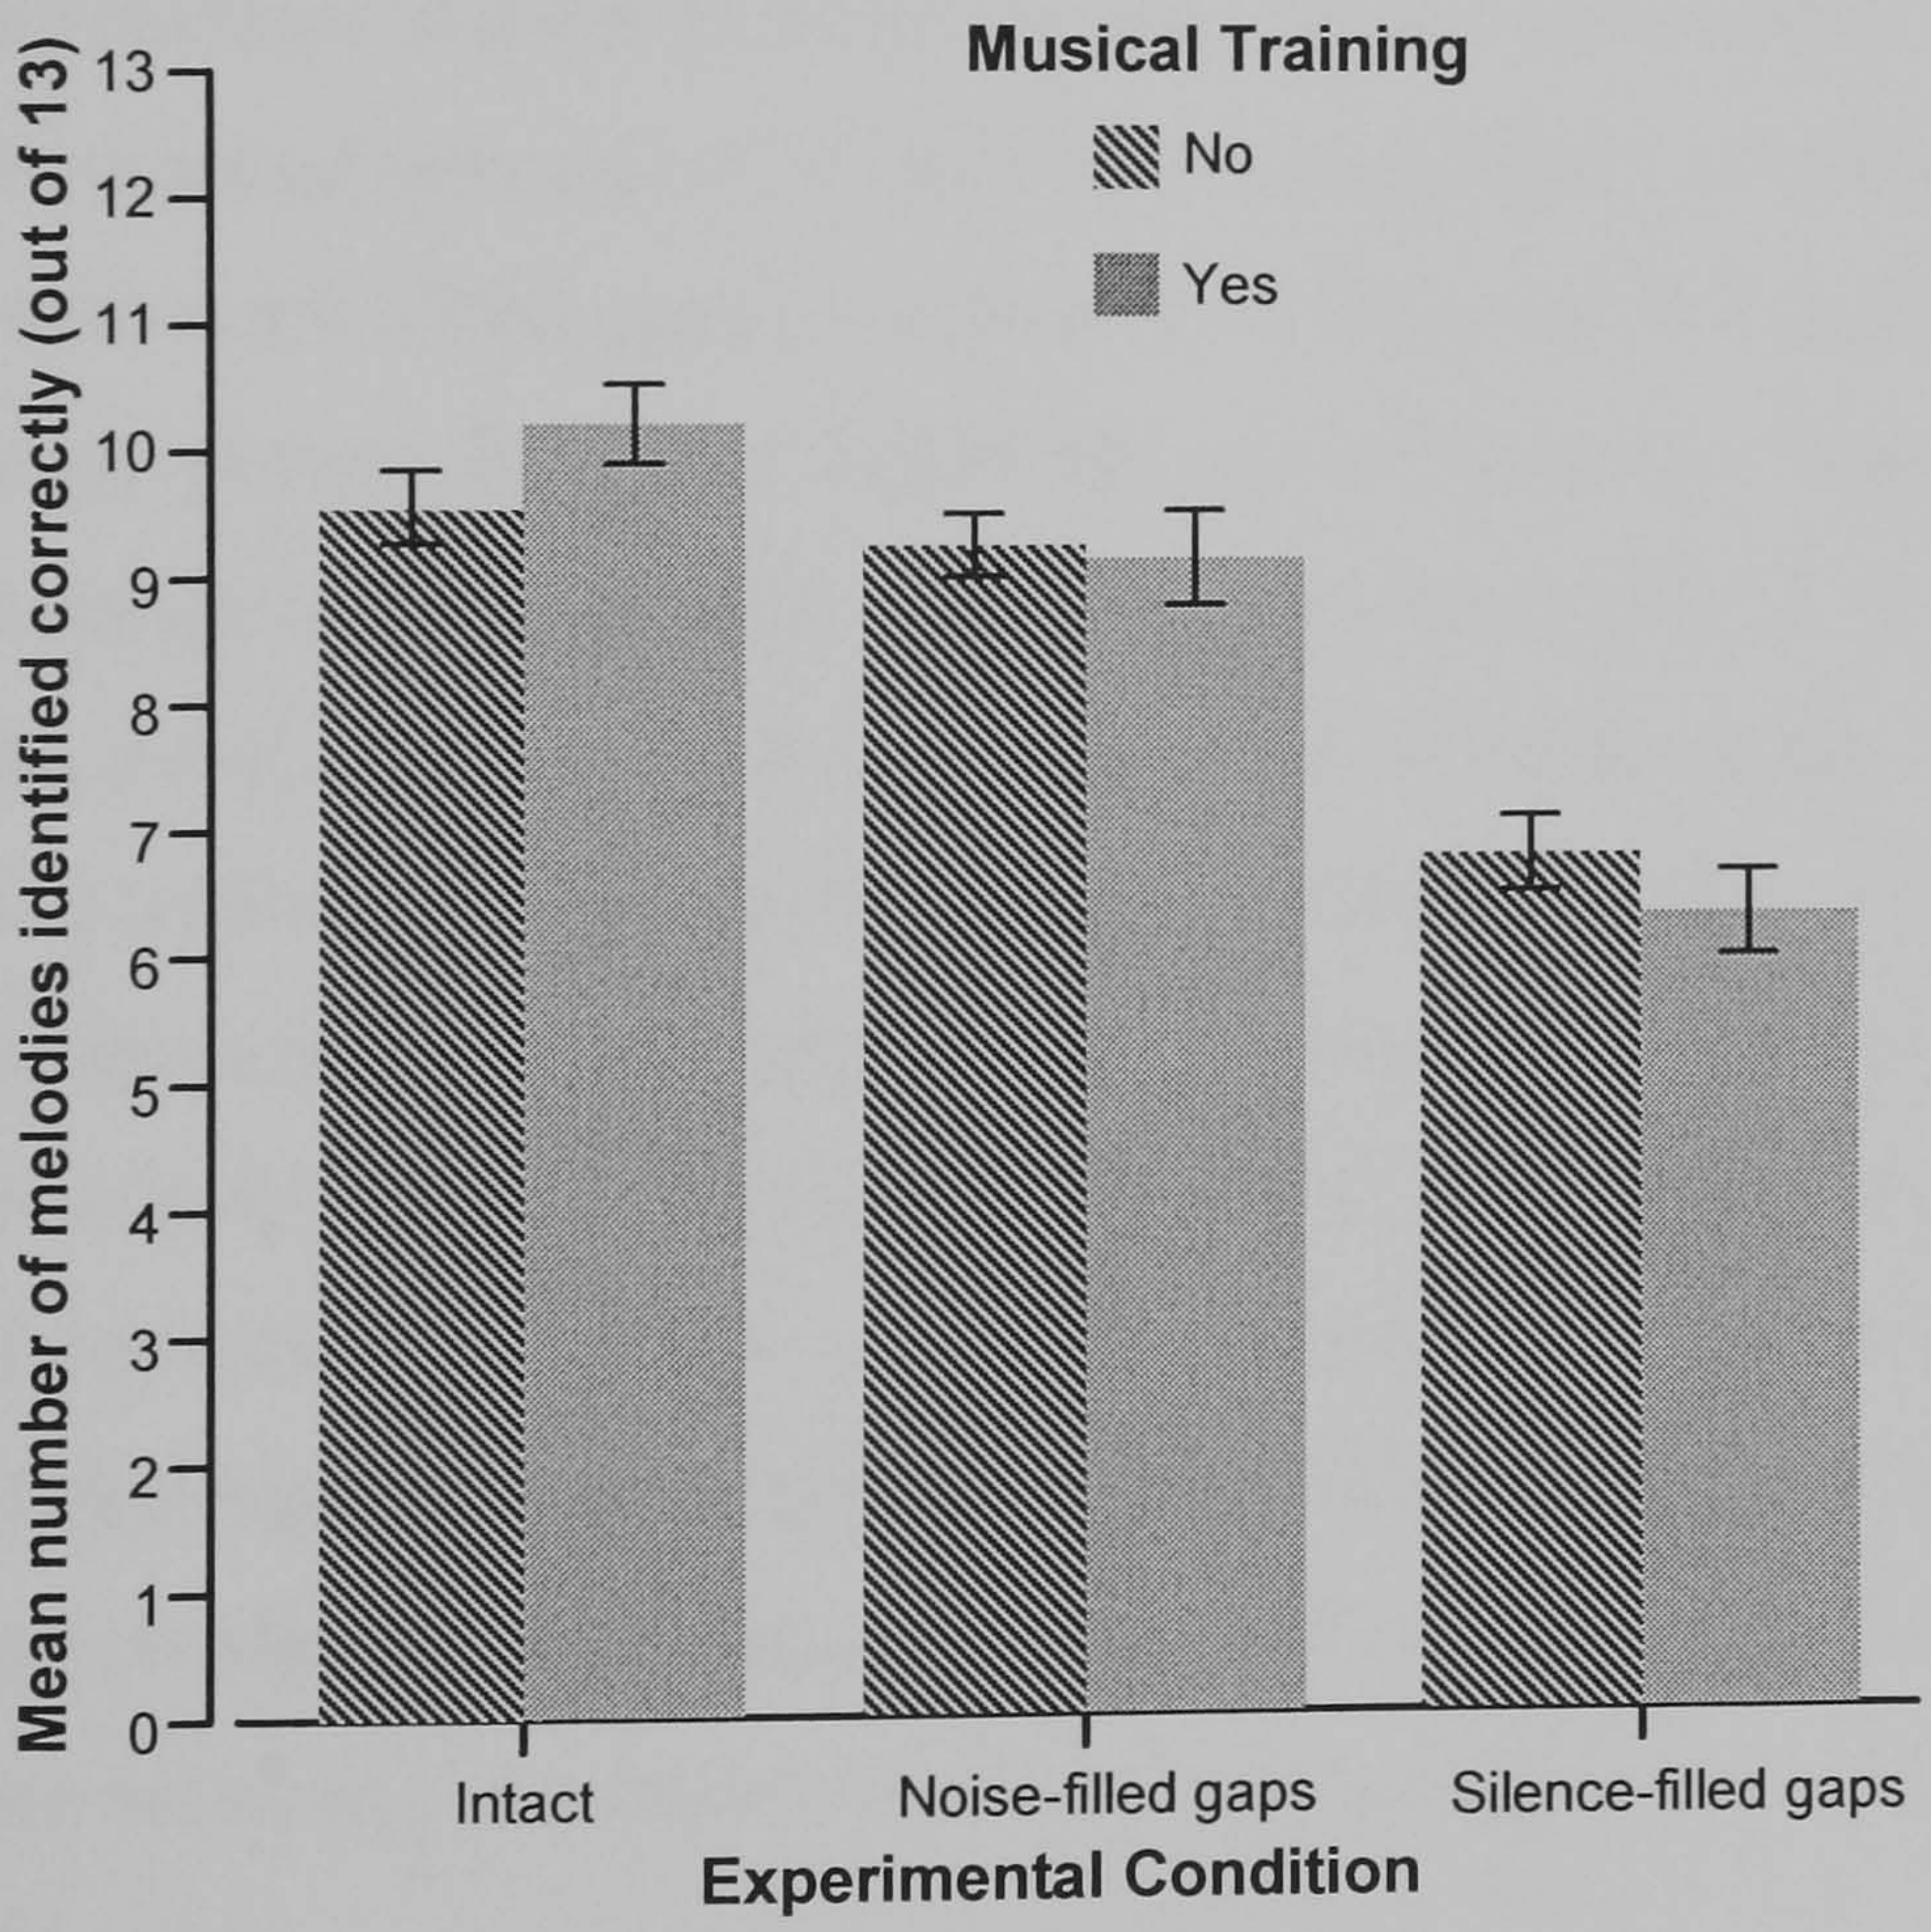

Figure 10.5. Mean number of melodies identified correctly (± *SE*) by participants in each experimental condition (intact, noise-filled gaps or silence-filled gaps) with and without formal musical training.

These data were analysed using a 3 (experimental condition: intact; noise-filled gaps; silence-filled gaps) × 2 (musical training: yes; no) between-subjects ANOVA. There was no effect of musical training on identification scores, $F(1, 84) = .10, p > .05$, and, most importantly, musical training did not interact with experimental condition, $F(2, 84) = 1.68, p > .05$. Whilst Figure 10.5 implies that musical training improves identification of intact melodies, it is evident that musical training has no effect on identification of melodies with noise-filled gaps, which would represent the most important consideration.

Discussion

The main aim of the present experiment was to investigate differences in the perceptual restoration of musical melody with increasing age. Whilst experiments 1 to 6 found evidence of musical perceptual restoration in 4- to 6-year-old children, their identification of stimuli with noise-filled gaps was not at ceiling hence it was expected that older children and adults might display greater proficiency in this ability.

Perceptual restoration, as shown by a significant identification advantage of stimuli with noise-filled gaps over those with silence-filled gaps, was demonstrated for all three age groups. However, whilst the size of the noise-filled gaps over silence-filled gaps advantage in terms of effect size was very similar for 4- to 6-year olds and 9- to 11-year-olds, a much larger effect was evident for the adult group. This suggests that perceptual restoration was much stronger for adults than for children, at least as far as this measure is taken to indicate strength of perceptual restoration. In contrast, Newman (2004) found that despite lower overall levels of repetition accuracy of sentences across all conditions for children when compared to adults, the size of the noise-filled gaps over silence-filled gaps advantage was roughly equivalent. The difference between Newman's findings and the results of the present experiment could represent a difference between developmental changes in perceptual restoration in music and speech domains.

With further reference to identification scores, there were developmental effects evident for all experimental conditions. When melodies were presented intact, adults were superior to children of both age groups in identifying the melodies. However, a larger difference between age groups emerged in the noise-filled gaps condition, which represents the conditions under which perceptual restoration can occur. Most striking is the considerably larger increase in melody identification scores in this condition between younger children and older children, than between older children and adults. This indicates that the ability to combine previous knowledge with incoming acoustic information develops significantly between the ages of 4 and 11. Crucially, whilst perceptual restoration does operate in 4- to 6-year-old children, this ability is still to undergo considerable developmental change.

Also important for interpreting perceptual restoration abilities in young children is the finding of a significant difference between intact and noise-filled gaps conditions for the young children, which was not evident for both the older children and the adults. This suggests that young children's perception is more greatly affected by disruptions to auditory signals than perception in older children and adults. A similar conclusion was made by Newman (2004) when comparing the perceptual restoration of speech between young children and adults. Investigation of perceptual restoration effects in young children, and finding possible ways to develop them, thus becomes a key concern based on the finding that noisy disruptions can have a significant effect on young children's auditory perception.

In the silence-filled gaps condition, there was an increase in identification scores between the young children and the older children, but no increase between the older children and the adults. This suggests that young children have particular difficulty interpreting musical signals where missing sections remain silent. Consequently, perceptual restoration emerges as a critical mechanism for these young children. If identification of disrupted input with silent gaps is particularly difficult, this suggests that young children are unable to make the same inferences that older children and adults can make in the process of perception. Thus, noise as a gap replacement plays a key role in allowing young children to make sense of disrupted input.

Taken together, the identification scores indicate that perceptual restoration does show clear developments with age. The largest improvements occur between young children and older children, rather than between older children and adults. In particular, young children appear to have particular difficulty interpreting melodies with silent gaps when compared to the same ability in older children and adults. These findings support previous studies that have found developmental increases in the perceptual restoration of speech, both between younger and older children (Ackroff, 1981) and between young children and adults (Walley, 1988).

Analysis of the identification time data revealed that young children were significantly slower at identifying the melodies than both the older children and the adults. This was true across all experimental conditions, and suggests that less input is needed for older children and adults to identify the melodies. This result may also indicate that younger children are slower at using what is present within the signal to access the stored representation of the melody in their musical lexicon.

Reports of the continuity of the signal made by the older children and adults suggest that when missing sections of the melodies were replaced by noise, the majority of responses made by participants reflect the hearing of a continuous melody. However, the frequency of responses broken down by age group suggests that adults were less consistent in their responses, across all conditions, than the older children. This could be a result of a greater analysis of the question by the adults, and perception of more complex task demands than was actually the case. Nevertheless, across both age groups there was a clear pattern of responses across conditions, and a clear difference between those that reported the melodies as continuous and those that reported the melodies as fragmented in terms of the number of correct identification responses.

Consistent with the musical perceptual restoration experiments reported by DeWitt and Samuel (1990), the present experiment found no effect of musical training on the ability to identify melodies with missing sections replaced by noise, and no significant effect on identification in intact and silence-filled gaps conditions. DeWitt and Samuel concluded, on the basis of their results, that the operation of musical perceptual restoration is dependent on abilities that are not facilitated by formal musical training. This could be a premature conclusion. An important similarity between DeWitt and Samuel's experiments and the present experiment is the use of highly and universally familiar musical representations, such as nursery rhyme melodies, as stimuli. It is possible that formal musical training, and the associated ability to anticipate upcoming acoustic events with reference to musical regularities, would be useful for perceptual restoration of more complex and less familiar musical signals. Musical training increases knowledge of musical regularities and the ability to predict the continuation of a musical excerpt according to these regularities (Bartlett & Dowling, 1980). Therefore, the most appropriate conclusion is that formal musical training appears to have little influence on the perceptual restoration of highly familiar, simple melody lines. Whether an influence would emerge if more complex musical signals were presented, represents a question that can be addressed through further research.

The operation of perceptual restoration for all age groups tested indicates that the principle of Gestalt closure and a preference for global coherence both apply to music perception from young childhood to adulthood. Whilst stimuli in both the noise-filled gaps and silence-filled gaps conditions contain disruptions at a local level, only by replacing missing sections with noise does the stimulus satisfy the requirements of a perceptually coherent whole. However, there is no way of knowing whether perceptual restoration in children and adults, and in children of different ages, is mediated by the same perceptual processes. It is possible that age effects signify a different trade-off between top-down and bottom-up processes in music perception, with adults emerging as less affected by signal disruptions than children because they rely more heavily on top-down processing. As a result, acoustic disruptions to the signal have less of a detrimental effect on their perception. Ultimately, however, the present series of experiments cannot address the problem of whether perceptual restoration is the result of comparable processing mechanisms in adults and children, and in speech and music domains. Addressing this level of analysis represents an important area for future work.

Whilst it can be assumed, on the basis of familiarity reports made by the adult participants, and by school teachers on behalf of the young and older children, that the nursery rhymes were familiar to all participants, it is likely that levels of exposure at the time of the experimental testing session would differ considerably. Whilst the young children are likely to be exposed to these nursery rhymes on a daily basis, adults may not have heard them for a very long time. This can probably account for the fact that even in the intact condition melody identification by adults was not at ceiling level. Alternatively, it is possible that the degree of familiarity with the stimuli was similar between children and adults, but that the way this familiarity is translated into expectations that could influence perceptual restoration could differ between age groups. In order to ensure equal levels of familiarity between all age groups tested, newly familiarised stimuli could be used that would control for differences in exposure directly. However, despite this possible limitation, all stimuli were reported as familiar which for present purposes leads to the assumption that differences in identification might represent differences in perceptual processing.

A key question emerging from analysis of the results of the present experiment is exactly what it is that develops with age that leads to superior performance on a perceptual restoration task, and whether these developments are responsible for improvements in perceptual restoration across domains. If such key developments were identified, there is a likelihood that, by training young children in these abilities, it would be possible to improve their ability to comprehend auditory signals such as music in the noisy listening conditions that they are faced with on a daily basis. It would also be important to investigate perceptual restoration abilities at the intervening developmental stages to those tested in the present experiment. A linear increase in perceptual restoration ability might not emerge, and investigating the developmental course of perceptual restoration would help to reveal the important abilities underlying the operation of a mechanism that represents a fundamental adaptation to noisy listening conditions.

To summarise, the present experiment has added to the emerging picture of musical perceptual restoration in children from Experiments 1 to 7 by revealing that the perceptual restoration of music does show developments with age. This in turn suggests that some of the proposed perceptual processes underlying perceptual restoration might show changes with development. It is also possible that different weightings are applied to different sources of information within the acoustic signal by listeners of different ages. The results from Experiments 1 to 7 with young children therefore do not under any circumstances represent optimum performance, and it is imperative that the results be interpreted as demonstrating perceptual restoration abilities in young children, and that important changes in the results obtained might emerge if different age groups were tested on the same task.

Taken together, these experiments suggest that the perceptual restoration of music in young children, in some respects, shows many parallels to a comparable mechanism in adults' speech perception (Experiments 2 and 3) and is influenced by similar factors, such as familiarity (Experiments 4 and 5). These experiments have also suggested that musical perceptual restoration in young children is highly context-dependent (Experiments 6 and 7). These abilities demonstrated by 4- to 6-year-old children are even more remarkable given that their perceptual restoration abilities are still to undergo considerable developmental change. That perceptual restoration operates for very young children represents a fundamental adaptation to their often noisy environments, mediated by perceptual biases that operate in both adults and children.

This series of experiments, in using a set of stimuli, has included some that are likely to be more familiar than others, at least in terms of degree of exposure (such as *Twinkle Twinkle Little Star* and *Happy Birthday*). These differences have been important in preventing floor and ceiling effects in identification thus allowing examination of group differences, and have not affected the analysis of relative differences between experimental conditions. However, it is possible that, apart from differences in familiarity, musical characteristics of the nursery rhymes (such as aspects of their melody, rhythm, and implied harmony) could make some more identifiable, and more importantly, more identifiable with missing sections filled with noise, than others. The chapter that follows reports an examination of identification scores by stimulus, taking into consideration the musical characteristics of each stimulus, in order to investigate whether any patterns seen in identification scores could be a result of musical factors influencing musical perceptual restoration.

Chapter 11:

An analysis of the effect of the musical characteristics of the stimuli on identification in Experiments 1 to 8

Introduction

Children's nursery rhymes and play songs are often described as similar in terms of melodic, harmonic and rhythmic features (Dowling, 1988). Experiments 1 to 7 presented either thirteen (Experiments 1-5; 8) or six (Experiments 6 & 7) nursery rhymes as stimuli for identification, across three main experimental conditions (intact, noise-filled gaps and silence-filled gaps). Within the stimuli presented, there were differences in the frequency of correct identification responses across the different nursery rhymes. This was seen as an important part of the design so that floor and ceiling effects in identification could be avoided. It is possible that these differences were due to overall differences in familiarity or level of exposure to the different nursery rhymes. However, it is also possible that differences in the musical characteristics of the stimuli (e.g. features of the rhythm, melody and tonality) could account for differences in identification, particularly in noise-filled gaps conditions where perceptual restoration occurred. The by-material analysis reported here aimed to determine whether differences in identification were most likely due to familiarity differences or whether some songs were more identifiable, and more "restorable", due to their musical characteristics.

A similar analysis was reported by Steinke et al. (2001), where a musical characteristics analysis was carried out to compare differences between identification of song and instrumental melodies presented as stimuli to an adult listener. In discussing the identification differences, they raised the possibility that the musical characteristics of the stimuli could be responsible for the differences in identification observed. This analysis considered factors such as the complexity of rhythm and melodic contour, the number and rate of presentation of notes, and the range of notes presented and the average interval size. Steinke et al. found no consistent evidence that these factors played a role in identification differences of the stimuli presented. Ruling out such a possibility is an important part of the analysis reported in this chapter. If such factors do consistently influence identification scores in the present series of experiments, and particularly influence identification of stimuli where missing sections are replaced with noise, this could have important implications for understanding how perceptual restoration operates in the specific case of music perception.

Method

Musical characteristics analyses were carried out on stimuli presented in each of Experiments 1 to 8, and once calculated, the relationship between the musical characteristics of each stimulus and identification scores in each experimental condition was investigated. The measures taken of the musical characteristics of the stimuli were adapted from an analysis reported by Steinke et al. (2001). The characteristics chosen, and the methods of calculation for each characteristic, are described in Table 11.1. Analysis was carried out on only those stimuli that were presented as experimental trials.

Results and Discussion

For each experiment (1-8), the values of the musical characteristics shown in Table 11.1 were calculated for each nursery rhyme stimulus. The number of correct identification responses for each nursery rhyme was then correlated with values of the musical characteristics using Kendall's Tau correlations, due to small sample sizes and a large number of tied ranks (as suggested by Field, 2005). These analyses were conducted separately by experimental condition. Appendix 16 reports all values of musical characteristics, correct response frequencies, and correlation analyses.

For Experiments 2 (gap duration), 3 (noise amplitude), 4 (initial/later phrases), 5 (songs/melodies) and 8 (age), there were no relationships between any of the musical characteristics and identification scores, in any experimental condition. However, significant relationships were evident in Experiments 1 (basic paradigm), 6 (melody versus lyrics) and 7 (pitch versus rhythm). These significant correlations will now be discussed in turn.

Table 11.1: Description of measures of musical characteristics

Musical Characteristic	Description	Calculation Method
Number of tone onsets	The number of tones, regardless of duration, within the excerpt	Total number of notes in the melody
Range	The distance in semitones between the lowest and highest pitch in the melody (a value of 12 indicates a distance of one octave)	Highest pitch – lowest pitch notes in semitones
Average interval size	The average distance in semitones between two successive notes	The distance between each pair of successive notes in semitones is calculated and the average value taken
Presentation Rate	The average number of notes per second	Number of notes ÷ length of excerpt
Average note duration	The average duration of a note in the melody (milliseconds)	Length of excerpt ÷ number of notes
Contour complexity	A measure of the complexity of pitch direction changes in the melody	The ratio of direction changes to the total number of notes in the melody. The lower the value, the simpler the contour of the melody
Number of different note durations	The number of different note durations in the melody (e.g. crotchet, quaver, etc.)	The number of different note durations appearing in the melody is counted.
Rhythmic complexity	The percentage of notes accounted for by the most common note durations	The two or three most common note durations are identified and the percentage of all notes belonging to these categories is calculated
Tonal strength	A measure of how well the melody conforms to the key in which it is written.	Calculated using the key-finding algorithm reported by Krumhansl (1990). The melody is analysed according to the number of beats that represent each note of the chromatic scale. This is then correlated with the "probe tone rating" for each note which is a measure of how well a particular note represents the key. The resulting correlation is a measure of tonal strength; a higher correlation indicates that the melody is a good representation of the key in which it is written[a]
Tonal strength (2)	A measure of how well the melody conforms to the key in which it is written	Each note in the melody (regardless of duration) is given the corresponding probe tone value. The average probe tone value is then calculated, with a higher value indicating increased tonal strength.

[a] It is questionable whether this method applies to very simple melodies like nursery rhyme tunes. Since there is value for each note in the chromatic scale in terms of probe tone ratings, a good correlation requires that there is also a good distribution of tones across all the notes of the key within the melody. However, the notes of a simple melody could be composed of only 2 or 3 different pitches but still conform well to the key in which it is written. This would result in a poor correlation in terms of this measure. The tonal strength (2) measure is an alternative way of measuring tonal strength that avoids this problem, and might be more appropriate for the simple melodies in the present series of experiments.

Experiment 1

In Experiment 1, children were presented with the nursery rhyme melodies either intact, or with alternating 100 ms fragments replaced with noise or silence. Analysis of the relationship between the musical characteristics of the stimuli and identification of the stimuli revealed significant correlations between identification scores and the number of different note durations, $\tau = -.59, p < .05$, and between identification scores and rhythmic complexity, $\tau = .61, p < .01$, in the silence-filled gaps condition. Melodies with fewer notes of different durations were better identified, and melodies with a high number of notes accounted for by the most common durations (i.e. the lower the rhythmic complexity), were better identified.

It is of course possible that musical characteristics and overall familiarity interact, as by coincidence it could be possible that the most familiar stimuli happen to be the most rhythmically simple ones. However, if this were the case, then significant correlations would be expected across all experimental conditions, not just the silence-filled gaps condition. It is therefore likely that children in the silence-filled gaps condition of this experiment found it easier to identify stimuli with simple rhythms.

Experiment 6

In Experiment 6, new melodies were written to accompany the lyrics of the original nursery rhymes, and new lyrics were written to accompany the melodies of the original nursery rhymes. In addition, mismatch stimuli were created, where a familiar melody and familiar lyrics were mismatched. Analysis of the relationship between identification scores and musical characteristics was carried out separately for each stimulus type. This is because old melody and mismatch stimuli contained the original melody of the nursery rhyme, as presented in Experiment 1. However, new melodies were composed for old lyrics stimuli; hence, the musical characteristics of these stimuli differ to those of the original melody (although not significantly; see Appendix 8).

No relationship between the musical characteristics of the stimuli and identification scores were observed in any experimental condition for old lyrics or mismatch stimuli. However, for old melody stimuli, where the melody was the same as that presented in Experiment 1, in the noise-filled gaps condition there was a relationship between stimulus identification and the tonal strength of the stimuli, $\tau = .93$,

$p < .05$. This would indicate that when a familiar melody is paired with new lyrics, the better the melody represents the key in which it is written, the better it can be identified when missing fragments are replaced by noise. Whilst this might implicate tonal expectations in this process, this particular measure of tonal strength, as outlined earlier, may not be appropriate for analysing simple melodies such as those of nursery rhymes. Therefore, conclusions based on this relationship need to be made with this in mind.

Experiment 7

For Experiment 7, new pitch contours were written to accompany the rhythms from the original melodies (old rhythm stimuli), new rhythms were written to accompany the pitch contours from the original melodies (old pitch stimuli), and original pitch contours and rhythms were mismatched to create mismatch stimuli. Because no stimulus here maintained the original pairing of pitch and rhythm that determined the musical characteristics of the original melodies, new values of musical characteristics were calculated for stimuli of each type.

For old pitch stimuli, a significant relationship was found between the second (more appropriate) measure of tonal strength and identification of old pitch stimuli in the silence-filled gaps condition, $\tau = .86, p < .05$. This suggests that where the pitch contour conforms well to the key in which it is written, the pitch is easier to identify.

For old rhythm stimuli, the analysis revealed a significant relationship between identification of the stimuli and the rhythmic complexity of the stimuli; this relationship was again significant for the silence-filled gaps condition only, $\tau = -.91, p < .05$.

Finally, for mismatch stimuli, a significant correlation emerged between identification and tonal strength in the intact condition, $\tau = .83, p < .05$. If such a relationship emerged in the intact condition, it is difficult to see why this did not also influence identification in the other conditions.

Harmony analysis

As a further test of the effect of the musical characteristics of the stimuli on identification, two further measures were taken from each stimulus. The first was the scale degree of the note the melody begins on: I (tonic) III (mediant) or V (dominant). The second was the type of cadence that ended the excerpt. A cadence is determined by harmonic progressions and is a form of musical punctuation. A perfect cadence (called a "full close") gives the impression that the phrase is complete (analogous to a full stop in punctuation), and an imperfect cadence (called a "half close"; analogous to a comma in punctuation) gives the impression that the phrase has yet to complete (Kennedy, 1996). Each melody was categorised according to the scale degree of the note it started on (only the first, third and fifth scale degrees were represented in the melodies presented), and the type of cadence the excerpt ended on (either perfect or imperfect).

The difference in identification by both of these factors was analysed using Kruskall-Wallis tests (scale degree) and Mann-Whitney tests (cadence). This was conducted separately for the beginning phrases (Experiment 1), the later phrases (Experiment 4) and the new melodies used in Experiments 6 and 7. The results of the analyses are shown in Appendix 16.

This analysis demonstrated that overall, the scale degree of the start note and the end cadence of the excerpt have little effect on identification of stimuli. However, there was a significant difference between identification of stimuli ending in a perfect and in an imperfect cadence in the silence-filled gaps condition of Experiment 1, $U = 6.50, p < .05$. Identification of stimuli ending with a perfect cadence ($M = 6.86$) was superior to that of stimuli ending with an imperfect cadence ($M = 3.33$).

This analysis has only revealed a few conditions where the musical characteristics of stimuli are related to identification of those stimuli. The majority of these relationships occur in silence-filled gaps conditions. This suggests that musical characteristics might have their greatest influence where other factors that might assist identification are not present. Overall, lack of consistency in the way that the musical characteristics affect identification indicates that for the most part, we can conclude that differences in identification between different stimuli are more likely to be due to familiarity differences than factors to do with the melody itself. Even where rhythmic complexity factors influence identification, this is not for the noise-filled gaps conditions where perceptual restoration occurs so it seems unlikely that melodic factors affect whether missing fragments of the melody can be restored when replaced by noise. It is also important to consider the possibility that the small number of significant correlations that emerged did so by chance. There is little evidence to support the possibility that the results seen across experiments are consistently and reliably accounted for by the musical characteristics of the stimuli, at least in the case of the highly familiar, simple melodies presented here.

Chapter 12: General Discussion

Overview

This final chapter summarises the results obtained in the series of experiments reported here and considers how these findings relate to the issues raised in the Introduction. Specifically, this chapter will discuss how the findings from each experiment contribute to our understanding of the perceptual restoration of musical input, and the perceptual restoration process in children. The factors that influence the operation of perceptual restoration in children's music perception will be discussed in the context of the possible auditory perceptual processes involved in perceptual restoration. The implications that perceptual restoration has for theories of auditory perception will be considered, before the chapter concludes by outlining possible directions for further research in this area.

Summary of Results

Experiment 1 formed the basis of the empirical work by first establishing the methods to be used for the investigation of musical perceptual restoration in young children. This was achieved by presenting children with a melody identification task, and comparing children's ability to identify familiar nursery rhyme melodies intact, with missing fragments replaced by noise, and with missing fragments replaced by silence. Evidence of perceptual restoration was taken to be a significant identification advantage for melodies with missing fragments replaced by noise when compared to the same melodies with missing fragments replaced by silence.

The results were clear in showing evidence of perceptual restoration; children's identification of melodies with noise-filled gaps was significantly superior to that of melodies with silence-filled gaps, even though the same amount of intact melody was presented in both cases. This indicates that children use previous knowledge of likely signal content in the process of perception, but importantly this knowledge operates through an interaction with bottom-up corroboration of the expected content of missing information. However, there was also a significant difference between identification of melodies with noise-filled gaps and intact melodies. This indicates that children's music perception is detrimentally affected by signal disruptions, even though replacing missing sections with noise measurably improves melody identification relative to its absence. Further, there were no differences in the time taken to identify the melodies across the three conditions. Thus, Experiment 1 formed the basis for subsequent empirical work by providing evidence of a perceptual restoration mechanism operating in children's music perception by the age of four years.

In the context of children's music perception, the operation of acoustic constraints on perceptual restoration is potentially important for two reasons. First, investigating constraints on the operation of perceptual restoration can identify conditions under which the mechanism cannot operate, or is reduced in the strength of its operation. Second, the effect of acoustic factors on the perceptual restoration of music has not previously been investigated systematically; hence, the results of such manipulations can be informative in comparing the operation of perceptual restoration in speech and music domains, and in adult and child listeners.

The aim of Experiment 2 was to test for the operation of acoustic constraints on the perceptual restoration of music in children, and to this end involved a manipulation of the duration of the missing fragments in the melodies. When combined with the data from Experiment 1, a comparison was made between identification of melodies with missing fragments of 100, 200 and 300 ms in duration, with these fragments either filled with noise or silence. Across noise-filled gaps and silence-filled gaps conditions, identification improved with decreasing gap duration, suggesting that it is easier to identify melodies when disruptions are brief, despite the fact that half of the signal was removed and half remained intact regardless of the gap duration. Furthermore, a significant noise-filled gaps over silence-filled gaps advantage was evident for all gap durations. Noise was able to improve identification most measurably over silence where identification with silent gaps was especially poor, when the longer gaps were removed from the melodies (300 ms).

The findings were interpreted with reference to the differential effect of increasing gap duration on disruption to the contour of the melody, where it is possible that longer interruptions disrupt the melodic contour to a greater degree than shorter gap durations. However, the duration of the missing fragments had no effect on the time taken to identify the melodies. Taken together, the findings from Experiment 2 built directly on those from Experiment 1 by demonstrating the influence of acoustic factors on the perceptual restoration of music in children.

Experiment 3 was a further test of the influence of acoustic factors on the perceptual restoration of music in children, and involved a manipulation of the amplitude of the noise used to replace missing fragments of the melodies. This experiment was a direct test of theories relating to the perceptual restoration of speech in adults whereby only noise that is higher in amplitude than the signal it is replacing can act as a potential masker of the missing information, and thus lead to perceptual restoration (Kashino, 2006). Replacing missing sections of the melodies with noise only improved intelligibility over melodies with silence-filled gaps when the amplitude of the replacement noise exceeded the peak amplitude of the signal, such that the noise could serve as a potential masker of the missing information. Noise, if lower in amplitude than the signal, provided no intelligibility advantage over silence as a gap replacement. In contrast, there was no such effect of replacement noise amplitude on the time taken to identify the melodies. These results were described as evidence of the adaptive nature of perceptual restoration, where perceptual restoration does not occur in situations where missing input could not be masked by the extraneous sound.

When children were asked whether they thought that little pieces of the tune were missing, or all of the tune was there, there was a striking consistency in the way that children in the different conditions responded. When missing sections were replaced with noise that was higher in amplitude than the signal, and hence was a potential masker of the missing information, the majority of children reported that all of the tune was there. This is suggestive of the perception of continuity of what in reality were fragmented signals. In contrast, when replacement noise was not a potential masker of the missing information and was lower in amplitude than the peak amplitude of the signal, most children reported that they thought little pieces of the tune were missing, as they were in reality.

Of particular interest was the way that children's reports of continuity corroborated the data from the identification task. Only in the conditions where identification of melodies with noise-filled gaps had a significant advantage over identification of melodies with silence-filled gaps did the majority of children report the signals to be continuous. Further evidence that the perception of continuity and the restoration of intelligibility are related came from findings that regardless of experimental condition, those children that reported the signals to be continuous scored significantly higher in terms of identification scores than children who reported the signals to be fragmented. Taken together, these findings suggest that the masking potential rule (Kashino, 2006) applies to the perceptual restoration of music in children. The findings also support propositions that perceptual restoration involves two processes, the restoration of continuity and the restoration of intelligibility (e.g. Bashford et al., 1992). Following directly from Experiment 2, the findings from Experiment 3 provided further parallels with the operation of perceptual restoration in the speech domain, and suggested that even in children acoustic factors influence the operation of perceptual restoration.

Experiment 4 represented a move from testing acoustic influences to testing familiarity influences on the perceptual restoration of music in children. The specific issue addressed was whether the operation of musical perceptual restoration in children is restricted to typical, highly familiar representations, or whether it also operates for less familiar and less typical sections of a musical representation. Children's ability to identify melodies intact, with missing sections replaced by noise and with the same sections replaced by silence was compared for the highly familiar opening phrases of children's nursery rhymes and for less familiar phrases occurring later in the piece. Perceptual restoration, as shown by a significant noise-filled gaps over silence-filled gaps advantage in identification, was evident for both the beginning and later sections from the nursery rhymes. The size of this advantage, however, was considerably reduced for the less typical later phrases. The most striking finding was that there was only a significant difference between identification of the most typical beginning phrases and less typical later phrases for the noise-filled gaps conditions, which represent the conditions under which perceptual restoration occurs.

This suggests that when top-down influences on perceptual restoration are reduced, such as when a less typical section from a musical representation is presented, perceptual restoration of missing information is not as strong as when top-down influences are stronger, such as when a highly familiar section of a musical representation is presented. This result also supports the interaction between top-down and bottom-up processes in perceptual restoration in children, as only where the additional top-down influences for a familiar section interacted with bottom-up confirmation of increased expectations was there an effect on identification scores.

Identification time differences were evident only for the later melodic phrases, where intact melodies were identified significantly faster than those with noise-filled and silence-filled gaps. This result suggests that when the musical excerpt presented is a less typical one, disrupted stimuli take longer than intact stimuli to identify. This difference was not evident for the more typical beginning phrases.

Data on children's reports of the continuity of the melodies only included the later melodic phrases, and half of the children in the noise-filled gaps condition reported the signals to be continuous. Nevertheless, across all conditions, those children that reported the signals to be continuous scored significantly higher in the identification task than those who reported the signals to be fragmented. The findings from Experiment 4 contributed to the emerging picture of the perceptual restoration of music in children by demonstrating the effect of familiarity on perceptual restoration in children. Thus, these data provided evidence that both low-level acoustic factors and higher-level contextual factors both influence the perceptual restoration of music in children, and that they interact in their operation.

The perceptual restoration of songs was tested in Experiment 5. This represented an important empirical endeavour as the perceptual restoration of vocal music had never previously been tested experimentally with either adult or child participants. In addition, this represented a further way to test for familiarity influences on the perceptual restoration of music in children, as a song contains two familiar components, both melody and lyrics, in comparison to the melodic component alone. Children were presented with excerpts from nursery rhymes which either consisted of the melody and lyrics sung together, or just the melody sung to the syllable /la/. In the same way as in previous experiments, this was compared for intact, noise-filled gaps and silence-filled gaps versions of the stimuli. Whilst overall identification of the songs was superior to that of the melodies, the size of the noise over silence advantage (hence the strength of perceptual restoration) was no greater for songs than melodies. This somewhat unexpected finding led to two possible interpretations of the result. First, it was considered that lyrics are crucial to perceptual restoration and the children were only using the lyrics for identification with little assistance from the melody across all conditions. Alternatively, it was considered that the effect of lyrics is purely an additive one, and that lyrics only provide extra familiarity cues for identification in the context of the expected melody. This represented a key issue to be addressed in subsequent experiments.

Songs were identified significantly faster than melodies, supporting propositions that lyrics help to access the stored representation of a piece of music (Peretz et al., 2004). Furthermore, there was no clear evidence in this experiment that filling missing sections with noise led to the subjective report that the signals were continuous.

Experiment 6 was designed to clarify the unexpected result from Experiment 5, by investigating the role that melody and lyrics individually play in musical perceptual restoration in children, and whether this differs for intact and disrupted stimuli. This was achieved by putting them into direct conflict in the identification task. Familiar melodies were presented with new lyrics, familiar lyrics were presented with new melodies, and familiar melodies and incongruent familiar lyrics were mismatched and presented together. Perceptual restoration, as shown by a significant noise-filled gaps over silence-filled gaps advantage in identification, only occurred when the melodic component was familiar. When the lyrics were familiar but presented out of their usual context, stimuli with noise-filled gaps were no easier to identify than stimuli with silence-filled gaps.

When familiar melodies and lyrics were mismatched, intact stimuli were identified primarily on the basis of the lyrics. However, a clear strategy shift was evident in the case of noise-filled gaps and silence-filled gaps stimuli: the melodic component was used as the primary basis for identification. As for Experiment 5, there was no consistent evidence that stimuli with noise-filled gaps were perceived by children to be continuous. The findings from Experiments 5 and 6 together were clear in demonstrating that the operation of perceptual restoration of music in children is highly context-dependent. They further indicate that when stimuli are disrupted, melody is more useful for identification than lyrics, and that the superior identification of songs over melodies observed in Experiment 5 might represent an additive effect of lyrics on top of a familiar melody.

With particular reference to perceptual restoration, lyrics are only supported by perceptual restoration if heard in the correct melodic content. On the contrary, melody is supported by perceptual restoration when heard in a new context. However, melody itself constitutes the components of pitch and rhythm. Hence, Experiment 7 was designed to determine if there is a difference in the contribution of these two components to musical perceptual restoration in children, or whether they are required to be congruent in order for perceptual restoration to occur.

The method of Experiment 7 was very similar to that of Experiment 6, where familiar pitch patterns were set to new rhythms, and familiar rhythms were set to new pitch patterns. In addition, familiar pitch and rhythm patterns from different nursery rhymes were mismatched. Strikingly, no evidence of perceptual restoration was obtained in this experiment for any stimulus type. This was evidence that pitch and rhythmic context must be congruent in order for expectations to be generated about the content of missing information, and contribute to musical perceptual restoration in young children. Pitch information predominated over rhythmic information in identification of intact stimuli, but pitch alone was not enough to allow perceptual restoration of missing fragments.

When missing fragments were replaced by noise, nearly all children reported that there were little pieces of the tune missing. This is particularly interesting because where no perceptual restoration occurred, neither were the signals reported as being continuous. This indicates that as well as noise being unable to improve intelligibility of the disrupted signals presented in this experiment, it was also unable to restore the continuity of the signals. Instead, it was clear that fragments were missing. There is some evidence throughout the experiments presented here that the two proposed processes in perceptual restoration, the restoration of continuity and intelligibility (e.g. Bashford et al., 1992; Verschuure & Brocaar, 1983), are related.

The findings from Experiment 7 indicate that the advantage of melody in musical perceptual restoration demonstrated in Experiment 6 was the result of the correct combination of pitch and rhythmic information. Taken together, the results of Experiments 6 and 7 raised important issues about what the concept of familiarity in the context of perceptual restoration is taken to represent, which will be discussed in greater detail later in this chapter.

The final stage of the empirical work, Experiment 8, addressed the developmental aspect of the perceptual restoration of music. The basic identification task presented in Experiment 1 was administered to a group of older children (9-11 years of age) and a group of adults. These data were compared to those from 4- to 6-year old children presented in Experiment 1. The results demonstrated clear developmental effects. Whilst perceptual restoration operated in all three age groups, the adult group showed a much larger noise-filled gaps over silence-filled gaps identification advantage than the two child groups. In adults and older children there was no difference between identification of stimuli with noise-filled gaps and intact stimuli, yet in young children there was significant difference between intact and noise-filled gaps stimuli, thus demonstrating greater effects of signal disruptions in general.

Furthermore, older children and adults did not differ in their identification of stimuli with silence-filled gaps, but young children seem to have particular difficulty identifying melodies on the basis of partial information when noise does not mask the missing sections. Identification scores were corroborated by identification time data; young children were significantly slower at identifying the melodies than the other groups. Finally, within the older children and adult groups, musical training did not significantly affect melody identification, and noise appeared to restore continuity as well as intelligibility of the fragmented signals.

The findings from Experiment 8 provided the final, and crucial, interpretation to the results obtained in the preceding experiments: that children's musical perceptual restoration as demonstrated in Experiments 1 to 6 is not in any respect adult-like, and is still to undergo considerable developmental change. It is likely that some of the perceptual processes underlying the perceptual restoration of music might form the basis of some of this change.

Over all experiments where it formed part of the design, the subjective measure of the continuity of the stimuli provided some interesting results, and as such merits further discussion. This measure was introduced in an attempt to measure children's perceptual experience of the stimuli presented. Because this measure did not form part of the design until Experiment 3, the only continuity reports from young children for the basic piano melodies presented in early experiments came from presentation of the melodies with variations in the amplitude of the replacement noise. In this experiment, where the amplitude of the replacement noise exceeded that of the signal, the majority of children reported the signal to be continuous. Whilst it is important to be aware of the fact that standard amplitude noise, not high amplitude noise, was used in subsequent experiments, the proportion of "continuous" responses in the noise-filled gaps conditions of Experiment 4 (later, less typical phrases) and Experiment 5 (sung lyrics and melodies) was much lower than that reported in Experiment 3. The lack of a clear majority of "continuous" responses in the noise-filled gaps condition of Experiment 4 is consistent with the weaker perceptual restoration observed for this manipulation; weaker expectations seem to affect both the continuity and intelligibility components of the effect.

However, in Experiment 5 strong perceptual restoration was demonstrated for both songs and melodies, hence a majority of "continuous" responses would be expected here in noise-filled gaps conditions. In fact, only half of the responses for the song stimuli (melody and lyrics) were that of a continuous signal, with only a small majority for "continuous" responses for the sung melodies (melody only). This result could possibly be due to that fact that the song stimuli were very easy to identify, and therefore more processing resources were available to notice that there were in fact disruptions to the signals. This possibility is consistent with the processing load explanation given by Samuel (1981a) to explain weaker perceptual restoration for highly familiar stimuli. In addition, this result could possibly reflect differences in signal complexity between piano and voice.

When melody and lyrics components, and pitch and rhythm components, were manipulated in Experiments 6 and 7, respectively, there was no clear majority for "continuous" responses in the noise-filled gaps conditions. For Experiment 7, the clearly evident majority for "fragmented" responses in the noise-filled gaps condition reflects the lack of perceptual restoration observed in this experiment. However, in Experiment 6, there was evidence of perceptual restoration, at least for the old melody and mismatch conditions. It is possible that these stimuli would have been reported as continuous, but since the continuity judgement was made across all stimuli, including the old lyrics stimuli for which perceptual restoration did not occur, it is not possible to determine whether this might indeed be correct. Finally, in Experiment 8, adults and older children showed a clear majority for "continuous" responses when missing fragments of the melodies were replaced by noise.

Taken together, it is evident that some parallels emerge between the strength of perceptual restoration demonstrated and the reports of continuity given in the noise-filled gaps condition. The consistent finding across experiments of improved identification scores for those participants that reported the signals as continuous further indicates that hearing a signal as continuous (regardless of whether or not this is the actual status of the signal) is related to improved identification of the signal.

As reported in the Introduction to this thesis, Bashford et al. (1992) proposed that the restoration of intelligibility of fragmented signals is dependent on the perception of continuity. The data presented here are not consistent with this representation of perceptual restoration, because in some experiments perceptual restoration of intelligibility occurred without the accompanying perception of a continuous signal. However, this theory does not consider the possibility that the relationship between intelligibility and continuity in the perceptual restoration process might be bi-directional. Indeed, a more recent model outlined by Shahin et al. (2009) represents the perception of a fragmented signal as either continuous or fragmented as the outcome, not the determinant, of a sensory repair process that restores the content of the signal.

The data presented in this thesis are consistent with this more recent model; where contextual or acoustic information was insufficient to contribute to the repair process, or reduced in efficiency, a "fragmented" response was more likely to be made by the majority of listeners when missing fragments were replaced by noise. It is possible that the operation of the repair process does not have to be complete, but can occur to differing degrees. Thus, the repair process can restore intelligibility of a signal to the extent where noise-filled gaps can have an advantage over silence-filled gaps in terms of intelligibility, without repairing the content of the signal to the degree where a "continuous" response is registered.

In summary, the results from this series of experiments build on one another in revealing a musical perceptual restoration mechanism operating in children (Experiment 1) that is influenced in its operation by low-level acoustic factors (Experiments 2 & 3) and higher-level contextual factors (Experiments 4 & 5). The operation of musical perceptual restoration in children is highly context-dependent (Experiments 6 & 7) and improves with development (Experiment 8).

The musical characteristics analysis reported in chapter 11 was carried out in order to determine whether the musical characteristics of the melodies were partially responsible for differences in identification of individual nursery rhymes within the stimulus set. No consistent effects were found, and in the vast majority of cases, no significant relationships were evident between melody identification and the musical characteristics of the melody. However, the majority of relationships that were evident were between musical characteristics and identification of stimuli with silence-filled gaps. This suggests that when no facilitation from replacement noise is able to assist identification of disrupted stimuli, factors such as the musical characteristics of the stimuli might have a greater effect on identification.

It is crucial to remember that as a musical genre, nursery rhyme and play song melodies are extremely similar in terms of the musical characteristics assessed here, having simple melody, rhythm and tonality (Dowling, 1988). Therefore, differences in identification across different stimuli are most likely due to differences in familiarity and exposure.

Factors affecting the perceptual restoration of music in children

The data presented in this thesis support the operation of perceptual restoration in children's music perception. Moreover, the data also demonstrate the influence of acoustic, contextual and developmental factors on the operation of perceptual restoration in this context. The section that follows discusses the influence of these factors and how they might reflect the perceptual processes involved in perceptual restoration, and then considers the theoretical implications of these findings.

Acoustic factors

Experiments 2 and 3 indicated that the perceptual restoration of music in children is affected by the acoustic properties of the signal and of the noise that replaces missing information. Acoustic factors have been found to exert an influence on the perceptual restoration of speech; gap duration, noise amplitude and the similarity between replacement noise and the sound it is replacing all affect the degree to which noise can aid reconstruction of missing phonemes (e.g. Eimas et al., 1996; Powers & Wilcox, 1977; Warren & Obusek, 1971). Even in young children gap duration influences intelligibility of speech with noise-filled gaps (Newman, 2004). The results obtained in the experiments reported here support the hypothesis that the role of noise in perceptual restoration is to provide bottom-up confirmation of contextually-generated expectations of the content of the missing information. Experiments 2 and 3 indicated that bottom-up confirmation is facilitated by a short duration of interruption and by noise that exceeds the amplitude of the signal it is replacing. These results also suggest that it is not just higher-level cognitive influences that operate on perceptual restoration in the music domain, as had been tested in previous research (e.g. DeWitt & Samuel, 1990). Furthermore, the results of Experiments 2 and 3 indicated that perceptual restoration is not an all-or-nothing phenomenon, but that there are acoustic constraints on its operation. That acoustic factors have an influence on perceptual restoration in children suggests that there are some important parallels between the way that the mechanism operates in adults and children.

In the speech domain, the effect of gap duration on perceptual restoration, where perceptual restoration is less likely to occur for missing fragments longer in duration than the average length of a word, is described as serving an adaptive function. This is because gaps in a speech stream often signal important events such as the end of a sentence of phrase. For this reason, whilst brief interruptions are likely to occur as a result of a disruption, hence restoring this missing content is adaptive, longer interruptions are less likely to be caused by some disruption and more likely to be an important aspect of the rhythm of spoken speech. Therefore filling in these sections would not be adaptive. In music, periods of silence can also be meaningful and therefore the effect of gap duration on the perceptual restoration of music might serve a similar purpose in preventing inappropriate restoration of gaps that are supposed to be present.

Whether children would have assimilated sufficient knowledge to know that gaps of silence in music can be an important aspect of musical structure, and hence be influenced by such a constraint, needs to be considered in the context of this interpretation of gap duration effects on the perceptual restoration of music in children. Whilst children have had sufficient exposure to the music of their culture to learn about the regularities of music by a young age (e.g. Schellenberg et al., 2005) the operation of a gap duration constraint on the perceptual restoration of music in children might instead, or in addition, represent the process of corroboration of expectations in the perceptual restoration process. If one takes as the explanation of perceptual restoration that restoration occurs provided there is no evidence that the signal is in fact discontinuous, it is plausible that it is harder for noise to hide the evidence that the signal is discontinuous, the longer the length of the missing fragment. Furthermore, with claims by Bregman (1990) suggesting that if noise is able to hide the absence of missing information the signal is projected into the noise, there could be a durational limit to how far into the noise the signal can be projected.

A further possibility based on Bregman's criteria for perceptual restoration to occur is that the sections of the signal before and after the missing fragment are more likely to be treated as parts of the same auditory object through noise if the duration of the interruption is brief. It is also plausible that continued neural activation through an interruption is more likely through a short as opposed to a longer interruption. Thus, this discussion indicates that many of the auditory perceptual processes involved in perceptual restoration are able to explain how and why a gap duration constraint operates on the perceptual restoration of music. With particular reference to children, this increases the likelihood that some, or all, of these processes are involved in perceptual restoration in children.

Experiment 3 demonstrated the operation of a further acoustic constraint on the perceptual restoration of music in children. When replacement noise was not higher in amplitude than the signal it was replacing, perceptual restoration did not occur and noise-filled gaps were not able to provide any improvement to the intelligibility of the melodies over melodies with silent gaps. In the speech domain, this constraint on the operation of the perceptual restoration mechanism is described as serving an adaptive function since restoring missing speech (thus inferring it continues behind an extraneous sound) if it could not have been masked by the extraneous sound would be inappropriate. This constraint would also apply to the perceptual restoration of music, where if musical sounds could not have been masked by an extraneous sound, inferring the continuation of the music through the interruption could be erroneous.

The operation of a noise amplitude constraint on the perceptual restoration of music in children also demonstrates the interaction between processes of expectation and corroboration in the operation of perceptual restoration. Regardless of the amplitude of the replacement noise, the same amount of expectation-generating intact portions of the melody were present. However, perceptual restoration only occurred where noise was a potential masker of the missing sections and provided bottom-up corroboration of contextually-generated expectations.

It was not necessarily expected that children would respond in a similar way to acoustic factors as do adults, nor that acoustic factors would influence musical perceptual restoration in the same way as they have been shown to influence phonemic restoration. That acoustic constraints influence both perceptual restoration in children and perceptual restoration in the music domain is an important finding, and might represent the adaptive nature of perceptual restoration in these contexts, in preventing restoration of gaps in auditory signals in conditions where this is not the most appropriate interpretation of the given auditory input.

Contextual factors

Whilst contextual factors have been manipulated in different ways, the present experiments corroborate findings in the speech domain by indicating that contextual factors (e.g. top-down influences of familiarity and expectation) influence the perceptual restoration of music, and perceptual restoration in children. Previous work in the speech domain found evidence of the influence of familiarity on perceptual restoration; restoration of missing phonemes was stronger in words with higher lexical frequency than lower lexical frequency, and in words when compared to pseudowords (Samuel, 1981a). This was explained as an effect of the greater expectations about the content of missing fragments generated for more familiar words, and as demonstrating that a top-down flow of information activates stored representations even when acoustic information is incomplete. Contextual factors also influence the perceptual restoration of music in adults (DeWitt & Samuel, 1990).

A similar effect of familiarity was also found in Experiment 4 for the perceptual restoration of music; missing fragments from more typical musical excerpts were restored to a stronger degree than missing fragments from less typical musical excerpts (these two types of stimulus can be seen as comparable to high and low lexical frequency words, respectively). This, together with the findings from the speech domain, supports the role of contextually-generated expectations in perceptual restoration, where these expectations operate in conjunction with the bottom-up confirmation of these expectations that the noise provides. Previous knowledge is clearly important in the perceptual restoration process. In order to fill in missing input, the listener needs to know what should be present in the missing sections, which involves a process of inference based on the content of intact sections.

The results obtained in Experiment 4 served as a clear example of the interaction between expectation and corroboration in the perceptual restoration of music in children. The data are consistent with an influence of expectation on perceptual restoration, where stronger expectations are generated for a more typical and familiar musical excerpt (Schulkind, 2004). Top-down influences based on increased familiarity and expectation were not on their own leading to the difference in strength of perceptual restoration between more typical and less typical melodic excerpts. Increased expectations operating for the more typical excerpts did not improve identification relative to less typical excerpts where the melodies contained silent gaps. Only where increased expectations were operating in conjunction with confirmation of these expectations provided by replacement noise did an effect of more typical versus less typical excerpt emerge on identification. This suggests that higher-level representations can exert an influence on lower-level perceptual analysis in children, and that children's stored representations of familiar music are established enough to assist in the perception of disrupted input.

Whilst the result of Experiment 5 suggested that songs, which contain two familiar components (melody and lyrics) were identified better than melodies alone, there was no accompanying evidence that perceptual restoration was stronger for songs than melodies. This result provided the first indication that the operation of familiarity effects on the perceptual restoration of music in children was far from simple. In the context of musical perceptual restoration, further experiments identified an important caveat in the use of familiarity effects to explain perceptual restoration, and an important difference between music and speech perception that has critical implications for perceptual restoration. Music, unlike speech, is "multiply implicative" (Narmour, 1977) in the respect that expectations operate simultaneously on many levels (e.g. pitch, rhythm, and lyrics in the case of vocal music; plausibly also harmony and tonality although these aspects were not tested in the present series of experiments). For this reason, the concept of familiarity as a unitary influence seems inappropriate in the case of the perceptual restoration of music.

As demonstrated in Experiments 6 and 7, exactly what is familiar within a musical signal determines whether or not perceptual restoration occurs. In the perceptual restoration of vocal music, melody alone is supported by perceptual restoration but lyrics alone are not. Neither pitch nor rhythmic information alone is enough to lead to perceptual restoration. With reference to the proposed hypothesis-generating and evidence-gathering processes involved in perceptual restoration, it is possible that noise can only confirm hypotheses about the likely content of missing information if the intact portions of the signal contain the right information in order to know what to expect. This has critical implications for understanding perceptual restoration, and has further implications for other areas of music perception under experimental study.

Expectations generated for a complete musical representation (e.g. the songs presented in Experiment 5) appear to be more complex than the sum of expectations generated by individual components (e.g. lyrics or melody when heard out of their usual context in Experiment 6), even when stimuli are completely intact. It appears that expectations generated on individual levels interact, such that a correct combination of components, such as rhythm and pitch in Experiment 1, results in superior identification to identification of components in isolation (Experiment 7). Furthermore, in some cases perceptual restoration can only operate where individual components, such as pitch and rhythm, are heard in the correct context. The dynamic, context-dependent nature of perceptual restoration as demonstrated in these experiments serves to illustrate that there are important top-down influences on the perceptual restoration of music in children, and that these influences are a critical part of the perceptual restoration process.

Developmental factors

In the speech domain, research has demonstrated age developments in perceptual restoration (e.g. Ackroff, 1981; Newman, 2004; Walley, 1988), thus supporting propositions that developmental changes occur in perceptual processing strategies (Warren & Warren, 1971). The results from Experiment 8 demonstrate that such age developments also occur in the perceptual restoration of music, strengthening parallels between the operation of the mechanism in both domains. Furthermore, evidence of considerable developments beyond the abilities demonstrated by 4- to 6-year-olds in the present series of experiments indicates that these abilities demonstrated do not represent an adult-like mechanism. It is highly likely that the use of context to infer the likely content of auditory signals, whether intact or disrupted, improves with development (Warren & Warren, 1971). Furthermore, the use of partial information for identification of melodies is also likely to improve, given the large differences in identification of melodies with silent gaps between young children and the older children and adults. In particular, using partial information to generate hypotheses about what is to come next in an auditory signal is an important part of the perceptual restoration process, and thus would represent an important developmental improvement.

Throughout the experiments reported here (with the exception of Experiment 7), young children did show evidence of musical perceptual restoration. However, the finding that this ability is significantly inferior to older children and adults also demonstrates that consideration should be given to noise levels in learning environments and the effect that they could have on children's learning and development. It is crucial to consider that demonstrating age developments in perceptual restoration does not reveal the processes responsible for such developments. This issue will be addressed later in this chapter, in a discussion of directions for further research in this area.

Newman (2004) interpreted the operation of phonemic restoration in 5-year-old children as demonstrating children's ability to combine their previous knowledge with incoming acoustic information in the process of perception. The results of the experiments reported here are consistent with this interpretation of children's perceptual restoration ability. When the findings of the present research are considered together with Newman's results, it appears that children do possess a perceptual mechanism that helps them to understand auditory input in noisy listening conditions, operating in both speech and music perception. Newman also reported that children's speech perception is more detrimentally affected by signal disruptions than adults' perception. Experiment 8 presents a similar finding in the context of children's music perception. Thus, whilst similarities between perceptual restoration in children's speech and music perception have been identified, an important difference has emerged between perceptual restoration in adults and children.

Theoretical implications

Taken together, the results from the present series of experiments demonstrate that there are constraints on the conditions under which both the expectation and corroboration processes in perceptual restoration can operate. The relationship between the two processes is likely to be a complex interaction, which further highlights the important conclusion that the operation of perceptual restoration of music in children is highly context-dependent.

The operation of perceptual restoration in the music domain for young children supports constructive theories of perception, where previous knowledge and expectations play a key role in the perceptual process. Whilst this is consistent with the recent paradigm shift in cognitive neuroscience from passive to constructive approaches to perception (Engle et al., 2001), it is important to acknowledge the particular sensitivity that perceptual restoration appears to demonstrate to the acoustic conditions in determining whether or not the process occurs. Therefore, the main theoretical implication of the research reported here is that both top-down and bottom-up processes are crucial to the operation of perceptual restoration; one process in isolation is not sufficient to explain the operation of musical perceptual restoration in children. Thus, the most appropriate representation of the musical perceptual system in young children is one that combines auditory stimulation with previous knowledge in the process of perception.

In demonstrating the influence of previous knowledge on perceptual processing, perceptual restoration is often taken as evidence against a modular perceptual system that has no access to listeners' knowledge in the process of perceptual analysis (e.g. Prinz, 2005). The data reported in this thesis also appear to be inconsistent with a modular perceptual system. However, discussing the relationship between modularity and perceptual restoration requires consideration of the difference between general background knowledge that the listener has, and stored knowledge within a domain, for example the music domain. The knowledge that contributes to perceptual restoration can be viewed as domain-general, or as knowledge within, and hence available to, the module. For this reason, the very operation of perceptual restoration can be interpreted as evidence against modularity, or as described by Fodor (1983), can be seen to be consistent with a modular perceptual system. The exact nature of the knowledge that contributes to perception is yet to be determined.

Moreover, the case of child listeners adds a further important dimension to the debate, since it has been claimed in the neuroconstructivist approach to development that strict modularity represents the end-state of development, emerging through interaction with both internal and external environments, rather than applying to early development (e.g. Karmiloff-Smith, 1992). Crucially, it cannot be denied that human behaviour in both adults and children is widely open to the influence of experience. Such experience influences a wide variety of cognitive processes (Piaget, 1971), including those involved in perception. Thus, regardless of whether or not the perceptual system is viewed as strictly modular in the Fodorian sense, what the operation of perceptual restoration in the studies presented here demonstrates is the need to include the influence of previous knowledge on perceptual processing, in both adults and children.

The operation of perceptual restoration in the musical domain in young children provides insights into methods of auditory perceptual processing. The operation of perceptual restoration certainly suggests that auditory perceptual processing in young children favours global configural features rather than local cues, and thus demonstrates a similar preference for maintaining coherence of perceptual objects wherever possible that is demonstrated in adult perception. Together with research evidence of early visual perceptual completion abilities (e.g. Craton, 1996), this aspect of perceptual processing appears to be a general property of sensory systems that is not specific to speech perception. Both music and speech are characterised by temporal organisation and as such are believed to be processed globally. This means that by forming temporal wholes, individual notes or phonemes are not required for accurate perception (Warren, 1983). In the present series of experiments, many notes or parts of notes were missing from the melodies but the global structure was not detrimentally affected where the missing notes were replaced by noise. Further, due to their temporal structure, each element of music and speech occurs in a context that allows both "what" and "when" to be predicted based on previous acoustic events. By maintaining global coherence, noise-filled gaps do not disrupt this crucial aspect of music and speech perception.

The results obtained from the present series of experiments support the operation of hypothesis-generating and evidence-gathering processes in perceptual restoration in children. Parts of these processes were manipulated in the experiments presented here, through manipulating the context that contributes to the generation of hypotheses, and through manipulating the acoustic evidence available to the perceptual system. Putting the results from the experiments together, explanations such as the "neural perseveration" account and the "no evidence for discontinuity" account of perceptual restoration could also be involved in the process.

The results from the present experiments support a perceptual restoration process where hypothesis-generating processes try to determine, using the intact information presented, the most likely content of missing information. This hypothesis is then evaluated against the information provided by the evidence-gathering process. This process determines whether anything is missing from the auditory signal, with evidence for neural perseveration and lack of evidence for discontinuity plausibly being important parts of this process. The information provided by the two processes is then combined, to determine whether the hypothesis generated is compatible with the evidence gathered from the signal.

Where noise replaces missing information, and satisfies acoustic requirements of hiding the absence of missing information, the hypothesis generated is corroborated by the evidence in the acoustic signal. If however the missing sections are not successfully hidden by noise, the contextually-generated expectation is not consistent with the available acoustic evidence and hence perceptual restoration cannot operate. The data presented are fully consistent with this account of perceptual restoration: where perceptual restoration was found to operate, there was sufficient contextual evidence to generate hypotheses about the most likely content of the missing information. When such contextual evidence was lacking, or was not sufficient to generate hypotheses about the likely content of missing fragments, perceptual restoration did not occur, or was at best considerably weaker. Similarly, when such hypotheses could be generated, perceptual restoration only occurred where sufficient evidence to confirm these hypotheses was present within the acoustic signal. When, for example, replacement noise was not a potential masker of the missing information (Experiment 3), there was not sufficient evidence to confirm the contextually-generated hypotheses about the content of missing fragments.

It is clear that the findings from the present series of experiments are consistent with the operation of some of the proposed perceptual processes that are involved in perceptual restoration. In the section that follows, an account of the contribution these findings make to understanding perceptual restoration in children, music perception in children, and perceptual restoration in the music domain will be given.

Perceptual restoration has only previously been demonstrated in children in the speech domain. Furthermore, the majority of studies on the perceptual restoration of speech in children have had the primary aim of demonstrating its operation, rather than investigating the conditions under which it can operate. Therefore, the present series of experiments demonstrate that a perceptual restoration mechanism in young children, at least for the music domain, is sensitive to the acoustic properties of the signal, but most importantly, is sensitive to the interaction between expectation and corroboration processes. This interaction has previously been discussed in the context of adult perceptual restoration (e.g. Samuel, 1981a).

The main implication of the findings from the present research for understanding children's music perception is that they demonstrate active music perception in children, where children do not passively process incoming acoustic input in the music domain, but can actively reconstruct missing input based on what is expected to be present. Thus, cognitive processes play a role in music perception in young children. Interactive processes also operate in children's music perception, where children can combine incoming acoustic information with their previous stored knowledge in the process of perception, and make and evaluate hypotheses about the most likely content of missing input. Music perception in young children also seems to operate in an adaptive way. Perceptual restoration only operates where filling in missing sections of a signal represents the most appropriate interpretation of the acoustic input, for example, where noise is a potential masker of the missing input. Early music perception is crucial for children's development in a variety of areas (e.g. Anvari et al., 2002; Bilhartz et al., 2000; Lamb & Gregory, 1993; Schellenberg, 2004, 2005; Schön et al., 2004; Thompson et al., 2004). The kind of active music perception demonstrated here, that can compensate for noisy disruptions, means that children can get the most from the kind of early music experiences that lead to such benefits for development.

It was argued in the Introduction that perceptual restoration in children represents a critical adaptation to noisy everyday environments. In evaluating the importance of perceptual restoration for adults, Kashino (2003) describes how life in noisy environments without perceptual restoration would be "quite inconvenient" (p.18). For adults, who possess many strategies for overcoming disruptions to auditory signals, this description does not represent any understatement of the utility of perceptual restoration. For children, however, the consequences of noisy environments on perception are much more severe. The data from Experiment 8 demonstrate that young children, in comparison to older children and adults, have much greater difficulty using partial acoustic information for melody identification (as shown by inferior identification of melodies with silence-filled gaps). This could be due to reduced experience in young children, a reduced effect of experience on perception, or due to developments in sensory systems still to take place. Without an understanding of what developments underlie age improvements in perceptual restoration, these age differences are potentially difficult to interpret. Nevertheless, reduced use of partial acoustic information in young children means that perceptual restoration can be seen as an even more important mechanism for children than adults. Where children are less adept at using the redundancy in auditory signals, perceptual restoration makes partial information useful for children and results in improvements to the intelligibility of auditory signals relative to its absence.

Several parallels between speech and music perception were outlined in the Introduction to this thesis. On the basis of the empirical findings presented here, it is now possible to draw a further parallel based on the similarities between the operation of perceptual restoration in speech and music domains in children. Despite differences in unit size and signal complexity between speech and music, perceptual restoration in the two domains responds in a similar way to comparable experimental manipulations. This suggests that perceptual restoration is an example of a shared processing resource that operates in different auditory domains from early childhood.

The present findings also further understanding of perceptual restoration in the music domain by building on previous research (e.g. DeWitt & Samuel, 1990) in demonstrating that perceptual restoration in the music domain is sensitive to the acoustic conditions of the stimulus as well as contextual influences, and relatively robust to different experimental manipulations. Perhaps the most important conclusion, however, relates to how the concept of familiarity is interpreted in the context of perceptual restoration, particularly in the music domain where this influence is highly complex in nature.

Finally, the Introduction to this thesis provided an account of how, whilst the majority of empirical effort has focused on the perceptual restoration of speech, the general process of filling-in missing sensory input occurs in humans in many different domains. This includes speech, music, vision, sign language, written language, and also in auditory communication systems of non-human animals. The present findings, in demonstrating the perceptual restoration of music in children, contribute to the conceptualisation of sensory filling-in processes as a general property of sensory processing in humans and non-human animals, and across domains, that represents a fundamental adaptation to environments that rarely present optimal conditions for perception.

Evaluation of methods and directions for further research

The Introduction to this thesis reported how a variety of different approaches have been taken to investigating perceptual restoration, both in terms of the methods employed and the measure taken as evidence of its operation. The identification task employed in the present series of experiments was chosen as, through the use of perceptual restoration in a secondary task, this was believed to best reflect the possible use of a perceptual restoration mechanism in children's everyday music perception. However, the use of an identification task raises the possibility that incorrect responses may have been caused by children not knowing a nursery rhyme rather than being unable to identify it on the basis of the acoustic conditions presented. Nevertheless, reports from the class teachers indicated that all of the nursery rhymes and play songs used as stimuli were sung in school on a regular basis therefore it is unlikely that they were unfamiliar to the children. Despite this, it is possible that children would differ in levels of exposure to the nursery rhymes and play songs, depending on the degree to which they were exposed to them in the home environment. There was no way to control for individual levels of exposure in the present series of experiments but by randomly allocating children to experimental conditions, differences in exposure could be controlled to a certain extent (Calvert & Billingsley, 1998).

In many cases, corroborating evidence for the identification measure was obtained from children's continuity judgements. This subjective measure was an attempt to investigate whether replacing missing sections of musical signals with noise not only improved intelligibility of the melodies but also restored the continuity of what in reality was a fragmented signal. Whilst this is a subjective measure and hence cannot be used alone as evidence of perceptual restoration, the consistency of children's responses in intact and silence-filled gaps conditions (where the "continuous" and "fragmented" response, respectively, represents the true status of the signal) indicates that children understood the question. Furthermore, the close correspondence between the identification data and the continuity judgements in most experiments increases the validity of this measure of perceptual restoration. Where noise had a large identification advantage over silence, in most cases the majority of children in the noise-filled gaps condition reported that they thought all of the tune was there, that is, that no fragments were missing. In contrast, where the strength of perceptual restoration was reduced, or where it did not occur at all, the consistency of children's judgements was reduced and in the latter case, the majority of children reported the signals to be fragmented. From this measure of perceptual restoration, it can be inferred that where children reported signals with noise-filled gaps to be continuous, missing musical sounds were recreated when replaced by noise.

The identification time data presented did not consistently reflect the pattern of identification scores of the participants in the different experimental conditions. The manipulations related to the familiarity of the stimulus (presenting a less typical musical representation and adding lyrics to the stimuli), did have a larger effect on identification time data than other manipulations, but overall the results suggest that whether or not perceptual restoration occurs has little effect on the time taken to identify stimuli. This in itself is an important finding, and it could be possible that children did identify the nursery rhymes earlier than they gave their response, but waited to check this with the remainder of the excerpt for fear of giving an incorrect answer. This possibility is supported by the significantly slower identification times of young children when compared to the adults and older children, who gave much faster identification responses.

Consistent with previous research on perceptual restoration in children (e.g. Newman, 2004), the measure of perceptual restoration in the present series of experiments was taken to be a significant advantage, in terms of identification, for stimuli with noise-filled gaps over stimuli with silence-filled gaps. This measure has provided important insights into the conditions under which perceptual restoration does and does not occur in music perception in young children. In this discussion, consideration can also be given to the results that using a different measure of perceptual restoration, namely comparing identification of stimuli with noise-filled gaps to intact stimuli, would have provided. First, this measure would present a very different picture. Conclusions based on this measure would suggest that perceptual restoration does not operate in the musical domain for children, as in most cases identification of intact stimuli was significantly superior to that of stimuli with noise-filled gaps.

However, this measure would have been highly inappropriate in the present context for the following reasons. First, there is no basis for the assumption that in children, perceptual restoration is complete; that it restores intelligibility of disrupted signals to the status of intact signals. Whilst many children reported signals with missing sections replaced by noise as intact (i.e. all of the tune was there), there are other stimulus factors which could affect the relationship between intact and noise-filled gaps stimuli (e.g. presence of noise, differing amounts of contextual information, etc.).

Second, where the key interest in perceptual restoration is in its ecological value, the operation of perceptual restoration is best demonstrated when replacing missing sections of an auditory signal with noise improves intelligibility relative to its absence. This is clearly what the data from the present experiments demonstrate. Crucially, a comparison between stimuli with noise-filled gaps and those with silence-filled gaps, as conducted in the present series of experiments, involves comparing stimuli containing exactly the same amount of intact contextual information. This is not the case where stimuli with noise-filled gaps are compared to those that are intact.

Finally, it is important to consider the effect of presenting white noise bursts to children. It is unlikely that children would have encountered such sounds before and hence the possible distraction that this might cause must be considered. Thus, making a comparison between noise-filled gaps stimuli and intact stimuli cannot be easily interpreted.

As a final issue related to the methods used in the present series of experiments, it is important to highlight that identification of stimuli with silence-filled gaps was not at floor level. This is likely to be a result of the highly familiar nature of the stimuli, and the fact that when using a method where multiple alternating interruptions are made to a stimulus, parts of individual notes can remain intact. However, comparison of the data from young children with those from older children and adults indicated that the young children did have comparative difficulty using fragmented input for identification when no replacement noise was added. Furthermore, this issue demonstrates the importance of the corroborating continuity report data: in most cases, children presented with stimuli with silence-filled gaps reported that they thought little pieces of the tune were missing. Thus, whilst in some cases stimuli can be identified without the use of perceptual restoration, this does not come with the accompanying perception that nothing is missing.

A key issue to be considered in the context of laboratory studies of children's cognitive development is the extent to which children's responses to experimental situations reflect the way in which they would respond in real life situations (Bryant, 1974). This consideration reflects recent debate regarding cognitive ethology (Kingstone et al., 2008). This approach emphasises the need to consider how a mechanism, such as perceptual restoration, would operate in real world situations and, more importantly, improve adaptation to everyday environments. Unless this level of explanation is addressed, it could be argued that laboratory studies have little value (Kinstone et al., 2008). Applying the experimental results obtained here to understanding children's perception of music in everyday noisy conditions, it appears that perceptual restoration can operate where sufficient contextually-based expectations can be generated and where there is bottom-up confirmation of these expectations.

Whilst it is important to acknowledge that cognition is not independent of the environment in which it usually occurs, experimentally breaking down a perceptual mechanism such as perceptual restoration is sometimes the only way that empirical questions about its operation can be adequately addressed (Bryant, 1974). Once experimental work has elucidated how such a mechanism operates, and the factors that influence its operation, observation of the mechanism in normal everyday conditions can provide further insights over and above those obtained through experimental study. Therefore, the following proposed extension of experimental methods in the context of perceptual restoration represents an endeavour that complements, rather than replaces, the need for thorough experimental work.

Following in the suggestions of a cognitive ethology approach, an important goal for perceptual restoration research in the future is to address the personal level of explanation, that is, how perceptual restoration can support perception in real-world listening conditions, and not just in a more arbitrary laboratory task. The methods adopted in the current series of experiments represent a first step in this direction, by considering how noise can enhance intelligibility of fragmented auditory signals over its absence. The melody/song identification task represented one way that perceptual restoration could provide an adaptation to everyday environments, in view of the fact that identification of auditory input is a task that would operate in such environments, and that would be supported by the operation of a compensatory mechanism such as perceptual restoration.

Furthermore, whilst Kingstone et al. (2008) argue that laboratory stimuli can be artificial and hence alter what is being measured, the stimuli presented in the present experiment were produced to be as ecologically valid as possible. One way in which this was achieved was to reject the method where a single complete note is removed, because in real-world conditions a noisy disruption will not perfectly map onto a complete note in a melody. However, this ecological approach needs to consider perceptual restoration not just as a processing system, but the relationship between perceptual restoration and other cognitive processes, such as attention (Kingstone et al., 2008). A better approach might not be to try to measure the operation of perceptual restoration outside the laboratory, but to consider what perceptual responses to experimental stimuli have in common with similar perceptual processes in the real world (Bryant, 1974).

With these issues in mind, some more explicit questions to be addressed by further research can be identified, and categorised into three main aims, which shall be addressed in turn:

-to find methods of experimentally investigating perceptual restoration that more closely represent everyday listening conditions;

- to address the possibility that there are likely to be considerable individual differences in perceptual restoration ability, and that perceptual restoration for some individuals might not be the automatic, unconscious process it is often described as;
- to identify the perceptual processes that contribute to perceptual restoration, to investigate developmental improvements in perceptual restoration from infancy to adulthood, to determine which processes might be responsible for developmental changes in perceptual restoration ability, and to establish developmental relationships between sensory filling-in processes across domains.

Investigating perceptual restoration in ways that more closely represent everyday listening conditions does not necessarily require the use of non-experimental methods. Instead, consideration needs to be given to the methods used to test perceptual restoration, such that the methods used represent a possible use of perceptual restoration in normal conditions. Furthermore, research can move beyond the use of white noise as a replacement for missing information to the kinds of environmental sounds that would act as maskers in everyday listening conditions (e.g. background chatter, machine noises, etc.). A further possibility involves assessing perceptual restoration through its use in a secondary task, such as the degree to which perceptual restoration can restore the missing content of a set of instructions needed to perform some further task.

The data reported in the present series of experiments, whilst reporting group averages, showed little variation within each condition. However, beyond the tasks presented here, it is highly likely that, for any number of reasons, some children will be able to make better use of perceptual restoration than others will. The educational implications of such a proposition are numerous. For example, difficulty in following classroom instructions in some children could be a result of difficulties in compensating for disrupted auditory input in a noisy environment such as classroom, using a mechanism such as perceptual restoration. If some children do have difficulties in using perceptual restoration, the reasons behind these difficulties can provide important insights into the perceptual processes involved in the effect. Ultimately, however, the aim of such an endeavour should be to find ways of improving children's ability to use compensatory processes such as perceptual restoration to understand auditory input in noisy listening conditions.

Perceptual restoration involves various aspects of perception and cognition. For this reason it is possible that developmental change in perceptual restoration could be a result of developmental change in any of these abilities. Such abilities could be domain-specific, such that developmental change in perceptual restoration in different domains is mediated by developments in different areas of perceptual processing and cognition. Alternatively, general developments in perceptual processing and cognition might be responsible for developments in perceptual restoration across domains. A full understanding of how perceptual restoration develops is currently lacking in the field, and could prove important for a number of reasons.

First, understanding the cause of development is not possible until the course of development is fully understood (Bryant, 1990). Identifying the possible developmental origins of perceptual restoration, and tracing the development of perceptual restoration throughout childhood thus represents an important part of this aim. Visual perceptual completion abilities have been demonstrated in infants (e.g. Craton, 1996), hence it is possible that perceptual filling-in will also occur in the auditory domain in infancy. The origins of perceptual restoration are crucial to understanding the processes underlying the mechanism and the way in which it develops.

Second, identifying the skills or processes important for perceptual restoration that are not fully developed in young children would reveal which aspect of the process they have difficulty with, which in turn means that help could be given to develop the use of these processes. The neurophysiological methods outlined in the introduction used to investigate perceptual restoration in adults might offer a potential way to address this deeper level of analysis. Perceptual restoration in adults and children and in different domains might not operate in the same way or be mediated by the same neural resources.

Finally, isolated research in the auditory domain is limited in scope; future research should address the common nature of perceptual filling-in mechanisms across visual, haptic, and auditory domains. In all sensory domains, the ability to compensate for missing input seems to result from a fundamental tendency to group things together so that they are coherent in the respect of possessing a beginning, a middle, and an end.

Whether this ability follows the same developmental course in all domains represents a crucial aspect of furthering understanding of what appears to be such a basic perceptual mechanism.

Conclusion

This programme of research involved eight related experiments designed to establish the operation of perceptual restoration in the music domain in young children, and some of the factors that influence this operation. Based on the findings, some intriguing answers can be provided to questions posed in the Introduction. The first two aims of this research were to establish whether perceptual restoration in children is restricted in operation to the speech domain, and whether the perceptual restoration of music is restricted in operation to adult listeners. Clearly, the data presented here indicate that perceptual restoration in children operates in the music domain. This is evidence that answers both questions: perceptual restoration in children is not restricted to the speech domain and the perceptual restoration of music does not only occur in adults. The ability to make inferences that compensate for missing sensory input is a basic perceptual ability that can operate in both musical and linguistic contexts, at least from early childhood.

The third aim of this research was to investigate the role of expectation and corroboration in the perceptual restoration of music in children. The results suggest that both contextually-generated expectations and the degree to which the acoustic input corroborates these expectations both influence the perceptual restoration of music in children. It has been further demonstrated that musical perceptual restoration operates where the amount of intact and missing content are equated, through the removal of multiple periodic fragments. Such evidence indicates that the perceptual restoration of music occurs for more than just a single musical note in a melody. The operation of perceptual restoration in children's music perception highlights clear parallels between mechanisms of auditory perception in speech and auditory domains, and in both children and adults.

The fourth aim of the present research was to examine more closely the role of familiarity in the context of perceptual restoration, and the data presented here indicate that the influence of familiarity on the perceptual restoration process is more complex than has been suggested previously, at least in the context of children's music perception.

Finally, the fifth aim of this research was to investigate developmental changes in musical perceptual restoration. Clear developmental changes were evident, which made the perceptual restoration abilities demonstrated by 4- to-6-year-old children in the preceding experiments even more impressive given that there are still improvements to be made in the ability. However, due to difficulties that young children have in using partial information for identification, it has been argued that perceptual restoration, which can assist this process considerably, is a critically important mechanism for young children.

Overall, this programme of research has provided evidence that whilst perceptual restoration does operate in music perception in young children, its operation is highly context-dependent and is influenced by both the nature of the acoustic input and the nature of the expectations that are generated by the intact portions of the signal. Furthermore, perceptual restoration abilities in 4- to 6-year-old children do not represent a fully developed mechanism; a key goal of future research is to determine exactly what processes undergo developmental change that might contribute to perceptual restoration.

Most importantly, the present findings contribute to the representation of a perceptual restoration mechanism for filling in missing sensory input that is not restricted in operation to a particular age group or a particular domain. At an even more general level, perceptual restoration might be a good example of recent shifts in theories of perception and perceptual development. Pomerantz (2003) has argued that global configural features, rather than local cues, can be considered as forms of perceptual primitives. The present results demonstrate just how salient Gestalt-like properties are to making sense of perceptual input.

It was stated in the Introduction that nearly all human behaviour takes place in environments that present redundant, irrelevant, and disrupted input. In the perceptual world of the child ambiguous, disrupted, and incomplete sensory input represents the norm rather than the exception. Thus, compensatory processes such as perceptual restoration represent a critical adaptation to such environments, where a fundamental preference for coherence of perceptual input means that in most cases perception is not detrimentally affected.

References

Ackroff, J. M. (1981). The interrelationship of verbal transformations, phonemic restorations and age. *Dissertation Abstracts International: Section B: The Sciences and Engineering, 42,* 2106.

Aiello, R. (1994). Music and Language: Parallels and Contrasts. In R. Aiello & J. Sloboda (Eds.), *Musical Perceptions* (pp. 40-63). New York: Oxford University Press.

Anderson, J. R. (1983). *The Architecture of Cognition.* Cambridge, MA: Harvard University Press.

Anvari, S. H., Trainor, L. J., Woodside, J., & Levy, B. A. (2002). Relations among musical skills, phonological processing, and early reading ability in preschool children. *Journal of Experimental Child Psychology, 82,* 111-130.

Aronoff, J. (2006). Investigating auditory induction without complete continuity. *Journal of the Acoustical Society of America, 119,* 3333.

Aronoff, J. (2008). The role of similarity in restoring missing notes in music. *Dissertation Abstracts International: Section B: The Sciences and Engineering, 68,* 5757.

Aslin, R. N., & Smith, L. B. (1988). Perceptual Development. *Annual Review of Psychology, 39,* 435-473.

Atterbury, B. W. (1985). Musical differences in learning-disabled and normal-achieving readers, aged seven, eight and nine. *Psychology of Music, 13,* 114-123.

Bartlett, J. C., & Dowling, W. J. (1980). Recognition of Transposed Melodies: A Key-Distance Effect in Developmental Perspective. *Journal of Experimental Psychology: Human Perception and Performance, 6,* 501-515.

Bartlett, J. C., & Snelus, P. (1980). Lifespan Memory for Popular Songs. *American Journal of Psychology, 93,* 551-560.

Bashford, J. A., Jr., Meyers, M. D., Brubaker, B. S., & Warren, R. M. (1988). Illusory continuity of interrupted speech: Speech rate determines durational limits. *Journal of the Acoustical Society of America, 84,* 1635-1638.

Bashford, J. A., Jr., Riener, K. R., & Warren, R. M. (1992). Increasing the intelligibility of speech through multiple phonemic restorations. *Perception and Psychophysics, 51,* 211-217.

Bashford, J. A., Jr., & Warren, R. M. (1979). Perceptual synthesis of deleted phonemes. In J. J. Wolf & D. H. Klatt (Eds.), *Speech Communication Papers* (pp. 423-426). New York: Acoustical Society of America.

Bashford, J. A., Jr., & Warren, R. M. (1987a). Effects of spectral alternation on the intelligibility of words and sentences. *Perception and Psychophysics, 42,* 431-438.

Bashford, J. A., Jr., & Warren, R. M. (1987b). Multiple phonemic restorations follow the rules for auditory induction. *Perception and Psychophysics, 42,* 114-121.

Bashford, J. A., Jr., Warren, R. M., & Brown, C. A. (1996). Use of speech-modulated noise adds strong "bottom-up" cues for phonemic restoration. *Perception and Psychophysics, 58,* 342-350.

Bashford, J. A., Jr., Warren, R. M., & Lenz, P. W. (2005). Enhancing intelligibility of narrowband speech with out-of-band noise: Evidence for lateral suppression at high-normal intensity. *Journal of the Acoustical Society of America, 117,* 365-369.

Başkent, D., Eiler, C., & Edwards, B. (2009). Effects of envelope discontinuities on perceptual restoration of amplitude-compressed speech. *Journal of the Acoustical Society of America, 125*(6), 3995-4005.

Bentley, A. (1966). *Musical Ability in Children and its Measurement.* London: Harrap.

Bergeson, T. R., & Trehub, S. E. (2007). Signature tunes in mothers' speech to infants. *Infant Behavior and Development, 30,* 648-654.

Besson, M., Faita, F., Peretz, I., Bonnel, A. M., & Requin, J. (1998). Singing in the brain: Independence of Lyrics and Tunes. *Psychological Science, 9,* 494-498.

Besson, M., & Schön, D. (2003). Comparison Between Language and Music. In I. Peretz & R. J. Zatorre (Eds.), *The Cognitive Neuroscience of Music* (pp. 269-293). New York: Oxford University Press.

Bharucha, J. J. (1994). Tonality and Expectation. In R. Aiello & J. Sloboda (Eds.), *Musical Perceptions* (pp. 213-239). New York: Oxford University Press.

Bigand, E. (1997). Perceiving Musical Stability: The Effect of Tonal Structure, Rhythm, and Musical Expertise. *Journal of Experimental Psychology: Human Perception and Performance, 23,* 808-822.

Bilhartz, T. D., Bruhn, R. A., & Olson, J. E. (2000). The Effect of Early Music Training on Child Cognitive Development. *Journal of Applied Developmental Psychology, 20,* 615-636.

Boltz, M. G. (1993). The generation of temporal and melodic expectancies during musical listening. *Perception and Psychophysics, 53,* 585-600.

Bonnel, A. M., Faita, F., Peretz, I., & Besson, M. (2001). Divided attention between lyrics and tunes of operatic songs: Evidence for independent processing. *Perception and Psychophysics, 63,* 1201-1213.

Bosch, L., & Sebastian-Galles, N. (1997). Native-language recognition abilities in 4-month-old infants from monolingual and bilingual environments. *Cognition, 65,* 33-69.

Bowers, J. S., & Davis, C. J. (2004). Is speech perception modular or interactive? *Trends in Cognitive Sciences, 8,* 3-5.

Braaten, R. F., & Leary, J. C. (1999). Temporal Induction of Missing Birdsong Segments in European Starlings. *Psychological Science, 10,* 162-166.

Bregman, A. S. (1990). *Auditory Scene Analysis: The Perceptual Organization of Sound.* Cambridge, MA: MIT.

Bregman, A. S., Calantonio, C., & Ahad, P. A. (1999). Is a common grouping mechanism involved in the phenomena of illusory continuity and stream segregation? *Perception and Psychophysics, 61,* 195-205.

Bregman, A. S., & Dannenbring, G. L. (1977). Auditory Continuity and Amplitude Edges. *Canadian Journal of Psychology, 31,* 151-159.

Bremner, J. G., Johnson, S. P., Slater, A., Mason, U., Cheshire, A., & Spring, J. (2007). Conditions for young infants' failure to perceive trajectory continuity. *Developmental Science, 10,* 613-624.

Broadbent, D. E. (1958). *Perception and Communication.* Elmsford, NY: Pergamon Press.

Brown, S. (2001). The "Musilanguage" Model of Music Evolution. In N. L. Wallin, B. Merker, & S. Brown (Eds.), *The Origins of Music* (pp. 271-300). Cambridge, MA: MIT.

Brown, S., Merker, B., & Wallin, N. L. (2001). An Introduction to Evolutionary Musicology. In N. L. Wallin, B. Merker, & S. Brown (Eds.), *The Origins of Music* (pp. 3-24). Cambridge, MA: MIT.

Brumm, H. (2006). Animal Communication: City Birds Have Changed Their Tune. *Current Biology, 16,* R1003-R1004.

Bryant, P. (1974). *Perception and Understanding in Young Children: An Experimental Approach.* London: Methuen.

Bryant, P. (1990). Empirical evidence for causes in development. In G. Butterworth & P. Bryant (Eds.), *Causes of Development: Interdisciplinary Perspectives* (pp. 33-45). London: Harvester Wheatsheaf.

Burnham, D., & Dodd, B. (1999). Familiarity and novelty preferences in infants' auditory-visual speech perception: problems, factors and a solution. In C. Rovee-Collier, L. P. Lipsitt & H. Hayne (Eds.), *Advances in infancy research, Vol. 12* (pp. 170-187). Stamford, CT: Ablex.

Butzlaff, R. (2000). Can Music Be Used to Teach Reading? *Journal of Aesthetic Education, 34,* 167-178.

Calvert, S. L., & Billingsley, R. L. (1998). Young Children's Recitation and Comprehension of Information Presented by Songs. *Journal of Applied Developmental Psychology, 19*(1), 97-108.

Cherry, C., & Wiley, R. (1967). Speech Communication in Very Noisy Environments. *Nature, 214,* 1164.

Clarke, E. F. (1999). Rhythm and Timing in Music. In D. Deutsch (Ed.), *The Psychology of Music* (2nd ed., pp. 473-500). San Diego, CA: Academic Press.

Coltheart, M. (1999). Modularity and Cognition. *Trends in Cognitive Sciences, 3,* 115-120.

Craton, L. G. (1996). The Development of Perceptual Completion: Infants' Perception of Stationary, Partially Occluded Objects. *Child Development, 67,* 890-904.

Cross, I. (2003). Music, Cognition, Culture and Evolution. In I. Peretz & R. J. Zatorre (Eds.), *The Cognitive Neuroscience of Music* (pp. 42-56). New York: Oxford University Press.

Crowder, R. G., Serafine, M. L., & Repp, B. H. (1990). Physical interaction and association by contiguity in memory for the words and melodies of songs. *Memory and Cognition, 18,* 469-476.

Dalla Bella, S., Peretz, I., & Aronoff, N. (2003). Time course of melody recognition: A gating paradigm study. *Perception and Psychophysics, 65,* 1019-1028.

Dannenbring, G. L. (1976). Perceived auditory continuity with alternately rising and falling frequency transitions. *Canadian Journal of Psychology, 30,* 99-114.

Davidson, L. (1994). Songsinging by Young and Old: A Developmental Approach to Music. In R. Aiello (Ed.), *Musical Perceptions* (pp. 99-130). New York: Oxford University Press.

Davis, M. H., & Johnsrude, I. S. (2007). Hearing speech sounds: Top-down influences on the interface between audition and speech perception. *Hearing Research, 229,* 132-147.

Demany, L. (1982). Auditory stream segregation in infancy. *Infant Behavior and Development, 5,* 261-276.

Department for Children, Schools and Families (2008). *Practice Guidance for the Early Years Foundation Stage.* Nottingham, UK: DCSF Publications.

DeWitt, L. A., & Samuel, A. G. (1990). The Role of Knowledge-Based Expectations in Music Perception: Evidence From Musical Restoration. *Journal of Experimental Psychology: General, 119,* 123-144.

Donaldson, M. (1978). *Children's Minds.* London: Fontana Press.

Dowling, W. J. (1978). Scale and Contour: Two Components of a Theory of Memory for Melodies. *Psychological Review, 85,* 341-354.

Dowling, W. J. (1988). Tonal structure and children's early learning of music. In J. A. Sloboda (Ed.), *Generative processes in music: The psychology of performance, improvisation and composition* (pp. 113-128). Oxford: Clarendon.

Dowling, W. J. (1999). The Development of Music Perception and Cognition. In D. Deutsch (Ed.), *The Psychology of Music* (2nd ed., pp. 603-625). San Diego, CA: Academic Press.

Drake, C., & McAdams, S. (1999). The auditory continuity phenomenon: Role of temporal sequence structure. *Journal of the Acoustical Society of America, 106,* 3529-3538.

Eimas, P. D., Tajchman, G., Nygaard, L. C., & Marcus, D. J. (1996). Phonemic restoration and integration during dichotic listening. *Journal of the Acoustical Society of America, 99,* 1141-1147.

Elfner, L., & Caskey, W. E. (1965). Continuity Effects with Alternately Sounded Noise and Tone Signals as a Function of Manner of Presentation. *Journal of the Acoustical Society of America, 38,* 543-547.

Elfner, L., & Homick, J.L. (1966). Some Factors Affecting the Perception of Continuity in Alternately Sounded Tone and Noise Signals. *Journal of the Acoustical Society of America, 40,* 27-31.

Elliott, L. L., Connors, S., Kille, E., Levin, S., Ball, K., & Katz, D. (1979). Children's understanding of monosyllabic nouns in quiet and noise. *Journal of the Acoustical Society of America, 66,* 12-21.

Ellis, D. (1999). Using knowledge to organize sound: The prediction-driven approach to computational auditory scene analysis, and its application to speech/nonspeech mixtures. *Speech Communication, 27,* 281-298.

Elman, J. L., & McClelland, J. L. (1984). Speech Perception as a Cognitive Process: The Interactive Activation Model. In N. Lass (Ed.), *Speech and Language, Vol. 10* (pp. 337-374). New York: Academic Press.

Elman, J. L., & McClelland, J. L. (1988). Cognitive Penetration of the Mechanisms of Perception: Compensation for Coarticulation of Lexically Restored Phonemes. *Journal of Memory and Language, 27,* 143-165.

Engle, A. K., Fries, P., & Singer, W. (2001). Dynamic predictions: Oscillations and synchrony in top-down processing. *Nature Reviews Neuroscience, 2,* 704-716.

Feierabend, J. M., Clark Saunders, T., Holahan, J. M., & Getnick, P. E. (1998). Song Recognition among Preschool-Age Children: An Investigation of Words and Music. *Journal of Research in Music Education, 46,* 351-359.

Fernald, A. (1989). Intonation and Communicative Intent in Mothers' Speech to Infants: Is the Melody the Message? *Child Development, 60,* 1497-1510.

Fernald, A. (1991). Prosody in speech to children: Prelinguistic and linguistic functions. *Annals of Child Development, 8,* 43-80.

Field, A. (2005). *Discovering Statistics Using SPSS* (2nd ed.). London: Sage.

Fitch, W. T. (2006). The biology and evolution of music: A comparative perspective. *Cognition, 100,* 173-215.

Fodor, J. A. (1983). *The modularity of mind.* Cambridge, MA: MIT Press.

Foxton, J. M., Nandy, R. K., & Griffiths, T. D. (2006). Rhythm deficits in 'tone deafness'. *Brain and Cognition, 62,* 24-29.

Fraisse, P. (1982). Rhythm and Tempo. In D. Deutsch (Ed.), *The Psychology of Music* (pp. 149-180). New York: Academic Press.

Frith, C., & Dolan, R. J. (1997). Brain mechanisms associated with top-down processes in perception. *Philosophical Transactions of the Royal Society London: Biological Sciences, 352,* 1221-1230.

Gardner, H. (1971). Children's Duplication of Rhythmic Patterns. *Journal of Research in Music Education, 19*(3), 355-360.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston: Houghton Mifflin.

Gregory, R. L. (1970). *The intelligent eye.* New York: McGraw-Hill.

Grosjean, F. (1980). Spoken word recognition processes and the gating paradigm. *Perception and Psychophysics, 28,* 267-283.

Grossberg, S. (2003). Resonant neural dynamics of speech perception. *Journal of Phonetics, 31,* 423-445.

Hannemann, R., Obleser, J., & Eulitz, C. (2007). Top-down knowledge supports the retrieval of lexical information from degraded speech. *Brain Research, 1153,* 134-143.

Hannon, E. E., & Trainor, L. J. (2007). Music acquisition: effects of enculturation and formal training on development. *Trends in Cognitive Sciences, 11,* 466-472.

Hansen, D., & Bernstorf, E. (2002). Linking Music Learning To Reading Instruction. *Music Educators Journal, 88*(5), 17-21.

Hargreaves, D. J. (1986). *The Developmental Psychology of Music.* Great Britain: Cambridge University Press.

Hawkins, S. (in press). Phonological features, auditory objects, and illusions. *Journal of Phonetics.*

Hébert, S., & Peretz, I. (1997). Recognition of music in long-term memory: Are melodic and temporal patterns equal partners? *Memory and Cognition, 25,* 518-533.

Heinrich, A., Carlyon, R. P., Davis, M. H., & Johnsrude, I. S. (2008). Illusory Vowels Resulting from Perceptual Continuity: A Functional Magnetic Resonance Imaging Study. *Journal of Cognitive Neuroscience, 20,* 1737-1752.

Holloway, C. M. (1970). Passing the Strongly Voiced Components of Noisy Speech. *Nature, 226,* 178-179.

Houston-Price, C., & Nakai, S. (2004). Distinguishing novelty and familiarity effects in infant preference procedures. *Infant and Child Development, 13,* 341-348.

Husain, F. T., Lozito, T. P., Ulloa, A., & Horwitz, B. (2005). Investigating the Neural Basis of the Auditory Continuity Illusion. *Journal of Cognitive Neuroscience, 17,* 1275-1292.

Johnson, S. P. (2004). Development of Perceptual Completion in Infancy. *Psychological Science, 15,* 769-775.

Jones, M. R., & Boltz, M. (1989). Dynamic Attending and Responses to Time. *Psychological Review, 96,* 459-491.

Jordan, T. R., Thomas, S. M., & Scott-Brown, K. C. (1999). The illusory-letters phenomenon: An illustration of graphemic restoration in visual word recognition. *Perception, 28,* 1413-1416.

Kaminska, Z., & Mayer, P. (1993). Transformation, migration and restoration: Shades of illusion in the perception of music. *Contemporary Music Review, 9,* 151-161.

Karmiloff-Smith, A. (1992). *Beyond Modularity: A Developmental Perspective on Cognitive Science.* Cambridge, MA: MIT Press.

Kashino, M. (2003). Human Auditory Mechanisms. *NTT Technical Review, 1,* 18-23.

Kashino, M. (2006). Phonemic restoration: The brain creates missing speech sounds. *Acoustic Science and Technology, 27,* 318-321.

Kennedy, M. (1996). *The Concise Oxford Dictionary of Music.* Oxford, UK: Oxford University Press.

Kidd, G., Boltz, M., & Jones, M. R. (1984). Some Effects of Rhythmic Context on Melody Recognition. *American Journal of Psychology, 97,* 153-173.

Kim, J., & Davis, C. (2003). Hearing foreign voices: does knowing what is said affect visual-masked-speech detection? *Perception, 32,* 111-120.

King, A. J. (2007). Auditory Neuroscience: Filling in the Gaps. *Current Biology, 17,* R799-R801.

Kingstone, A., Smilek, D., & Eastwood, J. D. (2008). Cognitive Ethology: A new approach for studying human cognition. *British Journal of Psychology, 99,* 317-340.

Kobayashi, M., Osada, Y., & Kashino, M. (2007). The effect of a flashing visual stimulus on the auditory continuity illusion. *Perception and Psychophysics, 69,* 393-399.

Koelsch, S., Gunter, T. C., von Cramon, Y., Zysset, S., Lohmann, G., & Friederici, A. D. (2002). Bach Speaks: A Cortical "Language-Network" Serves the Processing of Music. *NeuroImage, 17,* 956-966.

Koroleva, I. V., Kashina, I. A., Sakhnovskaya, O. S., & Shurgaya, G. G. (1991). Perceptual restoration of a missing phoneme: New data on speech perception in children. *Sensory Systems, 5,* 191-199.

Koroleva, I. V., Shurgaya, G. G., Kashina, I. A., & Sakhnovskaya, O. S. (1996). Sex Differences in Verbal Perception: Effect of Missing Phoneme Perceptual Restoration. *Human Physiology, 22,* 77-82.

Krumhansl, C. L. (1990). *Cognitive Foundations of Musical Pitch.* New York: Oxford University Press.

Krumhansl, C. L. (1995). Music Psychology and Music Theory: Problems and Prospects. *Music Theory Spectrum, 17,* 53-80.

Krumhansl, C. L. (2000). Rhythm and Pitch in Music Cognition. *Psychological Bulletin, 126,* 159-179.

Krumhansl, C. L., & Castellano, M. A. (1983). Dynamic processes in music perception. *Memory and Cognition, 11,* 325-334.

Krumhansl, C. L., & Keil, F. C. (1982). Acquisition of the hierarchy of tonal functions in music. *Memory and Cognition, 10,* 243-251.

Lamb, S. J., & Gregory, A. H. (1993). The relationship between music and reading in beginning readers. *Educational Psychology, 13,* 19-27.

Layton, B. (1975). Differential effects of two nonspeech sounds on phonemic restoration. *Bulletin of the Psychonomic Society, 6,* 487-490.

Lee, C. S. (1985). The Rhythmic Interpretation of Simple Musical Sequences: Towards a Perceptual Model. In P. Howell, I. Cross, & R. West (Eds.), *Musical Structure and Cognition* (pp. 53-69). London: Academic Press.

Longuet-Higgins, C., & Lee, C. S. (1982). The perception of musical rhythms. *Perception, 11,* 115-128.

Maclean, M., Bryant, P., & Bradley, L. (1987). Rhymes, Nursery Rhymes, and Reading in Early Childhood. *Merrill-Palmer Quarterly, 33*(3), 255-281.

Maess, B., Koelsch, S., Gunter, T. C., & Friederici, A. D. (2001). Musical syntax is processed in Broca's area: an MEG study. *Nature Neuroscience, 4,* 540-545.

Magnuson, J. S., McMurray, B., Tanenhaus, M. K., & Aslin, R. N. (2003). Lexical effects on compensation for coarticulation: the ghost of Christmash past. *Cognitive Science, 27,* 285-298.

Manlove, E. E., Frank, T., & Vernon-Feagans, L. (2001). Why Should We Care About Noise in Classrooms and Child Care Settings? *Child & Youth Care Forum, 30,* 55-64.

Marmel, F., Tillmann, B., & Dowling, W. J. (2008). Tonal expectations influence pitch perception. *Perception and Psychophysics, 70,* 841-852.

Marsh, D. H. (1973). Auditory Figure-Ground Ability in Children. *American Journal of Occupational Therapy, 27,* 218-225.

Marslen-Wilson, W. D., & Warren, P. (1994). Levels of perceptual representations and process in lexical access: Words, phonemes and features. *Psychological Review, 101,* 653-675.

Masataka, N. (2007). Music, evolution and language. *Developmental Science, 10,* 35-39.

Maxwell, L. E., & Evans, G. W. (2000). The effects of noise on pre-school children's pre-reading skills. *Journal of Environmental Psychology, 20,* 91-97.

McAdams, S., & Bregman, A. S. (1979). Hearing Musical Streams. *Computer Music Journal, 3,* 26-43+60.

McClelland, J. L., & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology, 18,* 1-86.

McClelland, J. L., Mirman, D., & Holt, L. L. (2006). Are there interactive processes in speech perception? *Trends in Cognitive Sciences, 10,* 363-369.

McDermott, J. H. (2009). What Can Experiments Reveal About the Origins of Music? *Current Directions in Psychological Science, 18*(3), 164-168.

McDermott, J. H., & Oxenham, A. J. (2008). Spectral completion of partially masked sounds. *Proceedings of the National Academy of Sciences, 105,* 5939-5944.

McMullen, E., & Saffran, J. R. (2004). Music and Language: A Developmental Comparison. *Music Perception, 21,* 289-311.

Micheyl, C., Carlyon, R. P., Shtyrov, Y., Hauk, O., Dodson, T., & Pullvermüller, F. (2003). The Neurophysiological Basis of the Auditory Continuity Illusion: A Mismatch Negativity Study. *Journal of Cognitive Neuroscience, 15,* 747-758.

Miller, C. T., Dibble, E., & Hauser, M. D. (2001). Amodal completion of acoustic signals by a nonhuman primate. *Nature Neuroscience, 4,* 783-784.

Miller, G. A., & Licklider, J. C. R. (1950). The intelligibility of interrupted speech. *Journal of the Acoustical Society of America, 22,* 167-173.

Mills, C. B. (1980). Effects of the match between listener expectancies and coarticulatory cues on the perception of speech. *Journal of Experimental Psychology: Human Perception and Performance, 6,* 528-535.

Mills, J. H. (1975). Noise and Children: A review of literature. *Journal of the Acoustical Society of America, 58,* 767-779.

Mithen, S. (2005). *The Singing Neanderthals: The Origin of Music, Language, Mind and Body*. London: Phoenix.

Morrongiello, B. A. & Roes, C. L. (1990). Children's Memory for New Songs: Integration or Independent Storage of Words and Tunes? *Journal of Experimental Child Psychology, 50,* 25-38.

Morrongiello, B. A., Trehub, S. E., Thorpe, L. A., & Capodilupo, S. (1985). Children's Perception of Melodies: The Role of Contour, Frequency, and Rate of Presentation. *Journal of Experimental Child Psychology, 40,* 279-292.

Narmour, E. (1977). *Beyond Schenkerism: The need for alternatives in music analysis.* Chicago: University of Chicago Press.

Nelken, I. (2004). Processing of complex stimuli and natural scenes in the auditory cortex. *Current Opinion in Neurobiology, 14,* 474-480.

Newman, R. S. (2004). Perceptual restoration in children versus adults. *Applied Psycholinguistics, 25,* 481-493.

Newman, R. S. (2005). The Cocktail Party Effect in Infants Revisited: Listening to One's Name in Noise. *Developmental Psychology, 41*(2), 352-362.

Newman, R. S. (2006). Perceptual restoration in toddlers. *Perception and Psychophysics, 68,* 625-642.

Norris, D. (1995). Signal Detection Theory and Modularity: On Being Sensitive to the Power of Bias Models of Semantic Priming. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 935-939.

Obusek, C. J., & Warren, R. M. (1973). Relation of the Verbal Transformation and the Phonemic Restoration Effects. *Cognitive Psychology, 5,* 97-107.

Parsons, L. M. (2003). Exploring the Functional Neuroanatomy of Music Performance. In I. Peretz & R. J. Zatorre (Eds.), *The Cognitive Neuroscience of Music* (pp. 247-268). New York: Oxford University Press.

Patel, A. D., Gibson, E., Ratner, J., Besson, M., & Holcomb, P. J. (1998). Processing Syntactic Relations in Language and Music: An Event-Related Potential Study. *Journal of Cognitive Neuroscience, 10,* 717-733.

Patel, A. D., & Iversen, J. R. (2007). The linguistic benefits of musical abilities. *Trends in Cognitive Sciences, 11,* 369-372.

Peretz, I. (1993). Auditory agnosia: a functional analysis. In S. McAdams & E. Bigand (Eds.), *Thinking in Sound: The Cognitive Psychology of Human Audition* (pp. 199-230). New York: Oxford University Press.

Peretz, I. (2006). The nature of music from a biological perspective. *Cognition, 100,* 1-32.

Peretz, I., & Coltheart, M. (2003). Modularity of music processing. *Nature Neuroscience, 6,* 688-691.

Peretz, I., & Hyde, K. L. (2003). What is specific to music processing? Insights from congenital amusia. *Trends in Cognitive Sciences, 7,* 362-367.

Peretz, I., & Kolinsky, R. (1993). Boundaries of Separability between Melody and Rhythm in Music Discrimination. *Quarterly Journal of Experimental Psychology, 46A,* 301-325.

Peretz, I., & Morais, J. (1989). Music and Modularity. *Contemporary Music Review, 4,* 279-293.

Peretz, I., Radeau, M., & Arguin, M. (2004). Two-way interactions between music and language: Evidence from priming recognition of tune and lyrics in familiar songs. *Memory and Cognition, 32,* 142-152.

Pessoa, L., Thompson, E., & Noë, A. (1998). Finding out about filling-in: A guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences, 21,* 723-802.

Petkov, C. I., O'Connor, K. N., & Sutter, M. L. (2003). Illusory Sound Perception in Macaque Monkeys. *Journal of Neuroscience, 23,* 9155-9161.

Petkov, C. I., O'Connor, K. N., & Sutter, M. L. (2007). Encoding of Illusory Continuity in Primary Auditory Cortex. *Neuron, 54,* 153-165.

Petzold, R. G. (1969). Auditory Perception by Children. *Journal of Research in Music Education, 17*(1), 82-87.

Piaget, J. (1971). *Biology and Knowledge.* Chicago: University of Chicago Press.

Piccirilli, M., Sciarma, T., & Luzzi, S. (2000). Modularity of music: evidence from a case of pure amusia. *Journal of Neurology, Neurosurgery and Psychiatry, 69,* 541-545.

Pickering, M. J., & Garrod, S. (2007). Do people use language production to make predictions during comprehension? *Trends in Cognitive Sciences, 11,* 105-110.

Pitt, M. A., & Samuel, A. G. (2006). Word length and lexical activation: longer is better. *Journal of Experimental Psychology: Human Perception and Performance, 32*(5), 1120-1135.

Plack, C. J., & White, L. J. (2000). Perceived continuity and pitch perception. *Journal of the Acoustical Society of America, 108,* 1162-1169.

Pomerantz, J. R. (2003). Wholes, holes and basic features in vision. *Trends in Cognitive Sciences, 7*(11), 471-473.

Powers, G. L., & Wilcox, J. C. (1977). Intelligibility of temporally interrupted speech with and without intervening noise. *Journal of the Acoustical Society of America, 61,* 195-199.

Prinz, J. J. (2005). Is the Mind Really Modular? In R. Stainton (Ed.), *Contemporary debates in cognitive science* (pp. 22-36). New York: Blackwell.

Racette, A., & Peretz, I. (2007). Learning lyrics: To sing or not to sing? *Memory and Cognition, 35,* 242-253.

Rainey, D. W., & Larsen, J. D. (2002). The Effect of Familiar Melodies on Initial Learning and Long-term Memory for Unconnected Text. *Music Perception, 20,* 173-186.

Recanzone, G. H., & Sutter, M. L. (2008). The Biological Basis of Audition. *Annual Review of Psychology, 59,* 119-142.

Remez, R. E., Rubin, P. E., Berns, S. M., Pardo, J. S., & Lang, J. M. (1994). On the Perceptual Organization of Speech. *Psychological Review, 101*(1), 129-156.

Remijn, G. B., Nakajima, Y., & Tanaka, S. (2007). Perceptual completion of a sound with a short silent gap. *Perception, 36,* 898-917.

Remijn, G. B., Pérez, E., Nakajima, Y., & Ito, H. (2008). Frequency modulation facilitates (modal) auditory restoration of a gap. *Hearing Research, 243,* 113-120.

Repp, B. H. (1992). Perceptual restoration of a "missing" speech sound: Auditory induction or illusion? *Perception and Psychophysics, 15,* 14-32.

Repp, B. H., Frost, R., & Zsiga, E. (1992). Lexical Mediation Between Sight and Sound in Speechreading. *Quarterly Journal of Experimental Psychology, 45A,* 1-20.

Riecke, L., Mendelsohn, D., Schreiner, C., & Formisano, E. (2009). The continuity illusion adapts to the auditory scene. *Hearing Research, 247,* 71-77.

Riecke, L., van Opstal, A. J., & Formisano, E. (2008). The auditory continuity illusion: A parametric investigation and filter model. *Perception and Psychophysics, 70,* 1-12.

Riecke, L., van Opstal, A. J., Goebel, R., & Formisano, E. (2007). Hearing Illusory Sounds in Noise: Sensory-Perceptual Transformations in Primary Auditory Cortex. *Journal of Neuroscience, 27,* 12684-12689.

Rookes, P., & Willson, J. (2000). *Perception: Theory, development and organisation.* London: Routledge.

Rumelhart, D. E. (1977). Toward an interactive model of reading. In S. Dornic (Ed.), *Attention and Performance VI* (pp. 575-603). Hillsdale, NJ: Lawrence Erlbaum.

Saberi, K., & Perrott, D. R. (1999). Cognitive restoration of reversed speech. *Nature, 398,* 760.

Saffran, J. R. (2003). Musical Learning and Language Development. *Annals of the New York Academy of Sciences, 999,* 397-401.

Samuel, A. G. (1981a). Phonemic Restoration: Insights From a New Methodology. *Journal of Experimental Psychology: General, 110,* 474-494.

Samuel, A. G. (1981b). The role of bottom-up confirmation in the phonemic restoration illusion. *Journal of Experimental Psychology: Human Perception and Performance, 7,* 1124-1131.

Samuel, A. G. (1987). Lexical Uniqueness Effects on Phonemic Restoration. *Journal of Memory and Language, 26,* 36-56.

Samuel, A. G. (1991). A Further Examination of Attentional Effects in the Phonemic Restoration Illusion. *Quarterly Journal of Experimental Psychology, 43A,* 679-699.

Samuel, A. G. (1996). Does lexical information influence the perceptual restoration of phonemes? *Journal of Experimental Psychology: General, 125,* 28-51.

Samuel, A. G. (1997). Lexical Activation Produces Potent Phonemic Percepts. *Cognitive Psychology, 32,* 97-127.

Samuel, A. G. (2001). Knowing a word affects the fundamental perception of the sounds within it. *Psychological Science, 12,* 348-351.

Samuel, A. G., & Ressler, W. H. (1986). Attention within auditory word perception: Insights from the phonemic restoration illusion. *Journal of Experimental Psychology: Human Perception and Performance, 12,* 70-79.

Sasaki, T. (1980). Sound restoration and temporal localization of noise in speech and music sounds. *Tohoku Psychologica Folia, 39,* 79-88.

Schellenberg, E. G. (2004). Music Lessons Enhance IQ. *Psychological Science, 15,* 511-514.

Schellenberg, E. G. (2005). Music and Cognitive Abilities. *Current Directions in Psychological Science, 14,* 317-320.

Schellenberg, E. G., Bigand, E., Poulin-Charronnat, B., Garnier, C., & Stevens, C. (2005). Children's implicit knowledge of harmony in Western music. *Developmental Science, 8*(6), 551-566.

Schiffman, H. R. (2001). *Sensation and Perception: An Integrated Approach* (5th ed.). New York: John Wiley and Sons, Inc.

Schmuckler, M. A. (1997). Expectancy Effects in Memory for Melodies. *Canadian Journal of Experimental Psychology, 51,* 292-305.

Schmuckler, M. A., & Boltz, M. G. (1994). Harmonic and rhythmic influences on musical expectancy. *Perception and Psychophysics, 56,* 313-325.

Schön, D., Magne, C., & Besson, M. (2004). The music of speech: Music training facilitates pitch processing in both music and language. *Psychophysiology, 41,* 341-349.

Schulkind, M. D. (2004). Serial processing in melody identification and the organization of musical semantic memory. *Perception and Psychophysics, 66,* 1351-1362.

Schultz-Westre, C. (1985). *A visual analog of phonemic restorations: Sign restoration in American Sign Language.* Unpublished dissertation, The University of Wisconsin-Milwaukee.

Seeba, F., & Klump, G.M. (2009). Stimulus Familiarity Affects Perceptual Restoration in the European Starling (*Sturnus vulgaris). PLoS ONE, 4*(6), 5974.

Sekuler, R., & Blake, R. (1994). *Perception* (3rd ed.). Singapore: McGraw-Hill.

Serafine, M. L. (1989). What Music Is. *Journal of Aesthetic Education, 23,* 31-37.

Serafine, M. L., Crowder, R. G., & Repp, B. H. (1984). Integration of melody and text in memory for songs. *Cognition, 16,* 285-303.

Serafine, M. L., Davidson, J., Crowder, R. G., & Repp, B. H. (1986). On the Nature of Melody-Text Integration in Memory for Songs. *Memory and Language, 25,* 123-135.

Shahin, A. J., Bishop, C. W., & Miller, L. M. (2009). Neural mechanisms for illusory filling-in of degraded speech. *NeuroImage, 44,* 1133-1143.

Shinn-Cunningham, B. G. (2008). Object-based auditory and visual attention. *Trends in Cognitive Sciences, 12,* 182-186.

Shinn-Cunningham, B. G., & Wang, D. (2008). Influences of auditory object formation on phonemic restoration. *Journal of the Acoustical Society of America, 123,* 295-301.

Shriberg, E. E. (1992). Perceptual Restoration of Filtered Vowels with Added Noise. *Language and Speech, 35,* 127-136.

Sivonen, P., Maess, B., & Friederici, A. D. (2006). Semantic retrieval of spoken words with an obliterated initial phoneme in a sentence context. *Neuroscience Letters, 408,* 220-225.

Sivonen, P., Maess, B., Lattner, S., & Friederici, A. D. (2006). Phonemic restoration in a sentence context: Evidence from early and late ERP effects. *Brain Research, 1121,* 177-198.

Slabbekoorn, H., & den Boer-Visser, A. (2006). Cities Change the Songs of Birds. *Current Biology, 16,* 2326-2331.

Slabbekoorn, H., & Peet, M. (2003). Birds sing at a higher pitch in urban noise. *Nature, 424,* 267.

Smith, J. A. (2000). Singing and songwriting support early literacy instruction. *The Reading Teacher, 53*(8), 646-649.

Smith, K. C., & Cuddy, L. L. (1989). Effects of Metric and Harmonic Rhythm on the Detection of Pitch Alterations in Melodic Sequences. *Journal of Experimental Psychology: Human Perception and Performance, 15,* 457-471.

Smith, N. A. (2005). The perceptual restoration of music. *Dissertation Abstracts International: Section B: The Sciences and Engineering, 65*, 5444.

Solley, C. M., & Murphy, G. (1960). *Development of the perceptual world.* New York: Basic Books.

Srinivasan, S., & Wang, D. (2005). A schema-based model for phonemic restoration. *Speech Communication, 45,* 63-87.

Steinke, W. R., Cuddy, L. L., & Jakobson, L. S. (2001). Dissociations among functional subsystems governing melody recognition after right-hemisphere damage. *Cognitive Neuropsychology, 18,* 411-437.

Sugita, Y. (1997). Neuronal correlates of auditory induction in the cat cortex. *NeuroReport, 8,* 1155-1159.

Sumby, W. H., & Pollack, I. (1954). Visual Contribution to Speech Intelligibility in Noise. *Journal of the Acoustical Society of America, 26,* 212-215.

Thackray, R. (1972). *Rhythmic Abilities in Children.* London: Novello.

Thompson, W. F., Schellenberg, E. G., & Husain, G. (2004). Decoding Speech Prosody: Do Music Lessons Help? *Emotion, 4,* 46-64.

Thurlow, W. (1957). An Auditory Figure-Ground Effect. *American Journal of Psychology, 70,* 653-654.

Thurlow, W. R., & Elfner, L. F. (1959). Continuity Effects with Alternately Sounding Tones. *Journal of the Acoustical Society of America, 31,* 1337-1339.

Tougas, Y., & Bregman, A. S. (1990). Auditory streaming and the continuity illusion. *Perception and Psychophysics, 47,* 121-126.

Trehub, S. (2001). Human Processing Predispositions and Musical Universals. In N. L. Wallin, B. Merker, & S. Brown (Eds.), *The Origins of Music* (pp. 427-448). Cambridge, MA: MIT.

Trehub, S. E. (2003). Towards a Developmental Psychology of Music. *Annals of the New York Academy of Sciences, 999,* 402-413.

Trehub, S. E., Bull, D., & Thorpe, L. A. (1984). Infants' Perception of Melodies: The Role of Melodic Contour. *Child Development, 55,* 821-830.

Trehub, S. E., & Trainor, L. J. (1998). Singing to Infants: Lullabies and Play Songs. In C. Rovee-Collier, L. P. Lipsitt & H. Hayne (Eds.), *Advances in infancy research, Vol. 12* (pp. 43-77). Stamford, CT: Ablex.

Trout, J. D., & Poser, W. J. (1990). Auditory and Visual Influences on Phonemic Restoration. *Language and Speech, 33,* 121-135.

Valenza, E., Leo, I., Gava, L., & Simion, F. (2006). Perceptual Completion in Newborn Human Infants. *Child Development, 77,* 1810-1821.

Vallar, S. (2006). An Enactive-Phenomenological Approach to Veridical Perception. *Journal of Consciousness Studies, 13,* 39-60.

Vernon, M. D. (1966). Perception in Relation to Cognition. In A. H. Kidd & J. L. Rivoire (Eds.), *Perceptual Development in Children* (pp.391-406). London: University of London Press.

Verschuure, J., & Brocaar, M. P. (1983). Intelligibility of interrupted meaningful and nonsense speech with and without intervening noise. *Perception and Psychophysics, 33,* 232-240.

Walley, A. C. (1988). Spoken Word Recognition by Young Children and Adults. *Cognitive Development, 3,* 137-165.

Warren, R. M. (1970). Perceptual restoration of missing speech sounds. *Science, 167,* 392-393.

Warren, R. M. (1983). Auditory Illusions and Their Relation to Mechanisms Normally Enhancing Accuracy of Perception. *Journal of the Audio Engineering Society, 31,* 623-629.

Warren, R. M. (1984). Perceptual restoration of obliterated sounds. *Psychological Bulletin, 96,* 371-383.

Warren, R. M. (1999). *Auditory Perception: A New Analysis and Synthesis.* Cambridge, UK: Cambridge University Press.

Warren, R. M., Bashford, J. A., Jr., & Healy, E. W. (1992). The subtractive nature of auditory continuity: Reciprocal changes in alternating sounds. *Journal of the Acoustical Society of America, 91,* 2334.

Warren, R. M., Bashford, J. A., Jr., Healy, E. W., & Brubaker, B. S. (1994). Auditory induction: Reciprocal changes in alternating sounds. *Perception and Psychophysics, 55,* 313-322.

Warren, R. M., Gardner, D. A., Brubaker, B. S., & Bashford, J. A., Jr. (1991). Melodic and Nonmelodic Sequences of Tones: Effects of Duration on Perception. *Music Perception, 8*(3), 277-290.

Warren, R. M., & Obusek, C. J. (1971). Speech perception and phonemic restorations. *Perception and Psychophysics, 9,* 358-362.

Warren, R. M., Obusek, C. J., & Ackroff, J. M. (1972). Auditory induction: Perceptual synthesis of absent sounds. *Science, 176,* 1149-1151.

Warren, R. M., Riener Hainsworth, K., Brubaker, B. S., Bashford, J. A., Jr., & Healy, E. W. (1997). Spectral restoration of speech: Intelligibility is increased by inserting noise in spectral gaps. *Perception and Psychophysics, 59,* 275-283.

Warren, R. M., & Sherman, G. L. (1974). Phonemic restorations based on subsequent context. *Perception and Psychophysics, 16,* 150-156.

Warren, R. M., & Warren, R. P. (1970). Auditory Illusions and Confusions. *Scientific American, 223,* 30-36.

Warren, R. M., & Warren, R. P. (1971). Some age differences in auditory perception. *Bulletin of the New York Academy of Medicine, 47*(11), 1365-1377.

Waterman, A. H., Blades, M., & Spencer, C. (2000). Do children try to answer nonsensical questions? *British Journal of Developmental Psychology, 18,* 211-225.

Welch, G., Himonides, E., Saunders, J., & Papageorgi, I. (2008). *The National Singing Programme for Primary Schools in England: An initial baseline study overview, February 2008.* London: Institute of Education.

White, B. W. (1960). Recognition of Distorted Melodies. *American Journal of Psychology, 73,* 100-107.

Zatorre, R. J., Belin, P., & Penhune, V. B. (2002). Structure and function of auditory cortex: music and speech. *Trends in Cognitive Sciences, 6,* 37-46.

APPENDICES

Appendix 1: Song familiarity questionnaire sent to participating schools

UNIVERSITY OF SURREY

The Perceptual Restoration of Music in Children
Song Familiarity Questionnaire

Thank you for agreeing to take part in this research. Before we visit, it would be very helpful if you could indicate whether the following children's songs are: a) taught in your school, and b) sung on a regular basis. Many thanks for your help.

Name of School____________________________

Name of Song	***Taught in School? (Please tick)***	***Sung regularly in School? (Please tick)***
Baa Baa Black Sheep		
The Grand Old Duke of York		
Twinkle Twinkle Little Star		
Happy Birthday to You		
The Wheels on The Bus		
The Hokey Cokey		
Row, Row, Row Your Boat		
Old Macdonald Had a Farm		
London Bridge		
Jack and Jill		
Incy Wincy Spider		
Humpty Dumpty		
Hickory Dickory Dock		
If You're Happy and You Know It		
Head, Shoulders, Knees and Toes		

Appendix 2: Results of analyses for gender effects

Experiment 1

Experimental Condition	**Correct identification scores**	**Identification time**
Intact	$t(20) = .58, p > .05$	$t(17) = .49, p > .05$
Noise-filled gaps	$t(20) = .33, p > .05$	$t(20) = .71, p > .05$
Silence-filled gaps	$t(20) = -.92, p > .05$	$t(20) = .48, p > .05$

Experiment 2

Experimental Condition	**Correct identification scores**	**Identification time**
Noise 200	$t(20) = -.65, p > .05$	$t(20) = -.80, p > .05$
Noise 300	$t(20) = -.88, p > .05$	$t(20) = -.19, p > .05$
Silence 200	$t(20) = .04, p > .05$	$t(19) = 1.26, p > .05$
Silence 300	$t(15) = -.16, p > .05$	$t(15) = 1.07, p > .05$

Experiment 3

Experimental Condition (noise amplitude and gap duration)	**Correct identification scores**	**Identification time**
High Amplitude Noise 100	$t(19) = 1.61, p > .05$	$t(19) = -1.09, p > .05$
High Amplitude Noise 300	$t(19) = 2.41, p < .05$	$t(19) = -1.13, p > .05$
Low Amplitude Noise 100	$t(19) = -.83, p > .05$	$t\ 19) = .13, p > .05$
Low Amplitude Noise 300	$t(19) = .67, p > .05$	$t(15) = 1.02, p > .05$
Extra Low Amplitude Noise 100	$t(19) = -.09, p > .05$	$t(18) = -.34\ p > .05$
Extra Low Amplitude Noise 300	$t(19) = -.18, p > .05$	$t(17) = 1.77\ p > .05$

Experiment 4

Experimental Condition	Correct identification scores	Identification time
Intact	$t(20) = .00, p > .05$	$t(20) = .84, p > .05$
Noise-filled gaps	$t(20) = .40, p > .05$	$t(20) = -.40, p > .05$
Silence-filled gaps	$t(20) = .46, p > .05$	$t(20) = -.68, p > .05$

Experiment 5

Experimental Condition	Stimulus Type	Correct identification scores	Identification time
Intact	Lyrics and Melody	$t(18) = 1.67, p > .05$	$t(18) = -1.33, p > .05$
	Melody Only	$t(18) = .21, p > .05$	$t(16) = -1.39, p > .05$
Noise	Lyrics and Melody	$t(18) = .06, p > .05$	$t(18) = -.38, p > .05$
	Melody Only	$t(18) = -.43, p > .05$	$t(18) = -1.12, p > .05$
Silence	Lyrics and Melody	$t(18) = -1.07, p > .05$	$t(14) = .52, p > .05$
	Melody Only	$t(18) = -.81, p > .05$	$t(18) = -2.27, p < .05$

Experiment 6

Experimental Condition	Correct identification scores	Identification time
Intact	$t(16) = -.42, p > .05$	$t(16) = -1.05, p > .05$
Noise-filled gaps	$t(16) = -.24, p > .05$	$t(16) = -1.20, p > .05$
Silence-filled gaps	$t(16) = .30, p > .05$	$t(16) = .49, p > .05$

Experiment 7

Experimental Condition	**Correct identification scores**	**Identification time**
Intact	$t(18) = -.35, p > .05$	$t(18) = -1.99, p > .05$
Noise-filled gaps	$t(18) = 1.47, p > .05$	$t(18) = -.57, p > .05$
Silence-filled gaps	$t(18) = .66, p > .05$	$t(17) = -.74, p > .05$

Experiment 8

Age group	**Experimental condition**	**Correct identification scores**	**Identification time**
Older children	Intact	$t(18) = -1.39, p > .05$	$t(18) = .41, p > .05$
	Noise-filled gaps	$t(18) = -.39, p > .05$	$t(18) = -1.43, p > .05$
	Silence-filled gaps	$t(18) = .00, p > .05$	$t(18) = -.06, p > .05$
Adults	Intact	no males	no males
	Noise-filled gaps	$t(8) = .72, p > .05$	$t(8) = 1.56, p > .05$
	Silence-filled gaps	$t(8) = 1.70, p > .05$	$t(8) = -.45, p > .05$

Appendix 3: Observed/ Expected frequencies for chi-square analyses

Experiment 3: Noise Amplitude

Perception of Continuity		High 100 ms	High 300 ms	Low 100 ms	Low 300 ms	Very Low 100 ms	Very Low 100 ms
Fragmented	observed	6	7	15	17	19	17
	expected	13.5	13.5	13.5	13.5	13.5	13.5
	%	28.6%	33.3%	71.4%	81.0%	90.5%	81%
Continuous	observed	15	14	6	4	2	4
	expected	7.5	7.5	7.5	7.5	7.5	7.5
	%	71.4%	66.7%	28.6%	19.0%	9.5%	19%

Experiment 4: Beginning/later phrases

Perception of Continuity		Intact	Noise-filled gaps	Silence-filled gaps
Fragmented	observed	1	11	17
	expected	9.7	9.7	9.7
	%	4.5%	50.0%	77.3%
Continuous	observed	21	11	5
	expected	12.3	12.3	12.3
	%	95.5%	50.0%	22.7%

Experiment 5: Melody & Lyrics

Perception of Continuity		Intact	Noise-filled gaps	Silence-filled gaps
Fragmented	observed	1	11	19
	expected	10.3	10.3	10.3
	%	5.0%	55.0%	95.0%
Continuous	observed	19	9	1
	expected	9.7	9.7	9.7
	%	95.0%	45.0%	5.0%

Experiment 5: Melody Only

Perception of Continuity		Intact	Noise-filled gaps	Silence-filled gaps
Fragmented	observed	2	8	19
	expected	9.7	9.7	9.7
	%	10.0%	40.0%	95.0%
Continuous	observed	18	12	1
	expected	10.3	10.3	10.3
	%	90.0%	60.0%	5.0%

Experiment 6: Lyrics/ melody

Perception of Continuity		Intact	Noise-filled gaps	Silence-filled gaps
Fragmented	observed	0	9	17
	expected	8.7	8.7	8.7
	%	0.0%	50.0%	94.4%
Continuous	observed	18	9	1
	expected	9.3	9.3	9.3
	%	100.0%	50.0%	5.6%

Experiment 7: Pitch/ rhythm

Perception of Continuity		Intact	Noise-filled gaps	Silence-filled gaps
Fragmented	observed	2	15	17
	expected	11.3	11.3	11.3
	%	10.0%	75.0%	85.0%
Continuous	observed	18	5	3
	expected	8.7	8.7	8.7
	%	90.0%	25.0%	15.0%

Experiment 8: Age

Perception of Continuity		**Intact**	**Noise-filled gaps**	**Silence-filled gaps**
Fragmented	observed	2	7	27
	expected	12.0	12.0	12.0
	%	6.7%	23.3%	90.0%
Continuous	observed	28	23	3
	expected	18.0	18.0	18.0
	%	93.3%	76.7%	10.0%

Appendix 4: Details of the stimuli presented in Experiment 4

Nursery rhyme	Duration of excerpt (s)	Time signature	Key signature	Phrase position
Baa Baa Black Sheep	5.90	4/4	C major	second
The Grand Old Duke of York	5.50	4/4	C major	second
Twinkle Twinkle Little Star	5.90	4/4	C major	second
Happy Birthday to You	5.30	3/4	C major	second
The Wheels on the Bus	5.00	4/4	C major	second
The Hokey Cokey	5.60	2/4	C major	second
Row, Row, Row Your Boat	3.80	4/4	C major	second
Old Macdonald had a farm	5.50	4/4	C major	third
London Bridge	4.00	4/4	C major	second
Jack and Jill	4.10	2/4	C major	second
Incy Wincy Spider	4.60	4/4	C major	second
Humpty Dumpty	7.50	2/4	C major	second
Hickory Dickory Dock	6.20	4/4	C major	second
If You're Happy and You Know it Clap Your Hands	6.70	4/4	C major	second
Heads, Shoulders, Knees and Toes	3.00	4/4	G major/C major	third

Appendix 5: Details of the stimuli presented in Experiment 5

The name, duration, time and key signatures refer to both melody and lyrics and melody only stimuli; the lyrics presented refer to the melody and lyrics stimuli only.

Nursery rhyme	Lyrics	Duration of excerpt (s)	Time signature	Key signature
Baa Baa Black Sheep	***Baa baa black sheep have you any wool? Yes, sir, yes sir, Three bags full***	5.90	4/4	C major
The Grand Old Duke of York	***Oh, the grand old Duke of York, he had ten thousand men***	5.25	4/4	C major
Twinkle Twinkle Little Star	***Twinkle, twinkle little star, How I wonder what you are!***	5.75	4/4	C major
Happy Birthday to You	***Happy Birthday to you, Happy Birthday to you***	4.75	3/4	C major
The Wheels on the Bus	***The wheels on the bus go round and round, Round and round, round and round.***	5.32	4/4	C major
The Hokey Cokey	***You put you left leg in, your left leg out, In, out, in, out, And shake it all about***	6.50	2/4	C major
Row, Row, Row Your Boat	***Row, row, row your boat Gently down the stream***	4.25	4/4	C major
Old Macdonald had a farm	***Old Macdonald had a farm, E-i-e-i-o***	4.50	4/4	C major
London Bridge	***London Bridge is falling down, Falling down, falling down***	4.50	4/4	C major

Jack and Jill	***Jack and Jill went up the hill*** ***To fetch a pail of water***	4.40	2/4	C major
Incy Wincy Spider	***Incy Wincy Spider climbed the water spout***	5.00	4/4	C major
Humpty Dumpty	***Humpty Dumpty sat on a wall,*** ***Humpty Dumpty had a great fall***	8.00	2/4	C major
Hickory Dickory Dock	***Hickory, dickory, dock,*** ***The mouse ran up the clock***	5.50	4/4	C major
If You're Happy and You Know it Clap Your Hands	***If you're happy and you know it, clap your hands,*** ***If you're happy and you know it, clap your hands***	6.75	4/4	C major
Heads, Shoulders, Knees and Toes	***Head, shoulders, knees and toes, knees and toes***	3.25	4/4	G major/C major

Appendix 6: The three stimulus types presented in Experiment 6, shown for *Twinkle Twinkle Little Star*

1.*Twinkle Twinkle Little Star* Old Lyrics, New Melody

Melody:
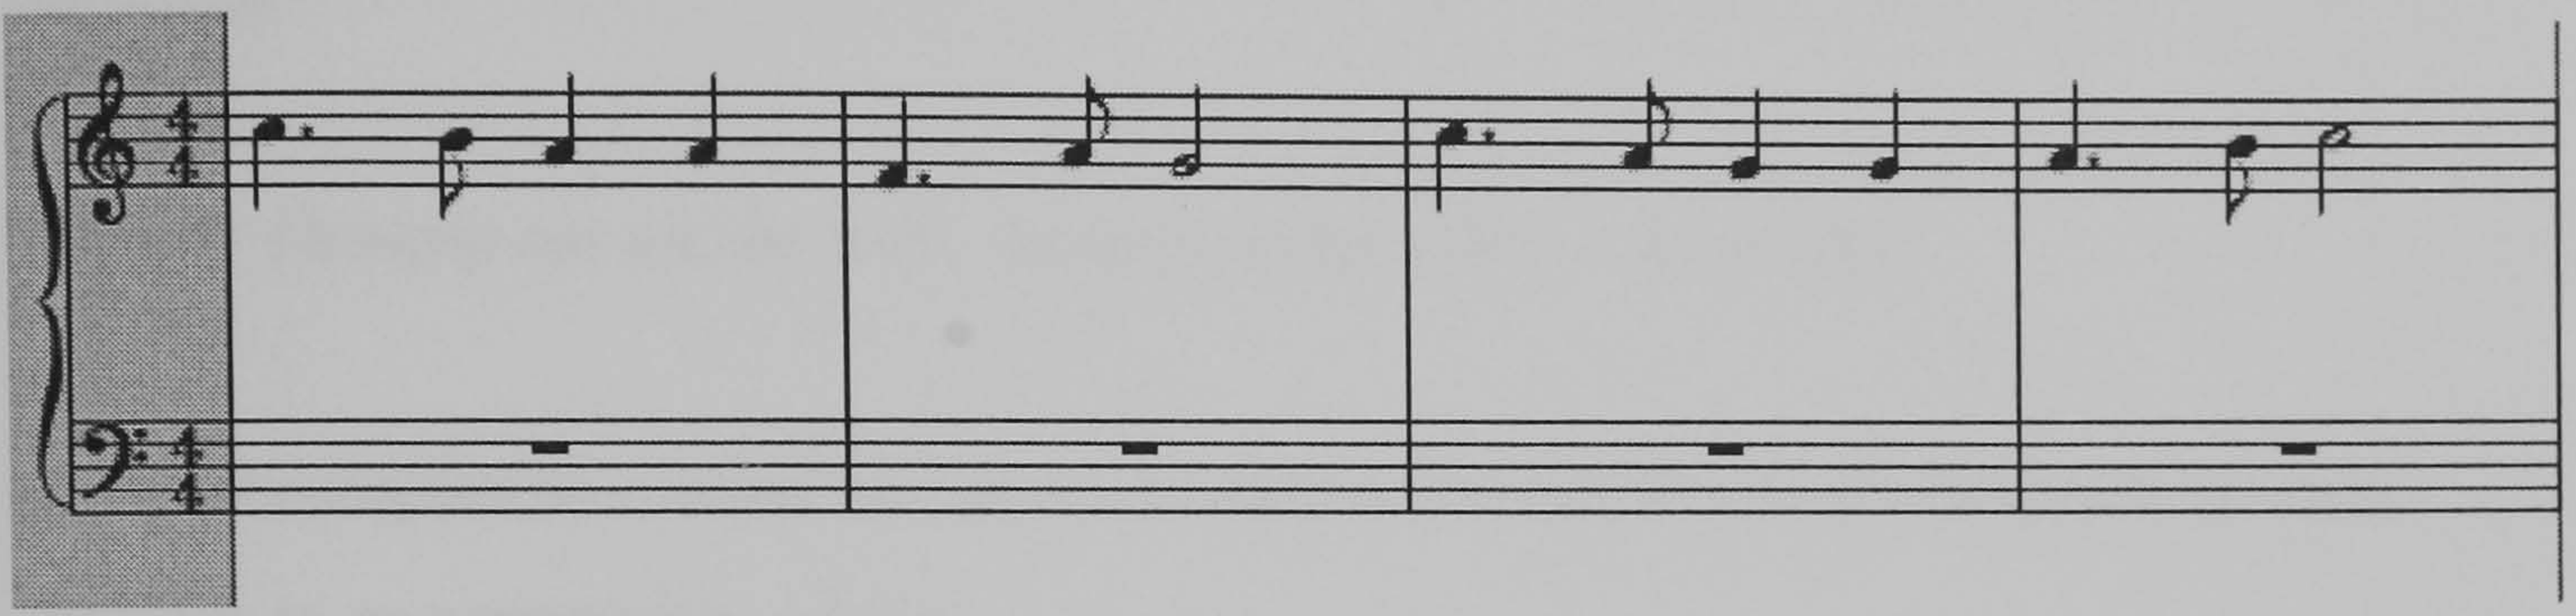
Lyrics:

Twinkle twinkle little star, how I wonder what you are

2.*Twinkle Twinkle Little Star* Old Melody, New Lyrics

Melody:

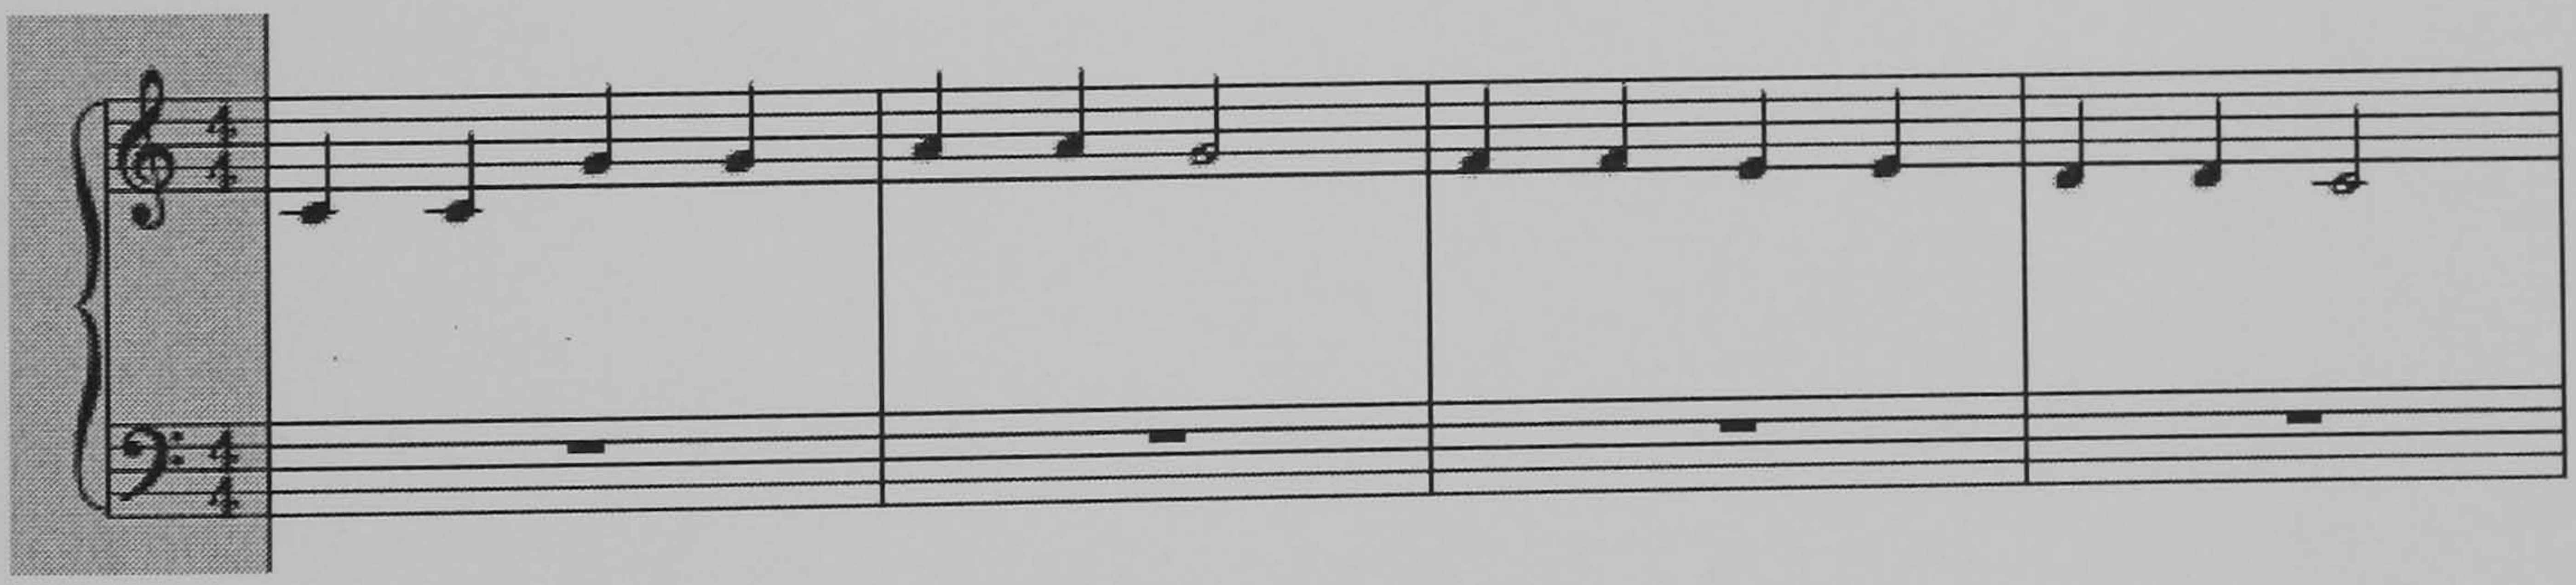

Lyrics:

Can you see the busy bee, buzzing round the apple tree?

3.*Twinkle Twinkle Little Star* mismatch

Melody:

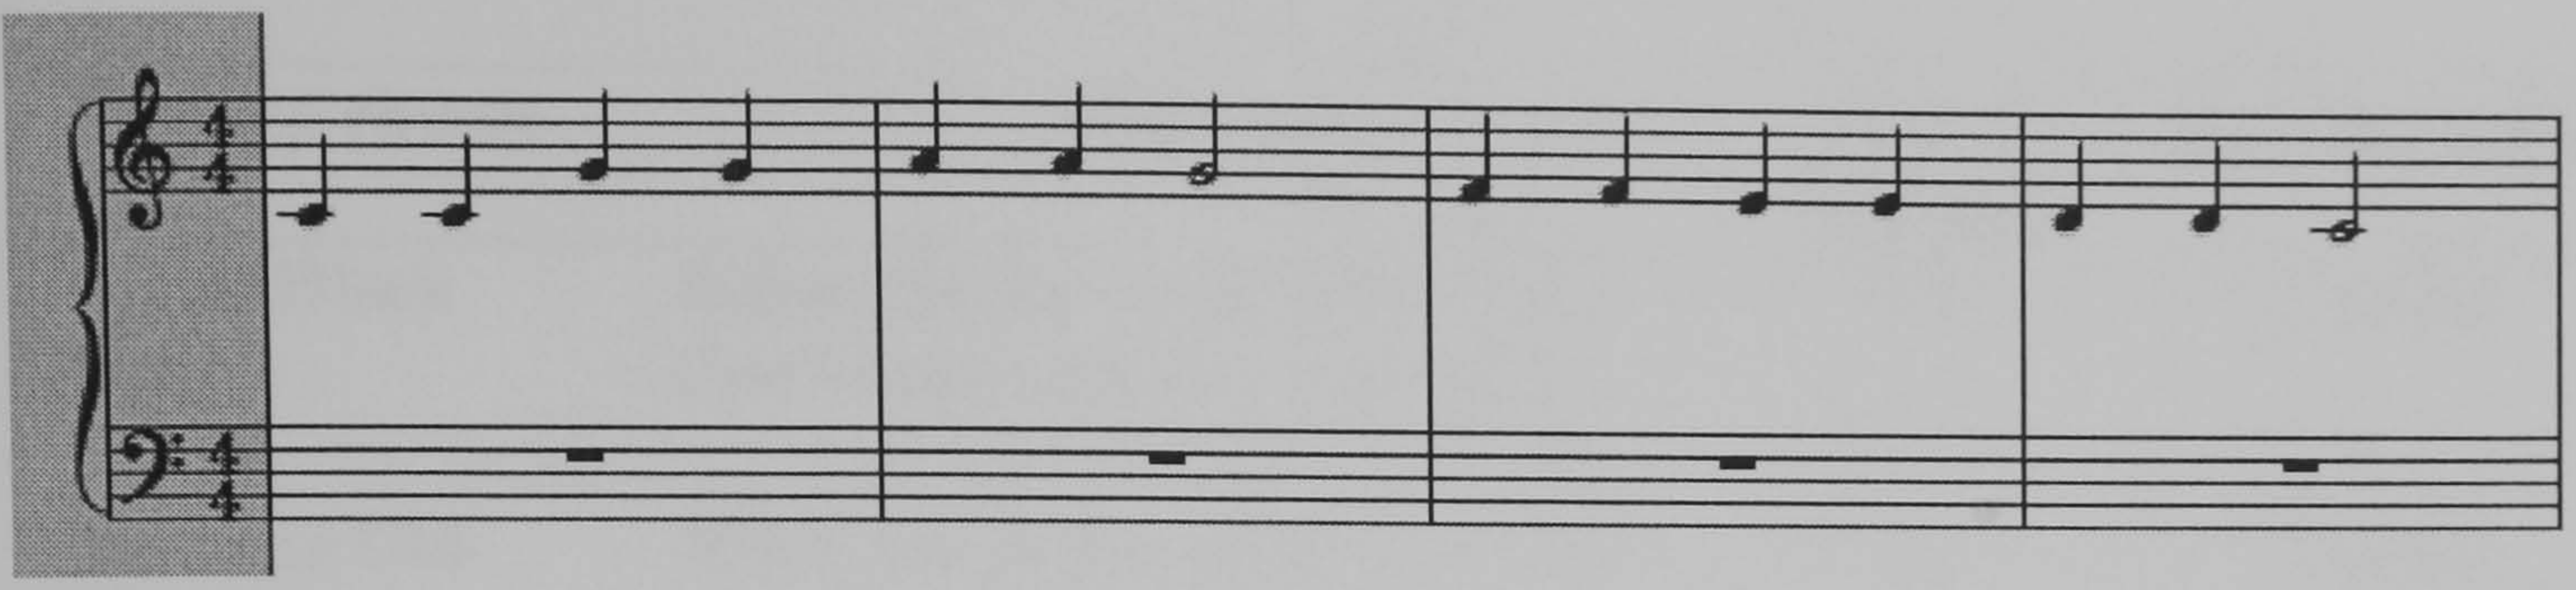

Lyrics:

Humpty Dumpty sat on the wall, humpty dumpty had a great fall

(Second crotchet in second bar and second crotchet in fourth bar changed to two quavers to accommodate lyrics)

Appendix 7: Details of the six stimuli presented in Experiment 6

Nursery rhyme	New Lyrics	Time signature	Key signature
Baa Baa Black Sheep	*Sammy Sailor on the ocean blue, Can I come and sail with you?*	4/4	C major
The Grand Old Duke of York	*When you look around the room, tell me the things you see*	4/4	C major
Twinkle Twinkle Little Star	*Can you see the busy bee, buzzing round the apple tree?*	4/4	C major
Happy Birthday to You	*Pick a flower for me, pick a flower for me*	3/4	C major
The Wheels on the Bus	*The duck on the pond made quite a splash, quite a splash, quite a splash*	4/4	C major
The Hokey Cokey	*The princess told the queen, I've lost my crown. You must find it the queen said with a frown*	2/4	C major
Row, Row, Row Your Boat	*Come down from the tree, naughty little Ben*	4/4	C major
Old Macdonald had a farm	*Clap your hands and stamp your feet, turn around and round*	4/4	C major
London Bridge	*Go to market, sell a cow, sell a cow, sell a cow*	4/4	C major

Appendix 8: Measures of musical characteristics for original melodies (old) and newly composed melodies (new) for Experiment 6

	Number of tone onsets[a]		range[b]		Average interval size[c]		Contour complexity[d]		Number of note durations[e]		% most common durations[f]		Tonal strength[g]	
	old	New	**old**	new	**old**	new	**old**	new	**old**	New	**old**	New	**old**	new
Baa Baa Black Sheep	**16**	16	**12**	12	**1.6**	2.33	**0.06**	0.38	**3**	3	**87.5**	93.75	**4.69**	4.62
Grand duke	**13**	13	**12**	12	**2.25**	2.38	**0.15**	0.23	**3**	2	**100**	100	**4.60**	4.87
Twinkle	**14**	14	**9**	7	**1.38**	2.15	**0.07**	0.29	**2**	4	**100**	85.71	**3.84**	4.48
Happy Birthday	**12**	12	**7**	7	**2.45**	1.81	**0.58**	0.42	**3**	3	**91.67**	100	**4.79**	4.79
Wheels on Bus	**15**	15	**12**	12	**2.14**	2.29	**0.20**	0.27	**3**	4	**100**	86.67	**4.87**	4.53
Hokey Cokey	**20**	20	**5**	5	**1.68**	1.3	**0.50**	0.30	**3**	3	**95**	90	**5.06**	4.21
Row Boat	**10**	10	**7**	5	**1.1**	1.56	**0.20**	0.30	**4**	2	**90**	100	**4.84**	4.92
Old Macdonald	**12**	12	**9**	9	**2.0**	2.09	**0.33**	0.25	**3**	4	**100**	91.67	**4.90**	4.75
London Bridge	**13**	13	**7**	7	**1.83**	1.5	**0.46**	0.23	**4**	4	**84.62**	100	**4.42**	4.57
Mean	**13.89**	13.89	**8.89**	8.44	**1.83**	1.96	**.28**	.30	**3.11**	3.22	**94.31**	94.20	**4.67**	4.64

[a] the number of notes in the melody

[b] the distance in semitones between the lowest and highest pitch notes in the melody (12 indicates a distance of one octave)

[c] the average distance in semitones between two successive notes

[d] a measure of the complexity of pitch direction changes in the melody. Calculated as the ratio of the number of direction changes to the number of notes. The lower the number, the simpler the contour of the melody

[e] the number of different note durations in the melody (e.g. crotchet, quaver, etc.)

[f] the percentage of notes accounted for by the two most frequent note durations; or between the three most frequent note durations if a three-way tie.

[g] a measure of how well the melody conforms to the key in which it is written. Each note in the scale is given a value; the measure here represents the average value. Note that the maximum value would be 6.35, representing a (hypothetical) case where each note of the melody was the tonic. The measures for a C major scale are:
C (tonic): 6.35; C# : 2.23#; D: 3.48; D#: 2.33; E: 4.38; F (subdominant): 4.09; F#: 2.52; G (dominant): 5.19; G#: 2.39; A: 3.66; A#: 2.29; B: 2.88

Appendix 9: Experiment 6 mismatch pairs and changes made to original songs in order to accommodate a mismatch

Melody	Lyrics	Changes made
The grand old duke of York*	Old Macdonald	One extra syllable added to lyrics
Twinkle Twinkle	Humpty Dumpty	Two crotchets changed to quavers to accommodate extra syllables
Happy birthday	Hickory Dickory Dock	One crotchet changed to quavers to accommodate extra syllable
The wheels on the bus	Baa Baa Black Sheep	Two crotchets tied to accommodate one syllable
Hokey cokey	The wheels on the bus	Extra syllable added and 3 syllables slurred across 5 notes
Old MacDonald	Row your boat	2 sets of 2 crotchets tied (i.e. 4 crotchets accommodate 2 syllables)
London bridge	Incy Wincy spider	1 syllable slurred across 2 crotchets

* mismatch demonstration trial

Appendix 10: The three stimulus types presented in Experiment 7, shown for *Twinkle Twinkle Little Star*

1.Twinkle Twinkle Little Star Old Pitch, New Rhythm
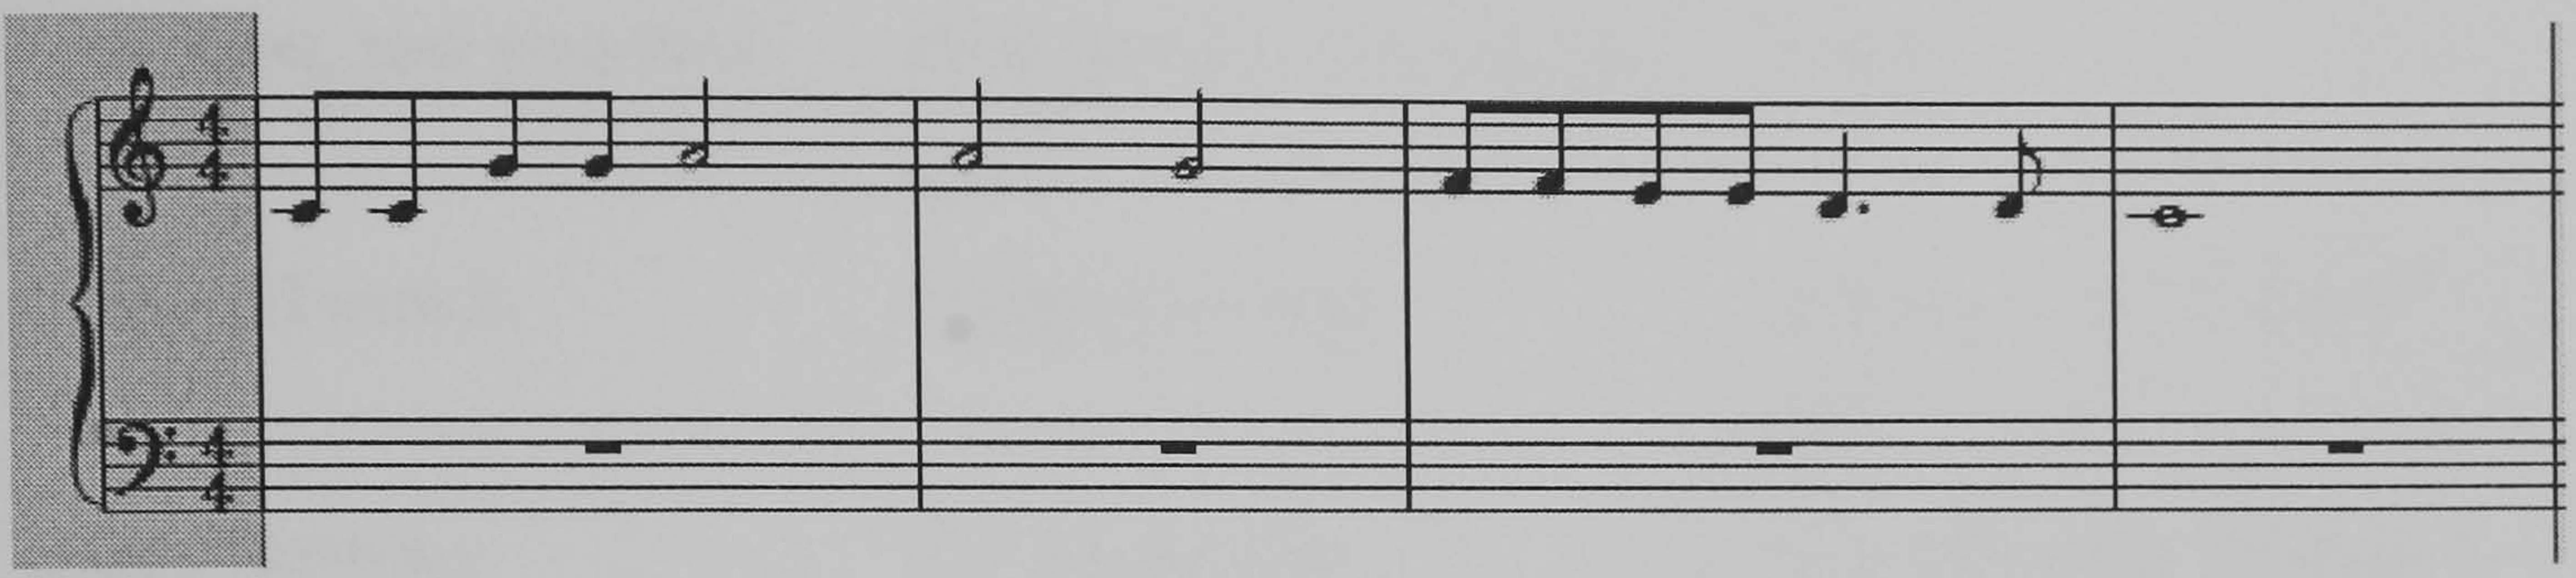
2.Twinkle Twinkle Little Star Old Rhythm, New Pitch

3.Twinkle Twinkle Little Star rhythm mismatch

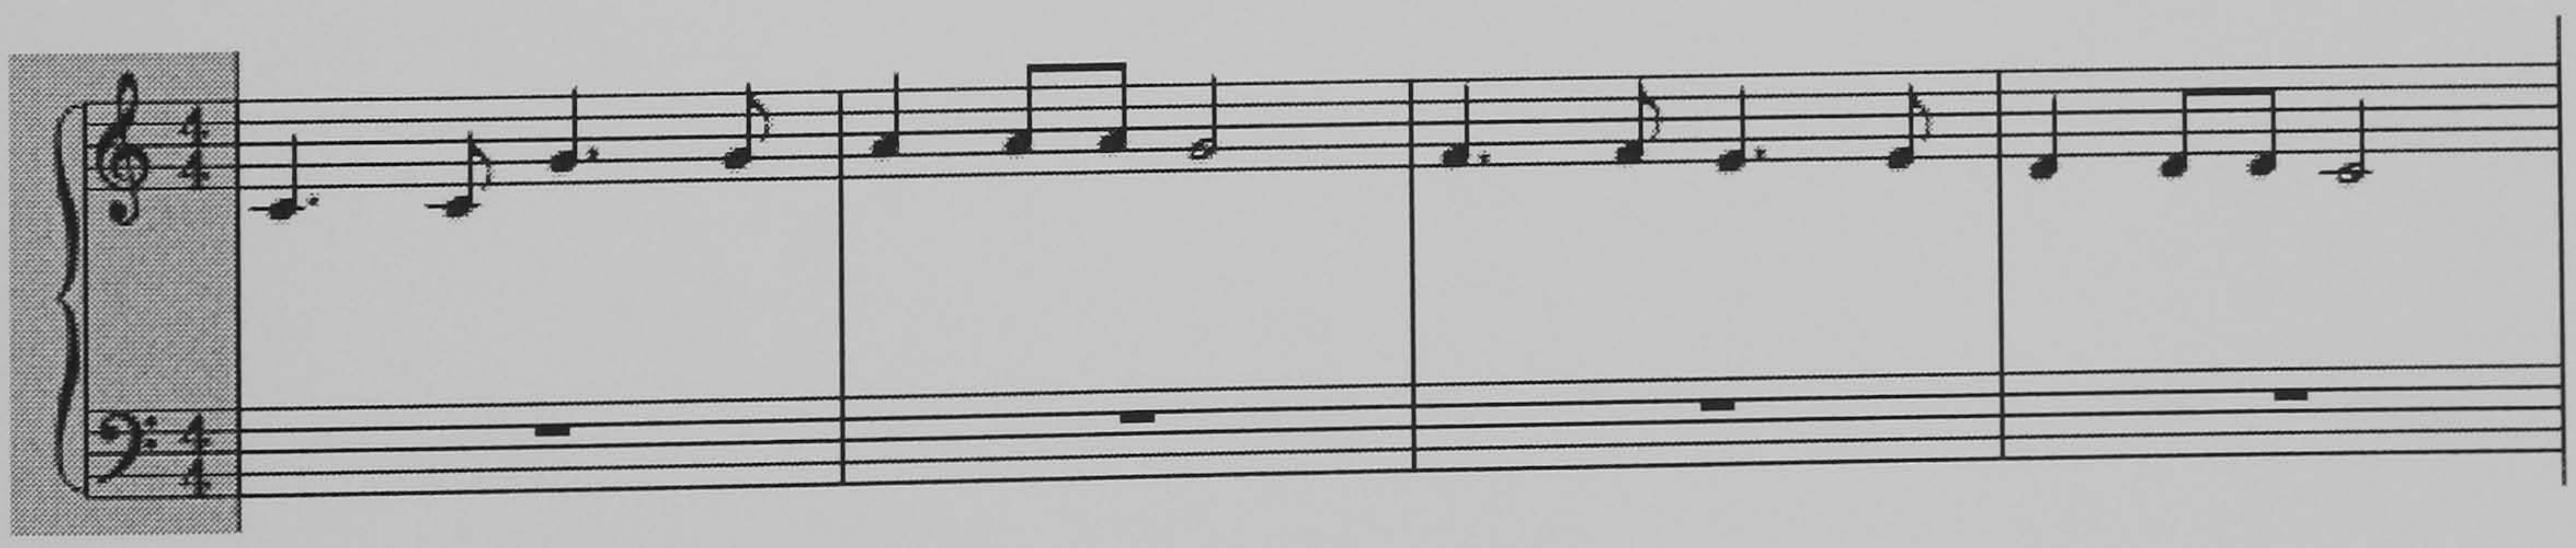

Pitch: *Twinkle Twinkle Little Star*

Rhythm: *Humpty Dumpty*

(extra note inserted in 2nd and 4th bars)

Appendix 11: Details of mismatch stimuli in Experiment 7

Nursery rhyme	**Mismatch partner**	**Changes made**
Baa baa black sheep	The wheels on the bus	None
The grand old duke of York	Hickory Dickory Dock	time signature changed from 4/4 to 3/4
Row, Row, row your boat	Head, shoulders knees and toes	None
Twinkle Twinkle	Humpty Dumpty	extra note in 2nd and 4th bars
Happy Birthday	Old Macdonald	time signature changed from 3/4 to 4/4
The Wheels on the Bus	Baa Baa black sheep	None
Old Macdonald	Happy Birthday	None
London Bridge	The grand old duke of York	None
Head, shoulders knees and toes	Row, row, row your boat	None

Appendix 12: Information sheet

Researcher: Naomi Worsfold
N.Worsfold@surrey.ac.uk
20AC04 ext: 6945
Supervisor: Dr. Alyson Davis

UNIVERSITY OF SURREY

Title of Research Project: The perceptual restoration of music

Information Sheet

Noise often accompanies the sounds we hear on a daily basis. Conversation can be masked by intermittent everyday sounds, and even music and other sounds can be disrupted. Research has demonstrated that the brain has a mechanism for overcoming the effect of noise; it is able to fill in the missing information based on what we expect to be present. This effect is called perceptual restoration.

Previous research has found evidence of the perceptual restoration of speech in both adults and children. If portions of a speech signal are removed and replaced with noise, both adults and young children are able to make sense of what is being said, even though parts are missing. However, when the same portions are removed and left silent, comprehension is extremely difficult.

The perceptual restoration of other sounds such as music has only been studied with adults. We know that in adults, the brain restores missing parts of a tune, but we want to compare this with restoration in children. We are particularly interested in how familiarity with a tune influences what is actually restored.

In this study, the session will be in the form of a music game, lasting approximately 5 minutes. Participants will be asked to listen to some children's nursery rhymes. Some of the music will have portions removed and replaced with white noise, and some with silence. Participants will then be asked if they can identify the music. The only piece of personal information required is participant age so that the effect of age can also be investigated. Each participant has the right to withdraw from the study at any point, without giving a reason. All information we collect will be kept confidential and securely stored. All information collected for the project will be destroyed when no longer needed.

This project conforms to the ethical guidelines of the British Psychological Society and the University of Surrey Ethics Committee. The researcher working with children has an up-to-date enhanced disclosure from the Criminal Records Bureau, and extensive experience of working with school-age children. I will be happy to answer any further questions regarding the study. Please contact Naomi Worsfold on 01483 686945, or N.Worsfold@surrey.ac.uk.

Appendix 13: Consent form

UNIVERSITY OF
SURREY

The perceptual restoration of music in adults and children

- I the undersigned voluntarily agree to take part in the study on the perceptual restoration of music.
- I have read and understood the Information Sheet provided. I have been given a full explanation by the investigators of the nature, purpose, location, and likely duration of the study, and of what I will be expected to do. I have been given the opportunity to ask questions on all aspects of the study and these have been answered to my satisfaction.
- I agree to comply with any instruction given to me during the study and to co-operate fully with the investigators. I shall inform them immediately if I experience any discomfort.
- I understand that all personal data relating to volunteers is held and processed in the strictest confidence, and in accordance with the Data Protection Act (1998). I agree that I will not seek to restrict the use of the results of the study on the understanding that my anonymity is preserved.
- I understand that I am free to withdraw from the study at any time without needing to justify my decision and without prejudice.
- I confirm that I have read and understood the above and freely consent to participating in this study. I have been give adequate time to consider my participation and agree to comply with the instructions of the study.

Name of volunteer (BLOCK CAPITALS______________________

Signed______________________

Date ______________________

Name of researcher (BLOCK CAPITALS)______________________

Signed ______________________

Date ______________________

Appendix 14: Debrief sheet

UNIVERSITY OF SURREY

The perceptual restoration of music
Debrief sheet

Thank you for taking part in this study. We have been looking at how children are able to reconstruct missing pieces of musical tunes when the absence of the missing sections is hidden with noise. We have compared children's ability to identify tunes when they are intact, when missing sections are replaced by noise, and when missing sections are replaced with silence.

The data we collect from adult participants will be compared to our data from child participants to see how the ability to reconstruct missing sections of music differs between adults and children.

You have been assigned a participant number and from this point on your data will only be identified by this number. We asked for your date of birth so that we can analyse the effect of age in more detail.

All data is stored securely and will be destroyed when no longer needed.

Please address any further questions or concerns about your participation in this study to Naomi Worsfold:

Email: N.Worsfold@surrey.ac.uk
Telephone: +44 (0)1483 686945

Appendix 15: Adult song familiarity questionnaire

Song Familiarity Questionnaire

UNIVERSITY OF SURREY

Please indicate whether you are familiar with the following
children's nursery rhymes by ticking the box if you know the words and/or the tune.

Name of Song	***Know the words (please tick)***	***Know the tune (please tick)***
Baa Baa Black Sheep		
The Grand Old Duke of York		
Twinkle Twinkle Little Star		
Happy Birthday to You		
The Wheels on The Bus		
The Hokey Cokey		
Row, Row, Row Your Boat		
Old Macdonald Had a Farm		
London Bridge		
Jack and Jill		
Incy Wincy Spider		
Humpty Dumpty		
Hickory Dickory Dock		
If You're Happy and You Know It		
Head, Shoulders, Knees and Toes		

Appendix 16: Results of the Musical Characteristics Analysis reported in Chapter 11

Experiment 1

The stimuli presented in Experiment 1 consisted of piano melodies of the first phrase of each nursery rhyme.

1.1 Values of musical characteristics for the thirteen experimental stimuli presented in Experiments 1-3, 5 & 8, and some stimuli in experiments 6 & 7.

Stimulus	Number of tone onsets	Presentation rate	Average note duration	Range	Average interval size	Contour complexity	Number of note durations	Rhythmic complexity	Tonal strength	Tonal strength (2)
Baa Baa black sheep	16	2.71	369	12	1.6	0.06	3	87.5	.862	4.69
The grand old duke of York	13	2.48	404	12	2.25	0.15	3	100	.569	4.60
Twinkle Twinkle	14	2.43	411	9	1.38	0.07	2	100	.816	3.84
Happy Birthday	12	2.53	400	7	2.45	0.58	3	91.67	.520	4.79
The wheels on the bus	15	2.82	355	12	2.14	0.20	3	100	.661	4.87
Hokey Cokey	20	3.08	325	5	1.68	0.50	3	95	.334	5.06
Row your boat	10	2.35	425	7	1.1	0.20	4	90	.716	4.84
Old Macdonald	12	2.67	375	9	2.0	0.33	3	100	.722	4.90
London Bridge	13	2.89	346	7	1.83	0.46	4	84.62	.569	4.42
Jack and Jill	16	3.64	274	7	0.8	0.06	3	87.50	.594	4.66
Incy Wincy spider	11	2.20	455	4	1.6	0.27	4	72.73	.494	5.03
Humpty Dumpty	16	2.00	500	10	2.67	0.56	4	75	.771	4.44
Hickory Dickory Dock	14	2.55	393	8	1.85	0.14	3	92.86	.766	4.61
Happy and you know it	22	3.26	307	9	1.36	0.27	2	100	.532	4.84
Head, shoulders knees and toes	9	2.77	361	5	1.5	0.44	2	100	.384	5.36
Mean	**14.2**	**2.69**	**380**	**8.2**	**1.75**	**0.29**	**3.07**	**91.79**	**.621**	**4.73**

1.2. The number of children giving a correct identification response (out of the 22 children in each condition) for each of the 13 experimental stimuli in Experiment 1.

Stimulus	**Number of children giving a correct response (out of 22)**		
	Intact	**Noise**	**Silence**
Twinkle Twinkle	21	22	16
Happy Birthday	21	17	7
Wheels on the Bus	14	13	6
Hokey Cokey	11	9	5
Row, Row Row Your Boat	16	15	2
Old Macdonald	21	19	9
London Bridge	23	9	5
Jack and Jill	9	3	0
Incy Wincy Spider	7	8	3
Humpty Dumpty	15	9	2
Hickory Dickory Dock	12	4	3
If You're Happy and You Know It	10	8	11
Head, Shoulders, Knees and Toes	12	6	6

1.3. Kendall's Tau correlation coefficients and significance values for identification and musical characteristics for Experiment 1.

		Identification score		
		Intact	**Noise**	**Silence**
Number of tone onsets	Correlation coefficient	-.128	-.148	.013
	Sig. (2-tailed)	.320	.496	.951
Presentation rate	Correlation coefficient	-.400	-.290	.079
	Sig. (2-tailed)	.063	.176	.712
Average note duration	Correlation coefficient	.400	.290	-.079
	Sig. (2-tailed)	.063	.176	.712
Range	Correlation coefficient	.357	.282	.226
	Sig. (2-tailed)	.113	.207	.313
Average interval size	Correlation coefficient	.374	.237	.079
	Sig. (2-tailed)	.083	.268	.712
Contour complexity	Correlation coefficient	.135	.173	.053
	Sig. (2-tailed)	.535	.422	.805
Number of different note durations	Correlation coefficient	-.048	-.016	-.585
	Sig. (2-tailed)	.841	.947	.014
Rhythmic complexity	Correlation coefficient	.243	.254	.606
	Sig. (2-tailed)	.278	.253	.006
Tonal strength	Correlation coefficient	.427	.290	.000
	Sig. (2-tailed)	.148	.176	1.000
Tonal strength (2)	Correlation coefficient	-.201	-.146	.040
	Sig. (2-tailed)	.353	.498	.853

Experiment 2

The stimuli presented in Experiment 2 were the same as those presented in Experiment 1, thus musical characteristics are shown in 1.1 above.

2.1. The number of children giving a correct identification response (out of the 22 children in each condition) for each of the 13 experimental stimuli in Experiment 2.

Stimulus	**Number of children giving a correct response (out of 22)**			
	Noise 200	**Noise 300**	**Silence 200**	**Silence 300**
Twinkle Twinkle	22	17	14	9
Happy Birthday	19	14	13	0
Wheels on the Bus	12	12	6	0
Hokey Cokey	4	8	0	1
Row, Row Row Your Boat	5	5	0	0
Old Macdonald	15	16	7	3
London Bridge	6	9	2	2
Jack and Jill	5	1	0	0
Incy Wincy Spider	8	7	2	2
Humpty Dumpty	4	4	1	0
Hickory Dickory Dock	1	5	0	1
If You're Happy and You Know It	11	11	5	3
Head, Shoulders, Knees and Toes	6	4	1	0

2.2. Kendall's Tau correlation coefficients and significance values for identification and musical characteristics for Experiment 2.

		Identification score			
		Noise 200	**Silence 200**	**Noise 300**	**Silence 300**
Number of tone onsets	Correlation coefficient	-.133	-.069	.040	.119
	Sig. (2-tailed)	.537	.754	.853	.602
Presentation rate	Correlation coefficient	-.092	-.162	-.052	.058
	Sig. (2-tailed)	.668	.455	.806	.795
Average note duration	Correlation coefficient	.092	.162	.052	-.058
	Sig. (2-tailed)	.668	.455	.806	.795
Range	Correlation coefficient	.140	.275	.278	.094
	Sig. (2-tailed)	.529	.225	.209	.690
Average interval size	Correlation coefficient	.065	.244	.234	-.146
	Sig. (2-tailed)	.759	.262	.270	.516
Contour complexity	Correlation coefficient	-.040	.082	.053	-.163
	Sig. (2-tailed)	.853	.708	.806	.474
Number of different note durations	Correlation coefficient	-.299	-.260	-.250	-.211
	Sig. (2-tailed)	.206	.281	.288	.400
Rhythmic complexity	Correlation coefficient	.294	.261	.389	.219
	Sig. (2-tailed)	.184	.248	.077	.349
Tonal strength	Correlation coefficient	.039	.162	.104	.117
	Sig. (2-tailed)	.854	.455	.624	.604
Tonal strength (2)	Correlation coefficient	-.026	-.150	-.065	-.162
	Sig. (2-tailed)	.902	.493	.759	.474

Experiment 3

The stimuli presented in Experiment 3 were the same as those presented in Experiment 1, thus musical characteristics are shown in 1.1 above.

3.1. The number of children giving a correct identification response (out of the 21 children in each condition) for each of the 13 experimental stimuli in Experiment 3.

Stimulus	**Number of children giving a correct response (out of 21)**					
	High amplitude noise		**Low amplitude noise**		**Very low amplitude noise**	
	100ms	**300ms**	**100ms**	**300ms**	**100ms**	**300ms**
Twinkle Twinkle	19	19	16	12	12	9
Happy Birthday	19	17	13	7	13	4
Wheels on the Bus	14	12	4	1	4	0
Hokey Cokey	17	15	2	0	6	1
Row, Row Row Your Boat	11	7	4	0	1	2
Old Macdonald	17	16	11	7	11	4
London Bridge	13	12	3	4	5	1
Jack and Jill	14	8	3	0	2	1
Incy Wincy Spider	6	9	4	2	3	0
Humpty Dumpty	13	6	0	0	2	1
Hickory Dickory Dock	6	6	5	1	0	1
If You're Happy and You Know It	10	8	6	0	7	5
Head, Shoulders, Knees and Toes	9	4	3	0	6	0

3.2. Kendall’s Tau correlation coefficients and significance values for identification and musical characteristics for Experiment 3.

		Identification Score*					
		HAN 100	**HAN 300**	**LAN 100**	**LAN 300**	**VLAN 100**	**VLAN 300**
Number of tone onsets	Correlation coefficient	.135	.080	-.122	-.266	.000	.159
	Sig. (2-tailed)	.536	.711	.576	.241	1.00	.482
Presentation rate	Correlation coefficient	.066	-.013	-.160	-.246	.078	.000
	Sig. (2-tailed)	.758	.951	.458	.271	.713	1.00
Average note duration	Correlation coefficient	-.066	.013	.160	.246	-.078	.000
	Sig. (2-tailed)	.758	.951	.458	.271	.713	1.00
Range	Correlation coefficient	.213	.070	.243	.140	.014	.318
	Sig. (2-tailed)	.344	.753	.281	.551	.950	.172
Average interval size	Correlation coefficient	.093	.170	.027	.275	.052	-.057
	Sig. (2-tailed)	.666	.425	.902	.218	.806	.799
Contour complexity	Correlation coefficient	.148	.119	-.257	.015	.263	-.029
	Sig. (2-tailed)	.497	.579	.238	.948	.219	.898
Number of different note durations	Correlation coefficient	-.223	-.141	.369	-.035	-.437	-.289
	Sig. (2-tailed)	.350	.549	.123	.889	.063	.243
Rhythmic complexity	Correlation coefficient	.213	.154	.372	.093	.334	.288
	Sig. (2-tailed)	.340	.486	.097	.689	.129	.213
Tonal strength	Correlation coefficient	.146	-.039	.267	.217	-.208	.283
	Sig. (2-tailed)	.498	.854	.216	.331	.327	.203
Tonal strength (2)	Correlation coefficient	-.133	-.053	-.148	-.248	.065	-.299
	Sig. (2-tailed)	.537	.806	.495	.270	.759	.181

*HAN= High amplitude noise; LAN = low amplitude noise; VLAN = very low amplitude noise

Experiment 4

Experiment 4 presented later phrases from each nursery rhyme.

4.1. Values of musical characteristics for the thirteen experimental stimuli presented in Experiment 4.

Song	number of tone onsets	Presentation rate	Average note duration	range	Average interval size	Contour complexity	Number of note durations	Rhythmic complexity	Tonal strength	Tonal strength (2)
Baa Baa black sheep	22	3.73	268	9	1.19	0.18	4	90.91	.675	4.47
The grand old duke of York	17	3.09	324	6	0.88	0.24	3	94.12	.557	4.30
Twinkle Twinkle	14	2.37	421	5	0.77	0.14	2	100	.479	4.40
Happy Birthday	13	2.45	408	12	3.17	0.38	3	92.31	.862	4.74
The wheels on the bus	12	2.40	417	12	3.00	0.33	4	75.00	.668	5.49
Hokey Cokey	18	3.21	311	5	1.33	0.61	4	83.33	.239	3.47
Row your boat	17	4.47	224	12	1.63	0.12	3	94.12	.770	5.31
Old Macdonald	21	3.82	262	5	0.75	0.10	2	100	.564	6.18
London Bridge	11	2.75	364	9	2.70	0.46	5	72.73	.781	4.65
Jack and Jill	16	3.90	256	8	1.13	0.31	3	87.50	.651	4.45
Incy Wincy spider	10	2.17	460	4	1.22	0.30	4	90.00	.405	4.54
Humpty Dumpty	20	2.67	375	13	1.89	0.40	3	95.00	.785	4.53
Hickory Dickory Dock	15	2.42	413	12	1.40	0.20	3	93.33	.727	4.57
Happy and you know it	27	4.03	248	9	1.33	0.22	2	100	.672	4.25
Head, shoulders knees and toes	9	3.00	.333	14	2.63	.333	3	88.88	.686	5.00
Mean	**14.2**	**2.69**	**380**	**8.2**	**1.75**	**0.29**	**3.07**	**91.79**	**.621**	**4.73**

4.2. The number of children giving a correct identification response (out of the 22 children in each condition) for each of the 13 experimental stimuli in Experiment 4.

Stimulus	**Number of children giving a correct response (out of 22)**		
	Intact	**Noise**	**Silence**
Twinkle Twinkle	20	11	12
Happy Birthday	21	12	10
Wheels on the Bus	14	9	2
Hokey Cokey	3	2	0
Row, Row Row Your Boat	7	7	1
Old Macdonald	13	6	7
London Bridge	14	6	3
Jack and Jill	11	4	1
Incy Wincy Spider	12	6	2
Humpty Dumpty	14	3	6
Hickory Dickory Dock	13	6	0
If You're Happy and You Know It	13	15	9
Head, Shoulders, Knees and Toes	15	6	4

4.3. Kendall's Tau correlation coefficients and significance values for identification and musical characteristics for Experiment 4.

		Identification score		
		Intact	**Noise**	**Silence**
Number of tone onsets	Correlation coefficient	-.320	-.055	.039
	Sig. (2-tailed)	.138	.801	.854
Presentation rate	Correlation coefficient	-.347	-.082	-.092
	Sig. (2-tailed)	.108	.705	.668
Average note duration	Correlation coefficient	.347	.082	.092
	Sig. (2-tailed)	.108	.705	.668
Range	Correlation coefficient	.343	.029	-.014
	Sig. (2-tailed)	.128	.897	.950
Average interval size	Correlation coefficient	.363	.180	.000
	Sig. (2-tailed)	.094	.412	1.000
Contour complexity	Correlation coefficient	.107	-.302	-.065
	Sig. (2-tailed)	.621	.166	.759
Number of different note durations	Correlation coefficient	-.215	-.284	-.391
	Sig. (2-tailed)	.360	.231	.092
Rhythmic complexity	Correlation coefficient	.109	.294	.400
	Sig. (2-tailed)	.618	.183	.064
Tonal strength	Correlation coefficient	.347	.137	.170
	Sig. (2-tailed)	.108	.529	.425
Tonal strength (2)	Correlation coefficient	.187	.192	.039
	Sig. (2-tailed)	.386	.378	.854

Experiment 5

In terms of musical characteristics, the stimuli presented in Experiment 5 were identical to those presented in Experiment 1.

5.1. The number of children giving a correct identification response (out of the 20 children in each condition) for each of the 13 experimental stimuli in Experiment 5.

Stimulus	**Number of children giving a correct response (out of 20)**					
	Lyrics & melody			**Melody Only**		
	Intact	**Noise**	**Silence**	**Intact**	**Noise**	**Silence**
Twinkle Twinkle	20	18	13	18	16	13
Happy Birthday	20	20	11	19	15	6
Wheels on the Bus	20	12	1	17	11	2
Hokey Cokey	12	9	1	13	10	2
Row, Row Row Your Boat	17	14	3	11	9	1
Old Macdonald	20	19	10	18	15	6
London Bridge	15	9	2	7	8	1
Jack and Jill	19	18	12	7	4	0
Incy Wincy Spider	20	15	7	10	6	1
Humpty Dumpty	20	11	5	10	6	3
Hickory Dickory Dock	15	12	3	4	4	1
If You're Happy and You Know It	18	19	12	7	8	4
Head, Shoulders, Knees and Toes	19	11	4	9	9	1

5.2. Kendall's Tau correlation coefficients and significance values for identification and musical characteristics for Experiment 5.

		Identification Score					
		Lyrics & Melody			**Melody Only**		
		Intact	**Noise**	**Silence**	**Intact**	**Noise**	**Silence**
Number of tone onsets	Correlation coefficient	-.163	-.027	.053	-.187	-.123	.196
	Sig. (2-tailed)	.474	.901	.805	.388	.575	.379
Presentation rate	Correlation coefficient	-.362	-.066	-.092	-.196	-.067	-.137
	Sig. (2-tailed)	.105	.758	.668	.357	.757	.532
Average note duration	Correlation coefficient	.362	.066	.092	.196	.067	.137
	Sig. (2-tailed)	.105	.758	.668	.357	.757	.532
Range	Correlation coefficient	.311	.142	.098	.000	.115	.338
	Sig. (2-tailed)	.185	.528	.660	1.00	.612	.141
Average interval size	Correlation coefficient	.275	-.119	-.196	.222	.094	.384
	Sig. (2-tailed)	.218	.579	.357	.297	.665	.080
Contour complexity	Correlation coefficient	.015	-.175	-.172	.146	.163	.306
	Sig. (2-tailed)	.948	.442	.423	.498	.456	.168
Number of different note durations	Correlation coefficient	-.017	-.255	-.346	-.094	-.242	-.347
	Sig. (2-tailed)	.944	.286	.143	.690	.315	.156
Rhythmic complexity	Correlation coefficient	.031	.142	.042	.196	.417	.353
	Sig. (2-tailed)	.894	.525	.849	.375	.064	.122
Tonal strength	Correlation coefficient	.217	.093	.222	.039	-.013	.165
	Sig. (2-tailed)	.331	.666	.297	.854	.951	.453
Tonal strength (2)	Correlation coefficient	.000	-.027	-.237	.132	.108	-.152
	Sig. (2-tailed)	1.00	.902	.269	.539	.620	.491

Experiment 6

Old lyrics, New Melody

The melodies for these stimuli were newly composed.

6.1. Values of musical characteristics for the six old lyrics stimuli presented in Experiment 6.

song	number of tone onsets	Presentation rate	Average note duration	range	Average interval size	Contour complexity	Number of note durations	Rhythmic complexity	Tonal strength	Tonal strength (2)
Twinkle Twinkle	14	2.37	421	7	2.15	.29	4	85.71	.606	4.48
Happy Birthday	12	2.53	396	7	1.81	.42	3	100.00	.752	4.79
The wheels on the bus	15	2.26	443	12	2.29	.27	4	86.67	.746	4.53
Hokey Cokey	20	3.45	290	5	1.37	.30	3	100.00	.439	4.21
Old Macdonald	12	2.61	383	9	2.09	.25	4	91.67	.870	4.75
London Bridge	13	2.00	500	7	1.50	.23	4	100.00	.557	4.57
Mean	**14.33**	**2.54**	**405.50**	**7.83**	**1.87**	**0.29**	**3.67**	**94.01**	**.662**	**4.56**

6.2. The number of children giving a correct identification response (out of the 18 children in each condition) for each of the 6 old lyrics stimuli in Experiment 6.

Stimulus	**Number of children giving a correct response (out of 18)**		
	Intact	**Noise**	**Silence**
Twinkle Twinkle	18	4	8
Happy Birthday	18	8	12
Wheels on the Bus	15	8	0
Hokey Cokey	7	0	0
Old Macdonald	16	7	8
London Bridge	13	6	2

6.3. Kendall's Tau correlation coefficients and significance values for identification and musical characteristics for old lyrics stimuli in Experiment 6.

		Identification score		
		Intact	**Noise**	**Silence**
Number of tone onsets	Correlation coefficient	-.357	-.071	-.445
	Sig. (2-tailed)	.330	.846	.234
Presentation rate	Correlation coefficient	.138	.138	.358
	Sig. (2-tailed)	.702	.702	.330
Average note duration	Correlation coefficient	-.138	-.138	.358
	Sig. (2-tailed)	.702	.702	.330
Range	Correlation coefficient	.000	.386	-.160
	Sig. (2-tailed)	1.000	.306	.677
Average interval size	Correlation coefficient	.276	.276	-.072
	Sig. (2-tailed)	.444	.444	.845
Contour complexity	Correlation coefficient	.414	.414	.845
	Sig. (2-tailed)	.251	.251	.173
Number of different note durations	Correlation coefficient	-.094	-.283	-.392
	Sig. (2-tailed)	.814	.481	.340
Rhythmic complexity	Correlation coefficient	-.386	-.077	.000
	Sig. (2-tailed)	.306	.838	1.000
Tonal strength	Correlation coefficient	.414	.414	.358
	Sig. (2-tailed)	.251	.251	.330
Tonal strength (2)	Correlation coefficient	.276	.276	.358
	Sig. (2-tailed)	.444	.444	.330

Old Melody, new lyrics

The melodies for these stimuli were identical to those in Experiment 1.

6.4. The number of children giving a correct identification response (out of the 18 children in each condition) for each of the 6 old melody stimuli in Experiment 6.

Stimulus	**Number of children giving a correct response (out of 18)**		
	Intact	Noise	**Silence**
Twinkle Twinkle	8	13	1
Happy Birthday	11	10	3
Wheels on the Bus	8	11	2
Hokey Cokey	6	5	1
Old Macdonald	8	13	1
London Bridge	8	11	3

6.5. Kendall's Tau correlation coefficients and significance values for identification and musical characteristics for old melody stimuli in Experiment 6.

		Identification score		
		Intact	**Noise**	**Silence**
Number of tone onsets	Correlation coefficient	-.713	-.296	-.322
	Sig. (2-tailed)	.067	.428	.404
Presentation rate	Correlation coefficient	-.602	-.501	.078
	Sig. (2-tailed)	.114	.173	.837
Average note duration	Correlation coefficient	.602	.501	-.078
	Sig. (2-tailed)	.114	.173	.837
Range	Correlation coefficient	.185	.615	-.084
	Sig. (2-tailed)	.642	.107	.832
Average interval size	Correlation coefficient	.602	-.215	.545
	Sig. (2-tailed)	.114	.559	.150
Contour complexity	Correlation coefficient	.086	-.645	.389
	Sig. (2-tailed)	.822	.079	.304
Number of different note durations	Correlation coefficient	.000	-.277	.603
	Sig. (2-tailed)	1.000	.485	.142
Rhythmic complexity	Correlation coefficient	-.096	.400	-.609
	Sig. (2-tailed)	.810	.298	.125
Tonal strength	Correlation coefficient	.086	.931	-.389
	Sig. (2-tailed)	.822	.011	.304
Tonal strength (2)	Correlation coefficient	-.430	-.358	-.389
	Sig. (2-tailed)	.260	.330	.304

Mismatch

The melodies for these stimuli were identical to those in Experiment 1.

6.6. The number of children giving a correct identification response (out of the 18 children in each condition) for each of the 6 mismatch stimuli in Experiment 6.

Stimulus	**Number of children giving a correct response (out of 18)**		
	Intact	**Noise**	**Silence**
Twinkle Twinkle	17	14	5
Happy Birthday	14	11	5
Wheels on the Bus	17	15	3
Hokey Cokey	16	11	2
Old Macdonald	14	12	4
London Bridge	18	15	6

6.7. Kendall's Tau correlation coefficients and significance values for identification and musical characteristics for mismatch stimuli in Experiment 6.

		Identification score		
		Intact	**Noise**	**Silence**
Number of tone onsets	Correlation coefficient	.222	.148	-.500
	Sig. (2-tailed)	.552	.692	.173
Presentation rate	Correlation coefficient	.215	.072	-.414
	Sig. (2-tailed)	.559	.845	.251
Average note duration	Correlation coefficient	-.215	-.072	.414
	Sig. (2-tailed)	.559	.845	.251
Range	Correlation coefficient	.077	.538	-.148
	Sig. (2-tailed)	.840	.158	.692
Average interval size	Correlation coefficient	-.358	-.072	.000
	Sig. (2-tailed)	.330	.845	1.000
Contour complexity	Correlation coefficient	-.358	-.501	.000
	Sig. (2-tailed)	.330	.173	1.00
Number of different note durations	Correlation coefficient	.185	.185	.178
	Sig. (2-tailed)	.642	.642	.647
Rhythmic complexity	Correlation coefficient	-.080	.160	-.386
	Sig. (2-tailed)	.835	.677	.306
Tonal strength	Correlation coefficient	.072	.358	.276
	Sig. (2-tailed)	.845	.330	.444
Tonal strength (2)	Correlation coefficient	-.358	-.358	-.690
	Sig. (2-tailed)	.330	.330	.056

Experiment 7

Old pitch, new rhythm

Whilst the pitch patterns of these stimuli are identical to those in experiment 1, the changed rhythmic patterns required new values of musical characteristics to be calculated.

7.1. Values of musical characteristics for the six old pitch stimuli presented in Experiment 7.

song	number of tone onsets	Presentation rate	Average note duration	range	Average interval size	Contour complexity	Number of note durations	Rhythmic complexity	Tonal strength	Tonal strength (2)
Twinkle Twinkle	14	2.33	429	9	1.38	0.07	4	78.57	.696	3.84
Happy Birthday	12	2.26	442	7	2.45	0.58	4	83.33	.507	4.79
The wheels on the bus	15	2.68	373	12	2.14	0.20	4	73.33	.661	4.87
Head, shoulders knees and toes	9	2.73	367	5	1.33	0.44	3	88.89	.304	5.36
Old Macdonald	12	2.40	417	9	2.00	0.33	4	75.00	.705	4.90
London Bridge	13	3.10	323	7	1.83	0.38	4	84.61	.569	4.42
Mean	**12.5**	**2.58**	**391.83**	**8.17**	**1.86**	**0.33**	**3.83**	**80.62**	**.574**	**4.70**

7.2. The number of children giving a correct identification response (out of the 20 children in each condition) for each of the 6 old pitch stimuli in Experiment 7.

Stimulus	**Number of children giving a correct response (out of 20)**		
	Intact	**Noise**	**Silence**
Twinkle Twinkle	9	2	1
Happy Birthday	11	5	2
Wheels on the Bus	8	8	5
Head, shoulders, knees and toes	10	4	5
Old Macdonald	12	6	5
London Bridge	2	2	2

7.3. Kendall's Tau correlation coefficients and significance values for identification and musical characteristics for old pitch stimuli in Experiment 7.

		Identification score		
		Intact	**Noise**	**Silence**
Number of tone onsets	Correlation coefficient	-.414	.071	-.322
	Sig. (2-tailed)	.251	.846	.404
Presentation rate	Correlation coefficient	-.467	-.138	.234
	Sig. (2-tailed)	.188	.702	.537
Average note duration	Correlation coefficient	.467	.138	-.234
	Sig. (2-tailed)	.188	.702	.537
Range	Correlation coefficient	-.072	.445	.000
	Sig. (2-tailed)	.845	.234	1.000
Average interval size	Correlation coefficient	.067	.414	.078
	Sig. (2-tailed)	.851	.251	.837
Contour complexity	Correlation coefficient	.200	.000	.078
	Sig. (2-tailed)	.573	1.000	.837
Number of different note durations	Correlation coefficient	-.115	.120	-.405
	Sig. (2-tailed)	.770	.766	.343
Rhythmic complexity	Correlation coefficient	-.067	-.552	-.078
	Sig. (2-tailed)	.851	.126	.837
Tonal strength	Correlation coefficient	-.200	.138	-.078
	Sig. (2-tailed)	.573	.702	.837
Tonal strength (2)	Correlation coefficient	.333	.414	.856
	Sig. (2-tailed)	.348	.251	.024

Old rhythm, new pitch

New pitch patterns were composed for these stimuli.

7.4. Values of musical characteristics for the six old rhythm stimuli presented in Experiment 7.

song	number of tone onsets	Presentation rate	Average note duration	range	Average interval size	Contour complexity	Number of note durations	Rhythmic complexity	Tonal strength	Tonal strength (2)
Twinkle Twinkle	14	2.41	414	7	2.15	0.29	2	100.00	.520	4.48
Happy Birthday	12	2.45	408	7	1.73	0.42	3	91.67	.722	4.79
The wheels on the bus	15	2.83	353	12	2.29	0.27	3	100.00	.615	4.53
Head, shoulders knees and toes	9	2.81	356	5	1.56	0.33	2	100.00	.506	4.49
Old Macdonald	12	2.40	417	9	2.62	0.42	3	91.67	.853	4.92
London Bridge	13	2.89	346	7	1.50	0.23	4	84.62	.520	4.58
Mean	**12.5**	**2.63**	**382.33**	**7.83**	**1.98**	**0.33**	**2.83**	**94.66**	**.623**	**4.63**

7.5. The number of children giving a correct identification response (out of the 20 children in each condition) for each of the 6 old rhythm stimuli in Experiment 7.

Stimulus	**Number of children giving a correct response (out of 20)**		
	Intact	**Noise**	**Silence**
Twinkle Twinkle	4	0	0
Happy Birthday	3	0	1
Wheels on the Bus	2	1	0
Hokey Cokey	2	0	0
Old Macdonald	2	0	2
London Bridge	1	0	2

7.6. Kendall's Tau correlation coefficients and significance values for identification and musical characteristics for old rhythm stimuli in Experiment 7.

		Identification score		
		Intact	**Noise**	**Silence**
Number of tone onsets	Correlation coefficient	.077	.598	-.161
	Sig. (2-tailed)	.838	.137	.676
Presentation rate	Correlation coefficient	-.596	.346	-.078
	Sig. (2-tailed)	.107	.380	.837
Average note duration	Correlation coefficient	.447	-.115	.234
	Sig. (2-tailed)	.227	.770	.537
Range	Correlation coefficient	-.083	.645	.174
	Sig. (2-tailed)	.830	.120	.661
Average interval size	Correlation coefficient	.298	.346	-.078
	Sig. (2-tailed)	.421	.380	.837
Contour complexity	Correlation coefficient	.386	-.359	.161
	Sig. (2-tailed)	.306	.373	.676
Number of different note durations	Correlation coefficient	-.609	.135	.727
	Sig. (2-tailed)	.125	.752	.074
Rhythmic complexity	Correlation coefficient	.435	.405	-.909
	Sig. (2-tailed)	.273	.343	.026
Tonal strength	Correlation coefficient	.231	.120	.322
	Sig. (2-tailed)	.539	.766	.404
Tonal strength (2)	Correlation coefficient	-.298	-.115	.701
	Sig. (2-tailed)	.421	.770	.064

Mismatch

New musical characteristics were calculated for the novel pairing of pitch and rhythm patterns.

7.7. Values of musical characteristics for the six mismatch stimuli presented in Experiment 7.

song	number of tone onsets	Presentation rate	Average note duration	range	Average interval size	Contour complexity	Number of note durations	Rhythmic complexity	Tonal strength	Tonal strength (2)
Twinkle Twinkle	16	2.00	500	9	1.38	0.07	4	75.00	.816	3.84
Happy Birthday	12	2.55	392	7	2.45	0.58	2	100.00	.557	4.79
The wheels on the bus	16	2.71	369	12	2.14	0.20	3	87.50	.661	4.87
Head, shoulders knees and toes	10	2.33	430	5	1.33	0.44	4	90.00	.359	5.36
Old Macdonald	12	2.53	396	9	2.00	0.33	3	91.67	.582	4.90
London Bridge	13	2.50	400	7	1.83	0.38	3	100.00	.569	4.42
Mean	**13.17**	**2.44**	**414.50**	**8.17**	**1.86**	**0.33**	**3.17**	**90.70**	**.591**	**4.70**

7.8. The number of children giving a correct identification response (out of the 20 children in each condition) for each of the 6 mismatch stimuli in Experiment 7.

Stimulus	**Number of children giving a correct response (out of 20)**		
	Intact	**Noise**	**Silence**
Twinkle Twinkle	11	9	7
Happy Birthday	13	11	4
Wheels on the Bus	14	5	7
Hokey Cokey	15	7	4
Old Macdonald	15	9	7
London Bridge	8	0	0

7.9. Kendall's Tau correlation coefficients and significance values for identification and musical characteristics for mismatch stimuli in Experiment 7.

		Identification score		
		Intact	**Noise**	**Silence**
Number of tone onsets	Correlation coefficient	-.445	-.148	.334
	Sig. (2-tailed)	.234	.692	.396
Presentation rate	Correlation coefficient	.138	.000	.078
	Sig. (2-tailed)	.702	1.000	.837
Average note duration	Correlation coefficient	-.138	.000	-.078
	Sig. (2-tailed)	.702	1.000	.837
Range	Correlation coefficient	.000	-.074	.669
	Sig. (2-tailed)	1.000	.843	.090
Average interval size	Correlation coefficient	-.138	.276	.078
	Sig. (2-tailed)	.702	.444	.837
Contour complexity	Correlation coefficient	.138	.138	-.545
	Sig. (2-tailed)	.702	.702	.150
Number of different note durations	Correlation coefficient	.161	-.322	.182
	Sig. (2-tailed)	.676	.404	.656
Rhythmic complexity	Correlation coefficient	-.071	.357	-.645
	Sig. (2-tailed)	.846	.330	.095
Tonal strength	Correlation coefficient	-.276	.000	.545
	Sig. (2-tailed)	.444	1.000	.150
Tonal strength (2)	Correlation coefficient	.828	-.276	.078
	Sig. (2-tailed)	.022	.444	.837

Experiment 8

The stimuli presented in Experiment 8 were identical to those in Experiment 1

8.1. The number of participants giving a correct identification response for each of the 13 stimuli in Experiment 8.

Stimulus	**Number of participants giving a correct response**					
	Older children (out of 20)			**Adults (out of 10)**		
	Intact	**Noise**	**Silence**	**Intact**	**Noise**	**Silence**
Twinkle Twinkle	20	19	20	10	10	7
Happy Birthday	20	20	19	10	10	9
Wheels on the Bus	14	17	7	7	10	6
Hokey Cokey	18	19	11	8	8	6
Row, Row Row Your Boat	17	17	8	10	10	3
Old Macdonald	20	19	16	10	9	10
London Bridge	19	16	13	9	10	6
Jack and Jill	6	11	2	7	7	1
Incy Wincy Spider	14	7	4	8	6	6
Humpty Dumpty	4	4	2	6	3	1
Hickory Dickory Dock	5	1	1	4	3	2
If You're Happy and You Know It	15	16	14	9	9	7
Head, Shoulders, Knees and Toes	15	9	11	10	6	4

8.2. Kendall's Tau correlation coefficients and significance values for identification and musical characteristics for stimuli in Experiment 8.

		Identification Score					
		Older children			**Adults**		
		Intact	Noise	**Silence**	**Intact**	Noise	**Silence**
Number of tone onsets	Correlation coefficient	-.203	-.027	-.079	-.387	-.057	-.083
	Sig. (2-tailed)	.352	.901	.712	.085	.799	.707
Presentation rate	Correlation coefficient	.013	.093	.052	-.098	.070	.108
	Sig. (2-tailed)	.951	.666	.806	.657	.751	.618
Average note duration	Correlation coefficient	-.013	-.093	-.052	.098	-.070	-.108
	Sig. (2-tailed)	.951	.666	.806	.657	.751	.618
Range	Correlation coefficient	-.057	.085	-.070	-.135	.165	.043
	Sig. (2-tailed)	.800	.705	.754	.559	.475	.848
Average interval size	Correlation coefficient	.013	.040	.026	-.183	-.042	.135
	Sig. (2-tailed)	.951	.853	.903	.410	.849	.534
Contour complexity	Correlation coefficient	.228	.134	.211	.128	-.057	.219
	Sig. (2-tailed)	.293	.536	.325	.567	.799	.318
Number of different note durations	Correlation coefficient	-.207	-.191	-.343	-.236	-.051	-.277
	Sig. (2-tailed)	.385	.423	.144	.337	.837	.252
Rhythmic complexity	Correlation coefficient	.270	.312	.362	.301	.226	.391
	Sig. (2-tailed)	.227	.162	.100	.191	.327	.083
Tonal strength	Correlation coefficient	-.119	-.013	-.078	-.126	.042	-.135
	Sig. (2-tailed)	.579	.951	.713	.568	.849	.534
Tonal strength (2)	Correlation coefficient	.053	.080	-.013	.141	-.170	.095
	Sig. (2-tailed)	.805	.711	.951	.526	.446	.662

Harmony analysis

*9.1.*Results of Kruskall-Wallis and Mann-Whitney tests for differences in stimulus identification as a result of scale degree of start note and end cadence of the melodies.

	Beginning phrases (Experiment 1)	**Later phrases (Experiment 4)**	**New melodies (experiment 6 & 7)**
Intact: scale degree	χ^2 (2) = .099, p=.952	χ^2 (2) = .199, p=.905	n/a*
Intact: end cadence	U=18.00, p=.731	U=25.50, p=.665	U=1.50, p=.267
Noise: scale degree	χ^2 (2) = 1.025, p=.599	χ^2 (2) = .249, p=.883	n/a*
Noise: end cadence	U=20.00, p=.945	U=28.00, p=.885	U=1.50, p=.267
Silence: scale degree	χ^2 (2) = 2.914, p=.233	χ^2 (2) = .891, p=.640	n/a*
Silence: end cadence	U=6.500, p=.035	U=24.50, p=.596	U=3.00, p=.800

*Analysis by scale degree of start note is not carried out for experiment 6/7 melodies as all began on the same note (tonic).