Loughborough
University

This item was submitted to Loughborough's Institutional Repository (https://dspace.lboro.ac.uk/) by the author and is made available under the following Creative Commons Licence conditions.

For the full text of this licence, please go to:
http://creativecommons.org/licenses/by-nc-nd/2.5/

Intelligence and negation biases on the Conditional Inference Task:

A dual-processes analysis

Nina Attridge[1][*] and Matthew Inglis[2]

[1]Centre for Pain Research,

Department for Health, University of Bath, UK

n.f.attridge@bath.ac.uk

[2]Mathematics Education Centre

Loughborough University, UK

m.j.inglis@lboro.ac.uk

Introduction: 2139

Method: 549

Results: 1757

Discussion: 1591

Word count: 6039

*Corresponding author

Abstract (134 words)

We examined a large set of conditional inference data compiled from several previous studies and asked three questions: how is normative performance related to intelligence, does negative conclusion bias stem from Type 1 or Type 2 processing, and does implicit negation bias stem from Type 1 or Type 2 processing? Our analysis demonstrated that a) rejecting denial of the antecedent and affirmation of the consequent inferences was positively correlated with intelligence, while endorsing MT inferences was not, b) the occurrence of negative conclusion bias was related to the extent of Type 2 processing, and c) that the occurrence of implicit negation bias was not related to extent of Type 2 processing. We conclude that negative conclusion bias is, at least in part, a product of Type 2 processing while implicit negation bias is not.

Keywords: Negative conclusion bias, implicit negation bias, affirmative premise bias, conditional inference, dual-processes.

Intelligence and negation biases on the Conditional Inference Task:

A dual-processes analysis

Dual-process theories suggest that reasoning occurs through two distinct types of processing – heuristic (or Type 1) processing and analytic (or Type 2) processing. Type 1 processing is fast and undemanding of working memory, while Type 2 processing is effortful and reliant on working memory (Evans, 2003; Evans, 2007; Evans, Handley & Bacon, 2009). Both types of processing can contribute to biased reasoning as well as normative reasoning (Evans & Stanovich, 2013). Biases thought to stem from Type 1 thinking include belief bias (Evans, Barston & Pollard, 1983; Sá, West & Stanovich, 1999) and matching bias (Evans, 1998; Wason and Evans, 1975), while those thought to stem from Type 2 thinking include an individual's tendency to rationalise their current mental model of a task as opposed to critically evaluating it (Evans, 2007). The Conditional Inference Task is a popular measure of deductive reasoning, but it is unclear how Type 1 and Type 2 processing contribute to performance on the task, and in particular, to the biases commonly shown on it.

Conditional Inference Behaviour

The Conditional Inference Task requires participants to decide whether conclusions drawn from conditional rules and premises are logically valid or invalid, with inferences following four forms: Modus Ponens (MP: if p then q; p; therefore q), Denial of the Antecedent (DA: if p then q; not p; therefore not q), Affirmation of the Consequent (AC: if p then q; q; therefore p) and Modus Tollens (MT: if p then q; not q; therefore not

p). The normative theory of conditional logic dictates that MP and MT inferences should be endorsed, whilst DA and AC inferences should be rejected. However, very few participants consistently show this pattern of responses. Evans, Handley, Neilens and Over (2007) found that while MP was almost universally endorsed (97.5%), MT was only endorsed around 50% of the time. AC, which should be rejected according to logical norms, was actually endorsed around 74% of the time, and DA around 47%.

The contributions of Type 1 and Type 2 processing to performance on the Conditional Inference Task are unclear. In Newstead, Handley, Harley, Wright and Farrelly's (2004) sample, AH4 intelligence scores correlated positively and significantly with MP (r=.35), DA (r=.50) and AC scores (r=.43), and negatively but non-significantly with MT scores (r=-.14). In other words, higher capacity for Type 2 processing was related to more normative performance on MP, DA and AC inferences, suggesting that Type 2 processing is involved in deciding their validity or invalidity. In Evans et al's (2007) sample, the number of normative responses to MP inferences was again positively correlated with AH4 scores (r=.194). The number of normative responses to DA and AC, on the other hand, were not correlated with intelligence scores, although participants in the highest quartile of intelligence scores endorsed them less than participants in the lowest quartile. MT responses were neither correlated with intelligence scores nor were different in the highest and lowest quartiles of intelligence scores. Overall this suggests that Type 2 processing is implicated in normative responses to MP inferences, that it may be implicated in normative responses to DA and AC inferences, and that it is not implicated in MT inferences.

However, MT inferences may be endorsed through two distinct routes, which may mask any contribution of Type 2 processing. Evans et al (2007) suggested that high-ability participants endorse MT through a difficult process of suppositional reasoning, using reductio ad absurdum, which requires a high level of working memory, i.e., Type 2 processing (Braine & O'Brien, 1998; Evans & Over, 2004). Low-ability participants, on the other hand, can draw MT by a simple equivalence strategy, whereby they expect p and q to go together, and not p and not q to go together, and this requires very little, if any, Type 2 processing. If this is the case, then endorsement of MT cannot reliably tell us anything about the extent of participants' on-task Type 2 processing. Furthermore, participants with a defective interpretation of the conditional statement (who are likely to be of high intelligence, Evans et al, 2007) may consistently reject MT inferences despite using Type 2 processing, further blurring the picture.

The first aim of the current paper was to add empirical data on the relationship between conditional inference endorsement rates and intelligence scores using a larger sample than previously analysed and with a different measure of intelligence. We measured intelligence using Raven's Advanced Progressive Matrices, which is non-verbal and highly g-loaded (Jensen, 1998).

Biases in Conditional Inference

Negative Conclusion Bias (NCB, Pollard & Evans, 1980) is commonly observed on the Conditional Inference task and theorists have proposed both Type 1 process based (e.g. Oaksford, Chater & Larkin, 2000) and Type 2 process based (e.g. Evans & Handley, 1999) accounts for it. NCB refers to the tendency to endorse inferences with negative

conclusions more often than inferences with affirmative conclusions. Take for example

the two modus tollens inferences below:

If the letter is D then the number is not 8

The number is 8

Conclusion: The letter is not D.

If the letter is not D then the number is 8

The number is not 8

Conclusion: The letter is D

The former inference leads to a negative conclusion, and is more often endorsed by

participants relative to the latter inference, which leads to an affirmative conclusion. NCB

is robustly observed on denial of the antecedent (DA) inferences and modus tollens (MT)

inferences, occasionally on affirmation of the consequent (AC) inferences, and never on

modus ponens (MP) inferences (Evans et al, 2007).

Two accounts of NCB have been proposed, one based on Type 1 processes and

one based on Type 2 processes. The Type 1 account proposes a "caution" heuristic or

probability-based decision (Evans, Newstead & Byrne, 1993; Oaksford, Chater & Larkin,

2000; Pollard & Evans, 1980), where affirmative conclusions are seen as being unlikely

and falsifiable in many ways (there are more not-blue things than there are blue things),

whereas negative conclusions are falsifiable in only one way (the assumption not-blue is falsified only by a blue example).

The Type 2 account proposes a "double negation" effect (Evans, Clibbens & Rood, 1995; Evans & Handley, 1999), which only applies to DA and MT items, where NCB is most commonly observed. DA and MT are both denial inferences with negative conclusions, except where there is a negative antecedent (in MT) or negative consequent (in DA), in which case the conclusion becomes affirmative. However, the affirmative conclusions are not straightforward; rather, they result from the conversion of a double negation into an affirmation (for example, converting 'not not p' into 'p'). This issue can be demonstrated with the DA inferences below:

If the letter is T then the number is 6

The letter is not T

Conclusion: The number is not 6

If the letter is T then the number is not 6

The letter is not T

Conclusion: The number is 6

There are three steps required to reach the conclusion in the first example, but four steps in the second:

If the letter is T then the number is not 6

The letter is not T

Conclusion: The number is not not-6

Conclusion: The number is 6

Reasoners may have difficulty in completing the extra logical step, possibly due to working memory constraints or lack of motivation, and thereby endorse the inference less often than the three-step equivalent.

There are arguments for both the Type 1 and Type 2 process based accounts of NCB. The double negation account can only explain NCB on DA and MT inferences (Evans & Handley, 1999), but the bias is also sometimes observed on AC inferences (Evans et al, 2007). On the other hand, NCB is never observed on MP inferences, and one might expect that it would be if the bias stemmed from heuristic processes (Evans, 1977). Furthermore, NCB is also not observed on 'p only if q' inferences (Evans, 1977), and again if it were a heuristic bias then it should be. Also in favour of the Type 2 processing account, Schroyens, Schaeken, Fias, & d'Ydewalle (2000) showed that response latencies were longer on problems that required a double negation than those that did not, suggesting that slow Type 2 processing was being used. Furthermore, NCB has been shown to be absent in children and to increase with development, as would be expected from the Type 2 processing account (Lott, 1999).

The second aim of the current paper is to add further empirical data that can differentiate between the two accounts of NCB. We analysed a large set of conditional inference data collated from several published and unpublished studies (Attridge & Inglis, 2011; Attridge & Inglis, 2013; Attridge & Inglis, in prep; Gilmore, Attridge & Inglis, 2011; Inglis & Simpson, 2008). We predicted that, if NCB is a Type 2 processing bias, then its occurrence should be related to patterns of individuals' overall responses – in other words, that the extent to which an individual's responses are normative (which we use as a proxy for Type 2 thinking – supported below by correlations with intelligence scores) should impact on the prevalence of NCB. If, on the other hand, NCB is a Type 1 processing bias, then its occurrence should be unrelated to the extent of normative responding.

To elaborate on these assumptions, we do not wish to argue that the number of normative responses is a direct measure of Type 2 processing. However, previous data and the analyses presented below suggest that normative responding to some inferences is related to intelligence scores, which in turn reflect capacity for Type 2 thinking. As such, normative responding may reflect the extent of Type 2 processing that is engaged on the task. We also report data on the relationship between measures of intelligence and NCB directly. Performance on intelligence tests is widely considered to measure the capacity and efficiency of Type 2 thinking, so although scores do not reflect Type 2 thinking on the Conditional Inference Task, they do reflect capacity for Type 2 thinking in general.

A relationship between the number of normative responses and NCB, if it existed, would likely take the form of an n-shaped quadratic curve. We predicted that, if NCB is a Type 2 processing bias, then as the number of normative responses increased so should NCB, but only up to a point. Low normative responding may indicate a lack of meaningful Type 2 engagement with the task, and if this is the case then NCB should also be low, as under the Type 2 hypothesis, NCB can only influence responses at the final stage of processing when a double negation is required. In other words, if NCB stems from Type 2 thinking and Type 2 thinking is not consistently present, then NCB should be low. At the other end of the scale, when responses are entirely normative, there would be no room for biases to be displayed – perfectly normative responding and NCB are mutually exclusive. In between these two extremes, we predicted that NCB would increase – when responding was based on Type 2 processing but did not perfectly conform with the normative model of the conditional there would be room for NCB to occur, and if it stemmed from Type 2 processing, then it would.

In addition to NCB, we also considered the implicit negation bias, previously considered to be an affirmative premise bias (Evans & Handley, 1999). Implicit negation bias (INB) refers to the tendency for reasoners to endorse more inferences with explicit negations in the minor premise than those with implicit negations in the minor premise. For example, the inference:

If the letter is D then the number is 9

The number is not 9

Conclusion: the letter is not D

would be endorsed more often that the inference:

If the letter is D then the number is not 9

The number is 7

Conclusion: the letter is not D

despite both being valid MT inferences. It has been proposed that this bias stems from the failure to recognise that implicitly negated premises are relevant to the rule. This is also thought to be the source of matching bias on the Wason Selection Task and Truth Table Task (Evans & Handley, 1999). Importantly, we are not aware of any arguments for INB (or affirmative premise bias) being the result of Type 2 processing. This makes it a useful comparison for our examination of NCB. We predicted that an implicit negation index (INI) on the conditional inference task would be unrelated to normative responding because it does not stem from Type 2 processing, while normative responding does.

To summarise, we predicted that INI would be unrelated to a Normative Responses Index (NRI) and that NCB would also be unrelated to the NRI if it were a Type 1 processing bias, and related to the NRI if it were a Type 2 processing bias.

Method

Participants

Data from 822 participants across five studies were collated (Attridge & Inglis, 2011; Attridge & Inglis, 2013; Attridge & Inglis, in prep; Gilmore, Attridge & Inglis, 2011; Inglis & Simpson, 2008). Participants were aged 16 – 48 and were either studying for A levels (the post-compulsory school-level qualification in England and Wales) or were staff or students of a university in the UK. None of the participants were recruited from courses with an explicit logic component (although some were recruited during the first week of an undergraduate mathematics course) and no participants had previously received formal training in conditional logic. Participants' native language was not recorded, but since they were all studying or working at a UK university they likely had a good familiarity with English.

Measures

Participants completed the abstract Conditional Inference Task (Evans, Clibbens & Rood, 1995). The task consists of 32 abstract items of four inference types: MP, DA, AC and MT. The inferences used are shown in Table 1; half of the problems used explicitly negated premises (e.g. not-4 was represented as 'not 4') and half used implicitly negated premises (e.g. not-4 was represented as, for example, '8'). The lexical content of the rules was generated randomly and the order of the problems was randomised for each participant. The instructions used were identical to those used by Evans et al. An example item which asks about the MT inference with the rule 'if not-*p*

then *q*' and an implicitly negated premise 'not-*q*' is as follows:

> This problem concerns an imaginary letter-number pair. Your task is to decide
>
> whether or not the conclusion necessarily follows from the rule and the premise.
>
> *Rule:* If the letter is not H then the number is 1.
>
> *Premise:* The number is 8.
>
> *Conclusion:* The letter is H.
>
> YES (it follows)   NO (it does not follow)

We measured the number of normative responses to MP, DA, AC and MT inferences (i.e. number of MP and MT inferences endorsed and number of DA and AC inferences rejected, each out of 8 items) and formed Normative Responses Indices (NRIs) both overall, and separately for each inference. A negative conclusion index (NCI) was calculated as the number of items with negative conclusions endorsed minus the number of items with affirmative conclusions endorsed and had a possible range of -16 to 16. Inference-specific NCIs were calculated as the number of inferences with negative conclusions endorsed minus the number of items with affirmative conclusions endorsed on each of the four inferences (MP, DA, AC, MT) separately. An implicit negation index (INI) was calculated as the number of items with an explicit negation in the minor premise endorsed minus number of items with an implicit negation in the minor premise endorsed. The INI had a possible range of -8 to 8.

A subset of 650 participants also completed an 18-item version of Raven's Advanced Progressive Matrices (RAPM, Sá, West & Stanovich, 1999). From this we took the number of items correctly answered out of 18 as a measure of fluid intelligence.

Procedure

Participants completed the tasks individually on paper, either alone or in groups of up to approximately 220. The RAPM task was administered first with a 15 minute time limit. The Conditional Inference Task was administered next, amongst various other paper-based tasks, without a time limit.

Results

The results are reported in four sections: descriptive statistics, the relationship between normative responding and intelligence, the relationship between normative responding and the biases, and the relationship between intelligence and the biases.

Descriptive statistics

Descriptive statistics for the measures taken are presented in Table 2. The MP NRI was high and with low variance (M=7.24, SD=1.00), while NRIs to the other inferences were lower and more varied (DA: M=3.91, SD=2.65; AC: M=2.97, SD=2.37; MT: M=5.35, SD=2.03).

NCI (M = 2.34, SD = 2.93) was significantly different from zero, $t(821)$ = 22.88, p < .001, and INI (M = .54, SD = 1.98) was also significantly different from zero, $t(821)$ = 7.77, p < .001. Furthermore, each of the inference-specific NCIs were significantly different from zero: MP NCI (M = -.08 SD = .66), $t(821)$ = 3.47, p = .001, d = -0.12, DA NCI (M = .88 SD = 1.27), $t(821)$ = 19.77, p < .001, d = 0.69, AC NCI (M = .64 SD = 1.19), $t(821)$ = 15.37, p < .001, d = 0.54, MT NCI (M = 1.18 SD = 1.38), $t(821)$ = 24.52, p < .001, d = 0.86. There is very rarely evidence of NCB on MP inferences, and in this case the effect is small and reversed (M = -.08, d = -0.12 indicating that MP inferences were more often endorsed with affirmative conclusions than with negative conclusions) and may only be significant due to the large sample size. In line with previous research, NCB was most prevalent on DA and MT inferences and present to lesser extent on AC inferences (Evans et al, 2007).

Inference endorsement and intelligence

For the subset of participants for whom RAPM data was available (N=650), we examined the relationships between NRIs for each inference and intelligence scores using Pearson's correlations. Given that MP scores had low variance compared to DA, AC or MT, we do not consider MP as a measure of individual differences in the analyses below.

The DA NRI showed a positive correlation with intelligence scores, $r(650)$=.243, p<.001, as did the AC NRI, $r(650)$=.290, p<.001. The MT NRI, however, was not related to intelligence scores, $r(650)$=-.012, p=.755. These relationships are displayed in Figure

1. This essentially replicates the findings of Newstead et al (2004), who found that AH4 intelligence scores correlated with DA scores at $r=.50$, AC scores at $r=.43$, and MT scores at $r=-.14$. Evans et al (2007) did not find significant correlations between AH4 scores and endorsement of DA ($r=-.149$) and AC ($r=-.150$) inferences, despite finding that the highest quartile of AH4 scorers endorsed them less. The magnitudes of the correlations found here were part way between those found by Newstead et al and Evans et al, and given that our sample was larger and our measure of intelligence non-verbal and highly g-loaded, we suggest that our correlations may be more representative of the true correlations between conditional inference behaviour and intelligence in the population.

The correlations between inference endorsement and intelligence were fairly small, which may suggest that Type 2 thinking is necessary but not sufficient for the production of normative responses. This is not surprising because a high capacity for Type 2 thinking is not the same as engaging that capacity - one needs to have the capacity but also the disposition to use it in order to engage it on a task (Stanovich, 2009). It may therefore be the case that normative responses to DA and AC inferences are better measures of *on-task* Type 2 processing than are responses to Raven's Matrices, despite the latter being a better reflection of capacity for Type 2 processing.

Curve fitting

Using normative responding to DA and AC inferences as a proxy for on-task Type 2 processing, we performed a curve fitting estimation on the relationship between the NRIs and NCI. If NCI were a Type 2 processing bias then we predicted that its occurrence would show an n-shaped relationship with NRIs on DA and AC inferences, and that this would be reflected by a significantly better fit to a quadratic curve than to a linear curve. In line with this prediction, the relationship between DA NRI and NCI was significantly fit by a quadratic model, $R^2 = .144$, $F(2,819) = 68.94$, $MSE = 7.35$ $p < .001$, and by a linear model, $R^2 = .026$, $F(1,820) = 21.89$, $MSE = 8.35$, $p < .001$, but the quadratic model provided a significantly better fit, $F(1,819) = 113.00$, $p < .001$. This was also the case for AC; the relationship between AC NRI and NCI was significantly fit by a quadratic model, $R^2 = .122$, $F(2,819) = 56.82$, $MSE = 7.54$ $p < .001$, and by a linear model, $R^2 = .007$, $F(1,820) = 5.92$, $MSE = 8.52$, $p = .015$, but the quadratic model provided a significantly better fit, $F(1,819) = 106.96$, $p < .001$. The data for these relationships are shown in Figure 2.

Next, a curve fitting estimation was performed on the relationship between the NRIs for the DA and AC inferences and INI. We predicted that both the quadratic and linear models would provide poor fits, given that INI is proposed to be a Type 1 processing bias which thus should be unrelated to Type 2 processing (proxied here by normative responding to DA and AC inferences). In line with this prediction, the relationship between DA NRI and INI had low fits for both the linear, $R^2 < .001$, and quadratic models, $R^2 = .009$. A similar pattern of results was observed for the relationship

between AC NRI and INI: linear $R^2 = .001$, quadratic $R^2 = .016$ (although in both cases the quadratic models provided a significantly better fit than the linear models, both ps < .01, we attribute this to the high power of these analyses: for NCI, both quadratic R^2s were above .12, for INI both were below .02, representing an essentially flat relationship). The data for these relationships are shown in Figure 2.

As we discussed above, normative responding to MT inferences is not straightforwardly related to Type 2 processing. Those using high levels of Type 2 processing may endorse MT via reductio ad absurdum, or reject it via a defective interpretation of the conditional. Others may endorse MT through a simple equivalence strategy which is not demanding on Type 2 processing. Following from this we would expect the same quadratic relationship for MT NRI and NCI as for the other inferences, but for a different reason.

With DA and AC a low NRI indicates low Type 2 processing and a high normative score indicates high Type 2 processing. We therefore predicted that higher NRIs would be related to higher Type 2 bias scores, within the constraint that biases cannot be shown alongside very high normative scores. This led to the prediction of an n-shaped curve, which we found. With MT, however, both low and high NRIs are consistent with high Type 2 processing, whilst mid-range NRIs are not. We therefore predict the same n-shaped curve, but at the bottom end this is not due to an absence of Type 2 processing but due to high consistency leaving no room for biases, as is the case at the top end.

To test this prediction, we performed a curve fitting analysis predicting NCI and INI from MT NRI scores. For NCI, the relationship with MT NRI was significantly fit by a quadratic model, $R^2 = .155$, $F(2,819) = 75.152$, $MSE = 7.26$ $p < .001$, and by a linear model, $R^2 = .031$, $F(1,820) = 26.03$, $MSE = 8.31$, $p < .001$, but the quadratic model provided a significantly better fit, $F(1,819) = 120.45$, $p < .001$. For INI scores, the relationship was not significantly fit by a linear model, $R^2 = .001$, $F(2,820) < 1$, $MSE = 3.94$ $p = .391$, nor by a quadratic model, $R^2 = .003$, $F(2,819) = 1.36$, $MSE = 3.93$, $p = .258$, and the quadratic model did not provide a significantly better fit than the linear model, $F(1,819) = 1.98$, $p = .160$.

Intelligence and biases

For the subset of participants for whom RAPM data was available (N=650), we calculated the correlations between RAPM and NCI and RAPM and INI. RAPM is a measure of intelligence, reflecting the capacity and efficiency of Type 2 processing (Stanovich, 2009a, 2009b). Therefore, we expected that if NCB were a Type 2 processing bias it would be correlated with RAPM, while INB would not be. Furthermore, we expected that higher RAPM scores would be related to *higher* NCI scores, in line with our hypothesis that greater engagement of Type 2 processing would result in more bias, not less. RAPM and NCI were significantly, albeit weakly, positively correlated, $r(650) = .142$, $p < .001$, while RAPM and INI were not, $r(650) = .062$, $p = .115$, however, these correlations were not significantly different, $t(647) = 1.44$, $p = .150$.

For completeness, we also performed a curve fitting estimation on the relationship between RAPM and NCI and RAPM and INI. However, it is unclear what should be predicted for this analysis. With the conditional inference scores we know the minimum and maximum possible, and that we have participants who attained the minimum and maximum, giving a full range of scores. With RAPM scores however, it is unclear where our participants lie on the full spectrum of intelligence (our measure was a subset of RAPM items, which cannot be standardized and compared to population norms). Since we are examining a university sample, it is likely that our data covers only a limited range of intelligence scores. If we looked at the relationship between intelligence and NCB in a wide range of participants, we would expect the same n-shaped curve as with DA, AC and MT scores, but here we do not know where on the full spectrum our sample lies, and therefore firm predictions are difficult.

The relationship between RAPM and NCI was significantly fit by a linear model, $R^2 = .020$, $F(2,648) = 13.43$, $MSE = 8.81$, $p < .001$, and by a quadratic model, $R^2 = .023$, $F(2,647) = 7.74$, $MSE = 8.79$, $p < .001$, but the quadratic model did not provide a significantly better fit than the linear model, $F(1,647) = 2.02$, $p = .156$. The relationship between RAPM and INI was not significantly fit by a linear model, $R^2 = .004$, $F(2,648) = 2.49$, $MSE = 4.21$, $p = .115$, nor by a quadratic model, $R^2 = .004$, $F(2,647) = 1.32$, $MSE = 4.22$, $p = .268$, and the quadratic model did not provide a significantly better fit than the linear model, $F < 1$.

Discussion

We investigated behavior on the Conditional Inference Task using data collected from five studies, and asked three questions: what is the relationship between intelligence and conditional inference behavior, is negative conclusion bias (NCB) a Type 1 process or Type 2 process bias, and is implicit negation bias a Type 1 process or Type 2 process bias? First, we showed that intelligence scores were correlated with performance on DA and AC inferences, but not MT inferences. Second, we found evidence that NCB stems from Type 2 processing while INB does not.

Our finding that the number of normative responses to DA and AC inferences is related to intelligence adds to the findings of Evans et al (2007) and Newstead et al (2004), showing a similar pattern of results but with a larger sample and a different measure of intelligence (a subset of Raven's Advanced Progressive Matrices as opposed to the AH4). We did not find evidence that the MT NRI was related to intelligence, which also concurs with previous authors. MT inferences may be endorsed through two distinct routes, one involving a high level of Type 2 processing, and one involving little Type 2 processing. High-ability participants may be capable of deducing MT through a process of suppositional reasoning, using reductio ad absurdum, which requires a high level of Type 2 processing (Evans et al, 2007). Low ability participants may also draw MT but through a simple equivalence strategy, whereby they expect p and q to go together, and not p and not q to go together (such a strategy makes 'if p then q, not q, therefore not p' seem obvious). Drawing MT through this route requires very little Type 2 processing.

Furthermore, participants with a defective interpretation of the conditional may consistently reject MT inferences despite being of high intelligence (Evans et al, 2007). Overall this makes the relationship between intelligence and MT very complex.

The magnitude of the correlations we found between the DA NRI and RAPM scores (r=.243) and the AC NRI and RAPM scores (r=.290) were fairly small, which suggests that a large portion of the variance in conditional reasoning behaviour is unrelated to intelligence. Some of this variance may due to Type 1 thinking or factors such as disposition for effortful thinking (Stanovich, 2009) and educational experiences over and above intelligence (Attridge & Inglis, 2013).

NCB is the tendency to endorse more inferences with negative conclusions than with affirmative conclusions on the Conditional Inference Task, and two accounts have been proposed for its occurrence. The first account proposes a heuristic that endorses inferences which are less easily falsified (Evans, Newstead & Byrne, 1993; Oaksford, Chater & Larkin, 2000; Pollard & Evans, 1980) – the conclusion 'The letter is not C' can only be falsified in one way (with a 'C') while the conclusion 'The letter is C' can be falsified in many ways (with any other letter of the alphabet). The second account proposes a difficulty with processing the double negation that results from an affirmative conclusion on DA and MT inferences (Evans et al, 1995; Evans & Handley, 1999).

In the current sample we found that NCB appeared to be related to Type 2 thinking – the relationship between the number of normative responses to DA and AC inferences (used here as a proxy for on-task Type 2 thinking, supported by their

significant positive relationships with intelligence scores) and NCB showed an n-shaped (quadratic) curve, while their relationship with a Type 1 bias, INB, was essentially flat. We also found a significant relationship between a measure of intelligence, RAPM, and prevalence of NCB. Those who were more intelligent showed greater extents of NCB, while intelligence was unrelated to INB (although the correlations did not differ significantly).

Although we found evidence that NCB was related to on-task Type 2 processing, this does not allow us to rule out a contribution of Type 1 processing to the occurrence of NCB, and indeed it is possible that, for some participants, the Type 1 account is correct. What we have shown is that accounts of NCB based on Type 1 processing alone (Evans, Newstead & Byrne, 1993; Oaksford, Chater & Larkin, 2000; Pollard & Evans, 1980) are not sufficient.

Our findings may reflect a pattern of individual differences (or even a developmental trend) in conditional reasoning, where reasoning with low levels of Type 2 processing results in low NCB at one end of the spectrum, with more systematic reasoning that allows for NCB to occur in the middle of the spectrum, and finally highly consistent high-Type 2 responding which does not leave room for NCB to occur at the other end of the spectrum. This is in contrast to the findings for INB, which is proposed to stem from a difficulty in seeing that implicitly negated premises (e.g. 'not 7' represented as '5') are relevant to the conditional rule (Evans & Handley, 1999). Such a bias would likely stem from Type 1 processing, and in line with this we found INB to be

largely unrelated to normative responding and intelligence. This both supports the generally accepted conclusion that INB stems from Type 1 thinking while simultaneously reinforcing the conclusion that negative conclusion bias stems from Type 2 thinking by means of contrast.

A potential limitation to our findings is that intelligence scores did not show a quadratic relationship with NCB scores. It might be expected that intelligence would be related to Type 2 biases in the same way that normative responding to DA and AC would be, where low intelligence scores reflect low Type 2 processing and low extents of bias, high intelligence scores reflect high type 2 processing and low extents of bias (due to being able to complete the double-negation in negative conclusions and avoid the bias), and mid-range intelligence scores reflect moderate Type 2 processing, but not enough to complete the double-negation, and therefore high extents of bias. However, there are two reasons why we may not expect to find this pattern, which indeed we did not.

The first reason is that intelligence scores reflect the capacity for Type 2 processing, but not the actual extent of Type 2 engagement on the Conditional Inference task. As Stanovich (2009) argued, engaging Type 2 processing on a task requires not only its availability but also the disposition to use it. Therefore, intelligence scores may not reflect on-task Type 2 processing as precisely as scores on DA and AC do. The second reason is that we do not know where our participants lie on the full spectrum of intelligence. Our sample was formed of staff and students from three high-ranking universities, meaning that the range of intelligence represented is likely to be narrow.

Since our measure was a subset of RAPM items, which cannot be standardized and compared to normative data, we cannot determine how restricted or high-ranking our sample was, and we therefore cannot know whether a prediction of a quadratic relationship would be reasonable.

A remaining issue is how to account for the presence of NCB on AC inferences as a product of Type 2 thinking. Schroyens, Schaeken & d'Ydewalle (2001) have attempted to account for this phenomenon in the context of mental models theory. Consider the following AC inference with an affirmative conclusion:

If the letter is A then the number is 3

The number is 3

Conclusion: The letter is A

When confronted with the conditional rule, the reasoner forms an initial model:

A, 3

…

The next step, if Type 2 thinking is engaged, is to search for counterexamples to the 'A, 3' model. This will include a search for instances where the antecedent is false and the consequent true, which in this case means searching for a 'not A, 3' example. If a 'not A,

3' example is found, then the conclusion will likely be rejected because A is not considered to be the only antecedent to 3. Now consider the following AC inference with a negative conclusion:

If the letter is not A then the number is 3

The number is 3

Conclusion: The letter is not A.

This time the initial model looks as follows:

Not A, 3

…

And the search for a false-antecedent true-consequent counterexample will mean searching for an instance of 'not not-A, 3'. At this stage the reasoner encounters the double negation elimination step that is proposed to be the source of problems with negative conclusions on DA and MT inferences. In order to find a counterexample, the reasoner must convert 'not not-A' into 'A', and the difficulty in doing so leads to a failure to produce a false-antecedent true-consequent counterexample. The failure to find such a counterexample thereby causes a failure to reject the negative conclusion (Schroyens, Schaeken & d'Ydewalle, 2001).

Our data are consistent with a double negation model of negative conclusion bias, but only if it is extended to include Schroyens et al's (2001) proposed mental models account of negative conclusion bias on AC inferences. Evans & Handley (1999) proposed that NCB stems from the difficulty of converting a double negative conclusion into a positive conclusion on DA and MT inferences, while Schroyens et al (2001) proposed that the same problem occurs in the search for counterexamples on AC inferences with negative antecedents. In each case, inadequate Type 2 processing leads to a preference for negative conclusions over affirmative conclusions.

Overall our data add to previous findings that rejection of DA and AC inferences is related to intelligence, and they support the argument that NCB stems from Type 2 processing and not Type 1 processing alone.

Acknowledgements

References

Attridge, N. & Inglis, M. (2011). *The development of conditional reasoning skills in undergraduate mathematics students.* Unpublished manuscript, Loughborough University, UK.

Attridge, N. & Inglis, M. (2013). Advanced mathematical study and the development of conditional reasoning skills. *PLoS ONE 8*(7): e69399. doi:10.1371/journal.pone.0069399

Attridge, N. & Inglis, M. (in preparation). Undergraduate education develops conditional reasoning skills.

Braine, M. D. S. & O'Brien, D. P. (1998). The theory of mental-propositional logic: Description and illustration. In M. D. S. Braine & D. P. O'Brien (Eds), *Mental Logic* (pp. 79-89). Mahwah, NJ: Erlbaum.

Evans, J. St. B. T. (1977) Linguistic factors in reasoning. *Quarterly Journal of Experimental Psychology, 29,* 297-306.

Evans, J. St. B. T. (1998). Matching bias in conditional reasoning: Do we understand it after 25 years? *Thinking & Reasoning*, *4*(1), 45-110.

Evans, J. St. B. T. (2003). In two minds: dual-process accounts of reasoning. *Trends in Cognitive Sciences*, *7*(10), 454-459.

Evans, J. St. B. T. (2007). *Hypothetical thinking: dual processes in reasoning and judgment.* Psychology Press: Hove, UK.

Evans, J. St. B. T., Barston, J. L., & Pollard, P. (1983). On the conflict between logic and belief in syllogistic reasoning. *Memory & cognition*, *11*(3), 295-306.

Evans, J. St. B. T., Clibbens, J., & Rood, B. (1995). Bias in conditional inference: Implications for mental models and mental logic. *The Quarterly Journal of Experimental Psychology*, *48*(3), 644-670.

Evans, J. St. B. T., & Handley, S. J. (1999). The role of negation in conditional inference. *The Quarterly Journal of Experimental Psychology: Section A*, *52*(3), 739-769.

Evans, J. St. B. T., Handley, S. J., & Bacon, A. M. (2009). Reasoning Under Time Pressure. *Experimental Psychology (formerly Zeitschrift für Experimentelle Psychologie)*, *56*(2), 77-83.

Evans, J. St. B. T., Handley, S. J., Neilens, H., & Over, D. E. (2007). Thinking about conditionals: A study of individual differences. *Memory & cognition*, *35*(7), 1772-1784.

Evans, J. St. B. T., Newstead, S. E., & Byrne, R. M. (1993). *Human reasoning: The psychology of deduction*. Psychology Press

Evans, J. St. B. T. & Over, D. (2004). *If*. Oxford: Oxford University Press.

Evans, J. St. B. T. & Stanovich, K. E. (2013). Dual-Process Theories of Higher Cognition: Advancing the Debate. *Perspectives on Psychological Science, 8,* 223-241.

Gilmore, C., Attridge, N. & Inglis, M. (2011). Measuring the approximate number system. *Quarterly Journal of Experimental Psychology, 64,* 2099-2109.

Inglis, M. & Simpson, A. (2008). Conditional inference and advanced mathematical study. *Educational Studies in Mathematics, 67,* 187-204.

Jensen, A., R. (1998). *The* g *factor*. Praeger: Westport, CT.

Lott, J. (1999). *Developmental trends in conditional reasoning.* Unpublished Masters thesis. Cited in Evans, J. St. B. T. (2007). *Hypothetical thinking: dual processes in reasoning and judgment.* Psychology Press: Hove, UK.

Newstead, S. E., Handley, S. J., Harley, C., Wright, H., & Farrelly, D. (2004). Individual difference in deductive reasoning. *Quarterly Journal of Experimental Psychology, 57,* 33-60.

Oaksford, M., Chater, N., & Larkin, J. (2000). Probabilities and polarity biases in conditional inference. *Journal of Experimental Psychology: Learning Memory and Cognition, 26(*4), 883-899.

Pollard, P., & Evans, J. St. B. T. (1980). The influence of logic on conditional reasoning performance. *The Quarterly Journal of Experimental Psychology*, *32*(4), 605-624.

Sá, W. C., West, R. F. & Stanovich, K. E. (1999). The Domain Specificity and Generality of Belief Bias: Searching for a Generalizable Critical Thinking Skill. *Journal of Educational Psychology, 91*, 497-510.

Schroyens, W., Schaeken, W., Fias, W., & d'Ydewalle, G. (2000). Heuristic and analytic processes in propositional reasoning with negatives. *Journal of experimental psychology. Learning, memory, and cognition*, *26*(6), 1713-1734.

Schroyens, W. J., Schaeken, W., & d'Ydewalle, G. (2001). The processing of negations in conditional reasoning: A meta-analytic case study in mental model and/or mental logic theory. *Thinking & Reasoning*, *7*(2), 121-172.

Stanovich, K. E. (2009). Distinguishing the reflective, algorithmic, and autonomous minds: Is it time for a tri-process theory. *In two minds: Dual processes and beyond*, 55-88.

Wason, P. C., & Evans, J. St. B. T. (1975). Dual processes in reasoning? *Cognition*, *3*(2),

141-154.

Table 1. The four inferences (modus ponens, denial of the antecedent, affirmation of the consequent and modus tollens) with and without negated premises (Prem) and conclusions (Con).

	MP		DA		AC		MT	
	Prem	Con	Prem	Con	Prem	Con	Prem	Con
if p then q	p	q	not-p	not-q	q	p	not-q	not-p
if p then not-q	p	not-q	not-p	q	not-q	p	q	not-p
if not-p then q	not-p	q	p	not-q	q	not-p	not-q	p
if not-p then not-q	not-p	not-q	p	q	not-q	not-p	q	p

Table 2. Descriptive statistics for the Conditional Inference Task.

Measure	Mean (SD)	Min	Max
MP NRI	7.24 (1.00)	0	8
DA NRI	3.91 (2.65)	0	8
AC NRI	2.97 (2.37)	0	8
MT NRI	5.35 (2.03)	0	8
INI	.54 (1.98)	-6	6
Overall NCI	2.34 (2.93)	-4	14
MP NCI	-0.08 (0.66)	-2	3
DA NCI	0.88 (1.27)	-3	4
AC NCI	0.64 (1.19)	-3	4
MT NCI	1.18 (1.38)	-4	4
RAPM	9.36 (3.28)	0	17

Figure 1. The relationships between RAPM score and DA, AC and MT (normative) scores. Error bars represent ±1 standard error of the mean.

Figure 2. The relationships between DA, AC and MT scores and NCI and INI scores. Error bars represent ±1 standard error of the mean.