THEORETICAL AND REVIEW ARTICLES

A signal detection analysis of the recognition heuristic

TIMOTHY J. PLESKAC
University of Basel, Basel, Switzerland

The recognition heuristic uses a recognition decision to make an inference about an unknown variable in the world. Theories of recognition memory typically use a signal detection framework to predict this binary recognition decision. In this article, I integrate the recognition heuristic with signal detection theory to formally investigate how judges use their recognition memory to make inferences. The analysis reveals that false alarms and misses systematically influence the performance of the recognition heuristic. Furthermore, judges should adjust their recognition response criterion according to their experience with the environment to exploit the structure of information in it. Finally, the less-is-more effect is found to depend on the distribution of cue knowledge and judges' sensitivity to the difference between experienced and novel items. Theoretical implications of this bridge between the recognition heuristic and models of recognition memory are discussed.

Simon (1990) observed that recognition is a natural mechanism for helping people to solve problems, such as those found in chess, medical diagnosis, or reading. Similarly, Axelrod (1985) postulated that recognition may be necessary for cooperation to be sustained in social interactions. The recent development of the recognition heuristic has added inferences to this list of indirect applications of recognition memory (Goldstein & Gigerenzer, 1999, 2002). According to the heuristic, recognition serves as a cue for making inferences about pairs of objects. The recognition heuristic can be quite accurate. These areas include, among others, making population inferences about German, U.S., and Swiss cities (Gigerenzer & Goldstein, 1996; Goldstein & Gigerenzer, 2002; Pohl, 2006), selecting stocks during a bull market (Borges, Goldstein, Ortmann, & Gigerenzer, 1999), and identifying National Hockey League players who have a high number of career points (Snook & Cullen, 2006).

Heuristics that use recognition as a predictor variable, such as the recognition heuristic, typically start with a judge's binary recognition decision (Gigerenzer & Goldstein, 1996, 1999; Goldstein & Gigerenzer, 1999, 2002). In contrast, memory research often focuses on the process that leads up to the recognition decision (Raaijmakers & Shiffrin, 2002). Research examining these memorial processes relies on laboratory experiments that have a defined learning phase in which participants study a list of items—sometimes novel, sometimes common—followed by a test phase in which their recognition memory is tested on the items that they learned. The results from the laboratory typically reveal that signal detection theory is useful for understanding why correct and incorrect recognition decisions are observed during memory experiments (Banks, 1970). In particular, two factors give rise to both decisions during memory experiments: (1) the ability to detect the difference between the familiarity of learned and unlearned items (sensitivity); and (2) various response factors (criterion location). Correct and incorrect recognition decisions, however, are more difficult to examine within the ecological framework of the recognition heuristic and the broader class of fast and frugal heuristics. The difficulty arises because their ecological framework requires heuristics to be examined with a representative sample of stimuli drawn from an environment or reference class experienced outside of the laboratory (see Gigerenzer, Todd, & the ABC Research Group, 1999). Furthermore, the environment should be selected in such a way that the heuristic could sensibly be used in the environment to make inferences (see Gigerenzer et al., 1999). Consequently, the researcher does not typically know the items a respondent has or has not experienced, making it difficult—if not impossible—to identify correct and incorrect recognition decisions.

The different methodology needed to study the recognition heuristic does not imply, however, that judges use a functionally different recognition process for items drawn from ecologically defined reference classes than for items learned in a laboratory setting. In fact, a recent fMRI study demonstrated a link between the recognition heuristic and recognition memory when areas of the medial parietal cortex—areas typically associated with recognition memory—were activated during the use of the recognition heuristic (Volz et al., 2006). A natural question is: How do these rec-

T. J. Pleskac, tim.pleskac@gmail.com

ognition processes affect the accuracy of recognition as a predictor variable? In this article, I formally address this question by integrating signal detection theory with the recognition heuristic to show the effect that correct and incorrect recognition decisions have when recognition is used as a predictor cue. In doing so, I rely on the generalizability of signal detection theory as a model of recognition memory examined in the laboratory (see, e.g., Banks, 1970; Ratcliff, Clark, & Shiffrin, 1990; Shiffrin, Huber, & Marinelli, 1995) to the inferential paradigm of the fast and frugal program of research (Gigerenzer et al., 1999). Next, to facilitate the integration, I introduce the recognition heuristic, the empirical evidence that supports its use, and the less-is-more effect that the recognition heuristic produces.

The Recognition Heuristic

The recognition heuristic is a single-variable decision rule that relies on recognition alone to make a judgment about an unknown target variable. When judges are confronted with a two-alternative, forced-choice question, such as choosing whether San Antonio or San Diego is more populous, the recognition heuristic states: If one of two objects is recognized and the other is not, then infer that the recognized object has the higher value with respect to the target variable.

In some domains, people who use this heuristic have a good chance of being correct. For example, Goldstein and Gigerenzer (2002) reported that the average Spearman correlation between American students' recognition of German cities with a population over 100,000 and their respective population size is .60 (the *recognition correlation)*. This correlation occurs because people learn city names from mediators in the environment, such as the news media. The *surrogate correlation* between American students' recognition of German city names and the number of mentions of those cities in the *Chicago Tribune* was .79. The mediators, in turn, can reflect other statistical relationships. The citation rates of the city names have an *ecological correlation* of .70 with the populations of German cities (Goldstein & Gigerenzer, 2002).[1] The recognition heuristic, by exploiting the structure of information diagrammed in Figure 1, is ecologically rational when inferring city populations or for any domain with the same information structure (Goldstein & Gigerenzer, 2002).

Judges with partial ignorance—those who have experience with some but not all of the objects in a specified reference class—benefit most from the recognition heuristic. For example, Goldstein and Gigerenzer (1999) reported that the recognition heuristic can explain why Germans answered the San Diego/San Antonio question quite accurately—100% in their sample—whereas only 62% of Americans answered the question correctly. This is because the recognition heuristic can be used only when just one of the items is recognized. American students recognized both cities and therefore resorted to other, less accurate cues to make an inference. When given a larger test bank of questions, partially ignorant judges can capitalize on their ignorance when the recognition heuristic is paired with a knowledge heuristic that follows the process diagrammed in Figure 2. Possible heuristics that follow
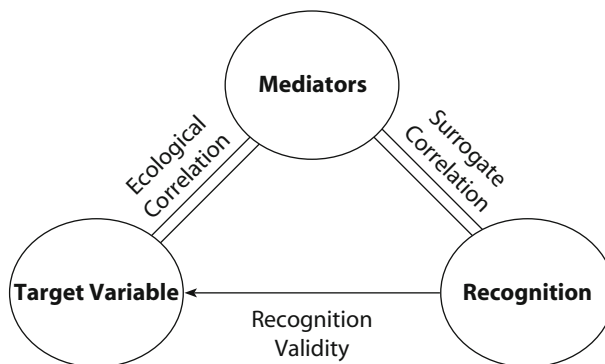

Figure 1. The ecological rationality of the recognition heuristic. In some domains, recognition of objects can be correlated with an unknown target variable (e.g., city population). This is because judges experience objects via mediators in the environment (e.g., newspapers), and the mediators reflect the target variable (e.g., more populous cities tend to be in the news more often). From "Models of Ecological Rationality: The Recognition Heuristic," by D. G. Goldstein and G. Gigerenzer, 2002, *Psychological Review*, *109*, p. 78. Copyright 2002 by the American Psychological Association. Adapted with permission.

this procedure include "take the best," "take the last," and "minimalist" (Gigerenzer & Goldstein, 1996). Together, these heuristics give rise to a less-is-more effect, in which less knowledge or experience can be beneficial (see Gigerenzer, 2004; Gigerenzer & Goldstein, 1996, 1999; Goldstein & Gigerenzer, 1999, 2002).[2]

To explain how and why the less-is-more effect can emerge, Goldstein and Gigerenzer (1999, 2002) derived the expected accuracy of judges who used the recognition heuristic when given an inferential test on pairs of objects. During the test, judges had to identify which object had a larger value on a target variable, such as population. The pairs were formed in such a way that the individual objects were a representative sample of a well-defined population of objects or reference class with N objects (e.g., German cities with over 100,000 inhabitants). A representative sample has the property that individual stimuli are equally likely to be selected (with replacement) from the reference class (Brunswik, 1955; Dhami, Hertwig, & Hoffrage, 2004; Gigerenzer, Hoffrage, & Kleinbölting, 1991).[3]

The recognition heuristic divides the N objects into two groups: n recognized objects and $N - n$ novel objects. If a person is confronted with a randomly drawn pair of items from $N(N - 1)/2$ of the possible pairs, then the recognition heuristic can be expected to be used for $2n(N - n)/[N(N - 1)]$ proportion of the pairs, where one item is recognized and the other is not.[4]

The *recognition validity* reflects the recognition correlation (see Figure 1) and is the proportion of times that the recognized object has a higher value on the target variable:

$$\alpha = \frac{R}{R + W}. \tag{1}$$

The variable R is the number of pairs that would lead to the correct inference, and W is the number of pairs that would lead to the incorrect inference, given a set of recog-
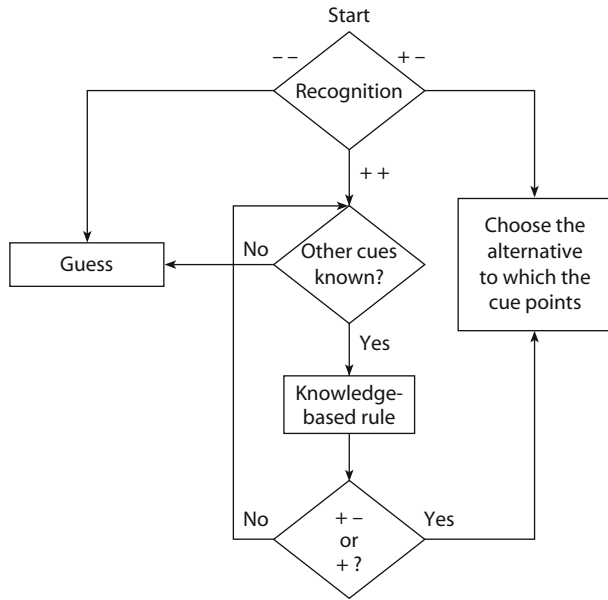
Figure 2. A flow diagram of knowledge heuristics that use the recognition heuristic in the first step. Typically, these algorithms include "take the best," "take the last," and "minimalist." The discrimination rule for these knowledge heuristics specifies that the judge choose the object with a "+" cue value when it is paired with an object whose cue value is "−" or "?". From "Reasoning the Fast and Frugal Way: Models of Bounded Rationality," by G. Gigerenzer and D. G. Goldstein, 1996, *Psychological Review*, 103, p. 653. Copyright 1996 by the American Psychological Association. Adapted with permission.

and the other does not; otherwise, it does not discriminate, and judges must guess. The expected accuracy for the pairs of items sent to the knowledge heuristic reflects this concept. It is the probability of a correct inference given that both items are recognized, β. Only in the limit when the knowledge heuristic has perfect discrimination is β the cue validity as defined in Equation 1. In all other cases, β is a weighted average between the cue validity and ½ in which the weights are determined by the discrimination rate of the knowledge heuristic.[5]

Summing the proportion of correct inferences for the recognition heuristic, guessing, and the knowledge heuristic produces the expected proportion of correct inferences, P, for the given reference class,

$$P = \frac{2n(N-n)}{N(N-1)}\alpha$$
$$+ \frac{(N-n)(N-n-1)}{N(N-1)} \cdot \frac{1}{2} + \frac{n(n-1)}{N(N-1)}\beta. \qquad (2)$$

Using Equation 2, P can be plotted as a function of n, the number of items a judge has recognized. Setting $N = 100$ and $\alpha = .8$, the curves in Figure 3 show this for four different levels of β. Goldstein and Gigerenzer (1999) illustrated the predictions with a story about three brothers with different levels of experience who take this inferential test. The youngest brother, who had no experience, recognized none of the objects ($n = 0$), had to guess on all the inferences, and was correct 50% of the time (see the dot on the left side of Figure 3). The middle brother had some experience and recognized half

nized and novel items. Goldstein and Gigerenzer (2002) suggested that this validity corresponds to Brunswik's (1955) *ecological validity*, in which the ecological cue validity is the true relative frequency of any object, p, having a larger value on the target variable than any other object, q, in a reference class in which p has a positive cue value on the particular cue and q does not (Gigerenzer et al., 1991). Typically, the ecological cue validity should be defined independently of any particular person (see Goldstein & Gigerenzer, 1999, note 1). Goldstein and Gigerenzer (2002) argued that recognition validity is different from other ecological validities in that the relationship between the target variable and the cue is mediated by the structure of the environment. A second difference that I will return to is that recognition validity is not defined independently of a person's psychological or memorial processes.

For the set in which neither object will be recognized, the judge must guess (see Figure 2). This is expected to happen for $[(N - n)][(N - (n - 1))/[N(N - 1)]$ proportion of the pair with an expected accuracy of ½. Finally, a knowledge-based rule can be used for the remaining pairs in which both items are recognized $[n(n - 1)/N(N - 1)]$ proportion of the time. This collapses across the subset of objects that the knowledge heuristic does and does not discriminate among (see Goldstein & Gigerenzer, 1999, note 1). According to Figure 2, the knowledge heuristic discriminates between objects when one has a positive cue

Figure 3. The less-is-more effect. The graph displays the proportion of correct inferences in solid lines when $\alpha = .8$ at four levels of β. When the knowledge validity β is .5, .6, or .7, a less-is-more effect occurs. The performance of the three brothers is indicated by the three points on the curve for $\beta = .6$. From "Models of Ecological Rationality: The Recognition Heuristic," by D. G. Goldstein and G. Gigerenzer, 2002, *Psychological Review, 109*, p. 79. Copyright 2002 by the American Psychological Association. Adapted with permission.

of the objects ($n = 50$). He also had extra cue knowledge so that when he recognized both cities, he was correct 60% of the time ($\beta = .6$). In this case, using Equation 2, the middle brother would score 68% on the test (see the middle dot in Figure 3). The oldest brother had extensive experience with the reference class ($n = 100$). Consequently, he could not use recognition and instead relied on cue knowledge. His cue knowledge was similar to that of the middle brother, so he would be correct 60% of the time and his score of 60% on his test reflected this (see the rightmost dot in Figure 3). Hence, the scores of the middle and oldest brothers exhibit the less-is-more effect: The middle brother benefited from partial ignorance and scored 8% better than the more knowledgeable brother. In general, to find the less-is-more effect, the recognition validity must be greater than the accuracy of the knowledge heuristics ($\alpha > \beta$); otherwise, the maximal performance will always be at $n = N$ (see Goldstein & Gigerenzer, 2002).

Do judges actually use the recognition heuristic? Goldstein and Gigerenzer (2002) reported that when making an inference about the more populous of two German cities, American college students chose the recognized city in 90% of all possible cases. Other studies (Bröder & Eichler, 2006; Newell & Shanks, 2004; Pachur & Hertwig, 2006; Pohl, 2006) have addressed whether recognition is used as a single variable predictor. Results from these studies indicate that when judges do not have direct knowledge about the target variable (i.e., judges use a probabilistic mental model; see Gigerenzer et al., 1991), they use recognition as a cue to make an inference. Furthermore, as the cue validity of recognition increases, the likelihood that recognition is used as a single variable predictor increases.

However, as knowledge of conflicting cues increases, the noncompensatory status of the recognition heuristics is also mitigated (Bröder & Eichler, 2006; Newell & Shanks, 2004; Oppenheimer, 2003; Pohl, 2006; Richter & Spath, 2006). Regardless, to use recognition as a binary predictor, whether in a single- or multiple-variable heuristic, judges must first generate a recognition decision. These decisions do not map directly onto past experience. Instead, they depend on judges' sensitivity to the difference between the familiarity of experienced versus novel items and their goals at the time of making their inference. Referring to Figure 1, these dependencies imply that memorial explanations other than an imperfect assessment of all possible mediators may explain why the surrogate correlation is not perfect. Integrating the recognition heuristic with signal detection theory reveals one possible explanation for the imperfect surrogate correlation and can explain how recognition memory processes affect inferential performance.

Integrating the Recognition Heuristic and Signal Detection Theory

The recognition heuristic makes an implicit assumption that recognition decisions perfectly reflect the split between experienced and novel items. Goldstein and Giger-enzer (2002) stated "Thus, with the term recognition, we divide the world into the novel and the previously experienced" (p. 77). However, judges are unlikely to both recognize all experienced objects and not recognize (reject) all novel objects.

In fact, the research Goldstein and Gigerenzer (1999, 2002) cited as evidence for the remarkable ability and accuracy of recognition memory also directly acknowledges that recognition is fallible (see Craik & McDowd, 1987; Jacoby, Woloshyn, & Kelley, 1989; Shepard, 1967; Standing, 1973). For example, Craik and McDowd and Jacoby et al. reported nonzero miss and false alarm rates for their respective recognition experiments. Shepard found that when respondents learned words and were later tested in his forced-choice paradigm on what they had learned, they responded correctly 88.4% of the time; with sentences, they were correct 89% of the time; and with pictures, they were correct 99.7% of the time.

Shepard (1967) showed that two classes of models might account for the observed data: (1) a signal detection framework where recognition is a function of a continuous construct and judges use an optimal criterion; and (2) a two-state threshold model where recognition is a function of a binary all-or-none process. Assuming no response bias, both models can account for the errors and the differences among stimuli as a function of the judges' sensitivity. Both models can be incorporated into the recognition heuristic.

Schooler and Hertwig (2005) have developed one possible high-threshold account with ACT–R. Global memory models, in comparison, embody the signal detection framework and have been successful in accounting for a range of different phenomena (Raaijmakers & Shiffrin, 2002). Examples include MINERVA-2 and MINERVA-DM (Dougherty, Gettys, & Ogden, 1999; Hintzman, 1988), REM (Shiffrin & Steyvers, 1997), SAM (Gillund & Shiffrin, 1984), and TODAM (Murdock, 1997). In general, these models presume that when a judge is shown a stimulus, a representation of the stimulus is formed in a memory probe. The probe is then compared with each item in memory, giving rise to a continuous level of activation or familiarity.[6] If the familiarity is above a criterion value, the judge decides the stimulus is old (he has recognized the item). If the familiarity is below the criterion, the judge decides the stimulus is new (he has not recognized the item).

From this perspective, the judge's familiarity with items from a reference class is correlated by means of the environmental mediators of the target variable (e.g., population size). During the inferential test, the judge must transform his familiarity into a binary decision in order to use recognition as a predictor variable. To make the model as general as possible, I do not use any specific global memory model but instead use the more basic Gaussian signal detection model. Furthermore, to keep as much of the recognition heuristic intact as possible, I assume that the use of memory is comparable to a detection task. Specifically, when judges are presented with two stimuli (e.g., San Antonio and San Diego) and want to make an inference about some target variable (e.g., city size) based on recognition, they first look at one stimu-

lus (San Antonio), decide whether it is old or new, and then turn to the remaining stimulus and decide whether it is old or new.[7] Each experienced item can be either correctly identified as old (hit) or incorrectly identified as new (miss). If the item is novel, then it can be either incorrectly identified as old (false alarm) or correctly identified as new (correct rejection).

These distribution-free, detection-based assumptions elicit novel predictions. When judges do not make recognition mistakes, there are three possible pair types. The recognition heuristic operates on the (hit, correct rejection) pairs. The guessing component is used on the (correct rejection, correct rejection) pairs, and the knowledge-based rules are used on the (hit, hit) pairs. As the rate of mistakes increases, the number of possible pair types increases from 3 to 10, and the distribution of pairs among these pair types changes.[8] The first column of Table 1 illustrates the 10 different possible pair types.

The specific component—recognition heuristic, guessing, or knowledge heuristic—that handles each of the 10 pair types is identified in the second column. To calcu-late the proportion of pairs for each of the 10 pair types using the signal detection model, the number of objects experienced, n_e, will be used, where the subscript "e" indicates the change from number recognized to number experienced. This is because the framework now predicts the recognition of objects on the basis of experience. The variables h and f represent the hit and false alarm rates, respectively. To illustrate the calculations, consider the pairing when both items are hits. There are hn_e items that are expected to be a hit for the first item and, once one item is removed from this set, $hn_e - 1$ items are left for the second item.[9] The expected number of pairs in which both items are hits, therefore, is $hn_e(hn_e - 1)$. Dividing that expression by the total number of possible pairs, $N(N - 1)/2$, produces the expected proportion of (hit, hit) pairs. The remaining expressions are found in a similar manner. The equations are shown in the third column of Table 1.

The expected accuracy associated with each pairing can also be derived. The accuracy of the recognition heuristic depends on which of the four pair types it is fed.

Table 1
The 10 Possible Pair Types That Judges' Recognition Decisions Can Produce

Pairs	Heuristic	Proportion of Pairs	Expected Proportion Correct
Hit, correct rejection	Recognition	$\dfrac{2hn_e(1-f)(N-n_e)}{N(N-1)}$	A
Miss, false alarm	Recognition	$\dfrac{2(1-h)n_e f(N-n_e)}{N(N-1)}$	$1-A$
Hit, miss	Recognition	$\dfrac{2(h-h^2)n_e^2}{N(N-1)}$	$\frac{1}{2}$
False alarm, correct rejection	Recognition	$\dfrac{2(f-f^2)(N-n_e)^2}{N(N-1)}$	$\frac{1}{2}$
Correct rejection, correct rejection	Guess	$\dfrac{(1-f)(N-n_e)[(1-f)(N-n_e)-1]}{N(N-1)}$	$\frac{1}{2}$
Miss, miss	Guess	$\dfrac{(1-h)n_e[(1-h)n_e-1]}{N(N-1)}$	$\frac{1}{2}$
Miss, correct rejection	Guess	$\dfrac{2(1-h)n_e(1-f)(N-n_e)}{N(N-1)}$	$\frac{1}{2}$
Hit, hit	Knowledge	$\dfrac{hn_e(hn_e-1)}{N(N-1)}$	B
Hit, false alarm	Knowledge	$\dfrac{2hn_e f(N-n_e)}{N(N-1)}$	$zA+(1-z)\cdot\frac{1}{2}$
False alarm, false alarm	Knowledge	$\dfrac{f(N-n_e)[f(N-n_e)-1]}{N(N-1)}$	$\frac{1}{2}$

Note—The second column identifies which heuristic judges would use to arrive at an answer for each pair. The third column gives the proportion of pairs of each type for a given reference class with N objects, n_e objects experienced by a judge, and hit and false alarm rates h and f, respectively. The fourth column gives the expected accuracy of each pair type.

However, the recognition validity, α, cannot be used; the cue validity must be independent of the judge's recognition ability. Consequently, I used the cue validity of the previously experienced and novel items, A. The value of A is the ecological validity of experience independent of any psychological or memorial process within the judge. It reflects the true environmental correlation between environmental mediators and the target variable (see Figure 1). Hypothetically, we could arrive at A if we could generate a list of German cities with more than 100,000 citizens that a typical judge might have encountered by reading the *Chicago Tribune*. More than likely, this list would not be a comprehensive list of cities from the reference class, and, according to past analyses, it would probably tend to include more of the larger cities (see Goldstein & Gigerenzer, 2002). Using this list, the number of pairs that would lead to a correct inference, R, and an incorrect inference, W, could be found. Like the cue validity of cues in the environment (e.g., whether a city has a soccer team or not), this cue validity is derived directly from the structure of the environment using Equation 1.

Of the four pairings that enter the recognition heuristic, only (hit, correct rejection) retains the expected accuracy of A. The second pairing, (false alarm, miss), has a different expected accuracy. If a judge accurately recognized both items, then this pair would also have gone into the recognition heuristic and retained A as the expected accuracy. However, now the novel item, false alarm, has been identified as recognized and the experienced item, miss, has been identified as not recognized, and the recognition heuristic would therefore pick the false alarm as the item with a higher target variable. Thus, referring to Equation 1, the W pairs used to calculate A would now lead to the correct inference, whereas the R pairs would lead to the incorrect inference. This results in an expected accuracy for (false alarm, miss) of $1 - A$.

The remaining two pairs, (hit, miss) and (false alarm, correct rejection), also have an expected accuracy different from A. Each of these two pairs comprises a correct and an incorrect detection. Consider first the (hit, miss) pair. A particular item is correctly recognized as old with a probability of h and is incorrectly identified as new with a probability $1 - h$. The recognition heuristic will produce the correct response for a (hit, miss) pairing only when the higher valued item is a hit and the lower valued item is a miss. This happens with a probability of $h(1 - h)$. The opposite can also happen. The higher valued item can be missed and the lower valued item can be correctly recognized, producing an incorrect inference. This particular pairing of a (hit, miss) also occurs with probability $(1 - h)h$. Thus, a (hit, miss) pairing produces a correct inference half of the time and an incorrect inference half of the time. Similar logic holds for the (false alarm, correct rejection) pair.

For the next three pairs (correct rejection, correct rejection), (correct rejection, miss), and (miss, miss), the guessing rule would be employed. For all three of these

pair types, the judge is expected to be correct half of the time.

Judges would use a knowledge-based heuristic for the remaining pairs when both items in a pair were recognized (see bottom three rows of Table 1). Although the cue validity of the knowledge cues—such as whether a city has a soccer team—is defined independently of the recognition ability of a judge, the probability of a correct inference given that both items are recognized (β), is not. To derive the expected accuracy of these items, I instead use B, the probability of a correct inference given that both items are experienced. In this case, only the original pairing (hit, hit) would retain the expected accuracy of B. The items in the pair with two false alarms are actually novel, so this pair has the lowest accuracy, $\frac{1}{2}$.

Interestingly, the remaining pair type, (hit, false alarm), can benefit indirectly from experience. It can have an expected accuracy as high as A or as low as $\frac{1}{2}$. To see why, consider Table 2, which models a hypothetical judge's experience with 20 items (a–t), recognition of the same 20 items, and knowledge of five cues for 20 items. The objects are displayed in descending order in terms of rank on the basis of a hypothetical target variable. For example, object a might be the city in a country with the largest population, and object t might be the city with the 20th largest population. The person has experienced 10 of the objects, as indicated by a "+" in the experience column. The "−" in the experience column indicates that the objects were novel. For these novel objects, the judge also has no binary cue

Table 2
A Hypothetical Person's Experience and Recognition of a Reference Class of 20 Objects (a–t) and Level of Knowledge of Five Knowledge Cues

| Rank | Object | Experience | Recognition | Cue | | | | |
				1	2	3	4	5
1	a	+	Hit	+	−	−	+	−
2	b	+	Hit	−	−	−	−	−
3	c	+	Miss	−	+	+	−	−
4	d	−	CR	?	?	?	?	?
5	e	+	Hit	−	+	−	−	−
6	f	+	Hit	−	−	−	−	−
7	g	−	CR	?	?	?	?	?
8	h	+	Hit	−	−	−	−	−
9	i	−	FA	?	?	?	?	?
10	j	+	Miss	−	−	+	−	−
11	k	+	Hit	−	+	−	−	−
12	l	−	CR	?	?	?	?	?
13	m	−	CR	?	?	?	?	?
14	n	+	Miss	−	−	−	+	−
15	o	+	Hit	−	−	−	−	−
16	p	−	CR	?	?	?	?	?
17	q	−	FA	?	?	?	?	?
18	r	−	CR	?	?	?	?	?
19	s	−	CR	?	?	?	?	?
20	t	−	FA	?	?	?	?	?

Note—The objects are ordered according to a hypothetical target variable. A "+" or "−" in the experience column indicates whether the object is experienced or novel, respectively. The Recognition column identifies the recognition decision at the time of the inferential test. CR, correction recognition; FA, false alarm. A "+," "−," or "?" in the cue columns indicates a positive, negative, or unknown cue value, respectively.

knowledge. This is indicated by a "?" in the cue columns. The recognition column identifies the classification of the person's recognition decision at the time of the inferential test. As Table 2 shows, there are two types of (hit, false alarm) pairs. Pair (a, i) is one. For this pair, the knowledge heuristic would infer item a as the larger item and be correct on the basis of the discrimination rule detailed in Figure 2, in which judges choose the positive cue over negative and unknown cue values. In fact, the hit item will always be chosen in a pair of this type whenever at least one positive cue value is associated with it in memory. Table 2 reveals that these (hit, false alarm) pairs are sampled from pairs that the experience cue would have differentiated among (i.e., one item has a "+" and the other a "−" in the experience column) and that the hit choices are consistent with the choices that the experience cue would have led to. As a result, this subset has an expected accuracy of A. However, some of the hit items (such as b and f) have no positive cue values. This can occur because of the structure of the environment or the limited cue knowledge of the judge. Either way, the discrimination rule (i.e., when both objects lack a positive cue value, guess) dictates that the judge would have to guess on this subset of (hit, false alarm) pairs. Thus, the expected accuracy for the (hit, false alarm) pair is a weighted average between A and ½, with the weight determined by z, the proportion of experienced items with at least one positive cue value.

To find the expected proportion of correct inferences for each pair type, multiply the expected proportion of pairs (column 3 of Table 1) by the accuracy rate (column 4 of Table 1). Table 3 provides a summary of the variables introduced for this derivation. To produce more precise predictions, I assume that novel and experienced stimuli give rise to a familiarity, t, that is normally distributed. For scaling purposes, the novel stimuli have a mean of 0, and experienced stimuli have a mean of d'. The parameter d' indexes the level of sensitivity to the difference between the familiarity of experienced versus novel items. As d' increases, judges become more sensitive to the difference between the two types of stimuli. Shepard's (1967) set of recognition experiments—using pictures, statements, and words—pro-

vides a convenient example of how sensitivity can vary according to item type. Interesting or important items can also induce higher levels of sensitivity within a domain (see, e.g., Gronlund, Ohrt, Dougherty, Perry, & Manning, 1998). Additionally, sensitivity can vary among domains on the basis of meaningfulness, similarity, and pleasantness of the objects, among other things (see also Glanzer & Adams, 1985). Finally, sensitivity can vary among people for a given domain. For example, repetitions and study time can improve discrimination. Consider two geography students, one diligent and one lackluster. The diligent student will study her list of cities more frequently and for longer periods of time, becoming more sensitive to differences between experienced and novel items, whereas the lackluster student will study the same list once and again only briefly the morning before the exam and would thus be less sensitive.

I also make a simplifying assumption that both distributions of familiarity have an equal variance, σ^2, set at 1.[10] At test, judges have a criterion, k, set at one point along the possible values that the familiarity could take. The rule for deciding whether a given item is old or new can now be formally stated as: Respond "old" if and only if $t > k$; if not, respond "new."

With the distributions specified, the probability of a hit for a given value of d' and k can be calculated as $h = 1 - \Phi(k - d')$, where the function $\Phi(\cdot)$ represents the standardized normal cumulative distribution function. The miss rate, m, is $m = (1 - h)$. The false alarm rate is $f = 1 - \Phi(k)$, and the correct rejection rate is $c = 1 - f$. The expected proportion of each type of pairings can now be calculated given a value for k and d'.

Besides changing the applicability of the recognition heuristic, experience can also change the recognition response process. Prior experience with the reference class increases the number of items the judge has experienced (n_e) prior to the test, thus making it a priori more likely during a representative test that the judge will be shown a previously experienced item. Within signal detection theory, a judge should capitalize on this by picking whichever decision (old or new) has the greater likelihood given his familiarity $p(\text{old} \mid \text{familiarity})$ or $p(\text{new} \mid \text{familiarity})$ and adjust his response criterion according to his level of prior experience. Thus, a judge with little to no experience would be fairly conservative in deciding whether he recognizes an item, but as his experience increased, his criterion would become more liberal.

To see this formally, the comparison can be stated in terms of the posterior odds,

$$\Omega = \frac{p(\text{old} \mid \text{familiarity})}{p(\text{new} \mid \text{familiarity})}$$
$$= \frac{p(\text{familiarity} \mid \text{old})}{p(\text{familiarity} \mid \text{new})} \frac{p(\text{old})}{p(\text{new})}, \quad (3)$$

where the first component on the far right side is the likelihood ratio, $p(\text{familiarity} \mid \text{old})/p(\text{familiarity} \mid \text{new})$, and the second component is the prior odds, $p(\text{old})/p(\text{new})$. The likelihood ratio of the data is found using the distributions specified for the experienced and novel items. The representative sampling of the stimuli in the inferential test—where

Table 3
Parameters and Variables Used for the Signal Detection Analysis

Variable	Description
N	Number of objects in a specified reference class
n	Number of objects recognized in the reference class
α	Validity of recognition in the reference class
β	Probability of a correct inference in the reference class, given that both items have been recognized
n_e	Number of objects experienced in the reference class
A	Ecological validity of experience in the reference class
B	Probability of a correct inference in the reference class given that both items have been experienced
h	Hit rate
m	Miss rate
f	False alarm rate
c	Correct rejection rate
t	Familiarity of an item
d'	Sensitivity to the difference between the familiarity of experienced and that of novel items

all objects are equally likely to occur in a pair (see Gigerenzer et al., 1991)—specifies the prior odds for the judge. For example, if an observer has experienced only 40 of 100 items in a reference class, the prior odds of an old item's appearing on the inferential test are .4/.6. More generally,

$$\frac{p(\text{old})}{p(\text{new})} = \frac{n_e}{N - n_e}.[11]$$ (4)

To maximize the probability of a correct detection, a judge's decision rule for recognition must be

$$\begin{cases} \text{If } \Omega \geq 1 \text{ or } n_e = N, \text{ then select old.} \\ \text{If } \Omega < 1 \text{ or } n_e = 0, \text{ then select new.} \end{cases}$$

Substituting Equation 4 into Equation 3 and taking the logarithm, the response can be reformulated in terms of a likelihood ratio observer who adopts a criterion with

$$k = -\log\left(\frac{n_e}{N - n_e}\right),$$

with the constraint that $0 < n_e < N$. The same decision rule can now be stated in terms of the likelihood ratio

$$\begin{cases} \text{If } R(\text{old:new}) \geq k \text{ or } n_e = N, \text{ then select old.} \\ \text{If } R(\text{old:new}) < k \text{ or } n_e = 0, \text{ then select new.} \end{cases}$$

Figure 5. The average α as a function of d' averaged across all levels of experience for four different levels of A (.6, .7, .8, and .9). Lower levels of sensitivity can bring about substantial decrements to the accuracy of their inferences.

The likelihood-ratio observer is consistent with recent empirical evidence from recognition experiments (see Glanzer, Adams, Iverson, & Kim, 1993; Shiffrin & Steyvers, 1997). This strategy can also be understood as adaptive. It maximizes the probability of correctly detecting whether an item is old or new and minimizes the interference that imperfect recognition memory would have on the recognition heuristic. To illustrate this, Figure 4 shows that the recognition error rate follows an inverted U centered around $n_e = 50$ for three different levels of d'.

An alternative strategy is to keep the criterion fixed. The dashed lines in Figure 4 show the effects this strategy would have on the error rate for three levels of a centered criterion (-1, 0, and 1) holding d' constant at 0.5.[12] The error rate for these strategies is a linear function of experience whose slope depends on how liberal the judge is in making a positive recognition decision. These error rates reveal that this decision rule is maladaptive and would lead to more recognition errors across all levels of experience, resulting in decreased inferential accuracy. A third class of criterion strategies, not shown in Figure 4, would be that of a respondent who grows more conservative with experience. This in effect would minimize the probability of a correct detection and would produce an upright U in Figure 4 (not shown). This does not seem very adaptive or plausible given the stated goal of inferential accuracy in the task.[13] Furthermore, the strategy would lead to odd predictions regarding the procedures outlined in Figure 2. For example, a complete expert ($n_e = N$) would completely neglect his knowledge and guess on 100% of the pairs. Taken together, this implies that the ecologically based response rule consistent with a Bayesian observer would allow a judge to fully exploit the ecological correlation in the environment and be ecologically rational (Gigerenzer, 2001).
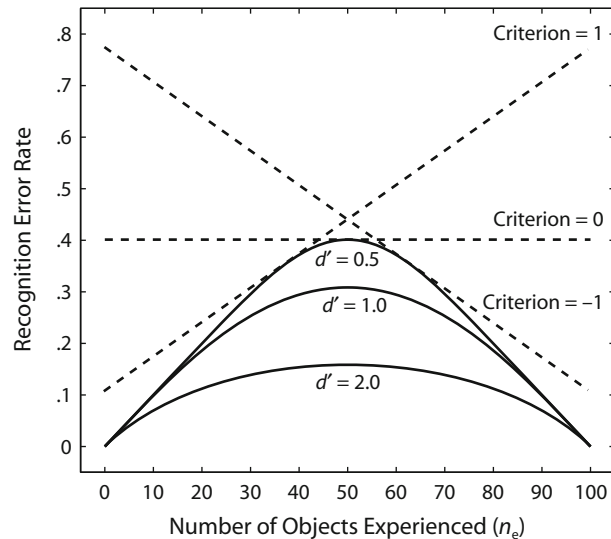

Figure 4. Hypothetical error rates for judges who adopt two different response strategies. The first strategy is to adjust the response criterion according to the level of experience, a Bayesian observer. The inverse U curves in the figure illustrate this strategy. At low levels of experience, judges are most conservative in deciding that they recognize an object; they grow more liberal with experience. This strategy minimizes the number of recognition errors. With increases in d', the error rate decreases. The second strategy is to fix the response criterion. The dashed lines show the predicted error rate for three different levels of criteria when $d' = 0.5$. The criterion values shown are centered so that a criterion of 0 indicates a criterion exactly between the two distributions. All possible fixed criterion strategies can be derived from the figure using the tangent of the inverted U of a given Bayesian observer. These fixed-criterion judges commit a large number of errors regardless of experience.

The Impact of Recognition Sensitivity on Inferential Accuracy

With this level of specification, the impact of d' on the judge's inferential accuracy can be assessed. Besides an imperfect assessment of all possible mediators that control a judge's experience in the environment, the model attributes sensitivity as a second psychological source for the imperfect surrogate correlation. To demonstrate this, I calculated the average predicted recognition validity, α, from the model using the four pairs that the judge would answer using the recognition heuristic (see Table 1). Specifically, I set $N = 100$ and calculated α for each level of n_e as a function of d' for four different levels of A (.6, .7, .8, and .9) when there was a nonzero probability of employing the recognition heuristic. Figure 5 shows the average α averaged across n_e. Lower levels of sensitivity can bring about substantial decrements to the accuracy of recognition-based inferences. For example, a d' with a value of 0.5 decreases an experience validity from $A = .8$

to a recognition validity of $\alpha = .59$; a d' of 1.0 decreases it to $\alpha = .63$; and a d' of 2.0 decreases it to $\alpha = .70$.

With its impact on both the recognition and knowledge heuristics, the recognition process also influences the less-is-more effect. With $N = 100$ and $A = .8$, the panels in Figure 6 plot the predicted proportion of correct inferences as a function of experience, n_e, when judges take the same inferential test described earlier. Three values of d' (0.5, 1.0, and 2.0) are varied across the columns, and three values of z (⅓, ⅔, or 1) are represented in the rows. Within each panel, there are four different levels of B, which is the probability of a correct inference when both items are experienced. Recall that judges are making recognition decisions to maximize the probability of a correct detection; therefore, the criterion, k, changes with each level of experience. As a result of this response rule, the oldest and youngest brothers for all levels of d' and z have the same scores as their counterparts in Figure 3. The bottom two rows of Figure 6 illustrate that decreasing sensitivity
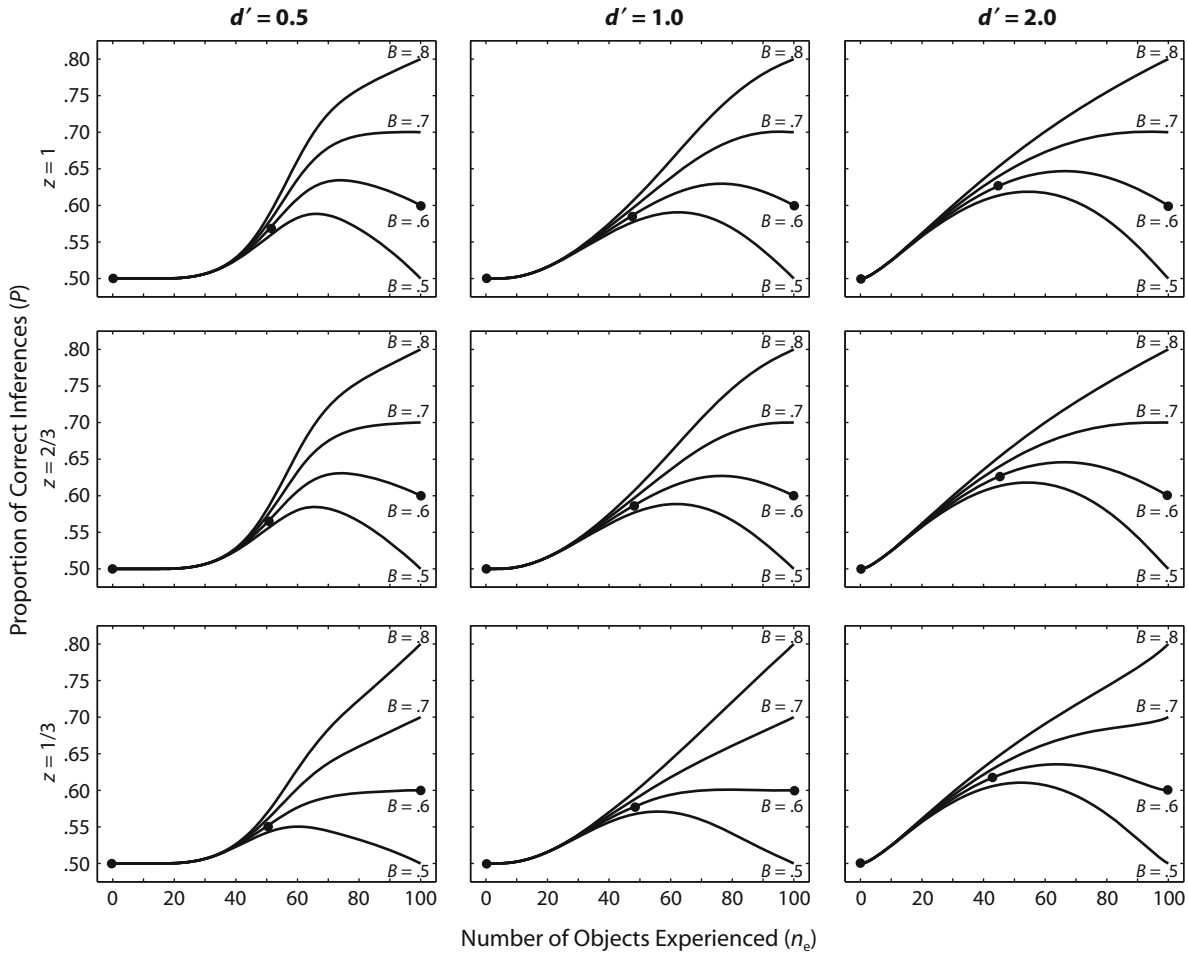

Figure 6. A reanalysis of the less-is-more effect assuming an equal-variance Gaussian signal detection model of the recognition process. The plots show different levels of sensitivity, d', and the proportion of experienced items with at least one positive cue, z. Three values of d' are varied across the columns and three values of z are varied up the rows. Within each panel, there are four different levels of B. Judges are making recognition decisions so as to maximize the probability of a correct detection; therefore, the criterion, k, changes with each level of experience. The bottom two rows of the figure illustrate that decreasing sensitivity (from right to left) can mitigate the less-is-more effect and give way to the less counterintuitive more-is-more effect. In the top row, the knowledge heuristics benefit indirectly from the judges' false alarms and help to maintain the less-is-more effect.

(from right to left) can mitigate the less-is-more effect and give way to the less counterintuitive more-is-more effect. The less-is-more effect is not existent for $B = .7$ for all six panels in the bottom two rows. Even for $B = .6$, where the less-is-more effect tends to persist, the magnitude of the effect is diminished. Consider, for example, the middle brother when $d' = 2.0$ and $z = \frac{2}{3}$. He is predicted to score 63%, a mere three-point advantage over his more experienced brother. Recall that, originally, partial ignorance gave the middle brother an eight-point advantage.

The signal detection model also reveals that the influence of sensitivity depends on the distribution of positive cue values across items. The top row of Figure 6 shows that when $z = 1$, the less-is-more effect is robust against lower levels of sensitivity. However, even here the range of experience for which the less-is-more effect is predicted to occur is reduced. In the figure, when $d' = 2.0$, the middle brother outscores his brother with an expected score of 64%. When $d' = 0.5$, he would get 58% correct, and when $d' = 1.0$ he would get 60% correct. In comparison, a person who has experienced $n_e = 75$ of the objects is predicted to score 66%, 65%, and 65% correct for d' values of 0.5, 1.0, and 2.0, respectively.

Why is the less-is-more effect predicted regardless of sensitivity and why does the range of experience where it is predicted become restricted when $z = 1$? The answer is that judges with intermediate levels of experience ($50 < n_e < 100$) have more false alarms and consequently more (hit, false alarm) pairs. When $z = 1$, the accuracy for these pairs is A, which protects judges from their mistaken recognition decisions. However, z and B are correlated. As the proportion of items with positive cue knowledge, z, decreases, the probability of a correct inference when both items are experienced, B, will also decrease. This is because the discrimination rule for the knowledge heuristic specifies that the judge guess when neither item has a positive cue value. Consequently, the very thing that can make the less-is-more effect more probable—less cue knowledge and therefore lower knowledge heuristic accuracy—also counteracts it to make it more susceptible to a judge's recognition sensitivity.

Discussion

When presented with a two-alternative forced choice task, judges who employ the recognition heuristic first look at one item in a pair, decide whether they recognize it or not, and then look at the other item in the pair and decide whether they recognize it or not (Goldstein & Gigerenzer, 2002). If one item is recognized and the other is not, then judges can use the recognition heuristic. Integrating a signal detection model of recognition memory with the recognition heuristic is a first step in understanding how recognition memory contributes to the recognition heuristic. Moreover, the integrated model makes it possible to assess how experience in the world translates to recognition serving as an accurate inferential predictor. To do so, the implicit assumption that recognition perfectly reflects experience was relaxed. That is, when recognizing objects, people make false alarms and misses, and these errors impact the inferences people make. For example,

a German professor using recognition to pick teams in the 2006 NCAA Division I basketball tournament might think that he recognizes Northwestern State—a 14th seed in the tournament—and pick it to win a game or two. Chances are, however, that this school, located in Natchitoches, Louisiana, has been mistaken for a school with a similar name located in Evanston, IL: Northwestern University.[14]

Expanding the analysis to an inferential test showed that as the error rate of recognition increased, the accuracy of the recognition heuristic fell. The errors also changed what tool or heuristic was used to answer the test questions. Furthermore, the ecologically rational goal for judges was shown to be one in which they adjusted their recognition response criteria according to their experience with the reference class. Finally, the less-is-more effect was shown to depend on judges' sensitivity to the difference between the familiarity of experienced versus novel items as well as the distribution of cue knowledge.

In the discussion that follows, I summarize how the signal detection framework makes it possible to assess how recognition sensitivity and response rules give rise to these errors and interact with the recognition heuristic to produce the observed inferential performance of the judge. I also discuss how the model can be used as a bridge between the recognition heuristic and other theories of recognition and how it can be informative for evaluating other knowledge heuristics. Finally, I will return to Simon's (1956) principle of bounded rationality and address how the signal detection framework can move the recognition heuristic closer to this principle.

Recognition sensitivity and response rules. From the perspective of theories of recognition memory, when there is an ecological correlation in the environment, the covert familiarity with objects from the environment can be correlated with the target variable of the inferential task. If judges want to exploit this correlation with the recognition heuristic, they have to transform their familiarity into a binary recognition decision (see Slegers, Brake, & Doherty, 2000, for an alternative framework to transform continuous knowledge cues into binary cues). The response rule that judges use to make this transformation depends on their goals and expectations during the task.

The ecologically rational goal in fast and frugal heuristics is "to exploit the structure of the information in the natural environment" (Goldstein & Gigerenzer, 2002, p. 76). Accordingly, to fully exploit the environmental correlation, judges need to adjust their criteria according to their level of experience with the reference class: When they have little to no experience, judges should be the most conservative in recognizing objects; as their experience increases, they should become more liberal. This Bayesian-observer response strategy (see Wickens, 2002) minimizes the error rate and allows judges to fully exploit the association between their familiarity and the inference's target variable (e.g., city population).

Given this response rule, a judge's sensitivity to the difference between experienced and novel items also influences his ability to exploit this correlation. With perfect sensitivity, the recognition validity reflects the validity of

the judge's experience, and with lower sensitivity, the recognition validity systematically decreases away from the validity of experience (A). As a result of this systematic change in the recognition validity, the original condition of the less-is-more effect still holds, $\alpha > \beta$ (Goldstein & Gigerenzer, 2002). Instead of changing this condition, the signal detection model parses this single condition into conditions related to structures of the environment and conditions related to the cognitive processes of judges. In the environment, the accuracy of experience has to be greater than the accuracy achievable when comparing objects that have both been experienced, $A > B$. At the same time, a judge's sensitivity needs to be high enough that recognition based on experience is still more accurate than the knowledge heuristics. For example, in this model, when $A \geq .8$ and $B \leq .6$, the less-is-more effect tends to persist when $d' > 1$, regardless of the distribution of positive cue knowledge, z. A final condition depends on both the environment and the judge: The more positive cue knowledge distributed among the objects in a reference class, the more robust the less-is-more effect against a judge's recognition errors. Recall, this can occur because of the distribution of positive cue values in the environment or because of the lack of cue knowledge on the part of the judge. In this model, values of z greater than approximately .8 tend to counteract lower values of sensitivity and maintain the less-is-more effect.

Admittedly, these are less precise criteria than the original derivation of $\alpha > \beta$. More precise conditions depend on the distributional assumptions of the models. However, disentangling the psychological and environmental contributions to the less-is-more effect continues to move questions about the recognition heuristic from empirical *what* questions (e.g., "What happens when the recognition heuristic . . . ?") toward more theoretically framed *why* questions (e.g.,"Why can recognition make accurate inferences?"). See Wallsten (1996) and Wallsten, Erev, and Budescu (2000) for a similar argument about investigating the cognitive processes involved in confidence judgments.

A bridge to theories of recognition memory. The framework can also serve as a bridge to larger and more expanded theories of recognition memory, like global memory models (Raaijmakers & Shiffrin, 2002). The REM global memory model (Shiffrin & Steyvers, 1997) has the most features in common with the signal detection framework that I have developed in this article. According to REM, an error-prone image or vector of feature values is stored in memory after an item has been studied. At retrieval, a probe is generated containing the features of the to-be-recognized test item, and this probe is matched with all images stored in episodic memory to produce a determination of the likelihood that the test item has been previously learned. In a manner similar to how this article describes the response strategy of ecologically rational judges, REM then calculates the posterior odds that the test item is old and makes a decision on the basis of this estimate, just like the response strategy of ecologically rational judges using the recognition heuristic.

Besides offering a competing recognition model for Schooler and Hertwig's (2005) ACT–R model of the recognition heuristic, an REM implementation of the recog-

nition heuristic has potential benefits for both the fast and frugal heuristics as well as REM. Although respondents can move the criterion to other values, the REM framework usually deals only with a default criterion set at equal odds (Shiffrin & Steyvers, 1997). Extending REM to encompass the ecological framework provides a natural prediction that judges adjust their criterion according to their prior experience with a reference class. In turn, REM can bring a more precise understanding to how judges learn sequentially about objects in the environment, develop a familiarity with items via encoding processes, and subsequently produce an estimate of an item's familiarity via retrieval processes.

Recognition's impact on noncompensatory and compensatory heuristics. Since most of the fast and frugal implementations of knowledge heuristics make use of the recognition principle, the signal detection model of recognition memory is also informative about their performance. Indeed, the analysis has shown that false alarms and misses systematically change the expected accuracy for noncompensatory rules like "take the best," "take the last," and "minimalist."

At the same time, there is recent debate about whether judges use noncompensatory single-variable inferential rules or more compensatory multiple-variable rules. This debate has occurred with regard to recognition (see Bröder & Eichler, 2006; Newell & Shanks, 2004; Oppenheimer, 2003; Pachur & Hertwig, 2006; Pohl, 2006; Richter & Spath, 2006) and the knowledge heuristic "take the best" (see Bröder, 2000, 2003; Lee & Cummins, 2004; Newell & Shanks, 2003). In terms of recognition, the debate is not on whether recognition is used as a predictor variable at all, but on how it is used to make inferences. If recognition is not used in a single variable manner, then other compensatory heuristics, such as tallying, weighted tallying, a unit-weight linear model, a weighted linear model, and multiple regression, can still use it (Gigerenzer & Goldstein, 1996).

Even with these heuristics, however, the use of recognition memory influences their performance if recognition is used as a predictor variable. This influence occurs because, for these heuristics, the response is a function of the sum of the cue values. For weighted tallying, the weighted linear model, and multiple regression, the cue values are typically weighted according to the ecological cue validity (see Gigerenzer & Goldstein, 1996). Consistently, if recognition is given a weight based on the ecological validity of experience, A, then recognition errors will have an increasingly detrimental impact on the accuracy of these heuristics as A increases. In contrast, tallying and the unit-weighting model ascribe equal weights to all of the cue values, deemphasizing the contribution of any one cue. Consequently, these heuristics will be the most robust against recognition errors or, more generally, errors in using any other inferential cue.

Conclusion. Using signal detection theory, this article has modeled judges' recognition of objects when they use the recognition heuristic to make an inference about the objects. The model shows that recognition ability plays a crucial role in the performance of the recognition heu-

ristic and the subsequent heuristics that use it. This is an important extension, because the recognition heuristic specifically, and fast and frugal heuristics in general, are supposed to be constructed according to Simon's principle of bounded rationality, which states: "To describe, predict, and explain the behavior of a system of bounded rationality we must both construct a theory of the system's processes and describe the environments to which it is adapting" (Simon, 1990, p. 6). This principle was developed in response to theories of economics being independent of the actor and solely a function of the environment. That is, theories such as expected utility assume that people make choices to maximize their own utility in a given environment, but ignore the cognitive abilities of the actor. The development of the recognition heuristic to date puts a great deal of emphasis on the characteristics of the environment that make it adaptive. Accounting for the recognition process within the heuristic better heeds the cognitive abilities of judges and moves the model closer to the principle of bounded rationality.

AUTHOR NOTE

REFERENCES

Axelrod, R. (1985). *The evolution of cooperation*. New York: Basic Books.

Banks, W. P. (1970). Signal detection theory and human memory. *Psychological Bulletin*, **74**, 81-99.

Borges, B., Goldstein, D. G., Ortmann, A., & Gigerenzer, G. (1999). Can ignorance beat the stock market? In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 59-72). New York: Oxford University Press.

Bröder, A. (2000). Assessing the empirical validity of the "take-the-best" heuristic as a model of human probabilistic inference. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **26**, 1332-1346.

Bröder, A. (2003). Decision making with the "adaptive toolbox": Influence of environmental structure, intelligence, and working memory load. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **29**, 611-625.

Bröder, A., & Eichler, A. (2006). The use of recognition information and additional cues in inferences from memory. *Acta Psychologica*, **121**, 275-284.

Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review*, **50**, 255-272.

Craik, F. I. M., & McDowd, J. M. (1987). Age-differences in recall and recognition. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **13**, 474-479.

Dhami, M. K., Hertwig, R., & Hoffrage, U. (2004). The role of representative design in an ecological approach to cognition. *Psychological Bulletin*, **130**, 959-988.

Dougherty, M. R. P., Franco-Watkins, A. M., & Thomas, R. (in press). The psychological plausibility of the theory of probabilistic mental models and the fast and frugal heuristics. *Psychological Review*.

Dougherty, M. R. P., Gettys, C. F., & Ogden, E. E. (1999). MINERVA-DM: A memory processes model for judgments of likelihood. *Psychological Review*, **106**, 180-209.

Gigerenzer, G. (2001). The adaptive toolbox. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality* (pp. 37-50). Cambridge, MA: MIT Press.

Gigerenzer, G. (2004). Fast and frugal heuristics: The tools of bounded rationality. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 62-88). Oxford: Blackwell.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, **103**, 650-669.

Gigerenzer, G., & Goldstein, D. G. (1999). Betting on one good reason: Take the best heuristic. In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 75-96). New York: Oxford University Press.

Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review*, **98**, 506-528.

Gigerenzer, G., Todd, P. M., & the ABC Research Group (1999). *Simple heuristics that make us smart.* New York: Oxford University Press.

Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, **91**, 1-67.

Glanzer, M., & Adams, J. K. (1985). The mirror effect in recognition memory. *Memory & Cognition*, **13**, 8-20.

Glanzer, M., Adams, J. K., Iverson, G. J., & Kim, K. (1993). The regularities of recognition memory. *Psychological Review*, **100**, 546-567.

Goldstein, D. G., & Gigerenzer, G. (1999). The recognition heuristic: How ignorance makes us smart. In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 37-58). New York: Oxford University Press.

Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, **109**, 75-90.

Gonzalez, R., & Nelson, T. O. (1996). Measuring ordinal association in situations that contain tied scores. *Psychological Bulletin*, **119**, 159-165.

Gronlund, S. D., Ohrt, D. D., Dougherty, M. R. P., Perry, J. L., & Manning, C. A. (1998). Role of memory in air traffic control. *Journal of Experimental Psychology: Applied*, **4**, 263-280.

Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*, **95**, 528-551.

Hirshman, E. (1995). Decision processes in recognition memory: Criterion shifts and the list-strength paradigm. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 302-313.

Jacoby, L. L., Woloshyn, V., & Kelley, C. (1989). Becoming famous without being recognized: Unconscious influences of memory produced by dividing attention. *Journal of Experimental Psychology: General*, **118**, 115-125.

Lee, M. D., & Cummins, T. D. R. (2004). Evidence accumulation in decision making: Unifying the "take the best" and the "rational" models. *Psychonomic Bulletin & Review*, **11**, 343-352.

Murdock, B. B. (1997). Context and mediators in a theory of distributed associative memory (TODAM2). *Psychological Review*, **104**, 839-862.

Nelson, T. O. (2003). Relevance of unjustified strong assumptions when utilizing signal detection theory. *Behavioral & Brain Sciences*, **26**, 351.

Newell, B. R., & Shanks, D. R. (2003). Take best or look at the rest? Factors influencing "one-reason" decision making. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **29**, 53-65.

Newell, B. R., & Shanks, D. R. (2004). On the role of recognition in decision making. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **30**, 923-935.

Oppenheimer, D. M. (2003). Not so fast! (and not so frugal!): Rethinking the recognition heuristic. *Cognition*, **90**, B1-B9.

Pachur, T., & Hertwig, R. (2006). On the psychology of the recognition heuristic: Retrieval primacy as a key determinant of its use. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **32**, 983-1002.

Pohl, R. F. (2006). Empirical tests of the recognition heuristic. *Journal of Behavioral Decision Making*, **19**, 251-271.

Raaijmakers, J. G. W., & Shiffrin, R. M. (2002). Models of memory. In H. Pashler & D. Medin (Eds.), *Stevens' Handbook of experimental psychology: Vol 2. Memory and cognitive processes* (3rd ed., pp. 43-76). New York: Wiley.

Ratcliff, R., Clark, S. E., & Shiffrin, R. M. (1990). List-strength effect: Data and discussion. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **16**, 163-178.

Ratcliff, R., Gronlund, S. D., & Sheu, C. F. (1992). Testing global memory models using ROC curves. *Psychological Review*, **99**, 518-535.

Richter, T., & Spath, P. (2006). Recognition is used as one cue among others in judgment and decision making. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **32**, 150-162.

Schooler, L. J., & Hertwig, R. (2005). How forgetting aids heuristic inference. *Psychological Review*, **112**, 610-628.

Shepard, R. N. (1967). Recognition memory for words, sentences, and pictures. *Journal of Verbal Learning & Verbal Behavior*, **6**, 156-163.

Shiffrin, R. M., Huber, D. E., & Marinelli, K. (1995). Effects of category length and strength on familiarity in recognition. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 267-287.

Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM—retrieving effectively from memory. *Psychonomic Bulletin & Review*, **4**, 145-166.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, **63**, 129-138.

Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, **41**, 1-19.

Slegers, D. W., Brake, G. L., & Doherty, M. E. (2000). Probabilistic mental models with continuous predictors. *Organizational Behavior & Human Decision Processes*, **81**, 98-114.

Snook, B., & Cullen, R. M. (2006). Recognizing national hockey league greatness with an ignorance-based heuristic. *Canadian Journal of Psychology*, **60**, 33-43.

Standing, L. (1973). Learning 10,000 pictures. *Quarterly Journal of Experimental Psychology*, **25**, 207-222.

Volz, K. G., Schooler, L., Schubotz, R. I., Raab, M., Gigerenzer, G., & von Cramon, D. Y. (2006). Why you think Milan is larger than Modena: Neural correlates of the recognition heuristic. *Journal of Cognitive Neuroscience*, **18**, 1924-1936.

Wallsten, T. S. (1996). An analysis of judgment research analyses. *Organizational Behavior & Human Decision Processes*, **65**, 220-226.

Wallsten, T. S., Erev, I., & Budescu, D. V. (2000). The importance of theory: Response to Brenner (2000). *Psychological Review*, **107**, 947-949.

Wickens, T. D. (2002). *Elementary signal detection theory.* Oxford: Oxford University Press.

NOTES

1. Spearman correlations are reported for continuity. A more appropriate and meaningful measure of association, given the structure of the data, may be Kendall's τ (see Gonzalez & Nelson, 1996). According to Kendall's τ, the recognition correlation is .43, the ecological correlation is .63, and the surrogate correlation is .53.

2. Other integration algorithms do not capitalize on the recognition heuristic per se but can use recognition as a cue within their frameworks. See Gigerenzer and Goldstein (1996) for more details.

3. An alternative sampling method would be proportional sampling, in which stimuli are sampled from a reference class on the basis of their relative frequency of occurrence in the environment (Dhami et al., 2004).

4. A necessary assumption for these calculations is that each object in the reference class has a unique value on a target variable.

5. The discrimination rate of a heuristic is the relative frequency with which the heuristic discriminates between any two objects from the reference class; it is directly related to the discrimination rate of a cue (see Gigerenzer & Goldstein, 1996).

6. This activation level can go under the guise of many labels, such as familiarity, strength, confidence, and activation.

7. An alternative use of memory would be consistent with a two-alternative forced-choice task (see Wickens, 2002) in which the inference problem involves determining which of two stimuli (San Antonio or San Diego) is more familiar. However, under these forced choice assumptions, the recognition heuristic, along with the entire flow of processes shown in Figure 1, breaks down. This is because familiarity, as a continuous predictor, would always discriminate between the two alternatives, and a judge would never guess or resort to a knowledge heuristic. As a result, a judge's expected accuracy would only be a function of the correlation between familiarity and the target variable, and the less-is-more effect would never be predicted (see Dougherty, Franco-Watkins, & Thomas, in press, for such an implementation).

8. More than 10 are possible if more than one distribution characterizes the familiarity of experienced objects.

9. An alternative and perhaps more appropriate way to make these derivations is in terms of relative experience, n/N. To be consistent with past work, however, I will continue to make derivations in terms of absolute experience, n.

10. Memory researchers typically find that the distributions tend to have unequal variances (see Nelson, 2003; Ratcliff, Gronlund, & Sheu, 1992). This has no substantial effect on the conclusions reached here.

11. In environments in which each item is not equally likely to occur (e.g., in a study using proportional sampling; see note 3), this expression would change to reflect the nonuniform nature of the distribution. Regardless, the general behavior of the response criterion developed here would still remain; that is, with increasing experience, a judge would adjust the response criterion to be more liberal in recognizing items. The adjustment of the criterion, however, would not move in equal intervals and would depend on the distribution of objects in the particular environment.

12. The criterion values are centered so that a criterion of 0 lies exactly between the two distributions.

13. Other goals are entirely possible and can give rise to different predictions. For instance, more conservative responding could go hand in hand with more experience if judges were increasingly punished for false alarms. Another instance that can lead to conservative responding occurs when the response criterion is set fixed but relative to the signal distribution, and the signal distribution shifts up in familiarity (see Hirshman, 1995).

14. Unfortunately for this article and this author, the German professor's error was a benefit. Northwestern State upset the third-seeded University of Iowa in their first-round game of the 2006 tournament.