Psychological Review
2011, Vol. 118, No. 3, 393–437

0033-295X/11/$12.00 DOI: 10.1037/a0024143

Cognitive Niches: An Ecological Model of Strategy Selection

Julian N. Marewski and Lael J. Schooler
Max Planck Institute for Human Development, Berlin, Germany

How do people select among different strategies to accomplish a given task? Across disciplines, the strategy selection problem represents a major challenge. We propose a quantitative model that predicts how selection emerges through the interplay among strategies, cognitive capacities, and the environment. This interplay carves out for each strategy a *cognitive niche*, that is, a limited number of situations in which the strategy can be applied, simplifying strategy selection. To illustrate our proposal, we consider selection in the context of 2 theories: the simple heuristics framework and the ACT–R (adaptive control of thought—rational) architecture of cognition. From the heuristics framework, we adopt the thesis that people make decisions by selecting from a repertoire of simple decision strategies that exploit regularities in the environment and draw on cognitive capacities, such as memory and time perception. ACT–R provides a quantitative theory of how these capacities adapt to the environment. In 14 simulations and 10 experiments, we consider the choice between strategies that operate on the accessibility of memories and those that depend on elaborate knowledge about the world. Based on Internet statistics, our model quantitatively predicts people's familiarity with and knowledge of real-world objects, the distributional characteristics of the associated speed of memory retrieval, and the cognitive niches of classic decision strategies, including those of the fluency, recognition, integration, lexicographic, and sequential-sampling heuristics. In doing so, the model specifies when people will be able to apply different strategies and how accurate, fast, and effortless people's decisions will be.

Keywords: strategy selection, memory, decision making, ACT–R, simple heuristics

Supplemental materials: http://dx.doi.org/10.1037/a0024143.supp

When you move to a new country you realize how important a sense of recognition is in making everyday decisions, such as what TV, yogurt, or car to buy. This is especially true when you are unable to read labels or talk with sales staff. Which car is likely to be of better quality, a Daewoo or a Volkswagen? Now, most Americans would opt for the Volkswagen, largely because it feels more familiar, and many of us follow the *heuristic*, or simple decision strategy, that highly recognizable, familiar brands are likely to be quality brands. Corporations spend billions of advertising dollars each year to exploit this tendency. Yet advertising influences not only a sense of recognition but also knowledge about products. For example, the classic Volkswagen ad featuring the German word *Fahrvergnügen* (joy of driving) reminded people that Volkswagens are German engineered. People often use a product's country of origin as an indication of quality, and, in fact, many of the other heuristics people use operate in this way; that is, they rely on knowledge rather than on recognition.

How and when people use different heuristics is a key question in the *simple heuristics research program* (Gigerenzer, Todd, & the ABC Research Group, 1999; Todd & Gigerenzer, 2000; see Marewski, Gaissmaier, & Gigerenzer, 2010a, 2010b, for a recent overview). This framework for studying decision making assumes that the mind is equipped with an "adaptive toolbox" consisting of a collection of heuristics. The thesis that people possess a repertoire of strategies to choose from has been similarly formulated in many areas of study, including choice (Einhorn, 1970; Fishburn, 1980; J. W. Payne, 1976; J. W. Payne, Bettman, & Johnson, 1988, 1993; Rapoport & Wallsten, 1972), social interactions (Erev & Roth, 2001; Stahl, 1999), mathematical skill (Siegler, 1988), probability judgments (Ginossar & Trope, 1987), categorization (Smith, Patalano, & Jonides, 1998), memory (Coyle, Read, Gaultney, & Bjorklund, 1998), word recognition (Eisenberg & Becker, 1982), and question answering (Reder, 1987). If one adopts this view, a major problem is then to determine how people select from these strategies to solve given tasks—the *strategy selection problem*.

Julian N. Marewski and Lael J. Schooler, Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Berlin, Germany.

Our thanks go to John R. Anderson, Nick Chater, Jakob Gierl, Gerd Gigerenzer, Kevin Gluck, Julian Heister, Bettina von Helversen, Andrew Howes, Konstantinos Katsikopoulos, Stefan Lindner, Magnus Persson, Carmi Schooler, Niels Taatgen, and the members of the ABC Research Group for many constructive comments as well as to Henry Brighton for pointing us to the literature on biological robustness. In particular, we would like to thank Anita Todd for editing the article, Gregor Caregnato and Britta Dittmann for helping to conduct the experiments, and Jakob Gierl and Julian Heister for programming the simulations. Finally, we would like to thank Jacqueline Schulze and Jürgen Rossbach for redrawing a portion of our figures for printing purposes. This work was first presented at the 2006 meetings of the Brunswik Society and the Society for Judgment and Decision Making in Houston, Texas.

Correspondence concerning this article should be addressed to Julian N. Marewski, who is now at University of Lausanne, Faculty of Business and Economics (HEC), Department of Organizational Behavior, Quartier UNIL-Dorigny, Bâtiment Internef, 1015 Lausanne, Switzerland. E-mail: marewski@mpib-berlin.mpg.de

The strategy selection problem represents a major theoretical challenge in the cognitive and decision sciences, but it is also a stumbling block for theories in other disciplines, including, for instance, economics and machine learning, where the problem presents itself in terms of the selection of actions, algorithms, operators, routines, or production rules (e.g., *for the cognitive and decision sciences*, Anderson & Betz, 2001; Beach & Mitchell, 1978; Bröder & Newell, 2008; Bröder & Schiffer, 2003, 2006; Busemeyer & Myung, 1992; Christensen-Szalanski, 1978; Cisek, 2007; Cokely, Kelley, & Gilchrist, 2006; Dror, Busemeyer, & Basola, 1999; Einhorn & Hogarth, 1981; Fu & Anderson, 2006; Gaissmaier, Schooler, & Rieskamp, 2006; Gigerenzer, Hoffrage, & Goldstein, 2008; Glöckner & Betsch, 2008; Gray, Sims, Fu, & Schoelles, 2006; Lovett & Anderson, 1996; Marewski, 2010; Marewski, Schooler, & Gigerenzer, 2010; McAllister, Mitchell, & Beach, 1979; Newell & Shanks, 2003; J. W. Payne et al., 1988; Rieskamp & Otto, 2006; Taatgen & Wallach, 2002; *for economics*, Haruvy & Stahl, 2007; Haruvy, Stahl, & Wilson, 2001; Stahl, 2000; Stahl & Wilson, 1994; *for machine learning*, Provost & Buchanan, 1995; Vilalta & Drissi, 2002; Walker, 2000).

In this article, we contribute to solving the strategy selection puzzle by proposing a quantitative model that predicts how selection emerges through the interplay among the strategies, the workings of basic cognitive capacities, such as memory, and the structure of the environment in which we humans live. To illustrate our proposal, we consider strategy selection in the context of two ecological theories, the aforementioned simple heuristics framework and the *adaptive control of thought—rational* theory of cognition (ACT–R; Anderson et al., 2004; Anderson & Lebiere, 1998). From the heuristics framework we adopt the notion of an adaptive toolbox of simple decision strategies. ACT–R provides a quantitative theory of the workings of memory, time perception, and other cognitive capacities essential to the strategies. By showing how the strategies and the capacities interact with the environment, we demonstrate how each strategy can be applied to a limited number of situations, simplifying strategy selection.

Approaches to Strategy Selection

In the cognitive and decision sciences, *cost–benefit* approaches have been used to explain how people select from a repertoire of strategies (Beach & Mitchell, 1978; Christensen-Szalanski, 1978; J. W. Payne et al., 1988, 1993). People trade a strategy's costs (e.g., the cognitive effort and time involved in adopting a strategy) against its benefits (e.g., its accuracy in making decisions) by, for instance, applying a metastrategy to master trade-offs between making accurate and effortless decisions. However, as has been argued by Rieskamp and Otto (2006; see also J. W. Payne et al., 1993), assuming a metastrategy can result in the recursive homunculus problem of deciding how to decide.

Most more recent approaches have abandoned the assumption of a metastrategy while still focusing on accuracy, cognitive capacity, or time and effort as determinants of strategy selection (e.g., Bröder, 2003; Bröder & Schiffer, 2006; Newell & Shanks, 2003; Pachur & Hertwig, 2006; Volz et al., 2006). For example, Rieskamp and Hoffrage (1999, 2008) showed that people use particularly fast and simple heuristics under time pressure (see also J. W. Payne et al., 1988). Understanding such determinants of strategy selection has proved useful; yet, at the same time, these approaches have also been criticized for being vague (e.g., Glöckner, Betsch, & Schindler, 2010), and in fact, most of them are not realized as quantitative models (see the various proposals for strategy selection in the simple heuristics framework; e.g., Mata, Schooler, & Rieskamp, 2007).

Many quantitative theories of strategy selection stress the role of learning (Busemeyer & Myung, 1992; Erev & Roth, 2001; Rieskamp, 2006). To illustrate this, according to Rieskamp and Otto's (2006) model, people choose among different strategies as a function of learning how accurate the strategies are. In ACT–R, similar reinforcement mechanisms favor the selection of cognitive processes (i.e., production rules) that implement successful strategies (Fu & Anderson, 2006; Lovett & Anderson, 1996). Although these models help us understand strategy choice in learning contexts, they do not explain how an adaptive, systematic selection of strategies emerges in the absence of feedback.

A Complementary Approach to Strategy Selection: The Cognitive Niche Framework

Explicitly or implicitly, most strategy selection theories assume that the strategies' accuracy or the time and effort involved in using them shape the selection. Yet there may be another way to select strategies—a complement to these theories—that follows the lead of Gibson's (1979) ecological approach to visual perception. Gibson hypothesized that, at every moment in time, an organism's environment poses opportunities and demands for different courses of action—food affords eating, predators require evasion, and caves offer protection. According to Gibson, the visual system allows organisms to perceive these functional meanings of the environment. He called them *affordances*. Gibson never developed his ideas into a computational model, and he might not have envisioned them applying beyond visual perception. The model we introduce below does this. It asks how cognitive capacities such as memory shape opportunities for applying different strategies—for instance, for inferring the quality of cars. According to our proposal, strategy selection emerges in two ways.

First, when faced with a decision task, people will not actually choose among all the strategies in their repertoire. Rather, the workings of the cognitive system will limit the number of strategies that can be executed, simplifying strategy selection by reducing the consideration set of applicable strategies. This limitation arises from how the strategies and cognitive capacities, such as memory or time perception, interact and represent regularities in the environment. Following Gibson (1979), who conceptualized an animal's environmental niche as a set of affordances, we use the term *cognitive niche* to refer to the situations in which a strategy is applicable, or afforded. Below, we model these niches in terms of landscapes that quantify the probability of an individual being able to apply a strategy as a function of the interplay between the cognitive system and the environment. We show that this interplay results in partially disjunctive probability landscapes for different strategies, that is, in *nonoverlapping* niches, which limits strategy selection to situations where the niches of two or more strategies overlap. As we demonstrate, this aspect of the selection process does not require the assumption of cost–benefit calculations, learning, or any of the other mechanisms of strategy choice that have populated the literature over the past decades.

Second, where different strategies' niches overlap, strategy selection depends on the cost–benefit, learning, and other selection mechanisms that assume accuracy, effort, and time as currencies of strategy choice. Our model maps out how these currencies are shaped by the interplay between the cognitive system and the environment. In doing so, the model refines an assumption made most notably by the cost–benefit approaches, namely, that people typically face trade-offs between making accurate, effortless, and quick decisions (but see Busemeyer, 1993; Gigerenzer & Goldstein, 1996). Although our model allows for such trade-offs, it also specifies the regions within a strategy's niche where they need not occur.

To illustrate our cognitive niche framework, we consider how people choose among classic strategies that can be used, for example, to infer which of two cars is likely to be of better quality. These strategies can be roughly divided into two types. The first leads to decisions based on *knowledge* about the world, say, about a manufacturer's country of origin. These decisions depend on the content of what is retrieved from memory. The second type depends on the characteristics of the retrieval. Such strategies are guided by the *accessibility* (e.g., Bruner, 1957; Higgins, 1996; Kahneman, 2003; Koriat, 1993; Tulving & Pearlstone, 1966) of memories, that is, the ease with which mental content comes to mind. This can take the form, for instance, of a sense of recognition of brand names.

Decisions From the Accessibility of Memories

Decades of research in decision making, social psychology, and memory have shown that the accessibility of memories guides judgment (e.g., Hertwig, Pachur, & Kurzenhäuser, 2005; Jacoby & Dallas, 1981; S. J. Payne, Richardson, & Howes, 2000; Reber, Schwarz, & Winkielman, 2004; Tversky & Kahneman, 1973; Whittlesea, 1993; Winkielman, Schwarz, & Belli, 1998). Across disciplines, several related concepts have been investigated: *recognition* (e.g., Goldstein & Gigerenzer, 2002; Pleskac, 2007), which we use here to distinguish between *objects*, such as brands people believe they have heard of before and those they have not; *familiarity* (e.g., Dougherty, Franco-Watkins, & Thomas, 2008; Mandler, 1980), which is frequently used to denote degrees of recognition; and a sense of *fluency* (e.g., Jacoby & Dallas, 1981; Oppenheimer, 2008) or *availability* (Tversky & Kahneman, 1973), which often refers to the ease or speed with which memories are retrieved.[1] To give a few examples, recognition and familiarity can influence consumer choice (Coates, Butler, & Berry, 2004, 2006). Fluent processing stemming from previous exposure can increase the perceived truth of repeated assertions (Begg, Anas, & Farinacci, 1992; Hertwig, Gigerenzer, & Hoffrage, 1997) and the perceived fame of names (Jacoby et al., 1989), and assessments of memories themselves can partly be based on a sense of fluency (Jacoby & Dallas, 1981).

Moreover, relying on the accessibility of memories can often lead us to make accurate *inferences* about the world, such as about uncertain future events or unknown quantities. For instance, our recognition of soccer teams and politicians can be used to forecast their future success in competitions and political elections, respectively (Gaissmaier & Marewski, 2011; Herzog & Hertwig, 2011; Pachur & Biele, 2007). Our recognition of universities and cities allows us to predict their quality and size, respectively (Hertwig & Todd, 2003; Reimer & Katsikopoulos, 2004), and the ease with which we retrieve the names of billionaires reflects their fortunes (Hertwig, Herzog, Schooler, & Reimer, 2008).

The Adaptive Toolbox: A Repertoire of Decision Strategies

There are ecological reasons for why the accessibility of memories can be so informative in inference: The press, the Internet, and other *environmental mediators* (Goldstein & Gigerenzer, 2002) make it likely that we will encounter objects (e.g., brands) that score high on a *criterion* of interest (e.g., quality) more frequently in our environment than those that score low. As a result, objects with high criterion values (e.g., high-quality brands) tend to be more accessible in memory. Thus, we can rely on a sense of accessibility to accurately infer which objects are likely to score higher on the criterion (see Figure 1).

A key tenet of the simple heuristics framework is that the people choose from heuristics that exploit how memory and other cognitive capacities reflect such environmental regularities. Many of these heuristics can be categorized according to whether they operate on knowledge and/or on the accessibility of memories. Table 1 provides an overview of the various heuristics considered in this article.

Knowledge-based heuristics use features of objects as *cues* to make inferences about the objects. For instance, to infer which of two cars is of better quality, a person could rely on one of the classic *integration strategies* (e.g., Dawes, 1979; Dawes & Corrigan, 1974). Here, we refer to two of them as *tally1* and *tally2*. Following tally1, a person would consider one or more cues, such as where the cars are produced, their reliability, or their price. An object's value on a cue is coded as *positive* (e.g., high price, high reliability), as *negative* (e.g., low price, poor reliability), or as *unknown* (e.g., unknown price, unknown reliability). For each car, the person adds up the number of positive cue values and subtracts the number of negative cue values and then infers that the car with the higher sum is better. Tally1 is akin to Gigerenzer and Goldstein's (1996) *unit-weight linear model*, which is a special case of the *equal-weight linear model* (Huber, 1989) and Einhorn and Hogarth's (1975) *Model 2*. Tally2, resembling the *tallying model* (Gigerenzer & Goldstein, 1996), sums only positive cue values, ignoring negative ones (see Alba & Marmorstein, 1987, and J. W. Payne et al., 1993).

[1] The term *recognition* is often used to refer to a person's ability to distinguish between stimuli that have been presented in an experiment (e.g., in a study list) and those that have not. Usually, a person is familiar with the stimuli before participating in the experiment (e.g., stimuli could be the names CARTER and BUSH). Here, we adopt Goldstein and Gigerenzer's (2002; see also Pleskac, 2007; Schooler & Hertwig, 2005) usage of the term to refer to a person's ability to discriminate between novel stimuli that have not been heard of before (e.g., the name SCHILKE) and those that have (e.g., CLINTON). For a discussion of the similarities between this notion of recognition and the notions of availability/fluency, see Schooler and Hertwig (2005) and Hertwig et al. (2008). See also Jacoby and Dallas (1981, p. 333) and Jacoby, Kelley, Brown, and Jasechko (1989, p. 328) for similarities between their fluency/familiarity concept and availability, and Hertwig et al. (2005), Schooler and Hertwig, and Sedlmeier, Hertwig, and Gigerenzer (1998) for differences in notions of availability.

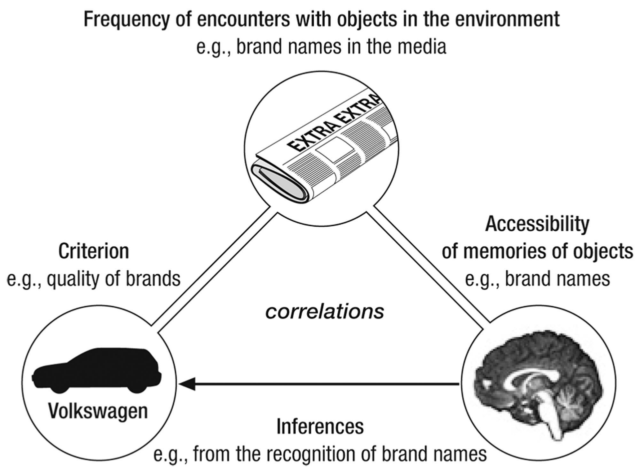

Figure 1. An unknown criterion (e.g., car quality) is reflected by a mediator (e.g., the press). The mediator makes it more likely that a person will encounter objects with larger criterion values than those with smaller ones (e.g., high-quality cars are mentioned in the news more often than low-quality cars). As a result, objects with larger criterion values are more accessible in memory (e.g., more easily recognized), and ultimately, objects' accessibility in memory can be relied upon to infer the criterion (e.g., recognition can be used to infer quality). These connections can be thought of in terms of three correlations: between the criterion and the mediator, between the mediator and a person's access to memories, and between the accessibility of memories and the criterion.

Lexicographic strategies (e.g., Tversky, 1969) search through cues sequentially, basing a decision, for instance on the first cue on which the two objects differ (e.g., where one car has a negative value and the other a positive). To illustrate this, *take-the-best* (Gigerenzer & Goldstein, 1996), which resembles Tversky's (1972) classic *elimination-by-aspects heuristic*, considers cues in the order of how likely they are to help a person make accurate inferences.

The *sequential-sampling* strategies *take-the-first-cue*, *take-the-first-value1*, and *take-the-first-value2*, in turn, exploit the order in which cue values come to mind, independent of the cues' accuracy. These three knowledge-based strategies thus take the accessibility of the knowledge into account. They are inspired by classic sequential-sampling models (Busemeyer & Townsend, 1993; Ratcliff & Smith, 2004), but the sampling of knowledge is based on the speed of memory retrieval rather than on random walks or other sampling procedures. Related models have been discussed by Bröder and Gaissmaier (2007), Dougherty et al. (2008), Gaissmaier (2007), Johnson and Raab (2003), Lee and Cummins (2004), and others.

An example of a purely accessibility-based strategy is the *fluency heuristic*, which has been defined in various ways (e.g., Jacoby & Brooks, 1984; Whittlesea, 1993). Here, we use Schooler and Hertwig's (2005) model. Building on these earlier definitions, their version of the fluency heuristic can be seen as a computational instantiation of Tversky and Kahneman's (1973) classic notion of availability. When inferring the quality of cars, for instance, the heuristic opts for the brand whose name is perceived as having been more quickly retrieved. This sense of *retrieval fluency* for brand names, such as Daewoo or Volkswagen, is modeled in terms of the time it takes to recognize the brand names, that is, in terms of a graded sense of recognition. Another accessibility-based strategy is Goldstein and Gigerenzer's (2002; see also Gigerenzer & Goldstein, 2011) *recognition heuristic*, which derives inferences from a binary, rather than a graded, sense of recognition, choosing recognized brands over those that have never been heard of before.

Structure of This Article

By modeling heuristics in accord with the ACT–R cognitive architecture, quantitative precision can be lent to memory and other capacities the heuristics exploit. At the same time, ACT–R provides a theory of how these capacities interact with the environment. Such an integrative theoretical framework is essential to modeling how a strategy's niche emerges.

Specifically, in what follows, we describe an ACT–R model of how the environment shapes memory. As the perception of the timing of memory retrieval is critical to the accessibility-based strategies, we also model time perception. Using environmental data, such as how often an object (e.g., Volkswagen) is mentioned in the media, the resulting *integrated model* of memory and time perception predicts different types of behavioral data. These data include (a) whether people recognize an object and (b) whether they additionally know something about it (e.g., that Volkswagen is a German company). These behavioral data also include (c) the speed of memory retrieval (e.g., recognition time of Volkswagen) as well as (d) people's perception thereof. The models' predictions about memory retrieval and time perception, in turn, enable detailed quantitative predictions (e) about the niches where different strategies are applicable and (f) in what regions of their niches the strategies help a person make accurate inferences, as well as (g) in what regions using the strategies requires little effort and time.

Owing to the complexity of these modeling efforts, we provide a roadmap in Figure 2. As is illustrated in the figure's Panels A and B, first we describe how we calibrate and test our ACT–R memory model, predicting memory retrieval from the environment. Second, we explain how we integrate our memory model with Taatgen, Van Rijn, and Anderson's (2007) ACT–R model of time perception (Panel C) and how we let the resulting integrated model orchestrate the various knowledge- and accessibility-based decision strategies from the adaptive toolbox (Panel D). As we show, this allows us to predict where the different strategies' niches overlap and where they do not, simplifying strategy selection (Panel E).

In the remainder of the article, we then use the integrated model to reveal how the mechanisms of strategy selection that have been discussed in the literature come into play, zooming in on situations where the niches of different strategies overlap (Panels F–I). In doing so, we turn to (a) accuracy, effort, and time as determinants of strategy choice and (b) those regions in the strategies' niches in which people can avoid the effort–accuracy trade-offs that have been so often assumed in the literature. Specifically, we describe how we use the integrated model to predict in which regions the fluency heuristic would help a person make accurate, effortless, and fast inferences (Panel F). Furthermore, we examine how such predictions about accuracy, effort, and time provide insight into the selection of the fluency heuristic where this heuristic's niche overlaps with those of knowledge-based strategies (Panel G). The article also turns to the niches of the recognition heuristic and the knowledge-based strategies, respectively, exploring the regions in

Table 1
Strategies Considered

Strategy	Description
Accessibility-based	
Recognition heuristic	Infers recognized objects to be larger than unrecognized ones; compares objects on a binary sense of recognition
Fluency heuristic	Infers objects that are perceived as more quickly recognized to be larger than those that are perceived as more slowly recognized; compares objects on a graded sense of recognition
Knowledge-based	
Tally1	Unit-weight linear integration strategy; adds up the number of positive cue values and subtracts the number of negative cue values for A and B,[a] inferring the object with the higher sum to be larger
Tally2	Unit-weight linear integration strategy; adds up the number of positive cue values for A and B, inferring the object with the higher sum to be larger
Take-the-best	Lexicographic strategy; considers cues in the order of their validities,[b] basing an inference on the first cue on which A has a positive and B a negative or unknown value, and inferring A to be larger than B
Knowledge-based, but sequentially samples knowledge as a function of its accessibility	
Take-the-first-cue	Sequential sampling strategy; first samples cues according to the speed with which positive cue values are retrieved, basing an inference on the first cue on which A has a positive value and B a negative or an unknown one, and inferring A to be larger than B; if this does not lead to an inference then cues are searched according to the speed with which negative cue values are retrieved, basing an inference on the first cue on which A has a negative value and B an unknown one, and inferring B to be larger than A
Take-the-first-value1	Sequential sampling strategy; relies on comparisons of objects by sampling cue values independently for each object according to the speed with which cue values are retrieved, basing an inference on the first comparison of objects in which A has a positive value on one cue and B a negative or an unknown one on the same or another cue, and inferring A to be larger than B
Take-the-first-value2	Sequential sampling strategy; relies on comparisons of objects by sampling cue values independently for each object according to the speed with which cue values are retrieved, basing inferences *either* on the first comparison of objects in which A has a positive value on one cue and B a negative or an unknown value on the same or another cue, *or* on the first comparison in which A has an unknown value and B a negative one, inferring A to be larger than B

[a] A and B are two objects. A and B's values on cues are coded as 1 (positive cue value), −1 (negative), or 0 (unknown). For example, when inferring the quality of cars, high price would be a positive cue value, poor price a negative cue value, and having no information about a car's price an unknown cue value. [b] The validity of a cue is the probability that A has a higher value on the criterion (e.g., car quality) to be inferred than B, given that A has a positive value on that cue and B a negative or unknown one.

which less effort and time are associated with greater accuracy (Panels H, I). We start by explaining how decision strategies can be modeled in accord with ACT–R.

ACT–R: An Integrative Framework

ACT–R is a unified, quantitative theory of cognition that accounts for many phenomena, ranging from memory performance (e.g., Anderson, Bothell, Lebiere, & Matessa, 1998) to time perception (Taatgen et al., 2007), driving (Salvucci, 2006), flying (Gluck, Ball, & Krusmark, 2007), and the learning of mathematics (Ritter, Anderson, Koedinger, & Corbett, 2007). It distinguishes between declarative memory (knowing *that*) and procedural memory (knowing *how*). The basic declarative representation is the *chunk*, which can simply be memories of objects' names, such as car brands, or elaborate knowledge, say, about attributes of a car. Procedural knowledge, in turn, is modeled with sets of *production rules* (i.e., if–then rules) whose conditions (i.e., the "if" parts of the rules) are matched against, for example, the contents of declarative memory. If the conditions of a production rule are met, then the rule is eligible to be selected. If selected, the actions specified in the "then" part of the rule are carried out.

Using ACT–R to Model Decision Strategies

We think of strategies such as tally1 or the fluency and recognition heuristics in terms of production rules. In doing so, we specify what it means for a strategy to be *applicable*, or afforded: A strategy is applicable if all the production rules implementing it have the potential to be selected.[2] The subsequent descriptions of the strategies can best be thought of as stylized, simpler versions of the production rules that are required to implement the strategies in ACT–R.

In modeling the fluency and recognition heuristics, we follow Anderson et al. (1998) and Schooler and Hertwig (2005) in assuming that (a) a chunk's retrieval implies recognizing the asso-

[2] Some strategies were originally defined as including guessing rules that fire when the strategy cannot be applied. For instance, following Gigerenzer and Goldstein's (1996) formulation of take-the-best, a person choosing between two objects would randomly guess if unable to retrieve knowledge about them. Likewise, according to Schooler and Hertwig's (2005) fluency heuristic, a person would randomly guess if unable to detect a difference in recognition times between two objects. Here, we conceptualize guessing as a separate decision strategy. That is, none of the strategies listed in Table 1 include guessing rules.

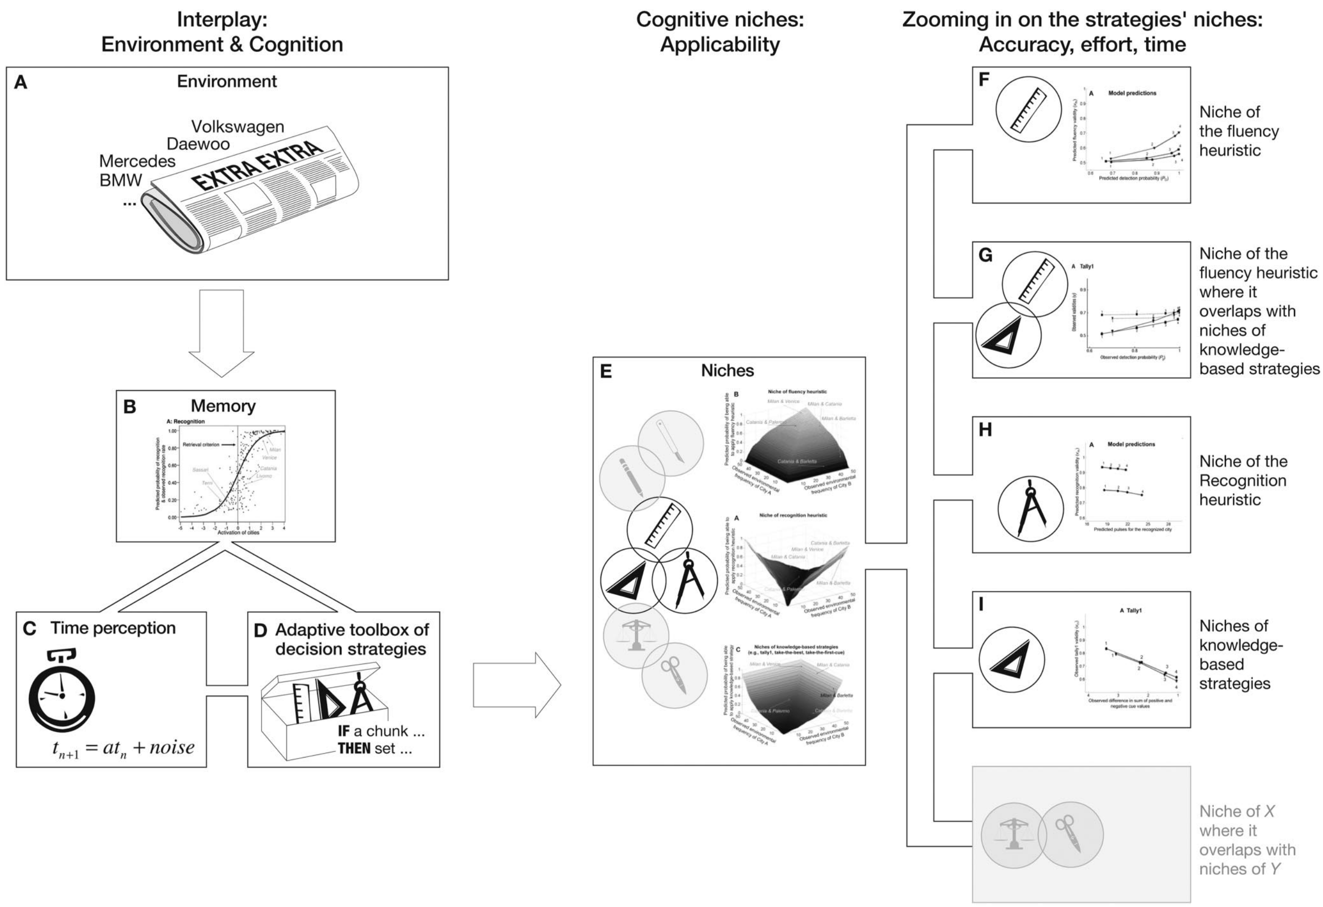

Figure 2. Roadmap of the modeling carried out. The environment shapes the contents of memory; memory, in turn, interacts with other cognitive capacities such as time perception and people's repertoire of decision strategies. This interplay carves out a cognitive niche for each strategy where it can be applied. The strategies' niches do not overlap completely, simplifying strategy selection. Where the niches overlap, the selection among such simultaneously applicable strategies depends on the cost–benefit, learning, and other selection mechanisms that assume accuracy, effort, and time as currencies of strategy choice. These currencies, too, depend on the interplay between the cognitive system and the environment. In the figure, the repertoire of strategies is symbolized by an adaptive toolbox, which contains different measurement tools, such as a ruler or a compass. Each tool pertains to a cognitive niche; niches partially overlap. The strategies' niches and their overlap are symbolized by Venn diagrams.

ciated object and (b) the more quickly the chunk is retrieved, the greater the sense of recognition and ease of retrieval. If a person's goal is to infer which of two objects, A or B, has a larger value on a given criterion, then the production rules that implement the recognition heuristic can be summarized as follows:

> If a chunk with the name A is retrieved and A is recognized, and a chunk with the name B cannot be retrieved,
>
> Then set the goal to respond that A has the larger value on the criterion.

For instance, when inferring which of two car brands, Volkswagen or Daewoo, is of higher quality, a person may recognize the name Volkswagen but not the name Daewoo. Following the recognition heuristic, the person would infer that the Volkswagen is of higher quality.

When both car brands are recognized, a person can use the fluency heuristic. This heuristic consists of production rules that fire when the memories of two objects are *both* available but are perceived to have been recalled with different speeds. A person using this heuristic would infer car brands that are more quickly recognized to be of higher quality:

> If two chunks with the names A and B are retrieved and recognized, and A is perceived to have been recognized faster than B,
>
> Then set the goal to respond that A has the larger value on the criterion.

Finally, the production rules implementing knowledge-based strategies can fire when a person remembers attributes of objects beyond just recognizing their names. For instance, tally2's production rules can be summarized as follows:

> If two chunks with the names A and B are retrieved and recognized, and chunks with knowledge representing more positive cue values for A can be retrieved than can be retrieved for B,
>
> Then set the goal to respond that A has the larger value on the criterion.

The sense of recognition, ease of retrieval, and knowledge on which the recognition, fluency, tally2, and other heuristics depend can be modeled with ACT–R's declarative memory.

Using ACT–R to Model Declarative Memory

In modeling declarative memory, ACT–R maps the frequency and temporal spacing of encounters with objects onto a mental currency called *activation*. Roughly speaking, the more often an object and/or information about it is encountered, the more strongly the corresponding chunks are activated. The more strongly a chunk is activated, (a) the more likely it is that it will be retrieved and (b) the less time it will take to retrieve it. The activation, A_i, of a chunk i is determined by a combination of the *base-level activation of the chunk*, B_i, and the S_{ji} *units of activation* the chunk receives from each of the j elements of the current context:

$$A_i = B_i + \sum_j S_{ji}. \tag{1}$$

The base-level activation B_i of a chunk i depends on the environmental pattern with which the corresponding object occurs, and the S_{ji} units of activation depend on the co-occurrence of the ith and jth element of the context. Where the object has been encountered n times in the past, with the kth encounter having occurred t_k time units in the past, and with the decay parameter, d, capturing the degree of forgetting in declarative memory,

$$B = \ln\left(\sum_{k=1}^{n} t_k^{-d}\right). \tag{2}$$

ACT–R assumes that there is stochastic variability in the activation level of a chunk. This variability is modeled as a logistic distribution, which adds noise to the retrieval process. With s representing noise in the retrieval process, the probability P that a chunk will be retrieved—that is, that its activation exceeds the *retrieval criterion*, τ—is modeled with a logistic function:

$$P = \frac{1}{1 + e^{-(A-\tau)/s}}. \tag{3}$$

The retrieval time, $T_{\text{retrieval}}$, of a chunk with activation A is

$$T_{\text{retrieval}} = F\, e^{-A}, \tag{4}$$

where F is a scaling parameter and $A > \tau$. In what follows, we explain which behavioral data we modeled using these and other components of ACT–R.

Overview of the Behavioral Data to Be Modeled With ACT–R

In the simple heuristics and related frameworks, a classic paradigm for studying people's repertoire of decision strategies is two-alternative choice. For instance, in the *cities problem*, the goal is to judge which of two cities is likely to have more inhabitants (e.g., Dougherty et al., 2008; Gigerenzer & Goldstein, 1996; Hogarth & Karelaia, 2006). In 10 experiments, we modeled strategy selection for this and related problems, examining what strategies people are able to apply to make inferences about cities' population and fame, countries' gross domestic product, companies' market capitalization, diseases' fame (i.e., notoriety), and politicians' fame. All experiments are described in detail in Appendix A. Descriptive statistics and a list of all stimulus materials are provided in the supplemental Online Materials A and B, respectively.

As summarized in Table 2, our experiments consisted of four modular tasks. First, in *inference tasks* we asked people to make inferences about objects, for example, which of two cities has more inhabitants. These tasks allowed us to examine when people rely on different strategies, for instance, when they use the fluency heuristic as opposed to knowledge-based strategies. Because we did not provide information about the objects, people had to make inferences by relying on their memory contents as they had been shaped by the environment outside of the laboratory, prior to participating in our study. Three additional tasks allowed us to examine these naturally acquired memory contents, offering a snapshot of people's declarative memories. Specifically, we used *recognition tasks* to assess people's recognition and recognition times of the objects. For example, we asked people whether they had heard of a city name prior to participating in our studies, measuring the time it took them to make this recognition judgment. Whether an object is recognized or not and how fast it is recognized is critical input to decision strategies such as the recognition and fluency heuristics, and as such is essential to modeling these strategies' niches. *General knowledge tasks* enabled us to separate out objects people knew something about from those they had merely heard of, recognizing the name but not knowing anything else about them. As we show below, this is a critical distinction for modeling the overlap between the niches of the fluency heuristic and knowledge-based strategies. *Cue-knowledge tasks* allowed us to further specify the knowledge people have about the objects, for instance, whether people know if cities such as Venice and Milan have international airports. Because strategies such as tally1 and take-the-best operate on such knowledge, assessing this information is essential for modeling them.

To provide an integrated picture of a person's declarative memory, all objects assessed in one task of an experiment were also assessed in all other tasks used in that experiment. As Table 3 shows, most experiments included recognition and general knowledge tasks, and a few of them additionally included inference and cue-knowledge tasks. As Table 3 also shows, our experiments differed in terms of the objects used as stimulus materials. Experiments 1–3 and 10 made use of cities as stimuli, and Experiments 4–9 used countries, companies, diseases, and politicians, enabling us to assess our model's ability to predict memory retrieval and the strategies' niches for different kinds of objects.

Table 2
Overview of the Main Experimental Tasks

Task	Schematic description of task	
Inference	The names of two objects (e.g., cities) appear on the computer screen, one on the left side and one on the right. Participants judge which of the two objects has a larger value on a given criterion (e.g., city size). Participants respond by pressing a key corresponding to either (a) the object on the left or (b) the object on the right.	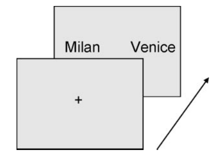
Recognition	The name of one object appears on the computer screen. Participants judge if they have heard or seen the object's name prior to participating in the study, that is, if they recognize the name. Participants respond with (a) "yes" or (b) "no" by pressing corresponding keys.	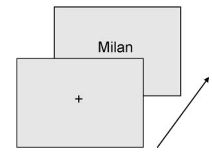
General knowledge	The name of one object appears on the computer screen. Participants indicate how much they know about each object. Participants respond for each object if they have (a) "never heard of it and never seen it before participating in the study"; (b) "heard of it or seen it before participating in the study but don't know anything else about it beyond recognizing its name"; or (c) "heard of it or seen it before participating in the study and know something about the city beyond simply recognizing it." Responses are made by pressing corresponding keys.	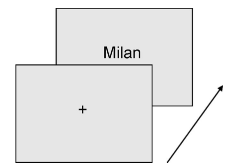
Cue-knowledge	The name of one city appears on the computer screen. The city name is combined with one of four cues,[a] namely, whether the city has an international airport (airport cue), has a university (university cue), is a significant industry site (industry cue), and is a world-famous tourist site (tourism cue). Participants respond to all combinations of cities and cues. Response options are (a) "yes" (e.g., the city does have an airport), (b) "no" (e.g., the city does not have an airport), or (c) "don't know" (e.g., do not know whether the city does have an airport or not). Responses are made by pressing corresponding keys.	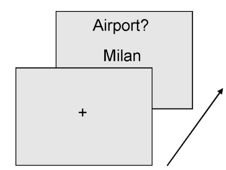

Note. All trials were preceded by a small fixation cross for 1,000 msec. We instructed participants to always fixate on this cross until it disappeared and to respond as quickly and accurately as possible upon stimulus onset. Positive responses were always made with the index finger of the right hand. In all tasks no feedback on the correctness of responses was given until after the experiment. Participants received a guaranteed minimum payment supplemented by a performance bonus for correct responses in the inference and cue-knowledge tasks. Within all tasks, the order of presentation of stimuli was randomized. All the objects included in one task of an experiment appeared in all the other tasks in that experiment. For detailed descriptions of all experiments, see Appendix A.
[a] In Experiment 3, the city name is combined with three cues.

How the Environment Shapes Memory: An Ecological ACT–R Memory Model

We built an ecological ACT–R memory model (throughout referred to as the *memory model*) that uses environmental data to predict what information is likely to have been stored in people's memories prior to their participation in a laboratory experiment (see also Schooler & Anderson, 1997). Specifically, we calculated activations of the chunks stored in participants' memories by estimating how often the participants would have encountered the associated objects' names in the environment outside the laboratory. To estimate the number of encounters, we obtained from the Yahoo search engine the objects' *web frequencies*, N, that is, the number of websites in which each object's name occurred. Appendix B explains how we derived the equations to calculate a chunk's activation from web frequencies. We used these activations to predict our participants' recognition and knowledge of objects as well as the associated recognition times in our experiments. In the recognition tasks, the *predicted recognition probability*, P_R, that an object will be recognized is

$$P_R = \frac{1}{1 + e^{-[(c_R + b_R \ln N) - \tau]/s}}, \tag{5}$$

where c_R is a constant and b_R a scaling parameter that together figure into an estimation of the unknown relation between how often we encounter an object in our environment and the object's web frequency, N. Similarly, in the general knowledge tasks, the *predicted knowledge probability*, P_K, of retrieving knowledge about an object is

$$P_K = \frac{1}{1 + e^{-[(c_K + b_K \ln N) - \tau]/s}}, \tag{6}$$

where c_K is a constant and b_K a scaling parameter that together relate an object's web frequency, N, to the activation of a person's knowledge about that object.

In the recognition tasks, the *predicted recognition time*, $T_{\text{recognition}}$, of an object is how long it takes to judge the object as recognized, measured from the object's presentation until a recognition response is made by pressing a key. We assume that an object's recognition time equals its retrieval time, $T_{\text{retrieval}}$, plus the time it takes to execute additional processes, $T_{\text{perceptual-motor}}$, such as encoding the object's name and pressing a key:

$$T_{\text{retrieval}} = T_{\text{recognition}} - T_{\text{perceptual-motor}} = Fe^{-(c_R + b_R \ln N)}. \tag{7}$$

Table 3
Overview of the Experiments

Experiment	Objects	Inference task	Recognition task	General knowledge task	Cue-knowledge task	Other tasks	No. of participants
1	240 cities[a]	X[b]	X	X	X		49
2	240 cities[a]	X[c]	X	X		X[e]	71
3	240 cities[a]	X[d]	X	X	X	X[e]	55
4	168 countries		X	X			20
5	80 companies		X	X			21
6	54 diseases		X	X			20
7	189 politicians		X	X			19
8	54 diseases		X				118
9	189 politicians		X				83
10	240 cities[a]					X[f]	32

Note. All experiments are described in detail in Appendix A.
[a] To obtain clean measures of recognition time, we only used city names of similar length (see Supplemental Online Materials B). [b] In Experiment 1, participants inferred which of two cities is larger. [c] In Experiment 2, participants inferred which of two cities would be recognized by a larger number of 100 randomly chosen students. Judging other people's recognition of names can be thought of as a judgment about fame. Using fame as a criterion follows the tradition of Jacoby et al.'s (1989) classic "fame" studies, where recognition processes were examined by asking people to judge other people's fame. In one of Experiment 2's two experimental groups, the *instruct group*, we instructed participants to always apply the fluency heuristic when inferring which of two cities is recognized by more students. In a second group, the *no-instruct group*, we did not give any instructions on strategy use. [d] In Experiment 3, we put participants under time pressure, giving them only 900 msec to infer which of two cities is larger. [e] Task assesses whether people consider fluency and/or various knowledge cues useful for making inferences about cities. [f] Task to measure the time it takes to press a key on a computer keyboard.

We use a normal distribution to model noise in perceptual-motor times. We estimated the mean of this distribution (0.59 sec) by using ACT–R to simulate the encoding of an object's name and the pressing of a key in the recognition tasks (see Online Materials C).[3]

As explained in Appendix B, we further assume that the *total retrieval noise*, s, in Equations 3, 5, and 6 can be divided into *criterion noise*, s_τ, attributable to the retrieval criterion and *activation noise*, s_A, attributable to permanent and momentary changes in the activation level of a chunk:

$$s = \sqrt{(s_\tau^2 + s_A^2)}. \tag{8}$$

Using the Ecological Memory Model to Predict Memory Performance From the Environment: Simulation 1

As the interplay between the environment and memory is critical for the emergence of cognitive niches, in a first simulation we examined whether environmental data enable our memory model to account for the probabilities of recognizing objects and retrieving knowledge about them, as well as for the associated recognition times (see Figures 2A and 2B). Specifically, we *calibrated* the memory model to Experiment 1, estimating eight parameters (see Table 4). We then used the memory model to *predict* memory performance in Experiments 2–9, leaving these parameters unchanged throughout this article.[4] Next, we describe this first simulation; details of this and all other simulations are reported in Online Materials D.

Observed data. We call our behavioral measure of the probability of recognizing an object the *observed recognition probability*, computed for each object as the proportion of participants who recognized the object in the recognition tasks. In addition, by pooling the time it took participants to recognize each object, we obtained an *observed recognition time distribution* for each object. Finally, we computed the *observed knowledge probability* for each object as the proportion of participants who not only recognized the object but in addition indicated having knowledge about it in the general knowledge tasks.

Model calibration. To calibrate the memory model to these observed data, we first fit Equation 5 to the observed recognition probabilities of Experiment 1 with a regression analysis, estimating the constant, c_R (−8.52), the total retrieval noise, s (.83), and the scaling parameter, b_R (.70). In doing so, we anchored the activation scale by setting the expected value of the retrieval criterion distribution, τ, to zero; an object with an activation of 0 would have a 50% chance of being retrieved. With these parameters fixed, in a second calibration step we then fit Equation 6 to the observed knowledge probabilities of Experiment 1, estimating the constant, c_K (–9.93), and the scaling parameter, b_K (.68).

In a final calibration step, we fit Equation 7 to the recognition time distributions of Experiment 1, striking a balance between accounting for their 25th, 50th, and 75th percentiles by informally searching the parameter space. Specifically, we estimated the criterion noise, s_τ (.60), the scaling parameter, F (.49), and the standard deviation of the perceptual-motor times (0.12 sec). As implied by Equation 8, the activation noise, s_A, is fixed (.58) once

[3] We also modeled the perceptual-motor times another way, running Experiment 10 to estimate them and feeding this estimated distribution into Simulation 1. The pattern of results remained similar. The experiment is described in Appendix A.

[4] Ideally, such predictions should be made without knowing how well the memory model accounts for the to-be-predicted data, given a set of parameter values. In one situation we did not live up to this ideal. On the basis of the quality of the predictions for the recognition time distributions shown in Simulation 1's Figure 4 (to be discussed later), we went back and set the parameters F and s_τ to the values reported in Table 4.

Table 4
Parameters Estimated for the Memory Model From the Observed Data in Experiment 1

Parameter	Value
Recognition probability, P_R	
Scaling parameter, b_R	.70
Constant, c_R	−8.52
Total retrieval noise, s	.83
Knowledge probability, P_K	
Scaling parameter, b_K	.68
Constant, c_K	−9.93
Retrieval time, $T_{retrieval}$, and perceptual-motor time distributions, $T_{perceptual\text{-}motor}$	
Scaling parameter, F	.49
Criterion noise, s_τ	.60
Standard deviation of perceptual-motor times	.12

Note. ACT–R = adaptive control of thought—rational.

s_τ is estimated, and the parameters c_R and b_R were left at the values estimated previously from the observed recognition probabilities.

Data predicted by the memory model. To assess how well the memory model captures memory retrieval in different experiments, we then used Equations 5–8 with the fixed parameters to predict the recognition and knowledge probabilities as well as the recognition times in Experiments 2–9.

Results. Figure 3A shows the predicted and the observed recognition probabilities for the cities of Experiments 2 and 3 as a function of the expected value of their activation distributions. The inflection point of the curve is at an activation of zero, dividing cities with activations above (e.g., Milan) and below (e.g., Barletta) zero into cities with a high and a low probability of being recognized. For the same cities, Figure 3B shows the predicted and observed knowledge probabilities. Figure 4 depicts the cities' predicted and observed recognition times. Finally, Table 5 reports the correlations between predicted and observed data for all objects and experiments. As can be seen, our ecological memory model gives us a good handle on how memory performance reflects the environment, capturing, for example, such complexities as how the median, spread, and skew of the cities' recognition time distributions change as a function of activation. Next, we exploit this ability of our model to predict how cognitive niches emerge (see Figures 2C–2E).

Quantifying Cognitive Niches: Applicability

When making judgments such as which cars are of good quality, which cities are large, or which politicians are trustworthy, the cognitive niches of applicable strategies are shaped by what is stored in people's memories about the cars, cities, and politicians. We distinguish between three levels of knowledge. First, a person may have never heard of an object, not even recognizing its name (*unrecognized object*). Second, this person may merely recognize the name and know nothing else about the object (*tartle object*; *tartle* is a Scottish verb: People tartle when they recognize someone but cannot recall anything else about that person; Goldstein & Gigerenzer, 2002). Third, the person may retrieve some knowledge about a recognized object (*knowledge object*). Pairing these three kinds of objects, we further distinguish between six memory states a person can be in when comparing two objects: *unrecognized pairs* (two unrecognized objects), *tartle–unrecognized pairs* (a tartle and an unrecognized object), *knowledge–unrecognized pairs* (a knowledge and an unrecognized object), *tartle pairs* (two tartle objects), *tartle–knowledge pairs* (a tartle and a knowledge object), and *knowledge pairs* (two knowledge objects).

With our memory model, we can predict the probabilities of a person being in these six memory states. For instance, Figure 5 shows these probabilities when comparing two cities, plotted as a function of the cities' *environmental frequency* (i.e., as

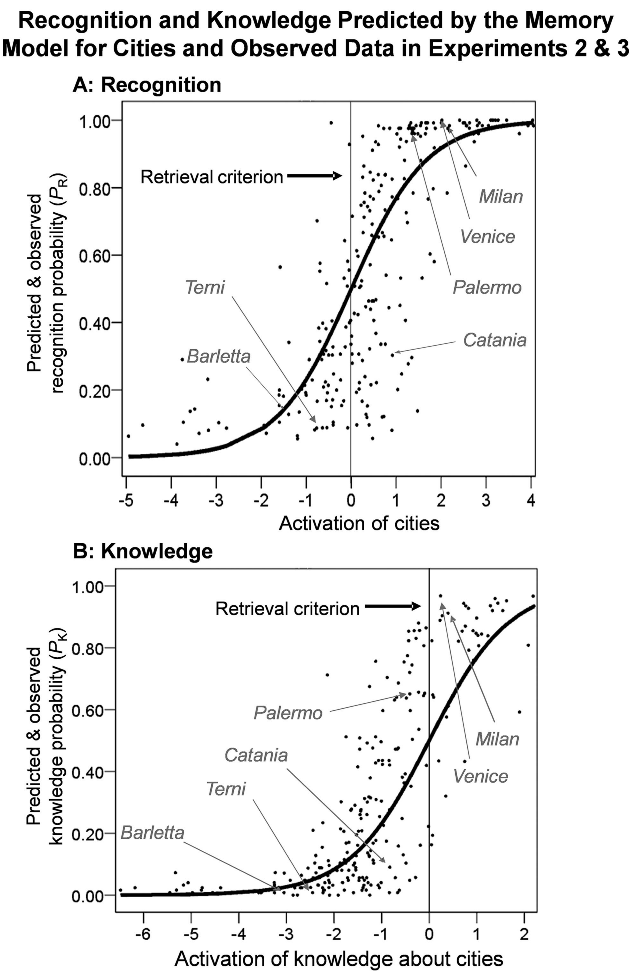

Figure 3. Simulation 1. (A) Predicted (Equation 5; curve) and observed (dots) recognition probabilities (P_R) for 240 cities computed across 126 participants in Experiments 2 and 3. Recognition probabilities are plotted as a function of the expected value of the cities' activation distributions. For instance, Milan has a mean activation of 2.18, the predicted recognition probability is .93, and the observed recognition probability is .98. (B) Predicted (Equation 6; curve) and observed (dots) knowledge probabilities (P_K) for 240 cities computed across 126 participants in Experiments 2 and 3. Knowledge probabilities are plotted as a function of the expected value of the knowledge's activation distributions. For instance, knowledge about Barletta has a mean activation of –3.05, the predicted knowledge probability is .03, and the observed knowledge probability is .01. In each panel, the vertical line shows the expected value of the retrieval criterion distribution.

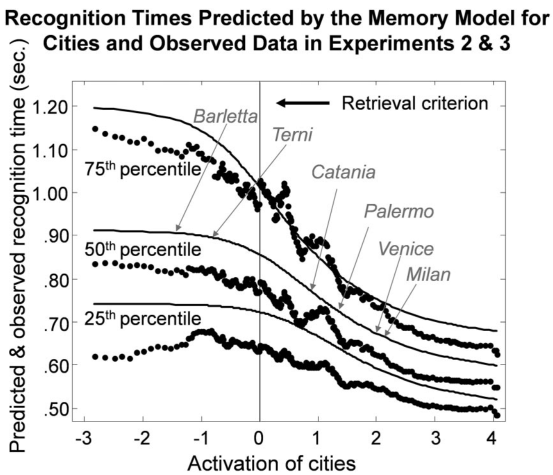

Figure 4. Simulation 1. Percentiles (25th, 50th, 75th) of recognition time distributions of cities predicted by our memory model ($T_{\text{recognition}}$; Equations 7–8; curves) and observed recognition time distributions of cities (dots) computed across participants in Experiments 2 and 3. Recognition times (in seconds) are plotted as a function of the expected value of the cities' activation distributions. Recognition times and activations are smoothed with a running window. The vertical line shows the expected value of the retrieval criterion distribution. For instance, Palermo has a mean activation of 1.35, the 50th percentile of its predicted recognition time distribution is 0.72 sec, and the 50th percentile of the observed recognition time distribution is 0.64 sec.

measured by web frequency). Importantly, each memory state may afford the application of a different set of strategies. To illustrate this, the recognition heuristic can be applied on tartle–unrecognized and knowledge–unrecognized pairs (see Figures 5B and 5C), but not on the other four types, because the heuristic relies on only one of two objects being recognized. Among the conditions for the applicability of the fluency heuristic is that both objects' names be recognized, which is the case for tartle, tartle–knowledge, and knowledge pairs (Figures 5D, 5E, 5F). Tally1, take-the-first-cue, and other strategies require knowledge, as is available in knowledge–unrecognized, tartle–knowledge, and knowledge pairs (Figures 5C, 5E, 5F). Figure 5 suggests that the applicability of all these strategies varies systematically as a function of the environment: The predicted probabilities of being in the six memory states are massed in different areas. Combining such predictions about the interplay between memory and the environment with predictions about time perception allows us to quantify the overlap among different strategies' niches.

Using ACT–R to Model Time Perception

Many strategies depend not on just one cognitive capacity but on several. The fluency, availability, and take-the-first-cue heuristics, for instance, draw on both memory, which influences retrieval and recognition times, and mechanisms for time perception that determine how a person will perceive these times. This is particularly important for carving out the fluency heuristic's niche: People will be able to apply this strategy only if they sense that one object has been recognized in less time than another.

To model time perception, we used Taatgen et al.'s (2007) ACT–R model (henceforth the *timing model*). According to it, a *pacemaker* generates *pulses* at certain time intervals, which are counted by an *accumulator*. The number of pulses accumulated serves the cognitive system as an estimate of time. To assess which of two objects is recognized more quickly, the system will compare the number of pulses that correspond to the recognition time of the first object with the number of pulses that correspond to the recognition time of the second. The system will detect a difference in recognition times if the numbers of pulses differ. Having detected a difference, the system will infer that objects that took fewer pulses to recognize have larger values on the criterion than those that took more pulses to recognize, when the fluency heuristic is used.

The timing model captures Weber's law, an important psychophysical property of time perception that previous studies (e.g., Hertwig et al., 2008; Schooler & Hertwig, 2005) on the fluency and related heuristics did not take into account. The rate at which pulses are generated gradually slows down, while at the same time the intervals between pulses become increasingly affected by noise. As a result, the ability to detect a difference in recognition times between two objects falls nonlinearly as the time it takes to recognize each object increases. For instance, whereas recognition times of 100 msec and 200 msec may correspond to the accumulation of 6 and 11 pulses, respectively, both 600 and 700 msec may well result in the accumulation of 20 pulses, making it possible to detect a difference between 100 and 200 msec, but not between 600 and 700 msec.

More formally, the time interval, t, separating one pulse from the next is a function of the time interval separating the previous two pulses:

Table 5
Correlations Between Predicted and Observed Data in Simulation 1

Data set	Objects	r_{P_R}	r_{P_K}	$r_{T_{\text{recognition}}}$
Experiment 1	Cities	.76	.82	.64
Experiment 2	Cities	.77	.83	.64
Experiment 3	Cities	.76	.81	.62
Experiments 2 & 3	Cities	.77[a]	.82[a]	.70[a]
Experiment 4	Countries	.63	.84	.74
Experiment 5	Companies	.58	.68	.61
Experiment 6	Diseases	.75	.77	.48
Experiment 7	Politicians	.81	.75	.48
Experiment 8	Diseases	.75		.51
Experiment 9	Politicians	.81		.43

Note. Pearson correlations between objects' observed and predicted recognition probabilities (r_{P_R}), between objects' observed and predicted knowledge probabilities ($r_{T_{\text{recognition}}}$), and between medians of objects' observed and predicted recognition time distributions ($r_{T_{\text{recognition}}}$) in Simulation 1 are shown. The correlations are computed across the objects. For computing $r_{T_{\text{recognition}}}$, only those objects that were recognized by at least two participants were included. The magnitude of the correlations lies in the range of those reported when using well-established corpora as a basis for calibrating models to response times in lexical decision tasks (cf. Adelman & Brown, 2008, who calibrated Bayesian, race, and other models to response times, using different experimental data sets and corpora).
[a] Observed recognition probabilities, observed knowledge probabilities, and observed recognition times, respectively, computed by collapsing the observed data across Experiments 2 and 3.

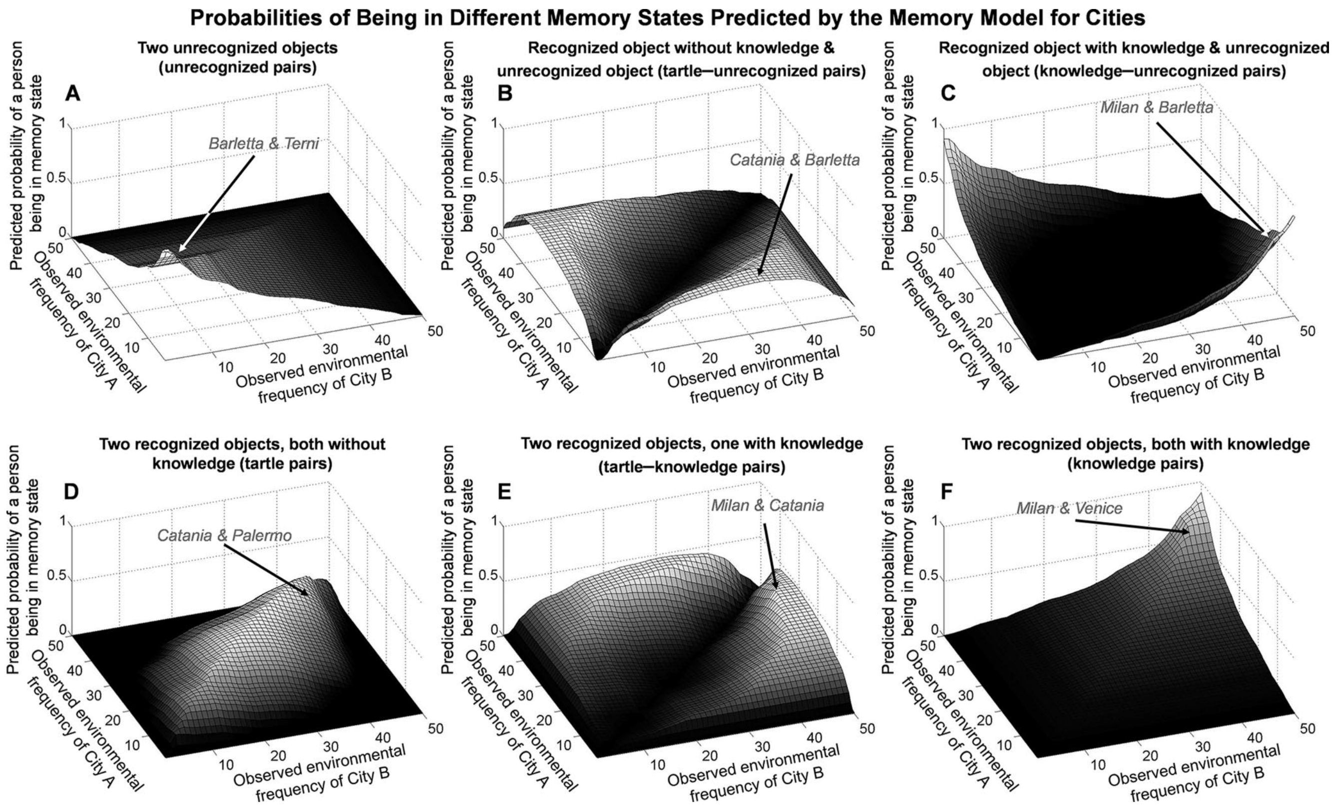

Figure 5. Memory states. The vertical axis shows the probability of a person being in six memory states (A–F) when comparing two cities. To generate the horizontal axes, we ranked the cities according to their environmental frequencies, as measured in terms of their web frequencies, and grouped them into 50 equally sized bins according to rank. These bins are shown on the horizontal axes. For instance, Panel A shows the probability of facing pairs of unrecognized cities and Panel C the probability of facing pairs consisting of an unrecognized city and a recognized city with knowledge. As can be seen, in Panel A the probability of facing two unrecognized cities (e.g., Barletta and Terni) is largest at the intersection of the first bins, where the cities' frequencies are the lowest. In Panel C, the probability of facing a recognized city with knowledge and an unrecognized city (e.g., Milan and Barletta) is highest where the 50th bin intersects the 1st bin, representing a highly frequent and an infrequent city, respectively. To predict these probabilities, the memory model was run 1,500 times, generating 1,500 *hypothetical persons'* recognition and knowledge responses (Equations 5–6). Specifically, according to the predicted recognition probability, P_R, we determined whether a hypothetical (i.e., simulated) person would recognize a city. If the city was recognized, then according to the predicted knowledge probability, P_K, we determined whether that person would additionally know something about it. For each hypothetical person, we then exhaustively paired the cities, grouping the pairs into equally sized bins, indexed by the ranks of each city's environmental frequency. In each bin, we computed the proportion of occurrence of the six different types of pairs. We averaged these proportions across hypothetical persons. These proportions represent the probability of a person being in each of the six memory states.

$$t_{n+1} = at_n + \text{noise}\ (M = 0,\ SD = b \cdot at_n), \qquad (9)$$

where n is the number of accumulated pulses, a and b are constants, and the noise is drawn from a logistic distribution. The duration t_1 is seeded with a start value, t_0. Taatgen et al. (2007) estimated $t_0 = 11$ msec, $a = 1.1$, and $b = 0.015$, stressing that these parameters should be left unchanged. We used these parameter values.

Quantifying Cognitive Niches: Simulation 2

To quantify the overlap among different strategies' cognitive niches, we integrated the timing model and the memory model (throughout referred to as the integrated model) by letting them interact in a simulation. Before turning to the data predicted by this integrated model, let us describe the observed data.

Applying the timing model to the observed data. For each participant of Experiments 2–7, objects were exhaustively paired and grouped into bins according to the ranks of each object's environmental frequency (i.e., as measured by web frequency). For each pair in each bin, we tested which strategy the participant would have been able to apply to make inferences about that pair.

Fluency heuristic. To test whether a participant would have been able to apply the fluency heuristic, we first checked if both objects in a pair were recognized. If so, the probability of a person being able to apply the fluency heuristic is the probability of detecting a difference in recognition times. To estimate this *detection probability*, P_D, we fed the participant's recognition times into the timing model. In each simulation run of this model, we let it count the number of pulses associated with the person's recognition time for each of the objects in a pair. Across simulation runs, for each pair, we counted how often a difference in pulses was detected (T_D) and how often it was not detected (T_{ND}). For each pair of recognized objects, the detection probability is

$$P_D = T_D/(T_D + T_{ND}). \tag{10}$$

If either one or both objects in a pair are unrecognized, by definition, the probability of a person being able to apply the fluency heuristic is zero. Across all pairs in a bin, we averaged the probabilities of the participant being able to apply the fluency heuristic. We averaged these probabilities across all participants.

Knowledge-based strategies. Similarly, for each participant, we estimated the probability that this individual would have been able to apply a knowledge-based strategy, assuming that knowledge based strategies, such as those listed in Table 1, are applicable when knowledge is available about at least one object in a pair. Across the pairs in a bin, we computed this probability as the proportion of pairs for which knowledge was available.

Recognition heuristic. Finally, across the pairs in each bin, we computed the probability that the participant would have been able to apply the recognition heuristic as the proportion of pairs in which the participant recognized one object but not the other.

Applying the timing model to the data predicted by the memory model. We use the terms *observed detection probabilities* and *observed pulses* to indicate that these timing measures have been estimated from participants' behavioral data, rather than from data predicted by our memory model, which we label *predicted detection probabilities* and *predicted pulses*. To generate the predicted data, we ran the memory model 1,500 times, creating 1,500 *hypothetical persons'* predicted recognition and knowledge responses.[5] Specifically, according to the predicted recognition probability, P_R, we determined whether a hypothetical (i.e., simulated) person would recognize an object and determined the object's recognition time, $T_{\text{recognition}}$, by drawing a sample from the object's predicted recognition time distribution. If the object was recognized, then according to the predicted knowledge probability, P_K, we determined whether that person would additionally indicate knowing something about it. To compute the probabilities of the hypothetical persons being able to apply the different decision strategies, we processed their predicted data in the same way as we processed the observed data. For instance, by feeding predicted recognition times into the timing model we predicted the probability of a hypothetical person being able to apply the fluency heuristic.

Results. For the same cities used in Figures 3–5, Figure 6 shows the probabilities that a person will be able to apply the fluency heuristic, recognition heuristic, and knowledge-based strategies in a comparison of two cities, plotted as a function of the cities' environmental frequencies. Figure 7 shows the same graph for countries, companies, politicians, and diseases. Regardless of the type of object, the integrated model of memory and time perception accurately predicts how very different probability distributions emerge for each of the three types of strategies, demonstrating how little the strategies' niches overlap. A person will most likely be able to apply knowledge-based strategies when comparing two objects that occur very frequently in the environment. For instance, when comparing the cities of Milan and Venice, which fall into the upper area of the environmental frequency axes of Figure 6C, a person has roughly a 76% chance of being able to rely on a knowledge-based strategy. In contrast, when comparing the slightly lower frequency cities of Catania and Palermo in Figure 6C, the probability that the same person can use a knowledge-based strategy is about 37%. Instead, for these cities, this person may rely on the fluency heuristic, which is likely to be applicable about 64% of the time (Figure 6B). Finally, when inferring whether a comparatively high-frequency city like Catania or a low-frequency city like Barletta is larger, the same person has a 60% probability of being able to rely on the recognition heuristic (Figure 6A) but only a 17% probability of being able to use the fluency heuristic (Figure 6B) and a 34% probability of being able to use knowledge-based strategies (Figure 6C).

In fact, Figures 6C, 6F, 7C, and 7F overestimate the niches of the six knowledge-based strategies listed in Table 1, because the aggregated data depicted in these figures do not distinguish between the niches of particular knowledge-based strategies, and the niches of these six knowledge-based strategies need not fully overlap. Table 6 summarizes when each of the strategies is applicable. For instance, as we elaborate below, tally2 requires the number of positive cue values to differ between two objects. Tally1, in contrast, requires sums of positive and negative cue values to differ, limiting the overlap between the two strategies' niches.

In short, Figures 6 and 7 show that the niches of different decision strategies do not completely overlap, simplifying the selection among them by constraining the consideration set of applicable strategies. This aspect of the strategy selection process does not assume cost–benefit calculations, reinforcement learning, or any of the other mechanisms of strategy choice that have been discussed in the literature over the past decades. In the remainder of this article, we report how we use the integrated model to examine those parts of the selection process in which these cost–benefit, learning, and other selection mechanisms are likely to come into play, focusing on accuracy, effort, and time as determinants of strategy choice in situations when the niches of different strategies overlap (see Figures 2F–2I).

Zooming in on the Fluency Heuristic's Niche: Regions of Accurate, Effortless Inferences

As explained above, in line with the literature on cost–benefit, learning, and other strategy selection mechanisms, we assume the choice among simultaneously applicable strategies depends on the strategies' accuracy or the effort and time

[5] In Simulations 2, 3, 9, C1, C2, and C3, each run of our memory model can be thought of as generating one hypothetical person's recognition and knowledge responses (including the recognition times) for each of the objects considered. Also, in all simulations where the memory model and the timing model interact, the timing model processes the data predicted by the memory model in the same way as the observed data.

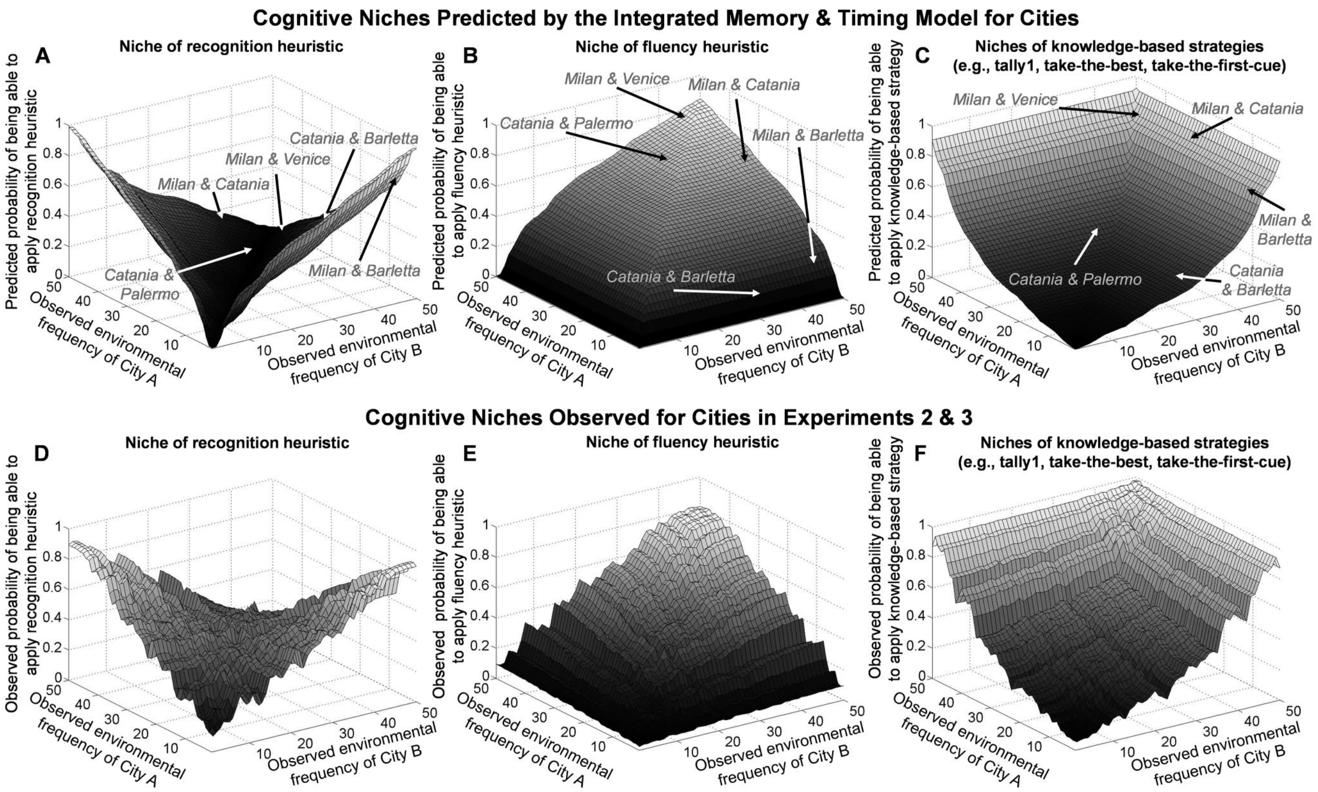

Figure 6. Simulation 2. The cognitive niches of different strategies, predicted by the integrated memory and timing model (Equations 5–10; e.g., P_R, P_K, $T_{recognition}$, P_D) for inferences about cities (A–C) and computed from the behavioral data observed in Experiments 2 and 3 (D–F). The strategies' applicability is plotted in terms of the rank of the cities' environmental frequency. The cities have been grouped into 50 equally sized bins according to their rank. Bin numbers are shown on the horizontal axes. For instance, in Panel B, the fluency heuristic is most likely applicable when comparing high to medium frequency cities (e.g., Milan and Venice or Catania and Palermo) that fall into the 30th to the 50th bin of the horizontal axes.

involved in using them. However, on the basis of this literature one might also believe that people are forced to master trade-offs among these currencies of strategy selection, for instance, between making accurate decisions or effortless ones. Our integrated model of memory and time perception maps out different strategies' niches directly from environmental data, (a) providing the currencies needed by the cost–benefit, learning, and related mechanisms to explain the choice among simultaneously applicable strategies and (b) enabling us to pinpoint areas in a strategy's niche where trade-offs between making accurate, fast, and effortless decisions need not occur.

To start, we use the integrated model to predict in which regions of its niche the fluency heuristic would help a person make accurate, effortless, and fast inferences (see Figures 2F and 2G). The subsequent sections are organized as follows. First we lay out the integrated model's predictions about the accuracy, effort, and time involved in employing the fluency heuristic. Second, using these predictions, we show how accuracy explains the selection between the fluency heuristic and knowledge-based strategies. Third, we turn to the role of accuracy in strategy choice when the fluency heuristic is applicable but knowledge is unavailable. Fourth, we further validate the integrated model's predictions about effort and time and explore how these predictions help to explain the role of effort and time in the selection between the fluency heuristic and knowledge-based strategies.

Before we begin, a comment is warranted. Although the integrated model provides the currencies required by the cost–benefit, learning, and related mechanisms, it does not on its own predict which strategy people will use when different strategies are applicable simultaneously. Deriving such predictions would require implementing one of the cost–benefit, learning, or other mechanisms in our cognitive niche framework, making such predictions dependent on the specifics of the mechanism implemented. For instance, different learning and cost–benefit mechanisms may weight accuracy and effort differently in strategy choice. As the integrated model's predictions about accuracy, effort, and time can, in principle, be used by any of the cost–benefit or learning mechanisms, we have not implemented any specific mechanism here but rather frame our explanations in general, illustrative terms.

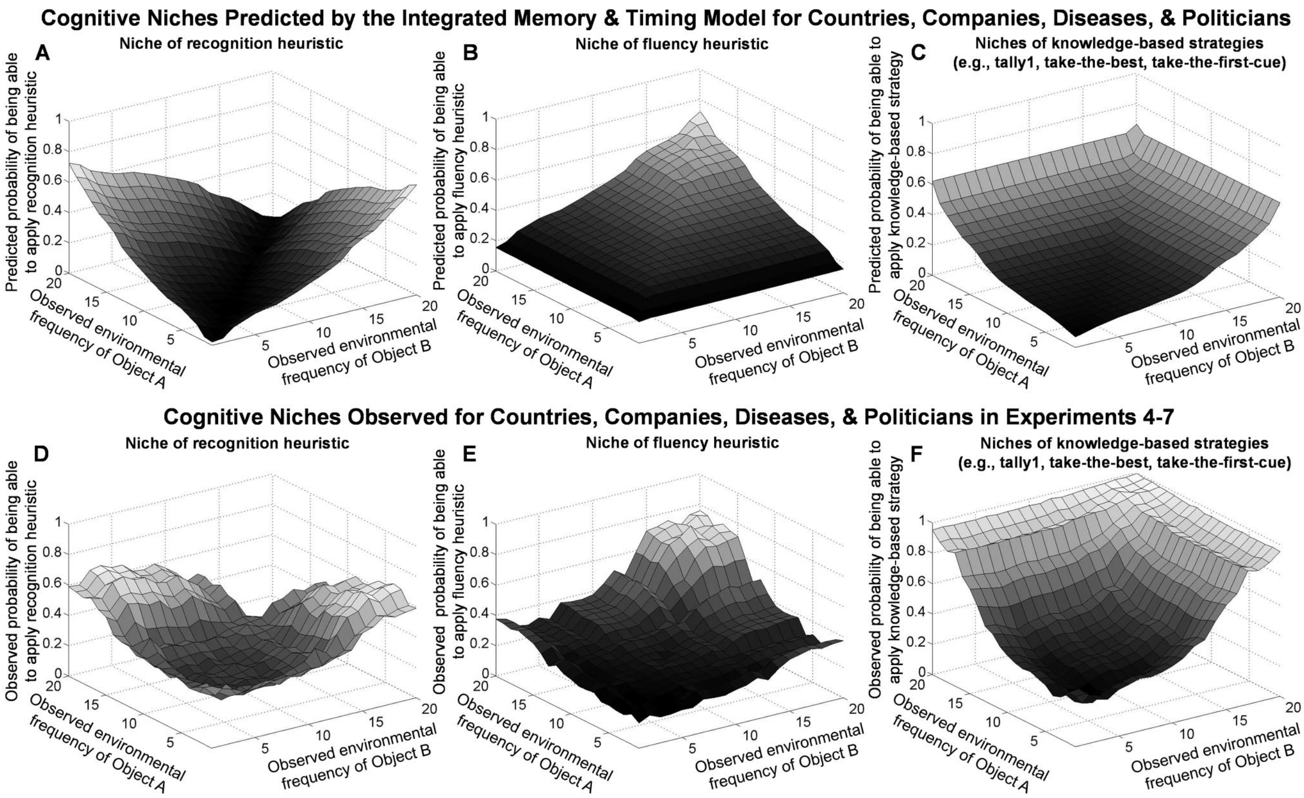

Figure 7. Simulation 2. The cognitive niches of different strategies, predicted by the integrated memory and timing model (Equations 5–10; e.g., P_R, P_K, $T_{recognition}$, P_D) for inferences about countries, companies, diseases, and politicians (A–C) and computed from the behavioral data observed in Experiments 4–7 (D–F). The data are collapsed across the four different types of objects. The strategies' applicability is plotted in terms of the rank of the objects' environmental frequency. The objects have been grouped into 20 equally sized bins according to their rank. Bin numbers are shown on the horizontal axes.

When Does the Fluency Heuristic Help Make Accurate Inferences? Simulation 3

We assume not only the applicability of the fluency heuristic but also the effort and time involved in applying it depend on a person's ability to detect differences in recognition times. For instance, all other things being equal, detecting larger recognition time differences may be easier than detecting smaller ones. Our memory model predicts that the magnitude of recognition time differences correlates with the availability of knowledge, and as such, with the applicability of knowledge-based strategies. Let us look at why this might be.

Encounters with an object are often associated with the acquisition of knowledge about the object. For example, the Volkswagen ad featuring *Fahrvergnügen* not only increases activation and, in turn, a sense of recognition of the brand name, but this German word in addition leads to the encoding or strengthening of the knowledge that Volkswagens are German engineered. This correlation between encountering an object and acquiring knowledge about it makes it likely that people can retrieve knowledge about objects they encounter more frequently than about those they encounter less often, resulting in knowledge objects being more strongly activated than tartle objects. As Figure 4 shows, our memory model predicts that the median and spread of recognition time distributions is larger for weakly activated objects (e.g., Barletta), which tend to be tartle objects, than for strongly activated ones (e.g., Milan), which tend to be knowledge objects. Because the median recognition time is larger for tartle objects than for knowledge objects, recognition times tend to differ more on tartle–knowledge pairs than on knowledge pairs. Recognition times tend to differ more on tartle pairs than on knowledge pairs, because the spread of the recognition time distributions is larger on the tartle pairs. Figure 8 illustrates this phenomenon, showing observed and predicted recognition time differences for tartle, tartle–knowledge, and knowledge pairs. To predict how the accuracy of inferences with the fluency heuristic depends on both (a) the ease of detecting such recognition time differences and (b) the retrieval of knowledge, we ran a simulation.

Applying the timing model to the observed data. The simulation consists of two parts. First, in Experiments 2–7, for each participant, we exhaustively paired the objects. We then used the timing model to compute for the resulting pairs the detection probability, P_D, that a participant would have been able to detect a difference in recognition times between two recognized objects. Recall, this is the probability that a person is capable of applying

Table 6
Applicability of Different Strategies

Strategy	Applicable . . .
Accessibility-based	
Recognition heuristic	If a person recognizes one object but not the other
Fluency heuristic	If two objects differ in terms of the numbers of pulses accumulated for them, such that a person would detect a difference in recognition times between the two objects
Knowledge-based	
Tally1	If the sum of positive and negative cue values, S_{pn}, differs between two objects
Tally2	If the sum of positive cue values, S_p, differs between two objects
Take-the-best	If ordering cues according to their validity yields a cue on which one object has a positive and the other a negative or an unknown value on the same cue
Knowledge-based, but sequentially samples knowledge as a function of its accessibility	
Take-the-first-cue	If sampling cues according to the number of pulses accumulated for positive cue values yields a comparison where one object has a positive value on one cue and the other a negative or an unknown value on the same cue, *or* if searching through cues according to the pulses with which negative cue values are retrieved results in a negative value on one cue for one object and an unknown value for the other object on the same cue
Take-the-first-value1	If sampling cue values independently for each object according to the cue values' pulses yields a comparison in which one object has a positive value on one cue and the other object a negative or an unknown value on another or the same cue
Take-the-first-value2	If sampling cue values independently for each object according to the cue values' pulses results in a comparison in which one object has a positive value on one cue and the other a negative or an unknown value on another or the same cue, *or* if one object has a negative value and the other an unknown one

Note. Just as we do for the fluency heuristic, we also implement the knowledge-based sequential sampling heuristics using the timing model (see Accordance rates section, para. 4, for details). In doing so, we assume that a person will have to decide which cue values have been retrieved most quickly, because the retrieval of cue values may happen in close succession. The timing model provides a mechanism for these discrimination decisions; pulses represent the model's currency for time.

the fluency heuristic; in addition, we assume that P_D also acts as a proxy for the effort and time required to apply it.[6]

Second, we ran another series of runs of the timing model. In each run of this second series, we let the timing model assess whether a person would have detected a recognition time difference between two objects in that run. On the pairs where the person would have detected a difference, we estimated the probability of the person making a correct inference if that person had used the fluency heuristic to infer quantities such as cities' size, countries' gross domestic product, companies' market capitalization, diseases' fame, or politicians' fame. This probability is called the *fluency validity*, v_{fh}, and is computed as

$$v_{fh} = C/(C + I), \tag{11}$$

where C is the number of correct inferences that can be made by using the fluency heuristic, and I is the number of incorrect inferences. Computing the fluency validity on the pairs where a person would have detected a difference yields the fluency validity *conditional* on a person having been able to apply the fluency heuristic. We computed such conditional validities separately for tartle, tartle–knowledge, and knowledge pairs. We binned the validities according to the previously computed detection probabilities, P_D. Bins were arranged by quartiles of the detection probabilities.

Applying the timing model to the data predicted by the memory model. As in Simulation 2, we let the memory model predict hypothetical persons' recognition and knowledge responses for the objects (P_R, P_K, $T_{recognition}$), pairing these objects into each hypothetical person's predicted tartle, tartle–knowledge, and knowledge pairs. We then processed each hypothetical person's predicted data in the same way as we processed our participants' observed data, using the timing model to compute both the predicted detection probability, P_D, of the hypothetical person being able to apply the fluency heuristic and the predicted conditional fluency validity, v_{fh}. For instance, calculating the validity on the pairs where a hypothetical person would have detected a difference in predicted recognition times yields the predicted fluency validity conditional on that person having been able to apply the fluency heuristic.

Results. Figure 9 shows the fluency validity, v_{fh}, plotted on the y-axis and the detection probability, P_D, on the x-axis. The integrated model correctly predicts three results. First, the fluency validity increases with rising detection probabilities. Second, these detection probabilities tend to be smaller in knowledge pairs than in tartle–knowledge and tartle pairs. This can be seen by comparing corresponding bins of the three types of pairs on the x-axis: Here the knowledge pairs tend to be the farthest to the left. For example, in the first bin of Figure 9B, the detection probability on knowledge pairs is .65; the corresponding probabilities on tartle and tartle–knowledge pairs are .71 and .70, respectively. This pattern is consistent with Fig-

[6] In Simulations 7 and 8, we provide experimental evidence suggesting that P_D reflects the effort and time required to use the fluency heuristic. That being said, we would like to add that our conclusions are equally supported by different proxies for effort and time (see Appendix C, Simulations C1 and C2).

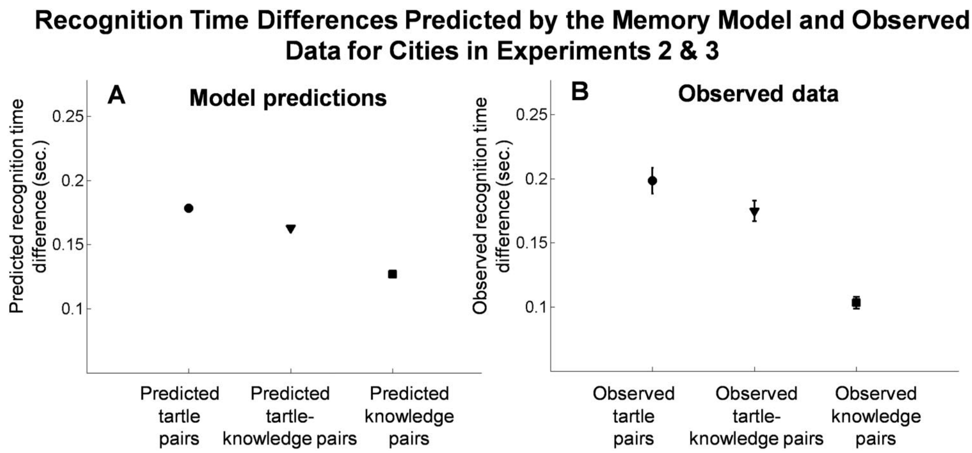

Figure 8. Recognition time differences (in seconds) between two cities in tartle, tartle–knowledge, and knowledge pairs, (A) predicted by the memory model (Equations 5–8; e.g., P_R, P_K, $T_{recognition}$) and (B) observed in Experiments 2 and 3. The data are collapsed across these experiments. Symbols in Panel A show mean (± 1 *SE*) predicted median time differences computed across simulation runs of our memory model; symbols in Panel B show mean (± 1 *SE*) observed median time differences computed across participants. In both panels, some of the error bars are obscured by the symbols. (A) To compute predicted recognition time differences, we let the memory model generate 1,500 hypothetical persons' recognition, knowledge, and recognition times, calculating the predicted recognition time differences for these hypothetical persons in an analogous way as we did for the experimental participants. (B) To compute the observed recognition time differences, we identified tartle and knowledge objects on the basis of participants' responses in the recognition and general knowledge tasks. For each participant, we exhaustively paired the objects into that participant's tartle, tartle–knowledge, and knowledge pairs, calculating the difference in recognition times between the objects in each pair. For each of these three kinds of pairs, we computed the median of recognition time differences for each participant, averaging the medians across participants.

ure 8, which shows that recognition time differences are smallest for knowledge pairs. The pattern that it is harder to detect recognition time differences for knowledge pairs is even more pronounced in Figures C1 and C2 (see Appendix C), where we plot the fluency validity in terms of other proxies of effort such as differences in pulses between two objects. Third, the model predicts that the fluency validity will be largest on tartle–knowledge pairs and smallest on tartle pairs, with the validity on knowledge pairs being somewhere in between.

Each of these three results has an important implication. First, as the fluency validity increases with the detection probability, it is easier to apply the fluency heuristic when using it helps a person make accurate inferences. There appears to be no effort–accuracy trade-off.[7]

Second, the detection probabilities show where the niches of the fluency heuristic and knowledge-based strategies overlap: On knowledge pairs, a person is least likely to be able to apply the fluency heuristic, because on these pairs recognition time differences are hardest to detect. On tartle pairs, differences are easier to detect, but in any event knowledge-based strategies cannot be applied. It is on tartle–knowledge pairs where the niche of the fluency heuristic and those of knowledge-based strategies overlap most, because on these pairs knowledge-based strategies are applicable and recognition time differences are easiest to detect.

Third, the fluency validity is largest on tartle–knowledge pairs—especially when recognition time differences are the easiest to detect (fourth bins in Figure 9). It is on such pairs that strategy selection between the fluency heuristic and its knowledge-based competitors is most likely to come into play based on currencies such as accuracy, effort, or time.

How Accuracy Explains Strategy Selection Between the Fluency Heuristic and Knowledge-Based Strategies

Next we pit the fluency heuristic against six knowledge-based strategies, examining how accuracy explains what strategy people

[7] Hertwig et al. (2008) reported a correlation between recognition time differences and the accuracy of inferences with the fluency heuristic. However, they lacked a memory model to predict the emergence of this correlation. They also had no model of time perception, which would have enabled them to predict *when* people detect differences in recognition times. Finally, they had no model to predict the knowledge people have about objects, and they also did not assess such knowledge. As we have learned when conducting our simulations, such models are warranted to understand the interplay between the fluency heuristic's accuracy and the effort required to use this heuristic, because this interplay turns out to be complex, varying as a function of people's recognition and knowledge, their recognition times, and their perception of such recognition times, and ultimately depending on the environment.

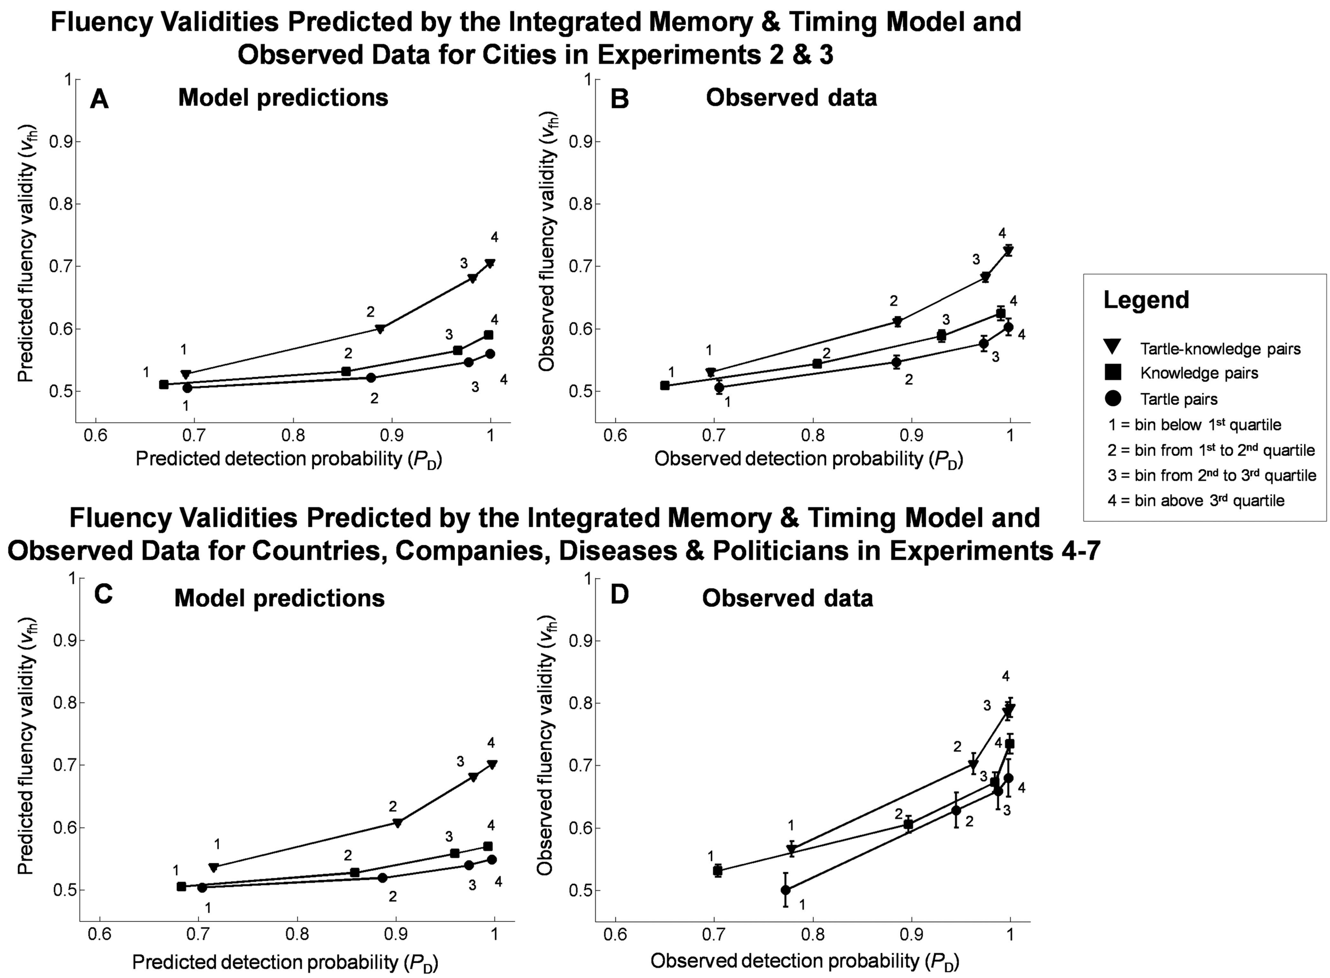

Figure 9. Simulation 3. Fluency validity (v_{fh}) as a function of the detection probability (P_D) that a person would detect a difference in recognition times between two cities. Predicted data were generated by the integrated memory and timing model (Equations 5–11; e.g., P_R, P_K, $T_{recognition}$, P_D, v_{fh}). The timing model was also applied to the observed data from Experiments 2–7. (A) Predicted and (B) observed data for inferences about cities' sizes. (C) Predicted and (D) observed data for inferences about countries' gross domestic product in 2006; companies' market capitalization on May 31, 2007; diseases' fame; and politicians' fame. We operationalized fame as the proportion of participants who recognized a disease in Experiments 6 and 8 and a politician in Experiments 7 and 9, respectively. In Panels C and D, the data are collapsed across the four different types of objects. For predicted data (A, C), symbols show mean (± 1 *SE*) predicted validities and mean predicted detection probabilities computed across simulation runs of the memory and timing model. For observed data (B, D), symbols show mean (± 1 *SE*) observed validities and mean observed detection probabilities computed across participants and simulation runs of the timing model. Some of the error bars are obscured by the symbols.

adopt when the strategies' niches overlap. Like the fluency heuristic, these knowledge-based strategies have been prominently discussed as descriptive and prescriptive models of behavior (e.g., Bröder & Gaissmaier, 2007; Dawes & Corrigan, 1974; Einhorn & Hogarth, 1975; Johnson & Raab, 2003; Nosofsky & Bergert, 2007; J. W. Payne et al., 1988; Rieskamp & Otto, 2006). The integration strategies tally1 and tally2 as well as the lexicographic heuristic take-the-best were introduced above; these models are summarized in Table 1. Other strategies that we have not yet described in detail are take-the-first-cue as well as take-the-first-value1 and take-first-value2. These three sequential-sampling strategies are explained in Figures 10 and 11.

Which Strategy Should People Rely on to Make Accurate Inferences?

As we have shown in Simulations 1–3, recognition times for objects as well as people's perception of them reflect noisy processes; they are sensitive, for example, to momentary changes in the activation levels of chunks, resulting in moderate levels of accuracy for the fluency heuristic in most regions of its niche (cf. Figure 9). At the same time, as we demonstrate next, even the accuracy of knowledge-based strategies that exploit the retrieval speed of cue values (e.g., take-the-first-cue) depends more on the cue values retrieved than on the type of noisy retrieval and time

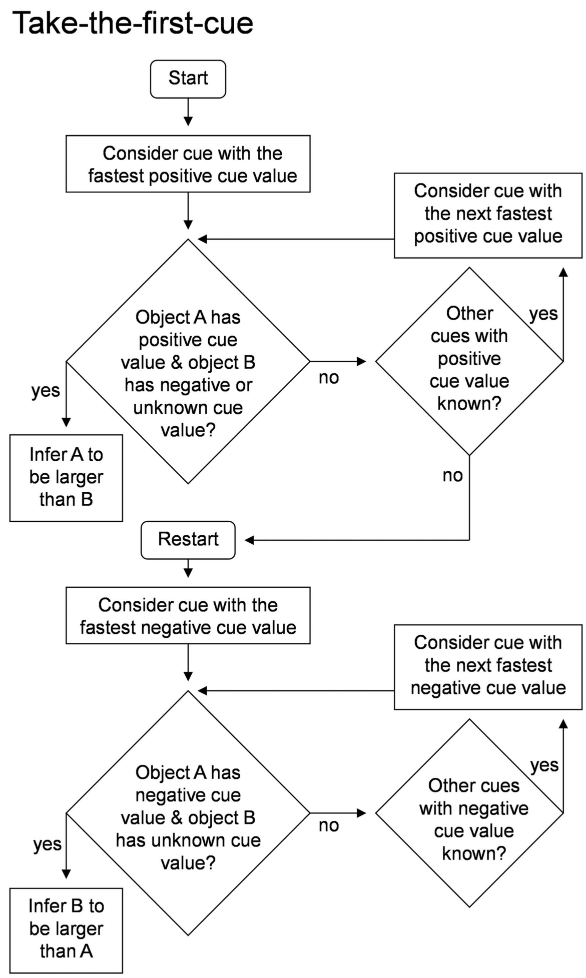

Figure 10. The sequential-sampling strategy take-the-first-cue considers cues in two search orders. First, it samples cues according to the speed with which the positive cue values are retrieved. For instance, when inferring whether Barcelona or Granada is larger, the fastest positive cue value may be that Barcelona has a world-famous tourist site. As a result, the tourism cue is the first cue on which the cities are compared. Among the remaining cues, the fastest positive cue value may be that Granada has a university. Thus, the university cue would yield the second comparison of the cities. Among the still remaining cues, the fastest positive cue value could be that Barcelona has an airport. Take-the-first-cue, in this case, would sequentially consult the tourism, university, and airport cues. It decides on the basis of the first cue in which Object A has a positive value and Object B has a negative or an unknown value. It infers A to be larger than B. In our example, take-the-first-cue would thus first compare the cities on the tourism cue. On this cue, Barcelona has a positive value. If Granada has no positive value on the tourism cue, then a decision can be made. If Granada also has a positive value on this cue, then the next cue would be considered, in this case, whether the cities have a university. If both have one, then the next cue would be considered, here, the airport cue. The algorithm continues until a decision can be made. Second, if sampling cues according to the speed with which positive cue values are retrieved does not allow making a decision, then the algorithm restarts, and take-the-first-cue considers the cues according to the speed with which negative cue values are retrieved. In our example, the fastest negative cue value may be that Granada has no major industrial site. As a result, take-the-first-cue would try to make a decision by comparing the cities on this industry cue. Take-the-first-cue relies on the first cue in which A has a negative value and B an unknown one, inferring B to be larger than A. The intuition is that knowing an object *does not have* a certain attribute (e.g., a city has no industry) makes it likely that this object is smaller than another for which one does not know if the object has the attribute or not. In our example, a person may not know whether Barcelona has a major industrial site. In this case, the comparison on the industry cue yields a decision and take-the-first would suggest that Barcelona is larger than Granada.

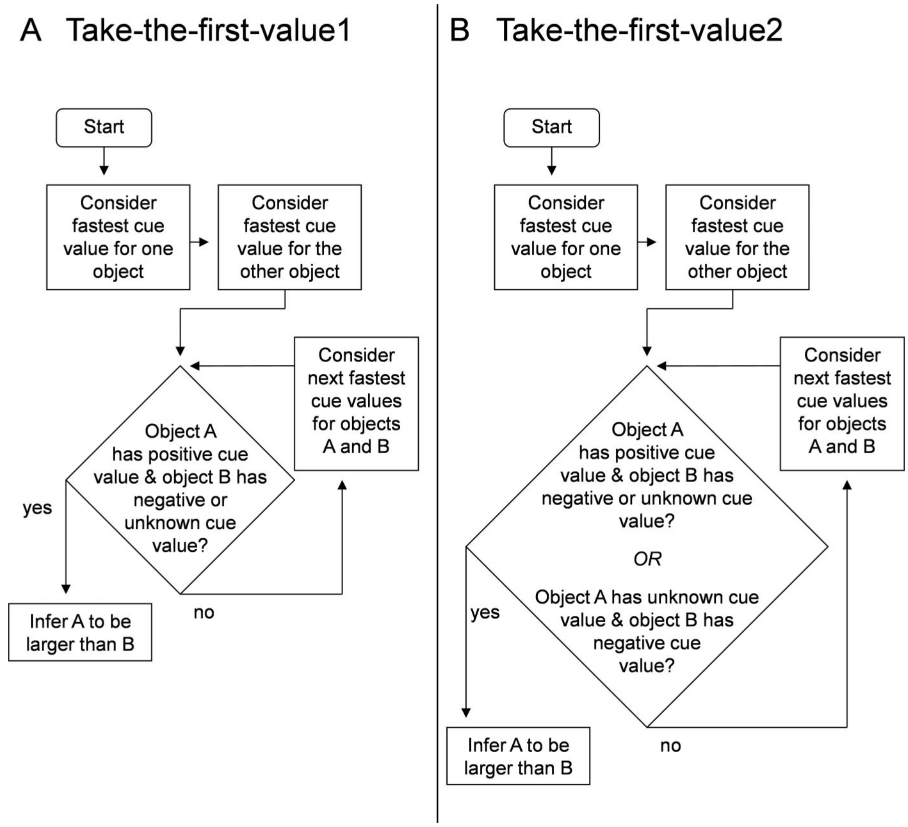

Figure 11. The sequential-sampling strategies take-the-first-value1 and take-the-first-value2. Take-the-the-first-cue compares objects on the same dimension, that is, on the *same* cue. In principle, however, objects could also be compared *across* cues. (A) Take-the-first-value1 relies on comparisons of objects by sampling cue values independently for each object according to the speed with which cue values are retrieved. To illustrate this, the fastest cue value generated for Barcelona may be that this city has an airport. The fastest cue value recalled for Granada may be that it is not an industry site. This is the first comparison of cities. For both cities, the second fastest cue value could be that they have universities, representing the second comparison. The third fastest cue value retrieved for Barcelona could be that it is an industry site. For Granada the third cue value could be that a person does not know whether Granada has an airport. Take-the-first-value1 consults these comparisons sequentially, relying on the first comparison in which Object A has a positive value on one cue and Object B a negative or an unknown value on the same or another cue. It infers A to be larger than B. In our example, this would happen in the first comparison of cities. (B) A variant of take-the-first-value1 is take-the-first-value2. It bases inferences *either* on the first comparison in which A has a positive value on one cue and B a negative or an unknown value on the same or another cue, *or* on the first comparison in which A has an unknown value and B a negative one. In both cases, take-the-first-value2 opts for A. As in take-the-first-cue, the intuition is that knowing an object lacks a certain attribute makes it likely the object will be small.

perception processes the fluency heuristic solely depends on. All else being equal, people would thus do well to rely on knowledge-based strategies when these strategies' niches and the fluency heuristic's niche overlap.[8]

Description of Observed Data and Measures

Observed data. To pit the fluency heuristic against knowledge-based strategies, in Experiment 1, we asked people to infer which city has more inhabitants. In addition, they took the recognition, general knowledge, and cue-knowledge tasks (see Tables 2 and 3). We used this recognition and knowledge data to test (a) whether relying on the fluency heuristic or on a knowledge-based strategy would help a person make more accurate inferences and (b) whether the former or the latter predicted their inferences best.

[8] On the basis of our earlier work (e.g., Hertwig et al., 2008), one might expect that people would rely on fluency rather than knowledge. Prior to developing the integrated model reported in the current article, we did not know how accurate the fluency heuristic would be compared with knowledge-based strategies, and we believed that retrieval fluency represents rapidly available, effortless information compared with knowledge. We had therefore thought that fluency would be prioritized over knowledge. At that time, we had not yet modeled people's recognition responses, knowledge responses, recognition times, and perception of recognition times, and, with perhaps one exception (Marewski, Gaissmaier, Schooler, Goldstein, & Gigerenzer, 2009), neither we nor anybody else had ever tested Schooler and Hertwig's (2005) fluency heuristic against knowledge-based strategies.

Accordance rates. To assess how well the fluency heuristic predicts people's inferences, for each participant we calculated the fluency heuristic *accordance rate*, k_{fh}, as the proportion of inferences consistent with the heuristic:

$$k_{\text{fh}} = E/(E + D), \tag{12}$$

where E is the number of times a more quickly recognized city was inferred to have more inhabitants than a more slowly recognized one, and D is the number of times the opposite happened. To model which city a participant would perceive as having been more quickly recognized, we used the timing model.

To assess how well tally1 predicts each participant's inferences, for each city we computed the *cue sum*, S_{pn}, of positive and negative cue values:

$$S_{\text{pn}} = C_1 + C_2 + C_3 + C_4, \tag{13}$$

where C is a participant's cue value for a given city on the airport cue (C_1), the university cue (C_2), the industry cue (C_3), and the tourism cue (C_4). Cue values could be either known, based on a participant's "yes" or "no" response in the cue-knowledge task, or unknown (i.e., a "don't know" response). Cue values were coded as 1, –1, and 0, respectively. The accordance with tally1 (k_{t1}) is computed analogously to the fluency heuristic accordance, using Equation 12, where E is the number of times a city with the larger cue sum, S_{pn}, was inferred to have more inhabitants than one with a smaller cue sum, and D is the number of times the opposite happened. To calculate the accordance with tally2 (k_{t2}), for each city we computed sums, S_{p}, of positive cue values, using Equations 12 and 13 as we did for tally1.

To implement take-the-best (see Table 1), we computed *cue validities* according to the actual attributes of the cities.[9] The validity of a cue is estimated as the proportion of times a city that has an attribute is larger than a city that does not have this attribute. The accordance with take-the-best (k_{ttb}) is computed for each participant individually, employing Equation 12 as for tally1 and tally2.

To implement take-the-first-cue, take-the-first-value1, and take-the-first-value2, we measured the speed of retrieving a city's cue value for each participant as the response time in the cue-knowledge task. We then let the timing model generate the number of pulses a person would have perceived while retrieving a cue value, sorting the cues (take-the-first-cue) and cue values (take-the-first-value1, take-the-first-value2), respectively, according to the number of pulses.[10] In doing so, we assumed that a person may have to decide which cue values have been retrieved most quickly, because the retrieval of multiple cue values may happen in close succession. The timing model provides a mechanism for these discrimination decisions and introduces additional noise in the retrieval-based sequential sampling of cues and cue values. The strategies' accordance rates (k_{ttfc}, k_{ttfv1}, k_{ttfv2}) are computed as for all other models.

The goal here is to test how well the strategies describe people's inferences *when their niches overlap*, that is, when they are applicable at the same time. For each model comparison, we selected each participant's pairs in which both the knowledge-based strategy to be compared and the fluency heuristic were applicable. Computing the accordance rates on these pairs yields accordance rates *conditional* on both the fluency heuristic *and* the respective competing knowledge-based strategy being applicable. Table 6 summarizes when the strategies are applicable. Next, we turn to the corresponding simulation.

Validities. To assess whether using the fluency heuristic or a knowledge-based strategy would be more helpful to a person inferring which of two cities is larger, we calculated the fluency validity and the validities for each of the knowledge-based strategies (v_{t1}, v_{t2}, v_{ttb}, v_{ttfc}, v_{ttfv1}, v_{ttfv2}). The knowledge-based strategies' validities are defined analogously to the fluency validity as the proportion of correct inferences a person can make by using a strategy. For each participant, the validities were computed conditional on both the fluency heuristic and the respective knowledge-based strategy being applicable. Next, we turn to the corresponding simulation.

The Fluency Validity and the Knowledge-Based Strategies' Validities: Simulation 4

Applying the timing model to the observed data. To examine which strategy would most help a person make accurate inferences, we ran a simulation consisting of two parts. First, we used the timing model to compute for each participant's pairs of two recognized objects the detection probability, P_{D}. Second, we reran the timing model, assessing in each run for each participant's pairs whether the participant would have been able to apply the fluency heuristic as well as a knowledge-based strategy in that run. In each run, we grouped each participant's tartle–knowledge and knowledge pairs where one of the knowledge-based strategies as well as the fluency heuristic were both applicable into four bins, arranged by quartiles, according to the previously computed detection probabilities, P_{D}. Across the pairs in each bin, we computed the strategies' validities conditional on a knowledge-based strategy and the fluency heuristic both being applicable.

Results. As Figure 12 shows, the knowledge-based strategies' validities tend to be larger than the fluency validity when the strategies' niches overlap. Hence, people would be better off relying on knowledge than on the fluency heuristic.

The Fluency Heuristic Accordance Rate and the Knowledge-Based Strategies' Accordance Rates: Simulation 5

Applying the timing model to the observed data. To assess how well the strategies account for people's inferences, we ran a simulation for Experiment 1. This simulation was similar to the previous one (i.e., No. 4), an exception being that we computed the strategies' accordance rates (rather than their validities) condi-

[9] That is, we looked up for each city whether it had at least one international airport and whether it had at least one university. To operationalize the industry cue, we looked up for each city whether it was the headquarters of at least one company that was registered in the major stock index of the country (i.e., CAC40, MIB30, IBEX, DAX, FTSE, Dow Jones, ATX). To operationalize the tourism cue, we looked up whether a city had a site included in the United Nations Educational, Scientific, and Cultural Organization (UNESCO) world heritage list.

[10] If for a city two cue values are retrieved with the same speed, then take-the-first-cue, take-the-first-value1, and take-the-first-value2 will decide at random which cue value will be considered first for this city.

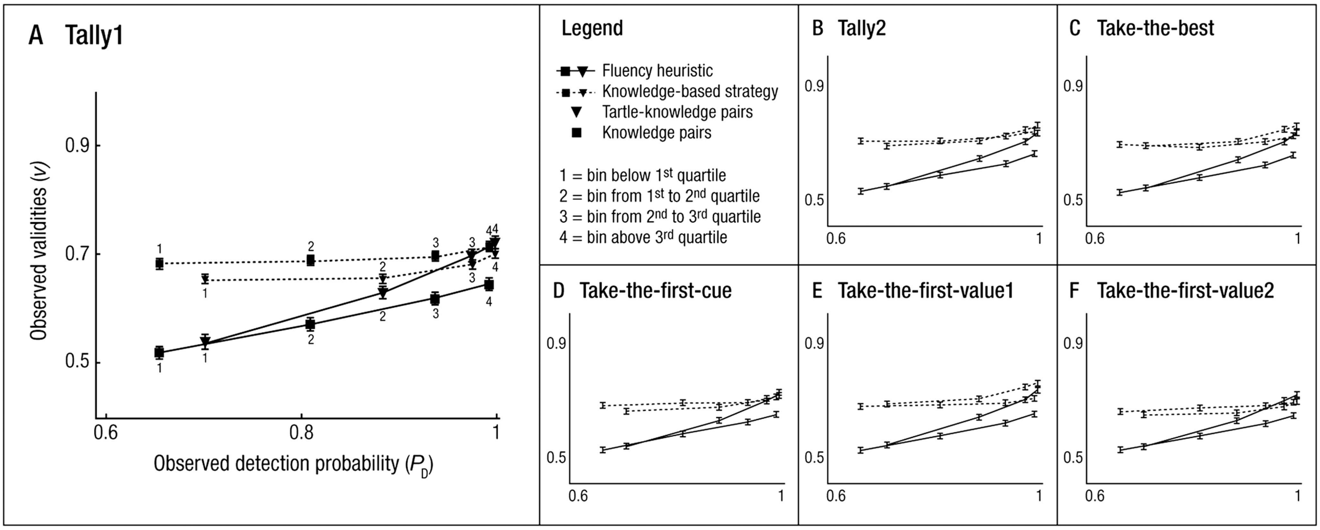

Figure 12. Simulation 4. Fluency validity (v_{fh}, solid lines) and validity of each of six knowledge-based strategies (dotted lines; A: tally1, v_{t1}; B: tally2, v_{t2}; C: take-the-best, v_{ttb}; D: take-the-first-cue, v_{ttfc}; E: take-the-first-value1, v_{ttfv1}; F: take-the-first-value2, v_{ttfv2}) for making inferences about cities' size on those pairs of cities where two strategies' niches overlap, plotted as a function of the detection probability (P_{D}) that a person would be able to detect a difference in recognition times between two cities. To make the graph easier to read, we enlarged the comparison between the fluency heuristic's and tally1's niches in Panel A; as can be seen in Panels B–F, similar results emerged when comparing the fluency validity to all other knowledge-based strategies' validities where the fluency heuristic's niche overlaps with each of these other strategies' respective niches. We modeled detection probabilities, participants' perception of recognition times, their perception of retrieval times of cues, and the validities by applying the timing model to each participant's observed data (Equations 9–11, 13). Data are collapsed across Experiments 1 and 3. In all panels, symbols show mean detection probabilities and mean (± 1 *SE*) validities, computed across participants and simulation runs of the timing model. Some of the error bars are obscured by the symbols.

tional on a particular knowledge-based strategy and the fluency heuristic both being applicable.

Results. The knowledge-based strategies account for people's inferences better than the fluency heuristic (see Figure 13). In short, people seem to rely on knowledge rather than on the fluency heuristic when these strategies' cognitive niches overlap. This allows them to make accurate inferences, because the knowledge-based strategies' validities are larger than the fluency validity.

How Accuracy Explains Strategy Choice When the Fluency Heuristic Can Be Used But Knowledge Is Not Available

Although people tend not to apply the fluency heuristic when they instead can rely on more accurate knowledge, our integrated model suggests that they would do well to use this heuristic when they cannot retrieve knowledge: As can be seen in Figure 9, the fluency validity rises up to between .55 and .68 on tartle pairs, depending on the objects considered. On tartle pairs, the fluency heuristic can thus help people make more accurate inferences than they would be able to achieve otherwise, for instance, by guessing. To compare, the accuracy obtainable by relying on a random guessing strategy is chance (i.e., .50). In fact, Experiment 1's participants correctly inferred which of two cities is larger in a mean proportion of .61 (*SE* = .05) of the tartle pairs (95% CI on the mean difference between .61 and .50 [.01, .20], $n = 45$), and as can be seen in Figure 14A, the fluency heuristic accordance rate rises to over .73 on tartle pairs, suggesting that people indeed relied on the fluency heuristic on the tartle pairs.

Do People Adopt the Fluency Heuristic When They Cannot Use Knowledge? Simulation 6

Observed data. To further validate that people rely on the fluency heuristic on tartle pairs, we created a benchmark. In Experiment 2, in an inference task, we asked students from Berlin to make judgments about the fame of cities, namely, to infer which of two cities would be recognized by a larger number of 100 randomly chosen students from Berlin. In one experimental condition (between subjects), the *instruct group*, we instructed participants to always apply the fluency heuristic on pairs of two recognized cities. We used the fluency heuristic accordance rate observed in this group to compare how well the fluency heuristic predicts behavior in a second group, the *no-instruct group*, in which we did not give any instructions on strategy use. In both groups, participants completed the same recognition and general knowledge tasks as in Experiment 1 (Tables 2 and 3).

Applying the timing model to the observed data. To model people's inferences with the fluency heuristic, we ran Simulation 6, which was identical to Simulation 5 with two relevant exceptions: First, we used the timing model to assess for each tartle pair (rather than tartle–knowledge and knowledge pairs as was the case

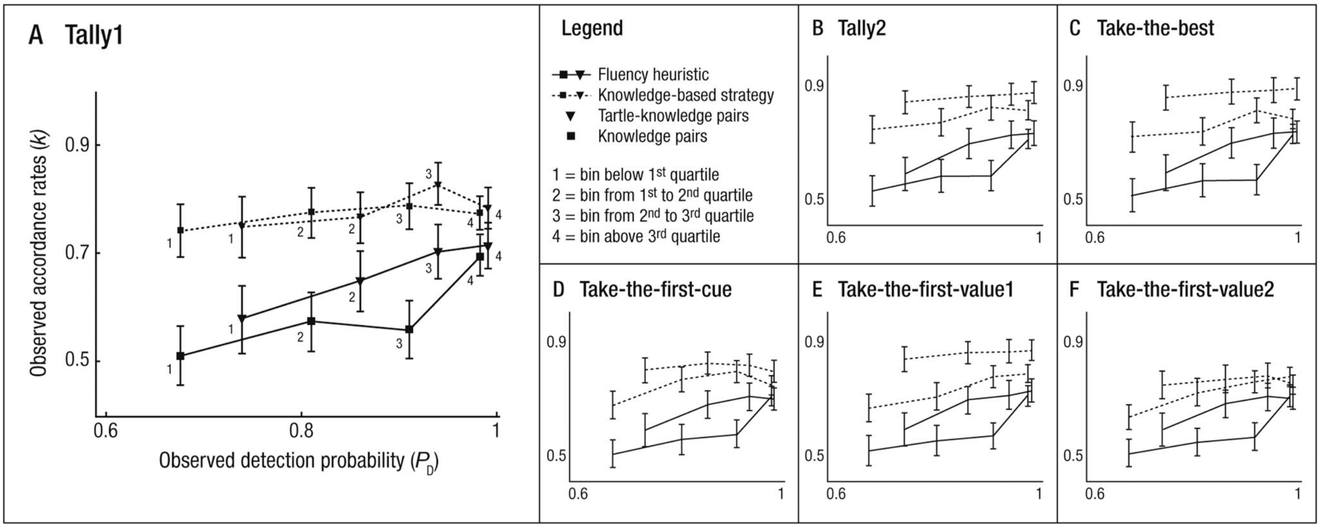

Figure 13. Simulation 5. Accordance rate for the fluency heuristic (k_{fh}, solid lines) and each of six knowledge-based strategies (dotted lines; A: tally1, k_{t1}; B: tally2, k_{t2}; C: take-the-best, k_{ttb}; D: take-the-first-cue, k_{ttfc}; E: take-the-first-value1, k_{ttfv1}; F: take-the-first-value2, k_{ttfv2}) in Experiment 1. Accordance rates are computed for making inferences about city size on those pairs of cities where two strategies' niches overlap and plotted as a function of the detection probability (P_D) that a person would be able to detect a difference in recognition times between two cities. To make the graph easier to read, we enlarged the comparison between the fluency heuristic and tally1 in Panel A; as can be seen in Panels B–F, similar results emerged when comparing the fluency heuristic accordance rate to all other knowledge-based strategies' accordance rates where the fluency heuristic's niche overlaps with each of these other strategies' respective niches. We modeled the detection probabilities, participants' perception of recognition times, their perception of retrieval times of cues, and the accordance rates by applying the timing model to each participant's observed data (Equations 9, 10, 12, 13). In all panels, symbols show mean detection probabilities and mean (± 1 *SE*) accordance rates, computed across participants and simulation runs of the timing model. Some of the error bars are obscured by the symbols.

in Simulation 5) the detection probability, P_D, that a participant would have been able to apply the fluency heuristic. Second, we computed the fluency heuristic accordance rate on tartle pairs conditional on the fluency heuristic being applicable (rather than conditional on both this heuristic and other strategies being applicable).

Results. The two groups did not differ in terms of the fluency heuristic accordance rate on tartle pairs (see Figures 14B and 14C). Rather, in both groups the accordance rate approximated that observed on tartle pairs in Experiment 1 (see Figure 14A). Importantly, the accordance rate increased with the detection probability, P_D, that a person would have been able to apply the fluency heuristic, which is consistent with the hypothesis that people indeed use this heuristic on tartle pairs.

In short, consistent with our integrated model's predictions about accuracy, when people cannot retrieve knowledge, they seem to employ the fluency heuristic. This helps them make more accurate inferences than they could otherwise make by simply guessing.

How Effort and Time Explain Strategy Selection Between the Fluency Heuristic and Knowledge-Based Strategies

Besides accuracy, the effort or time required to execute a strategy can also influence strategy selection. In Simulation 3 (see Figure 9), we illustrated that our integrated model allows us to predict when trade-offs among these currencies of strategy use need not occur. Next, we first validate the model's predictions further, showing that indeed less effort or time is required to employ the fluency heuristic when applying it is likely to help a person make accurate inferences. Second, we then use the model to examine the choice between the fluency heuristic and knowledge-based strategies, focusing on effort and time as currencies of strategy selection.

When Is the Fluency Heuristic Easy to Use? Simulation 7

Our integrated model predicts that a person using the fluency heuristic is most likely to make correct inferences when differences in recognition times are easy to detect (see Figure 9). To validate this prediction, we ran a simulation for the participants in the instruct group of Experiment 2, where participants had been instructed to always rely on the fluency heuristic when inferring which of two cities was recognized by more students (see Table 3). We expected that consistently using this heuristic would take them less time the more likely it was that the heuristic would help them make a correct inference.

Applying the timing model to the observed data. As in Simulation 5, in a first set of simulation runs, we used the timing model to compute for each participant's pairs of recognized cities

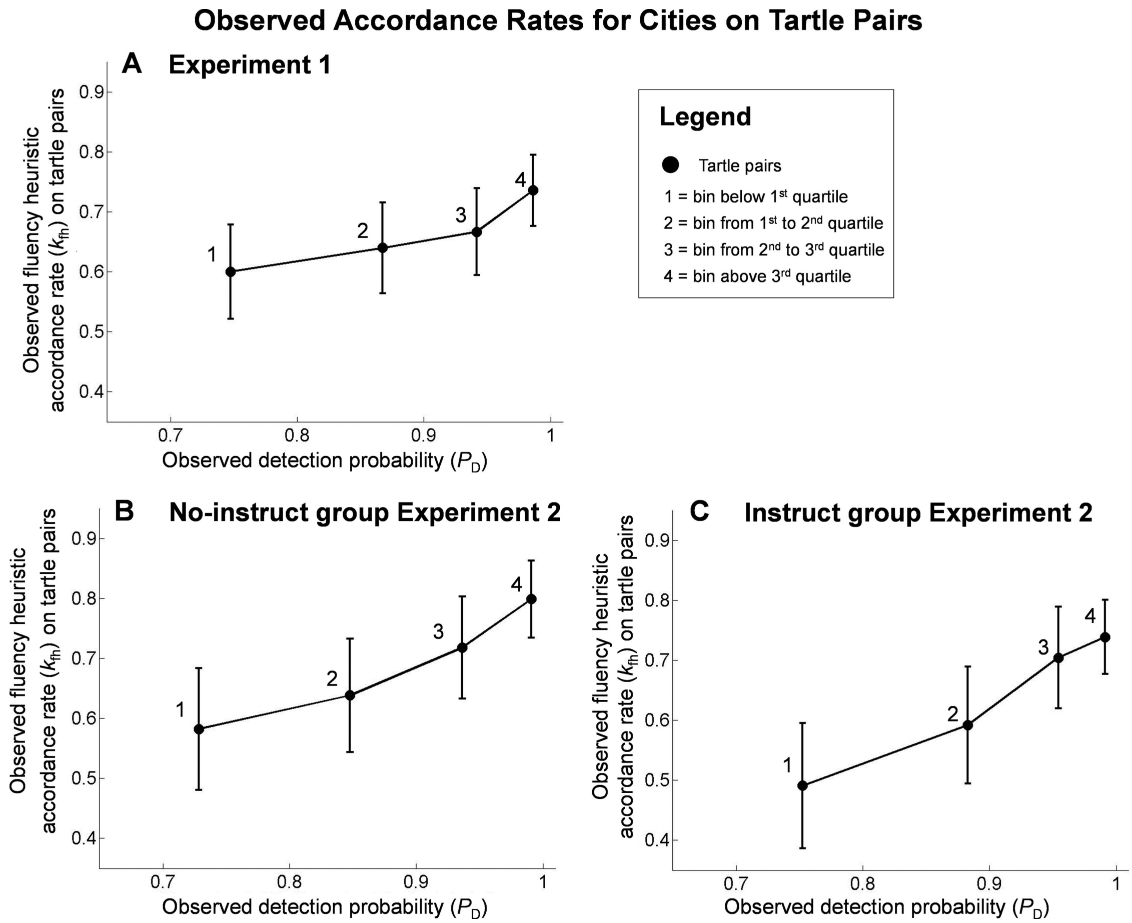

Figure 14. Simulation 6. Fluency heuristic accordance rate (k_{fh}) for making inferences about city size (A: Experiment 1) and city fame (B: Experiment 2, no-instruct group; C: Experiment 2, instruct group) on tartle pairs, plotted as a function of the detection probability (P_D) that a person would be able to detect a difference in recognition times between two cities. Participants' perception of recognition times, the detection probabilities, and the accordance rates were modeled by applying the timing model to each participant's observed data (Equations 9, 10, 12). In all panels, symbols show mean detection probabilities and mean (± 1 *SE*) accordance rates, computed across participants and simulation runs of the timing model. Some of the error bars are obscured by the symbols.

the detection probability, P_D, that the participant would have been able to detect a difference in recognition times and apply the fluency heuristic. Deviating from Simulation 5, in a second set of runs, we then computed three kinds of behavioral data conditional on a participant having been able to apply the fluency heuristic in that run. First, for each pair of recognized cities, we used the timing model to assess which city the participant would infer was recognized by more students, assuming the participant had used the fluency heuristic. Second and third, on those pairs where the participant had made an inference consistent with the fluency heuristic, we also examined how many correct inferences the participant had made and assessed the time it took the participant to make an inference (henceforth *inference time*). (To judge the inferences' correctness, we counted how many participants from Experiments 1–3, $N = 175$, had recognized each city.) As we did in Simulation 5, we then grouped the pairs into four bins, arranged by quartiles, according to the previously computed detection probabilities, P_D. We computed (a) the fluency heuristic accordance rate, (b) the proportion of correct inferences, and (c) the median inference time across the pairs in a bin.

Results. Recall that participants were instructed to *always* use the fluency heuristic. As Figure 15A shows, the fluency heuristic accordance rate increased the easier it was to detect a difference in recognition times and apply the heuristic. People were also more likely to make correct inferences about the cities when it was easier for them to detect a difference in recognition times (Figure 15B), and it also took them less time to make inferences the easier it was for them to detect a difference in recognition times (Figure 15C).

When Is the Fluency Heuristic Fast to Use? Simulation 8

Figure 15 indicates that people are more likely to make accurate inferences with the fluency heuristic the less time it takes them to apply it. This suggests people may do well to apply the fluency heuristic when forced to make quick decisions.

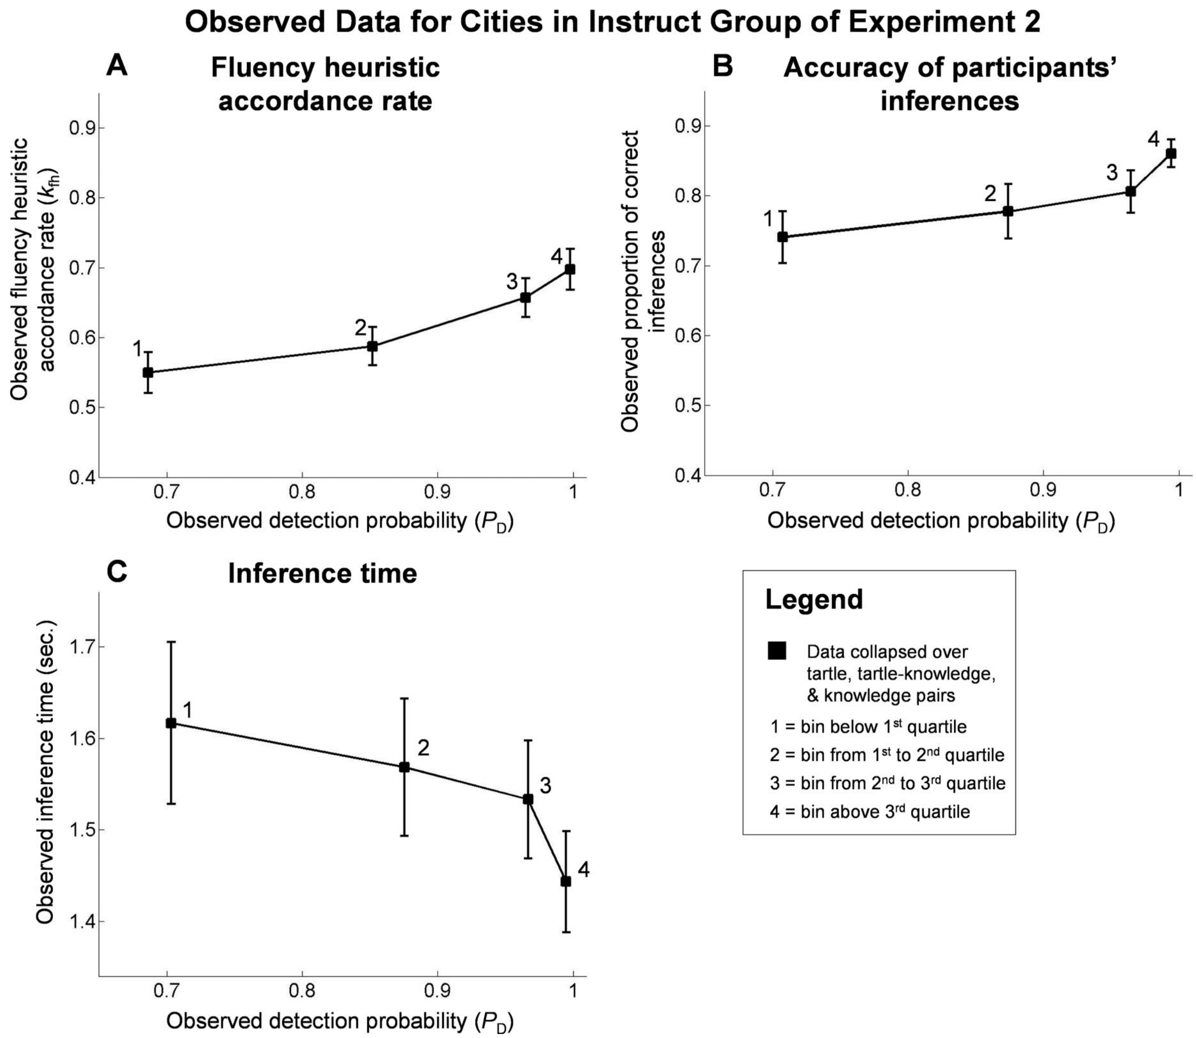

Figure 15. Simulation 7. (A) Fluency heuristic accordance rate (k_{fh}), (B) the proportion of correct inferences participants made, and (C) inference times (in seconds) for making inferences about the fame of cities in the instruct group of Experiment 2. Data are plotted as a function of the detection probability (P_D) that a person would be able to detect a difference in recognition times between two cities. Participants' perception of recognition times for cities, the detection probabilities, and the accordance rates were modeled by applying the timing model to each participant's observed data (Equations 9, 10, 12). In all panels, symbols show means (± 1 *SE*) computed across participants and simulation runs of the timing model. Some of the error bars are obscured by the symbols. Data are collapsed across tartle, tartle–knowledge, and knowledge pairs of cities.

There is evidence that people switch to accessibility-based rather than knowledge-based heuristics under time pressure (Pachur & Hertwig, 2006). Additionally, a sense of familiarity or fluency is often thought of as an automatic form of memory that is more quickly available than recollection, which entails effortful and intentional retrieval processes (Atkinson & Juola, 1974; Gronlund & Ratcliff, 1989; Hintzman & Curran, 1994; Jacoby, 1991; Mandler, 1980; McElree, Dolan, & Jacoby, 1999; Ratcliff & McKoon, 1989). Consistent with this literature, Hertwig et al. (2008) hypothesized the fluency heuristic would be faster to apply than knowledge-based strategies, which, in their view, depend on slow and effortful knowledge retrieval. One might thus expect that people will prefer the fluency heuristic over knowledge-based strategies in situations of time pressure.

Our integrated model does not predict whether such a preference will occur—such a prediction would depend, among other things, on the amount of time pressure and on the specifics of the mechanism (e.g., cost–benefit or learning) implemented in our cognitive niche framework—however, our model explains in what areas of its niche such a preference for the fluency heuristic is most likely to emerge. This is the case on tartle–knowledge pairs, where differences in recognition times are the easiest to detect (see Figure 9; when comparing corresponding bins, the tartle–knowledge pairs tend to be the farthest to the right). As Figure 12 shows, these are also the pairs where people are the most likely to make as accurate inferences by using the fluency heuristic as they can by relying on a knowledge-based strategy (see the fourth bins).

Observed data. To explore whether there is evidence that a preference for the fluency heuristic most likely emerges on tartle–knowledge pairs with easily detectable recognition time differences, we ran Experiment 3. This experiment was essentially identical to Experiment 1, the major exception being that we put participants under time pressure in the inference task, giving them only 900 msec to infer which of two cities is larger.

Applying the timing model to the observed data. We ran the same simulation (No. 5) for Experiment 3 as for Experiment 1. That is, we calculated the fluency heuristic's and each of six knowledge-based strategies' accordance rates conditional on the heuristic and a knowledge-based strategy both being applicable, plotting each model's accordance rate as a function of the detection probability, P_D, that a person would have been able to apply the fluency heuristic.

Results. On tartle–knowledge pairs, when people have almost a 100% probability of detecting a difference in recognition times, the fluency heuristic predicts people's inferences as well as any of its knowledge-based competitors (see fourth bins in Figure 16). To compare, in Experiment 1, where people were not under time pressure, the knowledge-based strategies predicted people's inferences better (see Figure 13). In short, consistent with our integrated model's predictions about effort and time, for situations of time pressure, there is evidence that if there were a preference for the fluency heuristic over knowledge-based strategies, this preference would most likely emerge when differences in recognition times are easy to detect and available knowledge is sparse. In this situation the heuristic also most likely helps a person make accurate inferences.

Zooming in on the Niches of the Recognition Heuristic and Knowledge-Based Strategies: Regions of Accurate, Effortless Inferences

Our integrated model not only maps out the accuracy, effort, and time entailed by using the fluency heuristic, it also makes corresponding predictions for other strategies, (a) providing the currencies the cost–benefit, learning, and related mechanisms use to explain the selection among simultaneously applicable strategies and (b) enabling us to pinpoint regions in the strategies' niches where effort–accuracy trade-offs need not occur. In the remainder of this article, we illustrate this by zooming in on the recognition heuristic's niche (see Figure 2H) before closing with a glimpse of the knowledge-based strategies' niches (see Figure 2I).

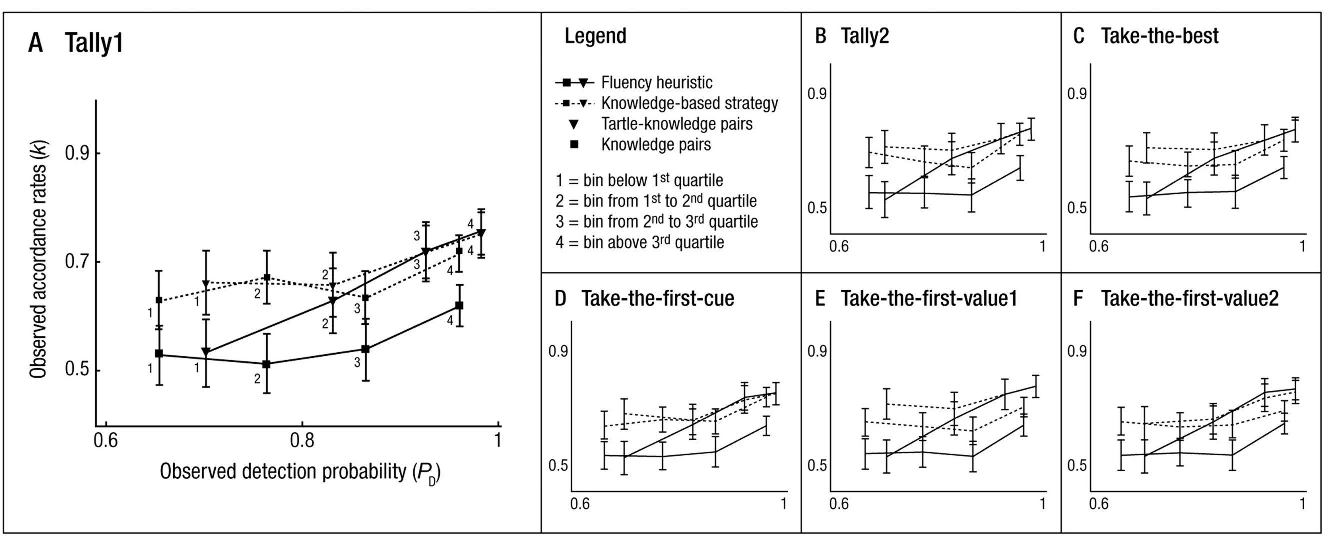

Figure 16. Simulation 8. Accordance rate for the fluency heuristic (k_{fh}, solid lines) and each of six knowledge-based strategies (dotted lines; A: tally1, k_{t1}; B: tally2, k_{t2}; C: take-the-best, k_{ttb}; D: take-the-first-cue, k_{ttfc}; E: take-the-first-value1, k_{ttfv1}; F: take-the-first-value2, k_{ttfv2}) in Experiment 3. Accordance rates are computed for making inferences about city size on those pairs of cities where two strategies' niches overlap and plotted as a function of the detection probability (P_D) that a person would be able to detect a difference in recognition times between two cities. To make the graph easier to read, we enlarged the comparison between the fluency heuristic and tally1 in Panel A; as can be seen in Panels B–F, similar results emerged when comparing the fluency heuristic accordance rate to all other knowledge-based strategies' accordance rates where the fluency heuristic's niche overlaps with each of these other strategies' respective niches. We modeled the detection probabilities, participants' perception of recognition times, and the accordance rates by applying the timing model to each participant's observed data (Equations 9, 10, 12, 13). In all panels, symbols show mean detection probabilities and mean (± 1 *SE*) accordance rates, computed across participants and simulation runs of the timing model. Some of the error bars are obscured by the symbols.

When Does the Recognition Heuristic Help Make Accurate Inferences?

To zoom in on the recognition heuristic's niche, we ran two simulations. As these simulations show, the integrated model predicts that the accuracy, effort, and time involved in applying this heuristic vary systematically as a function of a person's ability to retrieve knowledge about an object as well as a function of a person's perception of the object's recognition time. Before reporting the two simulations, let us provide an intuitive account.

As explained in Simulation 3, the correlation between encountering objects (e.g., Volkswagens) and acquiring knowledge about them (e.g., German engineered) results in people being more likely to retrieve knowledge about frequently encountered objects than about infrequently encountered ones. Due to being encountered more often, knowledge objects are likely to be more strongly activated, and more quickly recognized, than tartle objects (see Appendix C, Figure C3). In environments in which the probability of retrieving knowledge additionally correlates with the criterion to be inferred (e.g., car quality), knowledge objects also score on average higher on the criterion than do tartle or unrecognized objects. Because both knowledge and tartle objects are likely to have larger criterion values than unrecognized objects, knowledge–unrecognized pairs will tend to reflect larger differences on the criterion than tartle–unrecognized pairs. As a result, people are more likely to score accurate inferences with the recognition heuristic on knowledge–unrecognized pairs than on tartle–unrecognized pairs.[11] At the same time, because knowledge objects are more quickly recognized than tartle objects, applying the recognition heuristic on knowledge–unrecognized pairs is likely to be faster and less effortful than applying it on tartle–unrecognized pairs. Finally, in such environments, not only a person's ability to retrieve knowledge about objects but also the objects' recognition times themselves correlate with the criterion. Therefore, the recognition heuristic is most likely to help a person to make accurate, fast, and effortless inferences the more quickly an object is perceived as recognized.

Recognition Validity as a Function of Perceived Recognition Times: Simulation 9

Applying the timing model to the observed data. To model how the accuracy of inferences with the recognition heuristic correlates with the retrieval of knowledge and people's perception of recognition times, we first grouped each participant's objects into tartle–unrecognized and knowledge–unrecognized pairs, using their responses in the recognition and general knowledge tasks in Experiments 2–7. For each participant, we fed each tartle and knowledge object's recognition time into the timing model. In doing so, we computed each object's *perceived recognition time*, expressed in pulses. Second, according to the number of pulses, we further grouped each participant's pairs into four bins, arranged by quartiles. We calculated the *recognition validity*, v_{rh}, for each participant in each bin. This validity is the probability that a person will score a correct inference by using the recognition heuristic. It is estimated as the proportion of times a recognized object has a larger criterion value than an unrecognized one, using Equation 11 as for all other strategies.

Applying the timing model to the data predicted by the memory model. To generate the model predictions, we integrated the memory and the timing models in a simulation. Specifically, we first let the memory model predict hypothetical persons' recognition and knowledge responses, including the recognition times (P_{R}, P_{K}, $T_{\text{recognition}}$). We then processed these hypothetical persons' predicted data in the same way as we processed the observed data, feeding them into the timing model and calculating the predicted perceived recognition times (i.e., pulses) and the predicted recognition validity for each hypothetical person.

Results. The integrated model of memory and time perception correctly predicts three important results (see Figure 17). First, because knowledge objects tend to be recognized more quickly than tartle objects, more pulses tend to be accumulated for the recognized object in tartle–unrecognized pairs than in knowledge–unrecognized pairs (the bins of the knowledge–unrecognized pairs are the farthest to the left). Second, the fewer pulses that are accumulated the larger the recognition validity will be. Third, the recognition validity is larger on knowledge–unrecognized pairs than on tartle–unrecognized pairs. These results suggest that people should make the fastest and most accurate inferences on knowledge–unrecognized pairs. Correspondingly, on these pairs people would do well to rely most on the recognition heuristic, which is a hypothesis we test next.

When Is the Recognition Heuristic Easy and Fast to Use? Simulation 10

Applying the timing model to the observed data. To test whether people are most likely to make fast and accurate inferences in accordance with the recognition heuristic on knowledge–unrecognized pairs, in Simulation 10 we processed the observed data in a similar way to what we did in Simulation 9, a relevant exception being that we computed three behavioral measures across the pairs in each bin. First, for each participant, we calculated the recognition heuristic accordance rate (k_{rh}) in the inference task, which is the proportion of inferences the participant made consistent with the recognition heuristic. It is computed analogously to all other strategies' accordance rates, using Equation 12, and counting how often a participant infers a larger criterion value for a recognized city than for an unrecognized one. Second and third, on those pairs where the participant had made an inference consistent with the recognition heuristic, we also calculated the proportion of correct inferences the participant had made and

[11] Consistent with these intuitions, Hilbig and Pohl (2009), Marewski (2008), and Marewski, Gaissmaier, Schooler, Goldstein, and Gigerenzer (2010) reported longer recognition times for tartle objects than for knowledge objects. Marewski, Gaissmaier, Schooler, et al. also found a correlation between the accuracy of inferences with the recognition heuristic and people's knowledge. However, Marewski, Gaissmaier, Schooler, et al. and Hilbig and Pohl lacked models to predict these phenomena, and in fact, Marewski, Gaissmaier, Schooler, et al.'s reports of these phenomena were motivated exclusively by the simulations reported in the current article. With our integrated model we can now predict these and other phenomena directly from environmental data (e.g., Simulation 9; Simulation C3 in Appendix C).

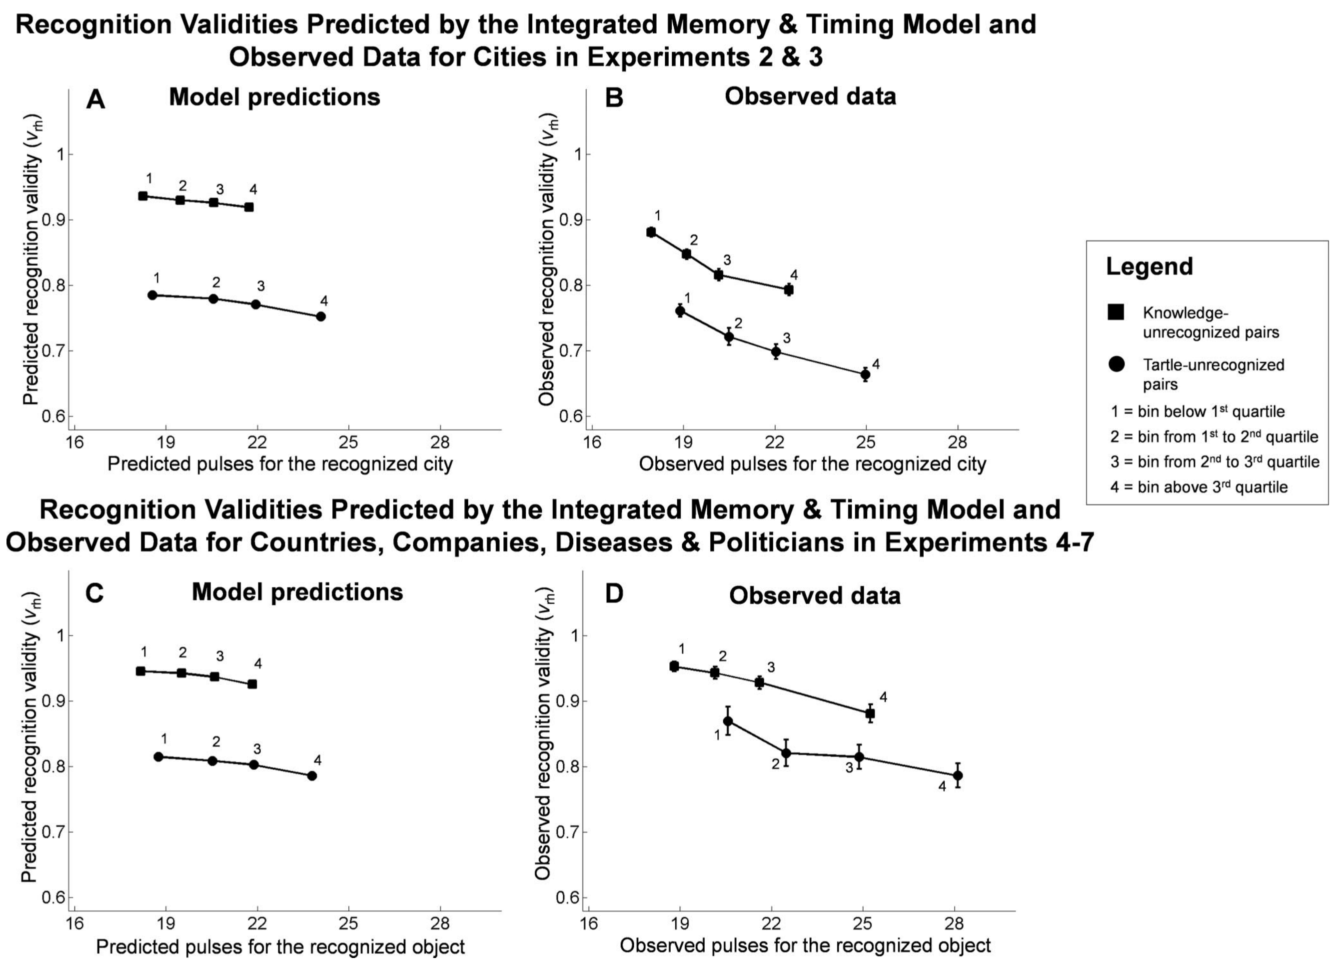

Figure 17. Simulation 9. Recognition validity (v_{rh}) as a function of the pulses a person would perceive when recognizing the tartle city and when recognizing the knowledge city in tartle–unrecognized and knowledge–unrecognized pairs, respectively. Predicted data were generated by the integrated memory and timing model (Equations 5–9, 11; e.g., P_R, P_K, $T_{recognition}$, v_{rh}). The timing model (Equations 9 and 11) was also applied to each participant's observed data. (A) Predicted and (B) observed data for inferring cities' size in Experiments 2 and 3. (C) Predicted and (D) observed data for inferring countries' gross domestic product in 2006; companies' market capitalization on May 31, 2007; diseases' fame; and politicians' fame in Experiments 4–7. In Panels C and D, the data are collapsed across the four different types of objects. (A, C) Mean (± 1 *SE*) predicted validities and mean predicted pulses computed across simulation runs of the integrated model. (B, D) Mean (± 1 *SE*) observed validities and mean observed pulses computed across participants and simulation runs of the timing model. Some of the error bars are obscured by the symbols.

computed the median time it took the participant to make each inference.

Results. Figure 18 shows the results for Experiment 1 and the no-instruct group of Experiment 2, which are the data sets without experimental manipulations (such as instructing people to use the fluency heuristic). Consistent with our integrated model's predictions, people were most likely to make inferences in line with the recognition heuristic on knowledge–unrecognized pairs (Panel A). On these pairs their inferences were also the most accurate (Panel B) and made the fastest (Panel C).

It has been debated whether the recognition heuristic is an adequate model of behavior (cf. Marewski, Pohl, & Vitouch, 2010, 2011). For instance, when both recognition and knowledge are acquired naturally, outside the laboratory, as is the case in the type of memory-based inference task we employ in our studies, recognition heuristic accordance rates have been found to vary as a function of knowledge. This and related findings have led several researchers to suggest that other (e.g., knowledge-based) decision strategies should account for people's inferences better than the recognition heuristic does (e.g., Hilbig & Pohl, 2009; Newell & Fernandez, 2006; Pohl, 2006). Yet, until recently no study has actually tested a computationally specified knowledge-based strategy (or any other model of decision-making processes) against the recognition heuristic in this type of memory-based inference task. The only study that conducted corresponding formal comparative model tests, for this type of task, found that neither knowledge- nor accessibility-based strategies predicted people's inferences as well as the recognition heuristic does (Marewski, Gaissmaier, Schooler,

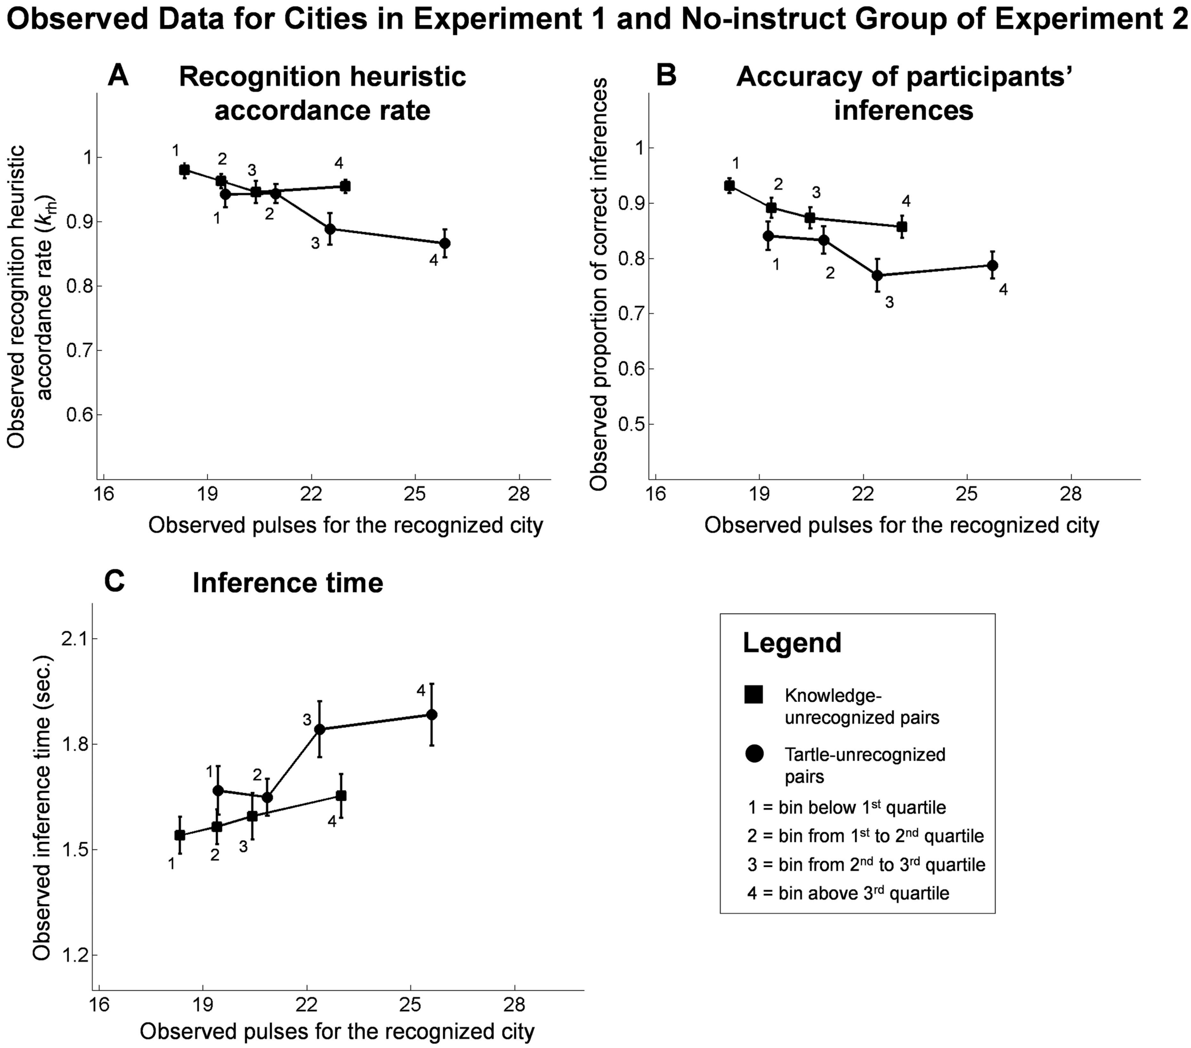

Figure 18. Simulation 10. (A) Recognition heuristic accordance rate (k_{rh}), (B) the proportion of correct inferences participants made, and (C) inference time (in seconds) for inferences about cities as a function of the pulses a person would perceive when recognizing the tartle city and when recognizing the knowledge city in tartle–unrecognized and knowledge–unrecognized pairs, respectively. Data are aggregated across Experiment 1 and the no-instruct group of Experiment 2. Participants' perception of pulses and the resulting accordance rates were modeled by applying the timing model to each participant's observed data (Equations 9 and 12). In all panels, symbols show means (± 1 *SE*) computed across participants and simulation runs of the timing model. Some of the error bars are obscured by the symbols. Note that participants made their inferences faster on knowledge–unrecognized pairs than on tartle–unrecognized pairs (C). This inference time difference is consistent with our memory model, which correctly predicts faster recognition times for knowledge objects than for tartle objects (see Figure C3, Appendix C).

et al., 2010). This also holds true for the knowledge–unrecognized and tartle–unrecognized pairs from our current Experiment 1, which Marewski, Gaissmaier, Schooler, et al. (2010) reanalyzed, and for which we reported the simulations of the integrated model above. These simulations show how the recognition validity, recognition times, and people's perception of recognition times covary with knowledge, providing a rationale for why one might expect recognition heuristic accordance rates and inference times to behave the way they do (see Figure 18; cf. Erdfelder, Küpper-Tetzel, & Mattern, 2011). In short, our integrated model predicts *when* people would do well to rely on the recognition heuristic, and Marewski, Gaissmaier, Schooler, et al.'s comparative model tests provide some evidence to suggest that they do rely on this heuristic in the situations specified by the integrated model.

When Are Knowledge-Based Strategies Easy to Use? Simulation 11

Several studies have examined the accuracy, effort, and time involved in using knowledge-based strategies such as take-the-best or tally2 (e.g., Gigerenzer & Goldstein, 1996; Hogarth & Karelaia, 2007; J. W. Payne et al., 1993). Yet, little emphasis has been placed on how these strategies depend on the way the environment interacts with memory and other cognitive capacities. With our integrated model it is now possible to make systematic predictions about the knowledge-based strategies' cognitive niches. For instance, we can quantify the overlap between niches (e.g., Simulation 2). We can also examine in which regions of their niches the knowledge-based strategies help people make accurate and effortless inferences. In what follows, for the purpose of illustration we briefly sketch out one such analysis of accuracy and effort. For simplicity, in this exploratory analysis, we focus on the data observed in our experiments, assessing how different strategies' validities behave as a function of how easy it is to apply these strategies.

Integration strategies such as tally1 or tally2 are likely to be easier to apply when the sums of cue values differ greatly between two objects. The intuition is that it is easier to tell which of two objects (e.g., two cars) scores higher on a criterion (e.g., quality) when one exhibits a large number of positive cue values (e.g., a Mercedes: high price, famous brand) and the other does not than when both objects exhibit many positive cue values (e.g., Mercedes and BMW: both are expensive, famous, etc.).

In lexicographic and sequential-sampling strategies, the effort involved in using them is likely to increase with the number of comparisons of cue values that need to be considered to make an inference. For instance, if the first pair of cue values considered is sufficient to make an inference, then this inference is likely to take less time than if lots of cue values have to be searched through. The number of cue values that need to be compared in a given situation varies among these strategies; recall that the lexicographic take-the-best heuristic sorts cues according to their validities and that sequential-sampling strategies such as take-the-first-cue or take-the-first-value1, in contrast, rely on people's perception of retrieval times of cue values (see Table 1).

Applying the timing model to the observed data. To explore in what regions of these knowledge-based strategies' niches effort–accuracy trade-offs need not be made, we ran a simulation for Experiments 1 and 3, which are the data sets where we assessed participants' knowledge about cues. First, we organized each participant's objects into tartle–knowledge and knowledge pairs and assessed how effortful it would have been for a participant to use a strategy. For the integration strategies, we computed differences between sums of cue values as a measure of effort, and for the lexicographic and sequential-sampling strategies the number of comparisons of cue values that need to be considered prior to making an inference. For the sequential-sampling strategies, this required running the timing model, using each participant's reaction times observed for each cue value in the cue-knowledge tasks as input for this model. Second, for each participant and each strategy, we selected those tartle–knowledge and knowledge pairs where a strategy was applicable. We grouped those pairs where a strategy was applicable into four bins, arranged by quartiles, according to the previously calculated effort involved in using the strategy. In each of the bins, for each participant, we calculated each strategy's validity (v_{t1}, v_{t2}, v_{ttb}, v_{ttfc}, v_{ttfv1}, v_{ttfv2}). This simulation procedure yields each strategy's validity conditional on the strategy being applicable.

Results. As Figure 19 shows, on tartle–knowledge pairs—except for take-the-first-cue, which always decides on the first comparison of cue values (Panel D)—each strategy's validity decreases the more effortful it is to rely on the strategy. This pattern holds regardless of whether a strategy sums cue values (e.g., tally1; Panel A) or consults them sequentially (e.g., take-the-first-value1; Panel E). In short, within these areas of their niches, it is likely to be easier for a person to apply a knowledge-based strategy when using such a strategy is also most likely to help that person make accurate inferences.

At the same time, Figure 19 demonstrates that the niches of the six knowledge-based strategies differ. For instance, the validities of the lexicographic and sequential-sampling strategies (Panels C–F) remain almost constant as a function of effort on the knowledge pairs. In contrast, the validities of two integration strategies (Panels A, B) fall sharply as effort increases on the knowledge pairs as well as on the tartle–knowledge pairs.

A comparison of Figures 19 and 12 illustrates a general methodological point: It is useful to examine both a strategy's niche separately (as in Figure 19) *and* those regions of the niche where it overlaps with the niches of other strategies (as in Figure 12). For instance, in Figure 12A, the tally1 validity increases only a little as a function of the probability, P_D, of a person being able to detect a difference in recognition times. In Figure 19A, in contrast, the tally1 validity increases as a function of increasing differences in sums of cue values, corresponding to its own currency of effort. From the perspective of the cognitive niche framework this is to be expected, because the ability to detect recognition time differences reflects the effort involved in applying the fluency heuristic, but not necessarily the effort involved in using tally1.

General Discussion

We join others in stressing the importance of studying how the cognitive system nestles into the environment (Anderson, 1990; Anderson & Milson, 1989; Dougherty et al., 2008; Gibson, 1979; Gigerenzer, Hoffrage, & Kleinbölting, 1991; Howes, Lewis, & Vera, 2009; Howes, Vera, & Lewis, 2007; Oaksford & Chater, 1998; Shepard, 2001; Simon, 1956; Stewart, Chater, & Brown, 2006; Vicente, 2003). Using the ACT–R cognitive architecture, we modeled how the natural environment, outside the laboratory, structures memory. Memory, in turn, interacts with time perception and other cognitive capacities, yielding the information on which decision strategies operate. Building an integrative, ecological model of the interplay between the environment, memory, time perception, and decision strategies, we were able to show how this interplay guides strategy selection by determining what strategies from our repertoire (a) are applicable, as well as (b) how accurate they will be and (c) how much effort and time will be involved in using them.

We use the term *cognitive niche* to characterize the situations in which a strategy can be applied, for instance, because sufficient knowledge can be retrieved or differences in recognition times detected. In particular, we demonstrated that the niches of the strategies in our repertoire may not fully overlap. That is, not all strategies will always be applicable simultaneously, thus simplifying strategy selection. At the same time, we showed what regions in the strategies' niches will allow people to make accurate, fast, and effortless inferences.

To elaborate our cognitive niche framework, we considered selecting among classic strategies, including knowledge-based in-

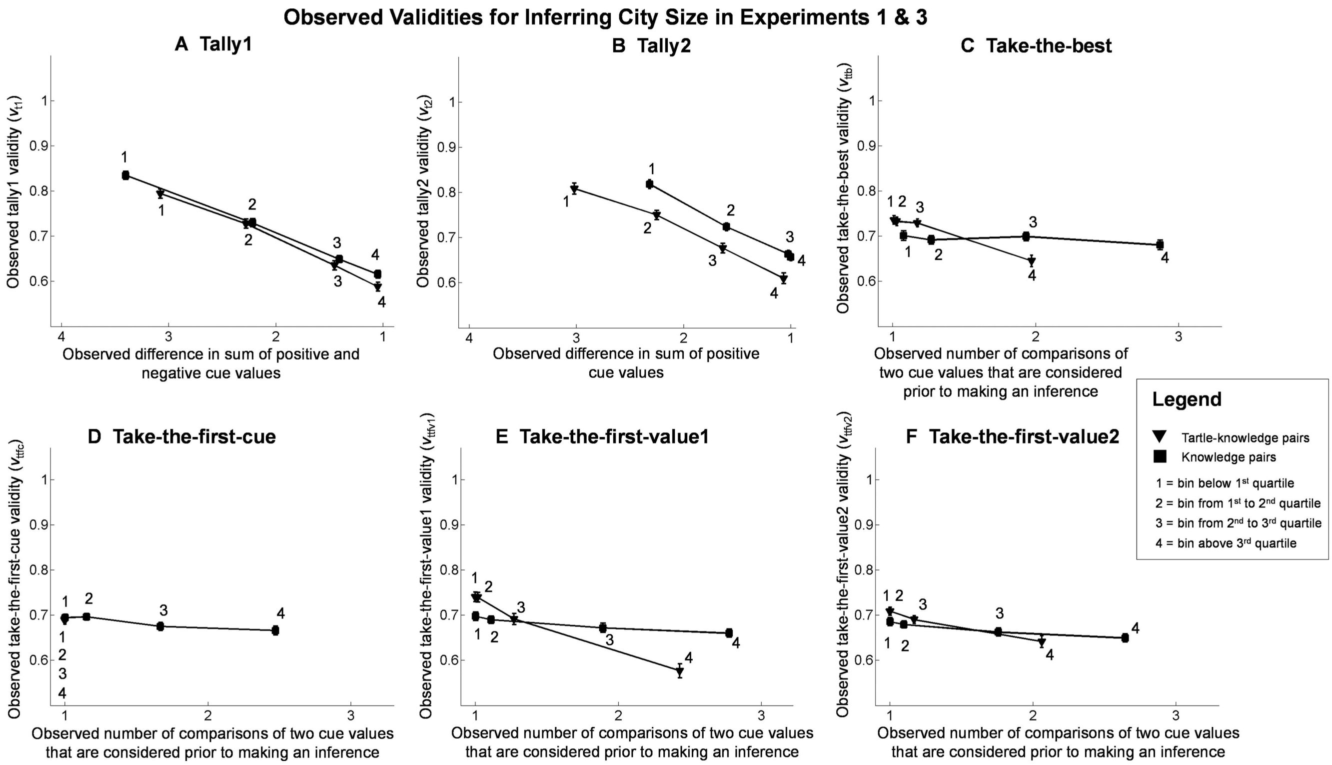

Figure 19. Simulation 11. (A) tally1, (B) tally2, (C) take-the-best, (D) take-the-first-cue, (E) take-the-first-value1, and (F) take-the-first-value2 validities (v_{t1}, v_{t2}, v_{ttb}, v_{ttfc}, v_{ttfv1}, v_{ttfv2}) for inferences about cities' size as a function of the effort entailed in applying each of the strategies. For all strategies, on the left-hand side of the *x*-axis the least effort is required to apply a strategy, on the right-hand side the most. For the sake of this illustrative analysis, for tally1 and tally2, effort is operationalized as differences between sums of cue values; for take-the-best, take-the-first-cue, take-the-first-value1, and take-the-first-value2, effort is the number of comparisons of two cue values that need to be considered. Participants' perception of retrieval times of cue values and the resulting validities were modeled by applying the timing model to each participant's observed data (Equations 9, 11, 13). Data are collapsed across Experiments 1 and 3. In all panels, symbols show mean effort and mean (± 1 *SE*) validities, computed across participants and simulation runs of the timing model. Some of the error bars are obscured by the symbols. Note that on tartle–knowledge pairs, take-the-first-cue is the most frugal strategy: The first comparison of cue values always consists of a positive or negative cue value for the knowledge object and an unknown cue value for the tartle object, allowing take-the-first-cue to make a decision on the first comparison of cue values (see Figure 10). Therefore, on tartle–knowledge pairs all four bins fall on top of each other (D). As it turns out, these four bins also fall exactly on the first bin of the knowledge pairs. As a result, the tartle–knowledge pairs' four bins and the knowledge pairs' first bin obscure each other.

tegration strategies, knowledge-based lexicographic strategies, knowledge- and accessibility-based sequential-sampling strategies, and purely accessibility-based strategies. For instance, we showed that the fluency heuristic is most likely applicable when a person has little or no knowledge about the objects in question. In such situations of limited knowledge, differences in recognition times tend to be easy to detect, favoring the applicability of the fluency heuristic. A person is less likely to be able to apply this heuristic when knowledge is abundant, because in this case differences in recognition times tend to be harder to notice. As it turns out, this is also a situation where knowledge-based strategies tend to help a person more to make accurate inferences than the fluency heuristic.

In the remaining sections of this article, we position our cognitive niche framework in broader theoretical and methodological contexts, discussing its connections to (a) complementary approaches to strategy selection, (b) the robustness of biological systems, (c) the ecological rationality of decision strategies, and (d) the development of memory models.

Cognitive Niches: A Complementary Approach to Strategy Selection

Much progress has been made in our understanding of strategy selection, for example, by assuming that people trade the accuracy of adopting a strategy against the effort and time required to use it (i.e., as in the cost–benefit approaches, e.g., J. W. Payne et al., 1988, 1993) or by describing the learning processes involved in the strategy's selection (e.g., Busemeyer & Myung, 1992; Erev & Roth, 2001; Rieskamp & Otto, 2006). By mapping out the niches where different strategies are applicable, our cognitive niche framework can be seen as a complement to these earlier approaches to strategy selection, rather than as an alternative.

For instance, neither Rieskamp and Otto's (2006) model nor other learning models, including ACT–R's mechanism for production rule selection, can easily explain how an adaptive, systematic selection occurs in the *absence* of feedback and learning processes. Moreover, like other approaches, Rieskamp and Otto's reinforcement learning model requires $N - 1$ free parameters for each of N strategies in a person's repertoire. To avoid an explosion in the number of parameters, Rieskamp and Otto suggested clustering the strategies into groups, such that the model selects between groups of strategies, rather than between the individual strategies themselves. The cognitive niche framework provides a theoretical basis for such grouping of strategies, because it allows modeling how the interplay of the environment and the cognitive system reduces the consideration set of strategies to those whose cognitive niches overlap. At the same time, for those situations where two or more strategies' niches do not overlap, the cognitive niche framework facilitates understanding how strategy selection emerges as a bottom-up process—in the absence of feedback and learning—solely through the interplay between the cognitive system and the environment.

Moreover, by mapping out how the interplay between the cognitive system and the environment shapes the regions in a strategy's niche where that strategy will help a person make accurate, fast, and effortless decisions, the cognitive niche framework provides the currencies needed by the cost–benefit, learning, and other earlier approaches to explain the selection among simultaneously applicable strategies. For instance, the cognitive niche framework enables us to model when positive feedback on the success of a strategy will be available for reinforcement learning and when not. To illustrate this, the fluency heuristic will most likely be reinforced when recognition time differences are easy to detect, because it is in this situation that a person using the heuristic is most likely to make accurate inferences.

Finally, with respect to the cost–benefit approaches, the cognitive niche framework specifies the regions within a strategy's niche where people may engage in the effort–accuracy trade-offs assumed by these approaches and where not. For example, the recognition heuristic most likely helps a person to make accurate, fast, and effortless inferences when a person has additional knowledge about an alternative (e.g., a car brand).

The Trade-Off Between Strategy Selection and Robustness in Biological Systems

The cognitive niche framework also connects to theories of human and animal cognition beyond the approaches to strategy selection discussed in this article. As is well known, a repertoire of strategies allows an organism to respond flexibly to environmental demands. At the same time, it makes an organism's cognitive system robust by creating redundancies in the system. To give an example in terms of the strategies discussed in this article, when a person's ability to perceive time is impaired, he or she may still be able to use knowledge-based strategies. In contrast, if the person has no knowledge, he or she may be capable of applying the fluency heuristic. This implies that there is a tension between having the cognitive niches of strategies overlap completely, which serves robustness, and having them be distinct, which simplifies strategy selection. Correspondingly, a well-functioning biological system will strike a balance by harboring a set of strategies that have partially overlapping cognitive niches.

Partial redundancies are often observed in biological systems, which can have multiple mechanisms that serve the same function. Because the mechanisms depend on different components, they are less susceptible to the failure of any single one (Kitano, 2004). Hammerstein, Hagen, Herz, and Herzel (2006) provided the metaphor of a carpenter's toolbox, which, rather than including several identical screwdrivers, is more likely to contain a set of different ones. Yet, if one is lost, another screwdriver will do reasonably well. To illustrate this, pigeons seem to rely on at least two mechanisms for navigation. The first depends on the sun. The second exploits the magnetic field of the earth. On overcast days, the first mechanism is hard to apply (Wiltschko & Wiltschko, 2003). Put in our terms, on sunny days, the niches of these mechanisms overlap; on overcast days they do not. Similarly, deer stags seem to make use of various mechanisms to judge a rival prior to engaging in a fight, including vision and hearing. At a distance, even when it cannot clearly see its rival, the stag can hear the deepness of its rival's roar, an indication of size. And as they close the distance to fight, the stag can then both see how large its rival is and still hear its roar (Hutchinson & Gigerenzer, 2005). Finally, we all often use our visual, auditory, sensory, and olfactory systems in parallel to explore our environment. For example, we can judge an unknown dish's taste by smelling it or by looking at it. However, except for those few with synesthesia, we will not "hear" how the dish tastes. That is, the niches of our sensory systems do not completely overlap.

An Integrative, Cognitive Approach to Ecological Rationality

In proposing the cognitive niche framework, we highlight that the study of decision making warrants detailed models of decision strategies, the cognitive capacities they draw on, and the environments in which both are embedded. For instance, within the simple heuristics framework, much research has focused on what is termed the heuristics' *ecological rationality*, that is, whether and how heuristics can exploit environmental structure (e.g., Gigerenzer et al., 1999; Hogarth & Karelaia, 2007; Katsikopoulos & Martignon, 2006; Martignon & Hoffrage, 1999). Yet, little research in this framework has implemented models of the cognitive capacities that are responsible for mapping these environmental structures into mental representations (see Dougherty et al., 2008; Tomlinson, Marewski, & Dougherty, 2011). Rather, environmental structure has often been characterized in terms of the objective properties of environments. To illustrate this, in their analyses of different heuristics' rationality, Gigerenzer and Goldstein (1996) and many others (e.g., Brighton, 2006; Chater, Oaksford, Nakisa, & Redington, 2003; Hogarth & Karelaia, 2006) studied how accurate the heuristics are for inferring objects' criterion values (e.g., a city's size) from the objects' attributes (e.g., whether a city has an airport). In doing so, these researchers relied on objective attributes, as they are listed in almanacs, rather than on what people know about the objects, or how easily they can recall what they know. Yet, as we have shown, it is the mental representations of the environment that define the niche of a strategy, influencing whether and when a person will be able to apply the strategy and how accurate, fast, and effortful using the strategy will be. If we had just studied cities' or other objects' objective attributes, rather than how they are mentally represented, then it would have been hard to discover a number of results.

Let us name just a few. For example, prior to carrying out our modeling efforts, we did not know (a) when a person using the fluency heuristic would be most likely to make accurate and effortless inference; (b) that a person using a knowledge-based strategy would be more likely to make accurate inferences than a person using the fluency heuristic; (c) that people's inferences would be best described by knowledge-based strategies; (d) that the fluency heuristic is most likely applicable, easy to use, and relied upon when there is little or no knowledge; and (e) that the recognition heuristic is most likely to help a person make fast and accurate inferences when knowledge is available, which is also (f) the situation when the recognition heuristic predicts people's inferences better than knowledge-based strategies (Marewski, Gaissmaier, Schooler, et al., 2010). In uncovering these results, we found some of our initial expectations, formed prior to our modeling efforts, did not hold up. For instance, on the basis of the recognition memory literature, we and others had expected that fluency-based information may have been prioritized over knowledge (see Hertwig et al., 2008; Pachur & Hertwig, 2006).

An Ecologically Grounded Approach for Studying Cognition

Since the publication of Woodworth's (1938) classic textbook, *Experimental Psychology*, the common practice in psychology of using *systematic design* has prescribed the isolation and manipulation of a few independent variables while all others are kept constant or randomly varied (Dhami, Hertwig, & Hoffrage, 2004; Hoffrage & Hertwig, 2006; Marewski & Olsson, 2009). The ecologically minded Brunswik (1955) argued that systematic design destroys the natural covariation of variables in the organism's habitat, making it hard to generalize from such controlled laboratory experiments to the conditions under which the organism actually performs in its environment.

In our investigation of the cognitive niches of decision strategies we were careful not to disrupt the natural covariation of memory variables. But as the strategies depend on the knowledge people bring with them to the experiment about, say, cities or companies, we needed to model in some detail how this knowledge is represented in memory. Using little more than web frequency data, we developed an ecological ACT–R memory model for predicting (a) how likely people are to recognize objects in the world, (b) how likely they are to know something more about these recognized objects, and (c) the associated recognition time distributions. The accessibility of these simulated memories reflects not only the natural environment, outside the laboratory, but also how easily a person in the laboratory can retrieve like memories. Our work thus aids developing quantitative, ecologically grounded models of accessibility, availability, familiarity, fluency, and recognition—notions that have been used, at least since Hume (1740/1978), to account for behavior (e.g., Goldstein & Gigerenzer, 2002; Jacoby & Dallas, 1981; Koriat, 1993; Tversky & Kahneman, 1973; Winkielman et al., 1998).

To conclude, in modeling strategy selection, by necessity, we contributed to research on how to build models of memory based on environmental data (e.g., Anderson & Schooler, 1991; Burgess & Lund, 1997; Griffiths, Steyvers, & Tenenbaum, 2007; Landauer & Dumais, 1997), additionally tying one such ecological memory model to models of decision strategies (e.g., Dougherty, Gettys, & Ogden, 1999; Juslin & Persson, 2002; Pleskac, 2007) as well as to models of other cognitive capacities like time perception. We hope that one day such detailed models of cognition will allow us to better understand how human decision making reflects the environment in which we act. Perhaps it is only then that we will begin to fully appreciate the subtle complexities of what we thought were among the simplest of decisions.

References

Adelman, J. S., & Brown, G. D. A. (2008). Modeling lexical decision: The form of frequency and diversity effects. *Psychological Review, 115,* 214–227. doi:10.1037/0033-295X.115.1.214

Alba, J. W., & Marmorstein, H. (1987). The effects of frequency knowledge on consumer decision making. *Journal of Consumer Research, 14,* 14–25. doi:10.1086/209089

Anderson, J. R. (1990). *The adaptive character of thought.* Hillsdale, NJ: Erlbaum.

Anderson, J. R. (1993). *Rules of the mind.* Hillsdale, NJ: Erlbaum.

Anderson, J. R., & Betz, J. (2001). A hybrid model of categorization. *Psychonomic Bulletin & Review, 8,* 629–647. doi:10.3758/BF03196200

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review, 111,* 1036–1060. doi:10.1037/0033-295X.111.4.1036

Anderson, J. R., Bothell, D., Lebiere, C., & Matessa, M. (1998). An integrated theory of list memory. *Journal of Memory and Language, 38,* 341–380. doi:10.1006/jmla.1997.2553

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought.* Mahwah, NJ: Erlbaum.

Anderson, J. R., & Milson, R. (1989). Human memory: An adaptive perspective. *Psychological Review, 96,* 703–719. doi:10.1037/0033-295X.96.4.703

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science, 2,* 396–408. doi:10.1111/j.1467-9280.1991.tb00174.x

Atkinson, R. C., & Juola, J. F. (1974). Search and decision processes in recognition memory. In D. H. Krantz, R. C. Atkinson, R. D. Lucek, & P. Suppes (Eds.), *Contemporary developments in mathematical psychology* (Vol. 1, pp. 243–293). San Francisco, CA: Freeman.

Beach, L. R., & Mitchell, T. R. (1978). A contingency model for the selection of decision strategies. *Academy of Management Review, 3,* 439–449. doi:10.2307/257535

Begg, I. M., Anas, A., & Farinacci, S. (1992). Dissociation of processes in belief: Source recollection, statement familiarity, and the illusion of truth. *Journal of Experimental Psychology: General, 121,* 446–458. doi:10.1037/0096-3445.121.4.446

Brighton, H. (2006). Robust inference with simple cognitive models. In C. Lebiere & R. Wray (Eds.), *AAAI Spring Symposium: Cognitive science principles meet AI-hard problems* (pp. 17–22). Menlo Park, CA: American Association for Artificial Intelligence.

Bröder, A. (2003). Decision making with the "adaptive toolbox": Influence of environmental structure, intelligence, and working memory load. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 611–625. doi:10.1037/0278-7393.29.4.611

Bröder, A., & Gaissmaier, W. (2007). Sequential processing of cues in memory-based multi-attribute decisions. *Psychonomic Bulletin & Review, 14,* 895–900. doi:10.3758/BF03194118

Bröder, A., & Newell, B. R. (2008). Challenging some common beliefs: Empirical work within the adaptive toolbox. *Judgment and Decision Making, 3,* 205–214.

Bröder, A., & Schiffer, S. (2003). Take the best versus simultaneous feature matching: Probabilistic inferences from memory and effects of representation format. *Journal of Experimental Psychology: General, 132,* 277–293. doi:10.1037/0096-3445.132.2.277

Bröder, A., & Schiffer, S. (2006). Stimulus format and working memory in fast and frugal strategy selection. *Journal of Behavioral Decision Making, 19,* 361–380. doi:10.1002/bdm.533

Bruner, J. S. (1957). On perceptual readiness. *Psychological Review, 64,* 123–152. doi:10.1037/h0043805

Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review, 62,* 193–217. doi:10.1037/h0047470

Burgess, C., & Lund, K. (1997). Modelling parsing constraints with high-dimensional context space. *Language and Cognitive Processes, 12,* 177–210. doi:10.1080/016909697386844

Busemeyer, J. R. (1993). Violations of the speed-accuracy tradeoff relation: Decreases in decision accuracy with increases in decision time. In O. Svenson & A. J. Maule (Eds.), *Time pressure and stress in human judgment and decision making* (pp. 181–193). New York, NY: Plenum Press.

Busemeyer, J. R., & Myung, I. J. (1992). An adaptive approach to human decision making: Learning theory, decision theory, and human performance. *Journal of Experimental Psychology: General, 121,* 177–194. doi:10.1037/0096-3445.121.2.177

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic–cognitive approach to decision making in an uncertain environment. *Psychological Review, 100,* 432–459. doi:10.1037/0033-295X.100.3.432

Chater, N., Oaksford, M., Nakisa, R., & Redington, M. (2003). Fast, frugal and rational: How rational norms explain behavior. *Organizational Behavior and Human Decision Processes, 90,* 63–86. doi:10.1016/S0749-5978(02)00508-3

Christensen-Szalanski, J. J. (1978). Problem solving strategies: A selection mechanism, some implications, and some data. *Organizational Behavior and Human Performance, 22,* 307–323. doi:10.1016/0030-5073(78)90019-3

Cisek, P. (2007). Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences, 362,* 1585–1599. doi:10.1098/rstb.2007.2054

Coates, S. L., Butler, L. T., & Berry, D. C. (2004). Implicit memory: A prime example for brand consideration and choice. *Applied Cognitive Psychology, 18,* 1195–1211. doi:10.1002/acp.1044

Coates, S. L., Butler, L. T., & Berry, D. C. (2006). Implicit memory and consumer choice: The mediating role of brand familiarity. *Applied Cognitive Psychology, 20,* 1101–1116. doi:10.1002/acp.1262

Cokely, E. T., Kelley, C. M., & Gilchrist, A. H. (2006). Sources of individual differences in working memory: Contributions of strategy to capacity. *Psychonomic Bulletin & Review, 13,* 991–997. doi:10.3758/BF03213914

Coyle, T. R., Read, L. E., Gaultney, J. F., & Bjorklund, D. F. (1998). Giftedness and variability in strategy processing on a multitrial memory task: Evidence for stability in gifted cognition. *Learning and Individual Differences, 10,* 273–290. doi:10.1016/S1041-6080(99)80123-X

Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist, 34,* 571–582. doi:10.1037/0003-066X.34.7.571

Dawes, R. M., & Corrigan, B. (1974). Linear models in decision making. *Psychological Bulletin, 81,* 95–106. doi:10.1037/h0037613

Dhami, M. K., Hertwig, R., & Hoffrage, U. (2004). The role of representative design in an ecological approach to cognition. *Psychological Bulletin, 130,* 959–988. doi:10.1037/0033-2909.130.6.959

Dougherty, M. R. P., Franco-Watkins, A. N., & Thomas, R. (2008). Psychological plausibility of the theory of probabilistic mental models and the fast and frugal heuristics. *Psychological Review, 115,* 199–211. doi:10.1037/0033-295X.115.1.199

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. doi:10.1037/0033-295X.106.1.180

Dror, I. E., Busemeyer, J. R., & Basola, B. (1999). Decision making under time pressure: An independent test of sequential sampling models. *Memory & Cognition, 27,* 713–725. doi:10.3758/BF03211564

Einhorn, H. J. (1970). The use of nonlinear, noncompensatory models in decision making. *Psychological Bulletin, 73,* 221–230. doi:10.1037/h0028695

Einhorn, H. J., & Hogarth, R. M. (1975). Unit weighting schemes for decision making. *Organizational Behavior and Human Performance, 13,* 171–192. doi:10.1016/0030-5073(75)90044-6

Einhorn, H. J., & Hogarth, R. M. (1981). Behavioral decision theory: Processes of judgement and choice. *Annual Review of Psychology, 32,* 53–88. doi:10.1146/annurev.ps.32.020181.000413

Eisenberg, P., & Becker, C. A. (1982). Semantic context effects in visual word recognition, sentence processing, and reading: Evidence for semantic strategies. *Journal of Experimental Psychology: Human Perception and Performance, 8,* 739–756. doi:10.1037/0096-1523.8.5.739

Erdfelder, E., Küpper-Tetzel, C. E., & Mattern, S. D. (2011). Threshold models of recognition and the recognition heuristic. *Judgment and Decision Making, 6,* 7–22.

Erev, I., & Roth, A. E. (2001). Simple reinforcement learning models and reciprocation in the prisoner's dilemma game. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 215–231). Cambridge, MA: MIT Press.

Fishburn, P. C. (1980). Lexicographic additive differences. *Journal of Mathematical Psychology, 21,* 191–218. doi:10.1016/0022-2496(80)90037-1

Fu, W.-T., & Anderson, J. R. (2006). From recurrent choice to skill learning: A reinforcement-learning model. *Journal of Experimental Psychology: General, 135,* 184–206. doi:10.1037/0096-3445.135.2.184

Gaissmaier, W. (2007). *The mnemonic decision maker: How search in memory shapes decision making* (Doctoral dissertation, Free University, Berlin, Germany). Retrieved from http://www.diss.fu-berlin.de/diss/receive/FUDISS_thesis_000000005913

Gaissmaier, W., & Marewski, J. N. (2011). Forecasting elections with mere recognition from lousy samples. *Judgment and Decision Making, 6,* 73–88.

Gaissmaier, W., Schooler, L. J., & Rieskamp, J. (2006). Simple predictions fueled by capacity limitations: When are they successful? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32,* 966–982. doi:10.1037/0278-7393.32.5.966

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston, MA: Houghton Mifflin.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review, 103,* 650–669. doi:10.1037/0033-295X.103.4.650

Gigerenzer, G., & Goldstein, D. G. (2011). The recognition heuristic: A decade of research. *Judgment and Decision Making, 6,* 100–121.

Gigerenzer, G., Hoffrage, U., & Goldstein, D. G. (2008). Fast and frugal heuristics are plausible models of cognition: Reply to Dougherty, Franco-Watkins, and Thomas (2008). *Psychological Review, 115,* 230–239. doi:10.1037/0033-295X.115.1.230

Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review, 98,* 506–528. doi:10.1037/0033-295X.98.4.506

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

Ginossar, Z., & Trope, Y. (1987). Problem solving in judgment under uncertainty. *Journal of Personality and Social Psychology, 52,* 464–474. doi:10.1037/0022-3514.52.3.464

Glöckner, A., & Betsch, T. (2008). Modeling option and strategy choices with connectionist networks: Towards an integrative model of automatic and deliberate decision making. *Judgment and Decision Making, 3,* 215–228.

Glöckner, A., Betsch, T., & Schindler, N. (2010). Coherence shifts in probabilistic inference tasks. *Journal of Behavioral Decision Making, 23,* 439–462. doi:10.1002/bdm.668

Gluck, K. A., Ball, J. T., & Krusmark, M. A. (2007). Cognitive control in a computational model of the Predator pilot. In W. Gray (Ed.), *Integrated models of cognitive systems* (pp. 13–28). New York, NY: Oxford University Press.

Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review, 109,* 75–90. doi:10.1037/0033-295X.109.1.75

Gray, W. D., Sims, C. R., Fu, W.-T., & Schoelles, M. J. (2006). The soft constraints hypothesis: A rational analysis approach to resource allocation for interactive behavior. *Psychological Review, 113,* 461–482. doi:10.1037/0033-295X.113.3.461

Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review, 114,* 211–244. doi:10.1037/0033-295X.114.2.211

Gronlund, S. D., & Ratcliff, R. (1989). The time course of item and associative information: Implications for global memory models. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 846–858. doi:10.1037/0278-7393.15.5.846

Hammerstein, P., Hagen, E. H., Herz, A. V. M., & Herzel, H. (2006). Robustness: A key to evolutionary design. *Biological Theory, 1,* 90–93. doi:10.1162/biot.2006.1.1.90

Haruvy, E., & Stahl, D. O. (2007). Equilibrium selection and bounded rationality in symmetric normal-form games. *Journal of Economic Behavior & Organization, 62,* 98–119. doi:10.1016/j.jebo.2005.05.002

Haruvy, E., Stahl, D. O., & Wilson, P. W. (2001). Modeling and testing for heterogeneity in observed strategic behavior. *Review of Economics and Statistics, 83,* 146–157. doi:10.1162/003465301750160117

Hertwig, R., Gigerenzer, G., & Hoffrage, U. (1997). The reiteration effect in hindsight bias. *Psychological Review, 104,* 194–202. doi:10.1037/0033-295X.104.1.194

Hertwig, R., Herzog, S. M., Schooler, L. J., & Reimer, T. (2008). Fluency heuristic: A model of how the mind exploits a by-product of information retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34,* 1191–1206. doi:10.1037/a0013025

Hertwig, R., Pachur, T., & Kurzenhäuser, S. (2005). Judgments of risk frequencies: Tests of possible cognitive mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 621–642. doi:10.1037/0278-7393.31.4.621

Hertwig, R., & Todd, P. M. (2003). More is not always better: The benefits of cognitive limits. In D. Hardman & L. Macchi (Eds.), *Thinking: Psychological perspectives on reasoning, judgment and decision making* (pp. 213–231). Chichester, England: Wiley.

Herzog, S. M., & Hertwig, R. (2011). The wisdom of ignorant crowds: Predicting sport outcomes by mere recognition. *Judgment and Decision Making, 6,* 58–72.

Higgins, E. T. (1996). Knowledge activation: Accessibility, applicability, and salience. In E. T. Higgins & A. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 133–168). New York, NY: Guilford Press.

Hilbig, B. E., & Pohl, R. F. (2009). Ignorance- vs. evidence-based decision making: A decision time analysis of the recognition heuristic. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35,* 1296–1305. doi:10.1037/a0016565

Hintzman, D. L., & Curran, T. (1994). Retrieval dynamics of recognition and frequency judgments: Evidence for separate processes of familiarity and recall. *Journal of Memory and Language, 33,* 1–18. doi:10.1006/jmla.1994.1001

Hoffrage, U., & Hertwig, R. (2006). Which world should be represented in representative design? In K. Fiedler & P. Juslin (Eds.), *Information sampling and adaptive cognition* (pp. 381–408). New York, NY: Cambridge University Press.

Hogarth, R. M., & Karelaia, N. (2006). Take-the-best and other simple strategies: Why and when they work "well" with binary cues. *Theory and Decision, 61,* 205–249. doi:10.1007/s11238-006-9000-8

Hogarth, R. M., & Karelaia, N. (2007). Heuristics and linear models of judgment: Matching rules and environments. *Psychological Review, 114,* 733–758. doi:10.1037/0033-295X.114.3.733

Howes, A., Lewis, R. L., & Vera, A. H. (2009). Rational adaptation under task and processing constraints: Implications for testing theories of cognition and action. *Psychological Review, 116,* 717–751. doi:10.1037/a0017187

Howes, A., Vera, A., & Lewis, R. L. (2007). Bounding rational analysis: Constraints on asymptotic performance. In W. D. Gray (Ed.), *Integrated models of cognitive systems* (pp. 403–413). New York, NY: Oxford University Press.

Huber, O. (1989). Information-processing operators in decision making. In H. Montgomery & O. Svenson (Eds.), *Process and structure in human decision making* (pp. 3–21). New York, NY: Wiley.

Hume, D. L. (1978). *A treatise of human nature.* Oxford, England: Clarendon Press. (Original work published 1740)

Hutchinson, J. M. C., & Gigerenzer, G. (2005). Simple heuristics and rules of thumb: Where psychologists and behavioural biologists might meet. *Behavioural Processes, 69,* 97–124. doi:10.1016/j.beproc.2005.02.019

Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language, 30,* 513–541. doi:10.1016/0749-596X(91)90025-F

Jacoby, L. L., & Brooks, L. R. (1984). Nonanalytic cognition: Memory, perception and concept learning. In G. H. Bower (Ed.), *Psychology of learning and motivation* (Vol. 18, pp. 1–47). New York, NY: Academic Press.

Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General, 110,* 306–340. doi:10.1037/0096-3445.110.3.306

Jacoby, L. L., Kelley, C. M., Brown, J., & Jasechko, J. (1989). Becoming famous overnight: Limits on the ability to avoid unconscious influences of the past. *Journal of Personality and Social Psychology, 56,* 326–338. doi:10.1037/0022-3514.56.3.326

Johnson, J. G., & Raab, M. (2003). Take-the-first: Option generation and resulting choices. *Organizational Behavior and Human Decision Processes, 91,* 215–229. doi:10.1016/S0749-5978(03)00027-X

Juslin, P., & Persson, M. (2002). PROBabilities from Exemplars (PROBEX): A "lazy" algorithm for probabilistic inference from generic knowledge. *Cognitive Science, 26,* 563–607. doi:10.1207/s15516709cog2605_2

Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist, 58,* 697–720. doi:10.1037/0003-066X.58.9.697

Katsikopoulos, K. V., & Martignon, L. (2006). Naïve heuristics for paired comparisons: Some results on their relative accuracy. *Journal of Mathematical Psychology, 50,* 488–494. doi:10.1016/j.jmp.2006.06.001

Kitano, H. (2004). Biological robustness. *Nature Reviews Genetics, 5,* 826–837.

Koriat, A. (1993). How do we know that we know? The accessibility model of feeling of knowing. *Psychological Review, 100,* 609–639. doi:10.1037/0033-295X.100.4.609

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review, 104,* 211–240. doi:10.1037/0033-295X.104.2.211

Lee, M. D., & Cummins, T. D. R. (2004). Evidence accumulation in decision making: Unifying the "take the best" and "rational" models.

Psychonomic Bulletin & Review, 11, 343–352. doi:10.3758/BF03196581

Lovett, M. C., & Anderson, J. R. (1996). History of success and current context in problem solving: Combined influences on operator selection. *Cognitive Psychology, 31,* 168–217. doi:10.1006/cogp.1996.0016

Mandler, G. (1980). Recognizing: The judgment of previous occurrence. *Psychological Review, 87,* 252–271. doi:10.1037/0033-295X.87.3.252

Marewski, J. N. (2008). *Ecologically rational strategy selection* (Doctoral dissertation, Free University, Berlin, Germany). Retrieved from http://www.diss.fu-berlin.de/diss/receive/FUDISS_thesis_000000019008?lang=en

Marewski, J. N. (2010). On the theoretical precision and strategy selection problem of a single-strategy approach: A comment on Glöckner, Betsch, and Schindler (2010). *Journal of Behavioral Decision Making, 23,* 463–467. doi:10.1002/bdm.680

Marewski, J. N., Gaissmaier, W., & Gigerenzer, G. (2010a). Good judgments do not require complex cognition. *Cognitive Processing, 11,* 103–121. doi:10.1007/s10339-009-0337-0

Marewski, J. N., Gaissmaier, W., & Gigerenzer, G. (2010b). We favor formal models of heuristics rather than loose lists of dichotomies: A Reply to Evans and Over. *Cognitive Processing, 11,* 177–179. doi:10.1007/s10339-009-0340-5

Marewski, J. N., Gaissmaier, W., Schooler, L. J., Goldstein, D. G., & Gigerenzer, G. (2009). Do voters use episodic knowledge to rely on recognition? In N. A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 2232–2237). Austin, TX: Cognitive Science Society.

Marewski, J. N., Gaissmaier, W., Schooler, L. J., Goldstein, D. G., & Gigerenzer, G. (2010). From recognition to decisions: Extending and testing recognition-based models for multi-alternative inference. *Psychonomic Bulletin & Review, 17,* 287–309. doi:10.3758/PBR.17.3.287

Marewski, J. N., & Olsson, H. (2009). Beyond the null ritual: Formal modeling of psychological processes. *Zeitschrift für Psychologie/Journal of Psychology, 217,* 49–60.

Marewski, J. N., Pohl, R. F., & Vitouch, O. (2010). Recognition-based judgments and decisions: Introduction to the special issue (Vol. 1). *Judgment and Decision Making, 5,* 207–215.

Marewski, J. N., Pohl, R. F., & Vitouch, O. (2011). Recognition-based judgments and decisions: Introduction to the special issue (II). *Judgment and Decision Making, 6,* 1–6.

Marewski, J. N., Schooler, L. J., & Gigerenzer, G. (2010). Five principles for studying people's use of heuristics. *Acta Psychologica Sinica, 42,* 72–87. doi:10.3724/SP.J.1041.2010.00072

Martignon, L., & Hoffrage, U. (1999). Why does one-reason decision making work? A case study in ecological rationality. In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 119–140). New York, NY: Oxford University Press.

Mata, R., Schooler, L. J., & Rieskamp, J. (2007). The aging decision maker: Cognitive aging and the adaptive selection of decision strategies. *Psychology and Aging, 22,* 796–810. doi:10.1037/0882-7974.22.4.796

McAllister, D. W., Mitchell, T. R., & Beach, L. R. (1979). The contingency model for the selection of decision strategies: An empirical test of the effects of significance, accountability, and reversibility. *Organizational Behavior and Human Performance, 24,* 228–244. doi:10.1016/0030-5073(79)90027-8

McElree, B., Dolan, P. O., & Jacoby, L. L. (1999). Isolating the contributions of familiarity and source information to item recognition: A time course analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25,* 563–582. doi:10.1037/0278-7393.25.3.563

Newell, B. R., & Fernandez, D. (2006). On the binary quality of recognition and the inconsequentiality of further knowledge: Two critical tests of the recognition heuristic. *Journal of Behavioral Decision Making, 19,* 333–346. doi:10.1002/bdm.531

Newell, B. R., & Shanks, D. R. (2003). Take the best or look at the rest? Factors influencing "one-reason" decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 53–65. doi:10.1037/0278-7393.29.1.53

Nosofsky, R. M., & Bergert, F. B. (2007). Limitations of exemplar models of multi-attribute probabilistic inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 33,* 999–1019. doi:10.1037/0278-7393.33.6.999

Oaksford, M., & Chater, N. (Eds.). (1998). *Rational models of cognition.* Oxford, England: Oxford University Press.

Oppenheimer, D. M. (2008). The secret life of fluency. *Trends in Cognitive Sciences, 12,* 237–241. doi:10.1016/j.tics.2008.02.014

Pachur, T., & Biele, G. (2007). Forecasting from ignorance: The use and usefulness of recognition in lay predictions of sports events. *Acta Psychologica, 125,* 99–116. doi:10.1016/j.actpsy.2006.07.002

Pachur, T., Bröder, A., & Marewski, J. N. (2008). The recognition heuristic in memory-based inference: Is recognition a non-compensatory cue? *Journal of Behavioral Decision Making, 21,* 183–210. doi:10.1002/bdm.581

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, and Cognition, 32,* 983–1002. doi:10.1037/0278-7393.32.5.983

Payne, J. W. (1976). Task complexity and contingent processing in decision making: An information search and protocol analysis. *Organizational Behavior and Human Performance, 16,* 366–387. doi:10.1016/0030-5073(76)90022-2

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 534–552. doi:10.1037/0278-7393.14.3.534

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker.* New York, NY: Cambridge University Press.

Payne, S. J., Richardson, J., & Howes, A. (2000). Strategic use of familiarity in display-based problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26,* 1685–1701. doi:10.1037/0278-7393.26.6.1685

Pleskac, T. J. (2007). A signal detection analysis of the recognition heuristic. *Psychonomic Bulletin & Review, 14,* 379–391. doi:10.3758/BF03194081

Pohl, R. (2006). Empirical tests of the recognition heuristic. *Journal of Behavioral Decision Making, 19,* 251–271. doi:10.1002/bdm.522

Provost, F. J., & Buchanan, B. G. (1995). Inductive policy: The pragmatics of bias selection. *Machine Learning, 20,* 35–61. doi:10.1007/BF00993474

Rapoport, A., & Wallsten, T. S. (1972). Individual decision behavior. *Annual Review of Psychology, 23,* 131–176. doi:10.1146/annurev.ps.23.020172.001023

Ratcliff, R. (2002). A diffusion model account of response time and accuracy in a brightness discrimination task: Fitting real data and failing to fit fake but plausible data. *Psychonomic Bulletin & Review, 9,* 278–291. doi:10.3758/BF03196283

Ratcliff, R., & McKoon, G. (1989). Similarity information versus relational information: Differences in the time course of retrieval. *Cognitive Psychology, 21,* 139–155. doi:10.1016/0010-0285(89)90005-4

Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review, 111,* 333–367. doi:10.1037/0033-295X.111.2.333

Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience? *Personality and Social Psychology Review, 8,* 364–382. doi:10.1207/s15327957pspr0804_3

Reder, L. M. (1987). Strategy selection in question answering. *Cognitive Psychology, 19,* 90–138. doi:10.1016/0010-0285(87)90005-3

Reimer, T., & Katsikopoulos, K. V. (2004). The use of recognition in group decision-making. *Cognitive Science, 28,* 1009–1029. doi:10.1207/s15516709cog2806_6

Rieskamp, J. (2006). Perspectives on probabilistic inferences: Reinforcement learning and an adaptive network compared. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32,* 1355–1370. doi:10.1037/0278-7393.32.6.1355

Rieskamp, J., & Hoffrage, U. (1999). When do people use simple heuristics, and how can we tell? In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 141–167). New York, NY: Oxford University Press.

Rieskamp, J., & Hoffrage, U. (2008). Inferences under time pressure: How opportunity costs affect strategy selection. *Acta Psychologica, 127,* 258–276. doi:10.1016/j.actpsy.2007.05.004

Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General, 135,* 207–236. doi:10.1037/0096-3445.135.2.207

Ritter, S., Anderson, J. R., Koedinger, K. R., & Corbett, A. (2007). Cognitive tutor: Applied research in mathematics education. *Psychonomic Bulletin & Review, 14,* 249–255. doi:10.3758/BF03194060

Salvucci, D. D. (2006). Modeling driver behavior in a cognitive architecture. *Human Factors, 48,* 362–380. doi:10.1518/001872006777724417

Schooler, L. J., & Anderson, J. R. (1997). The role of process in the rational analysis of memory. *Cognitive Psychology, 32,* 219–250. doi:10.1006/cogp.1997.0652

Schooler, L. J., & Hertwig, R. (2005). How forgetting aids heuristic inference. *Psychological Review, 112,* 610–628. doi:10.1037/0033-295X.112.3.610

Sedlmeier, P., Hertwig, R., & Gigerenzer, G. (1998). Are judgments of the positional frequencies of letters systematically biased due to availability? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 754–770. doi:10.1037/0278-7393.24.3.754

Shepard, R. N. (2001). Perceptual-cognitive universals as reflections of the world. *Behavioral and Brain Sciences, 24,* 581–601.

Siegler, R. S. (1988). Strategy choice procedures and the development of multiplication skill. *Journal of Experimental Psychology: General, 117,* 258–275. doi:10.1037/0096-3445.117.3.258

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review, 63,* 129–138. doi:10.1037/h0042769

Smith, E. E., Patalano, A. L., & Jonides, J. (1998). Alternative strategies of categorization. *Cognition, 65,* 167–196. doi:10.1016/S0010-0277(97)00043-7

Stahl, D. O. (1999). Evidence based rule learning in symmetric normal-form games. *International Journal of Game Theory, 28,* 111–130. doi:10.1007/s001820050101

Stahl, D. O. (2000). Rule learning in symmetric normal-form games: Theory and evidence. *Games and Economic Behavior, 32,* 105–138. doi:10.1006/game.1999.0754

Stahl, D. O., & Wilson, P. W. (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization, 25,* 309–327. doi:10.1016/0167-2681(94)90103-1

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology, 53,* 1–26. doi:10.1016/j.cogpsych.2005.10.003

Taatgen, N. A., Van Rijn, H., & Anderson, J. R. (2007). An integrated theory of prospective time interval estimation: The role of cognition, attention, and learning. *Psychological Review, 114,* 577–598. doi:10.1037/0033-295X.114.3.577

Taatgen, N., & Wallach, D. (2002). Whether skill acquisition is rule or instance based is determined by the structure of the task. *Cognitive Science Quarterly, 2,* 163–204.

Tobii Technology. (2005). *Tobii Eye Tracker & ClearView Analysis Software: User manual.* Danderyd, Sweden: Author.

Todd, P. M., & Gigerenzer, G. (2000). Précis of *Simple heuristics that make us smart. Behavioral and Brain Sciences, 23,* 727–780. doi:10.1017/S0140525X00003447

Tomlinson, T., Marewski, J. N., & Dougherty, M. R. (2011). Four challenges for cognitive research on the recognition heuristic and a call for a research strategy shift. *Judgment and Decision Making, 6,* 89–99.

Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior, 5,* 381–391. doi:10.1016/S0022-5371(66)80048-8

Tversky, A. (1969). Intransitivity of preferences. *Psychological Review, 76,* 31–48. doi:10.1037/h0026750

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review, 79,* 281–299. doi:10.1037/h0032955

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology, 5,* 207–232. doi:10.1016/0010-0285(73)90033-9

Vicente, K. J. (2003). Beyond the lens model and direct perception: Toward a broader ecological psychology. *Ecological Psychology, 15,* 241–267. doi:10.1207/S15326969ECO1503_4

Vilalta, R., & Drissi, Y. (2002). A perspective view and survey of meta-learning. *Artificial Intelligence Review, 18,* 77–95. doi:10.1023/A:1019956318069

Volz, K. G., Schooler, L. J., 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. doi:10.1162/jocn.2006.18.11.1924

Wagenmakers, E.-J., & Brown, S. (2007). On the linear relation between the mean and the standard deviation of a response time distribution. *Psychological Review, 114,* 830–841. doi:10.1037/0033-295X.114.3.830

Wagenmakers, E.-J., Zeelenberg, R., Steyvers, M., Shiffrin, R. M., & Raaijmakers, J. G. W. (2004). Nonword repetition in lexical decision: Support for two opposing processes. *Quarterly Journal of Experimental Psychology, 57*(A), 1191–1210.

Walker, M. (2000). An application of reinforcement learning to dialogue strategy selection in a spoken dialogue system for email. *Journal of Artificial Intelligence Research, 12,* 387–416.

Whittlesea, B. W. A. (1993). Illusions of familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 1235–1253. doi:10.1037/0278-7393.19.6.1235

Wiltschko, R., & Wiltschko, W. (2003). Avian navigation: From historical to modern concepts. *Animal Behaviour, 65,* 257–272. doi:10.1006/anbe.2003.2054

Winkielman, P., Schwarz, N., & Belli, R. F. (1998). The role of ease of retrieval and attribution in memory judgments: Judging your memory as worse despite recalling more events. *Psychological Science, 9,* 124–126. doi:10.1111/1467-9280.00022

Woodworth, R. (1938). *Experimental psychology.* New York, NY: Holt.

(*Appendices follow*)

Appendix A

Methods of Experiments Conducted

All experiments were conducted in the Center for Adaptive Behavior and Cognition at the Max Planck Institute for Human Development in Berlin, Germany. Descriptive statistics for the main experiments are reported in Online Materials A, Tables OA1–OA5.

Experiment 1

Participants and Materials

Forty-nine right-handed persons participated in the study (43% female; mean age 24.0 years, SD = 3.1). Participants received a guaranteed minimum payment of €13 ($17) supplemented by a performance bonus. We chose the 70 largest Austrian, British, French, German, Italian, Spanish, and U.S. cities as candidate stimuli for the experimental tasks. To obtain a clean measure of recognition times we included only one-word city names of similar word length (i.e., 5–8 letters). This resulted in 240 cities (see Online Materials B, Table OB1), which we also used as stimuli in Experiments 2, 3, and 10.

Procedure

In an *inference task* (cf. Table 2), we presented two cities on a computer screen (one on the left side and the other on the right) and asked participants to infer which city had more inhabitants. Pairs of cities were randomly drawn for each country without replacement such that each city could appear only once throughout the task. In addition to the guaranteed minimum payment of €13, participants received €0.04 ($0.05) for each correct response. Four cents were subtracted from this additional payment for each incorrect response. (No feedback on the correctness of the responses was given until after the experiment.)

In a subsequent computerized *recognition task*, we gave participants the name of one city at a time and asked them to judge whether they had seen or heard of the city prior to participating in the study. The response times for positive responses represented our measure of recognition time.

After a 60-sec break, we presented all cities again in a computerized general knowledge task. Here, we asked participants how much they knew about each city. There were three possible answers: (a) "never heard of it and never seen it before participating in the study"; (b) "heard of it or seen it before participating in the study but do not know anything else about it beyond recognizing its name"; or (c) "heard of it or seen it before participating in the study and know something about the city beyond simply recognizing it."

Last, in a cue-knowledge task, we asked participants to indicate for each city whether it (a) has an international airport (airport cue), (b) has a university (university cue), (c) is a significant industry site (industry cue), and (d) is a world-famous tourist site (tourism cue). Participants could respond with "yes," "no," or "don't know." Previous participants from the subject pool of the Center for Adaptive Behavior and Cognition considered these knowledge cues as the most useful ones for inferring city size (Pachur, Bröder, & Marewski, 2008). Participants received €0.04 for each correct response on top of the guaranteed minimum payment. For incorrect responses, €0.04 was subtracted from this additional payment. (No feedback on the correctness of the responses was given until after the experiment.) Participants did not receive payment nor did they lose money for "don't know" responses.

Within all tasks, the order of presentation of stimuli was randomized. Each trial was preceded by a small fixation cross for 1,000 msec. We instructed participants to always fixate on this cross until it disappeared and to respond as quickly and accurately as possible upon stimulus onset. All responses were made on a standard PC keyboard. Positive responses were always made with the index finger of the right hand. Participants took a 20-sec break every other 12 trials (inference, recognition, and general knowledge tasks) and every other 12 question blocks (cue-knowledge task), respectively.

Experiment 2

Participants and Materials

Seventy-one right-handed persons participated in the study (62% female; mean age 24.9 years, SD = 2.8). Participants were required to be enrolled in a university in Berlin. They received a flat payment of €8 ($11). We assigned participants randomly to two experimental groups (between subjects), labeled the *instruct group* (n = 37) and the *no-instruct group* (n = 34), respectively. Participants in the instruct group had been told to always use the fluency heuristic. Participants in the no-instruct group received no such instructions. Stimuli were the same cities used in Experiments 1, 3, and 10 (see Online Materials B, Table OB1).

Procedure

In the computerized experiment, participants were presented first with an inference task, second a recognition task, third a general knowledge task, and last an estimation task. The procedure differed from Experiment 1 only in the following respects.

First, in the inference task, we asked participants to infer which of two cities would be recognized by more of 100 randomly chosen students from Berlin (instead of which city had more inhabitants). In the instruct group, the instructions to use the fluency heuristic were as follows:

(*Appendices continue*)

> In pairs of cities in which you recognize both city names, that is, in which you have heard of or seen both city names before, please always use the following strategy: In such pairs, always decide that the city name you recognize faster is also recognized by more students.

In both groups, participants were not given additional payment for correct responses, nor were their payments reduced for incorrect ones. The inference task used in Experiment 2 follows the tradition of Jacoby et al.'s (1989) classic "fame" studies, where fluency and recognition processes were examined by asking people to judge other people's fame. Inferring other people's recognition of city names can be thought of as a judgment about the cities' fame.

Second, we replaced the cue-knowledge task with an estimation task. This task served to assess whether participants in both experimental groups considered retrieval fluency to be an equally valid piece of information. Participants estimated the fluency validity for inferring which of two randomly drawn cities from a country would be recognized by more students. The estimation task was presented in a frequency format (e.g., "Imagine 100 randomly drawn pairs of cities. . . . In how many pairs is the city you recognize faster the larger one. . .?"). As it turns out, the two experimental groups did not differ in terms of participants' estimates of the fluency validity (both groups $M = .72$, $SE = .01$).

Experiment 3

Participants and Materials

Fifty-five right-handed persons participated in the study (66% female; mean age 25.1 years, $SD = 3.5$). Participants received a minimum payment of €8 ($11), which they could raise on the basis of their performance in the experimental tasks. We used the same set of cities as in Experiments 1, 2, and 10 as stimuli (see Online Materials B, Table OB1).

Procedure

The procedure of the computerized experiment was identical to that of Experiment 1, with the following exceptions in the inference task and the cue-knowledge task: In the inference task, in which people inferred under time pressure which of two cities is larger, each presentation of a pair of cities started with a blank screen. After 2,000 msec, a first acoustic signal (Tone 1) sounded, followed by a second signal 900 msec later (Tone 2) that coincided with the presentation of a fixation cross. Again 900 msec later, upon the sound of a third signal (Tone 3), the fixation cross was replaced by a pair of cities. Participants were instructed to respond as quickly and accurately as possible but not later than a fourth, imaginary signal (Tone 4, imagined), 900 msec after the third signal and stimulus onset. (We asked participants to imagine the fourth signal to avoid possible interference of a real signal with the ongoing processing of the city pair. A similar procedure is employed in lexical decision tasks—see Wagenmakers, Zeelenberg, Steyvers, Shiffrin, & Raaijmakers, 2004—and has been used by Pachur & Hertwig, 2006, to study how people use recognition when making inferences under time pressure.) If a response was markedly delayed, that is, 1,200 msec after the third tone, an aversive tone sounded and the message "too late" appeared on the screen. On top of the guaranteed flat fee of €8, participants were paid €0.04 ($0.05) for each correct response. Incorrect responses resulted in a subtraction of 4 cents from this additional payment. (No feedback on the correctness of the responses was given until after the experiment.) Regardless of whether the response was correct, responses that were followed by the message "too late" always resulted in a subtraction of 4 cents.

In the cue-knowledge task, we asked three of the same questions as in Experiment 1, that is, whether a city has (a) an international airport, (b) a university, and (c) significant industry. We dropped the question about whether a city is a world-famous tourist site. (In another study [Experiment 16 in Marewski, 2008], participants gave the tourism cue the lowest validity rating of the knowledge cues.)

To make sure that our participants considered fluency and the knowledge cues to be valid for inferring city size, in a questionnaire we additionally asked them to estimate the fluency validity as well as the validity of the cues from the cue-knowledge task (using a frequency format; see Experiment 2). The order of tasks (i.e., estimates for fluency vs. for cues) was randomized. Both fluency and the knowledge cues were considered to be valid for inferring city size (range of mean estimated validities was .72 to .74, $SE = .01$).

Eye-Tracking and Data-Cleaning Procedures in Experiments 1–3

For the purpose of further analyses to be reported elsewhere, we monitored participants' eye movements in Experiments 1–3 during the inference tasks, the recognition tasks, and the general knowledge tasks. To allow for comparing results between the current analyses and future ones, in our initial analyses (i.e., in Simulation 1 as well as in reanalyses of our data published elsewhere; e.g., Marewski, Gaissmaier, Schooler, et al., 2010) we used the data as they were recorded and filtered by eye-tracking software and algorithms. In addition, we excluded from the analyses those trials for which the total duration of gazes on the screen computed from stimulus onset to stimulus offset did not differ by more than 15% from the response time chunked for the trial. In Experiment 1, this resulted in an average exclusion of 2.96% of the trials per participant ($Mdn = 0.83\%$, $SE = 0.98\%$). In Experiment 2, this resulted in an average exclusion of 2.59% of the trials per participant ($Mdn = 0.00\%$, $SE = 0.96\%$) and in Experiment 3 in an average exclusion of 0.93% of the trials per participant ($Mdn = 0.00\%$, $SE = 0.36\%$). The eye-tracking system we used (Tobii 1750 system) is completely noninvasive, and the experimental setup is hardly distinguishable from a setup without an eye tracker. The frame rate of the system is 50 Hz, and its accuracy has been tested to 0.5° (Tobii Technology, 2005). Assuming an average distance of 60 cm from the screen, in the inference task cities were separated by a 19° visual angle between the midpoints of the two city

(*Appendices continue*)

letter strings (20 cm). We used the font Courier New (18 pixels per letter) for the strings. Within each string, 4° separated letters from each other. Due to a programming error, in later analyses of the data, including all simulations except Simulation 1, we did not exclude trials on the basis of the eye tracking.

Experiments 4–9

Experiments 4–7: Participants and Materials

We assigned 80 right-handed persons (55% female; mean age 24.6 years, $SD = 3.9$) at random to four experiments. In two experiments, stimuli were either the names of 168 countries (Experiment 4, $N = 20$; Table OB2 in Online Materials B; including four filler stimuli) or the names of 80 companies (Experiment 5, $N = 21$; Table OB3 in Online Materials B). Countries were essentially those about which the German Federal Statistical Office provides data on the gross domestic product (see Online Materials B for details). Companies were all 80 firms listed in the DAX and MDAX, which are important German stock indices. In Experiment 6 ($N = 20$), stimuli were the names of 54 infectious diseases that are subject to registration in Germany (Table OB4 in Online Materials B). In Experiment 7 ($N = 19$), stimuli were the 189 names of all past and current federal ministers and chancellors of the Federal Republic of Germany from 1949 to 2007 (Table OB5 in Online Materials B). Countries, companies, diseases, and politicians' names differed in word length (Tables OB2–OB5). Participants' payment differed between the experiments as a function of the length of the experiment (countries: €7 [$10]; companies: €6 [$8]; diseases: €5 [$7]; politicians: €8 [$11]).

Experiments 4–7: Procedure

Participants completed computerized recognition and general knowledge tasks that were identical to those used in Experiments 1–3, except that the stimuli were countries, companies, diseases, and politicians instead of cities. Participants took a 10-sec break after 20 to 27 trials depending on the experiment.

Experiments 8 and 9: Participants, Materials, and Procedure

One hundred eighteen and 83 participants completed only the recognition task for diseases (Experiment 8; 60% female; mean age 42.4 years, $SD = 23.1$) and politicians (Experiment 9; 59% female; mean age 34.1 years, $SD = 18.6$), respectively. With these data we also operationalized politicians' and diseases' fame, namely, as the proportion of participants who recognized a disease in Experiments 6 and 8 and a politician in Experiments 7 and 9, respectively.

Experiment 10

Participants, Materials, and Procedure

Thirty-two right-handed persons (63% female; mean age 24.2 years, $SD = 3.5$) completed a *reaction task*. In this task, we presented participants one of 240 cities (Table OB1 in Online Materials B) at a time and asked them to press a key with their right index finger immediately upon stimulus onset (i.e., without reading the city's name). All other features of this task were identical to those of the recognition task employed in Experiments 1–3.

Appendix B

How to Predict Memory Retrieval From the Internet: Derivations of the Model Equations

Activation of Objects

Anderson (1993) showed that Equation 2 can be approximated with

$$B = k + \ln n - d \ln T_{\text{creation}}, \tag{B1}$$

where d is a decay parameter that captures the degree of forgetting in declarative memory, k is a constant, n is the number of times the object represented by the chunk has been encountered, and T_{creation} is how long it has been since the chunk was first created. We assume that T_{creation} is the same for all objects such that a constant, c_0, can absorb $k + -d \ln T_{\text{creation}}$:

$$B = c_0 + \ln n. \tag{B2}$$

We assume that the base-level activation, B_W, as estimated from the web counts, N, produced by the search engine Yahoo, can be written as $B_W = c_1 + \ln N$. It is related to the base-level activation, B, as produced by a person's true past encounters with objects:

$$B = b_0 B_W + c_2 = b_0 c_1 + c_2 + b_0 \ln N, \tag{B3}$$

where the scaling parameter, b_0, and the constant c_2 capture the relation between the two measures of base-level activation. We do not model the influence of contextual information as represented by the second term in Equation 1 (i.e., the S_{ji} units of activation a chunk i receives from each of the j elements of the current context,

(*Appendices continue*)

$\sum_j S_{ji}$). Rather, we assume that it contributes to the overall variance in activation. Correspondingly, $\sum_j S_{ji}$ can be replaced with a constant c_3 in Equation 1:

$$A = B + c_3. \tag{B4}$$

Combining Equations B3 and B4 gives

$$A = b_0\, c_1 + c_2 + b_0 \ln N + c_3. \tag{B5}$$

The terms $b_0\, c_1$, c_2, and c_3 can be represented by a new constant, c_4, which gives our estimate of the activation, A, of the chunk representing an object:

$$A = c_4 + b_0 \ln N. \tag{B6}$$

In the main text's Equation 5, we refer to c_4 as c_{R} and b_0 as b_{R} to signal that they are used to model the recognition of objects. Inserting Equation B6 into Equation 3 gives the recognition probability, P_{R}, that the chunk representing an object will be retrieved and the object recognized:

$$P_R = \frac{1}{1 + \mathrm{e}^{-[(c_4 + b_0 \ln N) - \tau]/s}}. \tag{B7}$$

Activation of Knowledge

Each encounter with an object can result in a person acquiring knowledge about the object, which can be information co-occurring with the object. Each encounter can increase both the activation of the object and the activation of knowledge. We assume that the activation, A_{R}, of a chunk representing an object is related to the activation, A_{K}, of a chunk representing knowledge about the object with

$$A_{\mathrm{K}} = c_5 + b_1\, A_R, \tag{B8}$$

where the constant, c_5, and the scaling parameter, b_1, capture the relation between the two levels of activation. In addition c_5 also captures the effects of interference from competing chunks. Combining Equations B6 and B8 gives

$$A_{\mathrm{K}} = c_5 + b_1\, c_4 + b_1\, b_0 \ln N, \tag{B9}$$

where the term $c_5 + b_1\, c_4$ can be replaced with a new constant, c_6, and the term $b_1\, b_0$ with a new scaling parameter, b_2, yielding the final equation to model the activation of knowledge:

$$A_{\mathrm{K}} = c_6 + b_2 \ln N. \tag{B10}$$

In the main text's Equation 6, we refer to c_6 as c_{K} and b_2 as b_{K} to signal that they are used to model knowledge about objects. Inserting Equation B10 into Equation 3 gives the knowledge probability, P_{K}, that knowledge about an object will be retrieved:

$$P_{\mathrm{K}} = \frac{1}{1 + \mathrm{e}^{-[(c_6 + b_2 \ln N) - \tau]/s}}. \tag{B11}$$

Adding Noise to the Retrieval Criterion

ACT–R assumes that retrieval variability depends on stochastic variability in the activation level of a chunk. However, in principle, this variability could instead be attributed to noise in the retrieval criterion, τ, as in the original *rational analysis of memory* (Anderson, 1990; Anderson & Schooler, 1991). Moreover, there could even be variability in activation and the retrieval criterion. Then, the *total retrieval noise*, s, in Equations 3, 5, and 6 can be divided into *criterion noise*, s_τ, attributable to the retrieval criterion, and *activation noise*, s_{A}, attributable to the activation. We define

$$s = \sqrt{(s_\tau^2 + s_{\mathrm{A}}^2)}. \tag{B12}$$

As Figure 4 illustrates, Equation B12 helps us model recognition time distributions that are characterized by an increase in the mean, median, and spread as a function of activation (cf. Ratcliff, 2002; Wagenmakers & Brown, 2007). The shape of the distributions predicted by our model also results from computing retrieval times for each activation level conditional on a chunk being retrieved. In Figure 4, we imposed this condition on the retrieval times by estimating the portion of each object's activation distribution that falls above the retrieval criterion, and we computed a retrieval time distribution from it (i.e., using Equation 7). This is the retrieval time distribution given that the object is retrieved (cf. Equation 4). Recall that we assume there is a probability distribution of retrieval criteria, rather than just one criterion. Therefore, we split the total retrieval noise into criterion noise and activation noise, which determines the shape of the retrieval criterion and activation distributions, respectively. We then computed retrieval time distributions across the retrieval criterion distribution, generating the corresponding recognition time predictions using Equation 7.

(*Appendices continue*)

Appendix C

Additional Simulations

Appendix C consists of three additional simulations we ran to test our model. As for all other simulations, details of the simulation procedures are reported in Online Materials D.

Robustness of Model Predictions Across Proxies for Effort: Simulations C1 and C2

In Simulation 3 (see Figure 9), we used the detection probability, P_D, as a proxy for how effortful it would be for a person to apply the fluency heuristic. The detection probability is derived from the timing model. To make sure that our conclusions from Simulation 3 do not depend on the proxy for effort chosen (i.e., P_D), in Simulation C1 we modified the design of Simulation 3. We used the memory model alone, that is, without the timing model, to predict how the fluency validity changes as a function of the raw recognition time differences. In doing so, we assumed that the fluency heuristic would be applicable when two objects differ in recognition times. The intuition is that it should be easier to apply the heuristic the larger the recognition time differences are.

Besides the detection probability P_D, another measure of how easy it would be for a person to apply the fluency heuristic is the difference in pulses the person would notice when comparing the objects. The intuition is that it is easier to tell two objects' recognition times apart the larger their differences in pulses. Although pulse differences are likely to correlate with detection probabilities, these proxies are not identical. To further test whether our results depend on the proxy chosen, in Simulation C2 we modified the design of Simulation 3, computing the fluency validity as a function of pulse differences.

Results

Larger differences in raw recognition times and pulses are both associated with a larger fluency validity (see Figures C1 and C2), demonstrating that our results hold independently of the proxy chosen.

Recognition Times as a Function of Knowledge: Simulation C3

We let the memory model predict recognition times for tartle and knowledge objects.

Observed Data

We used participants' responses in the recognition and general knowledge tasks of Experiments 2–3 to identify tartle and knowledge objects. For each participant, we calculated the median recognition times for these objects, averaging the medians across participants.

Data Predicted by the Memory Model

We used our memory model to predict hypothetical persons' recognition, knowledge, and recognition times (P_R, P_K, $T_{recognition}$), computing medians of predicted recognition times separately for predicted tartle and knowledge objects. We averaged the medians across simulation runs.

Results

Figure C3 shows the results for the cities of Experiments 2 and 3. As can be seen, the model predicts larger recognition times for tartle objects than for knowledge objects.

(*Appendices continue*)

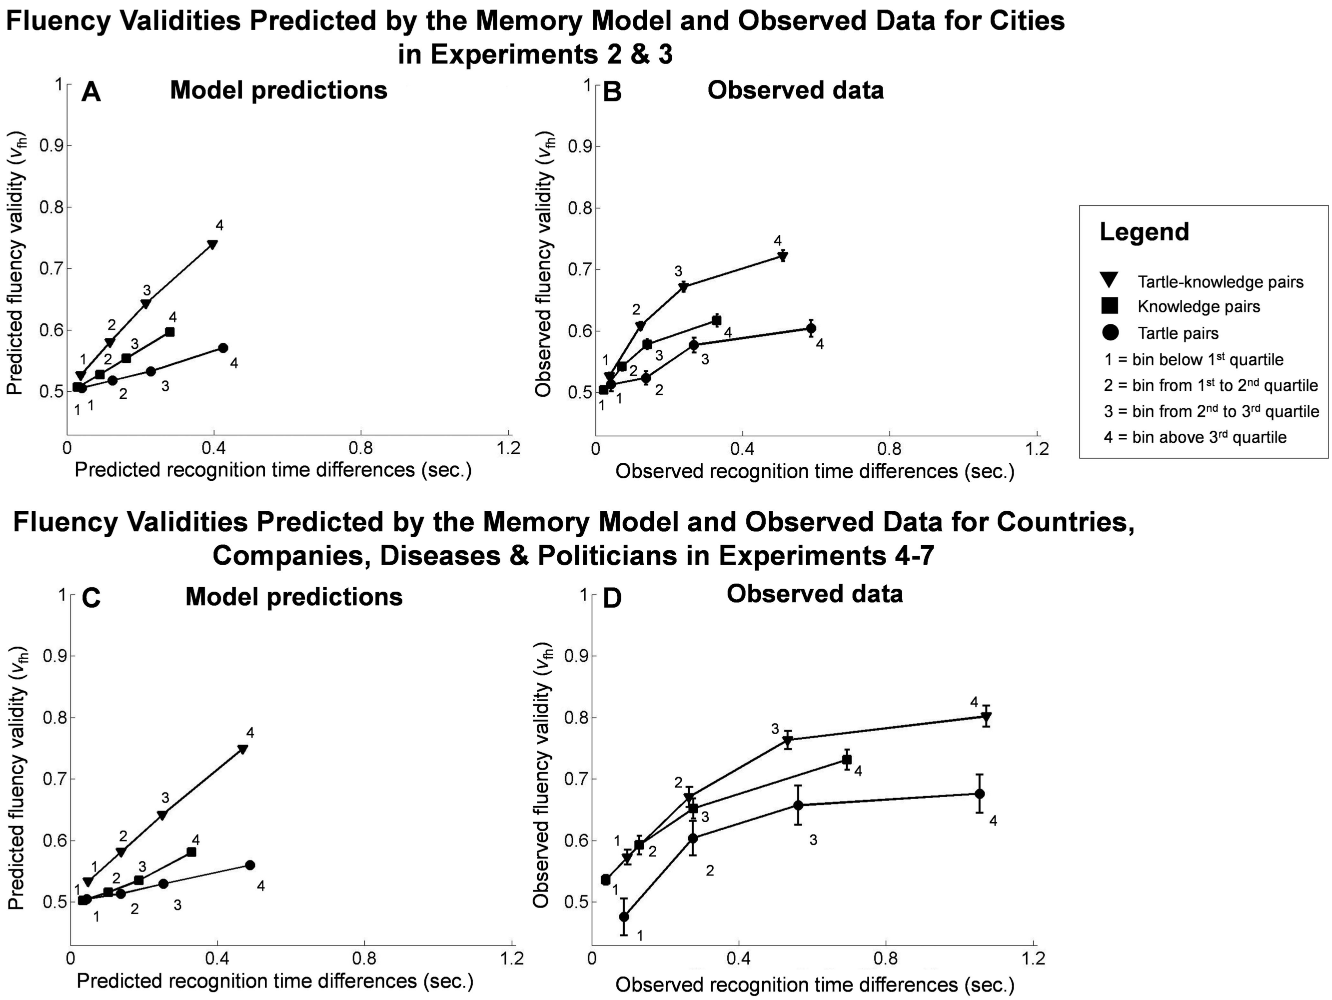

Figure C1. Simulation C1. Fluency validity (v_{fh}) as a function of the recognition time difference (in seconds) between two objects. Predicted data were generated by the memory model (Equations 5–8, 11; e.g., P_R, P_K, $T_{recognition}$, v_{fh}). (A) Predicted and (B) observed data for inferences about cities' size. (C) Predicted and (D) observed data for inferences about countries' gross domestic product in 2006; companies' market capitalization on May 31, 2007; diseases' fame; and politicians' fame. We operationalized fame as the proportion of participants who recognized a disease in Experiments 6 and 8 and a politician in Experiments 7 and 9, respectively. In Panels C and D, the data are collapsed across the four types of objects. For predicted data, symbols show mean (±1 *SE*) predicted validities and mean predicted median time differences computed across simulation runs of the memory model. For observed data, symbols show mean (±1 *SE*) observed validities and mean observed median time differences computed across participants. Some of the error bars are obscured by the symbols.

(*Appendices continue*)

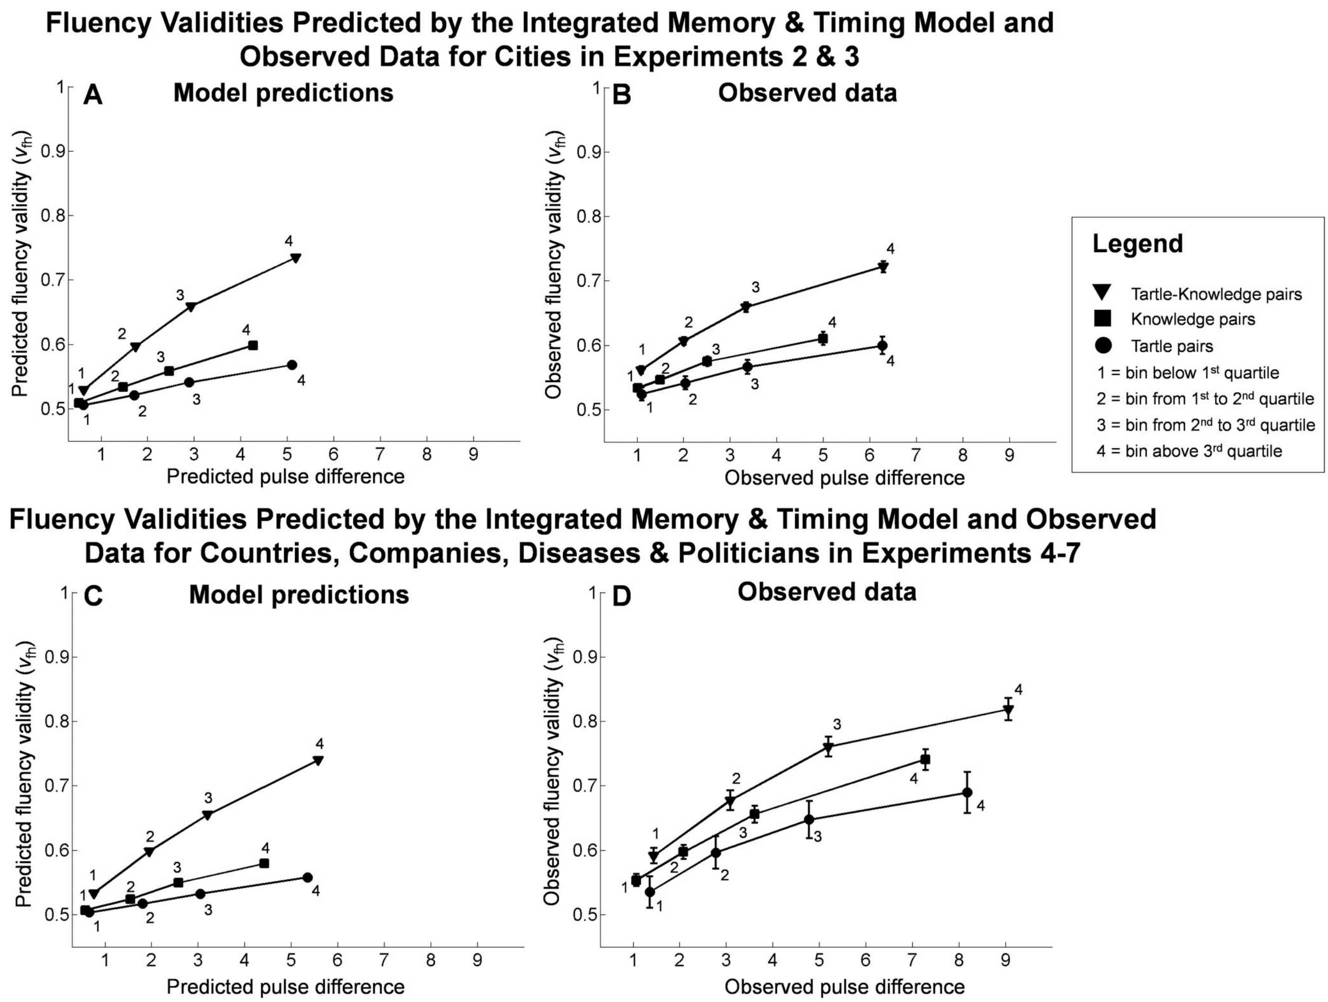

Figure C2. Simulation C2. Fluency validity (v_{fh}) as a function of the pulse difference between two objects. Predicted data were generated by the integrated memory and timing model (Equations 5–9, 11; e.g., P_R, P_K, $T_{recognition}$, v_{fh}). The timing model was also applied to the observed data in Experiments 2–7. (A) Predicted and (B) observed data for inferences about cities' size. (C) Predicted and (D) observed data for inferences about countries' gross domestic product in 2006; companies' market capitalization on May 31, 2007; diseases' fame; and politicians' fame. In Panels C and D, the data are collapsed across the four types of objects. For predicted data, symbols show mean (± 1 *SE*) predicted validities and mean predicted pulse differences computed across simulation runs of the memory and timing model. For observed data, symbols show mean (± 1 *SE*) observed validities and mean observed pulse differences computed across participants and simulation runs of the timing model. Some of the error bars are obscured by the symbols.

(*Appendices continue*)

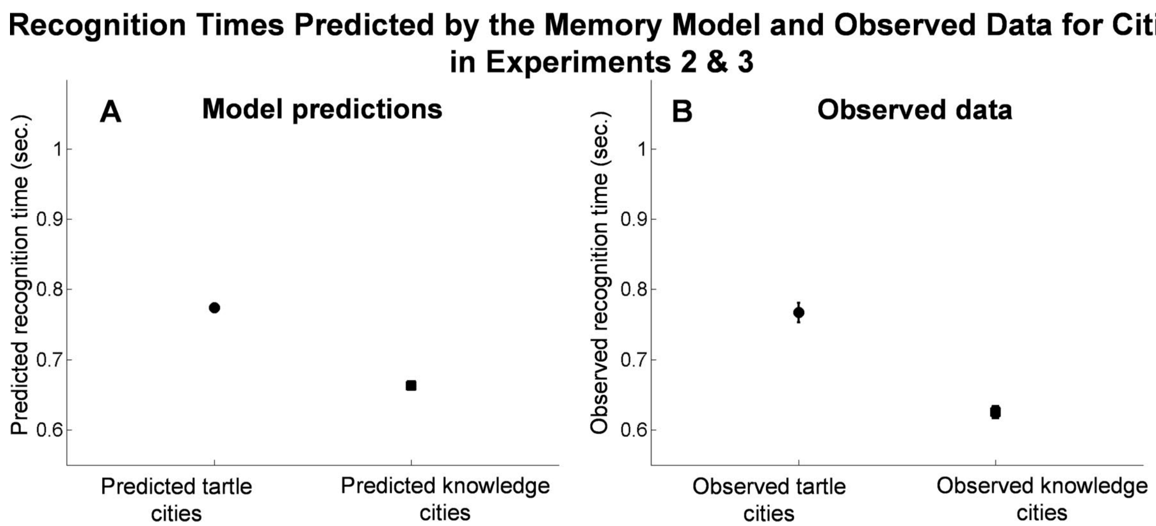

Figure C3. Simulation C3. Recognition times (in seconds) for tartle cities and knowledge cities, (A) predicted by the memory model (Equations 5–8; e.g., P_R, P_K, $T_{recognition}$) and (B) observed in Experiments 2 and 3. The observed data are aggregated across these experiments. (A) Means (±1 *SE*) of predicted median times computed across simulation runs of the memory model; (B) means (±1 *SE*) of observed median times computed across participants. Some of the error bars are obscured by the symbols.

Received July 20, 2008
Revision received March 7, 2011
Accepted March 9, 2011 ■