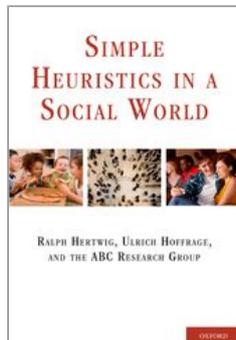


The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

University Press Scholarship Online

Oxford Scholarship Online



Simple Heuristics in a Social World

Ralph Hertwig, Ulrich Hoffrage, and ABC Research Group

Print publication date: 2012

Print ISBN-13: 9780195388435

Published to Oxford Scholarship Online: January 2013

DOI: 10.1093/acprof:oso/9780195388435.001.0001

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

Thorsten Pachur

Ralph Hertwig

Jörg Rieskamp

DOI:10.1093/acprof:oso/9780195388435.003.0009

Abstract and Keywords

Inferring latent event frequencies in the environment is a key cognitive function. Models of frequency judgments commonly assume that for such inferences, people rely on instance knowledge (e.g., people in a person's social network) processed in a compensatory fashion. The chapter examines the possible contribution and accuracy of noncompensatory processing of instance knowledge. For that purpose, the chapter extends the notion of ordered and limited search—which has received much attention in cue-based inference—to instance-based inference and propose the boundedly rational social-circle heuristic as one possible model. Unlike the common assumption of compensatory processing in models of instance-based inference, the social-circle heuristic is noncompensatory: it searches the social circles of a person's

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

network for relevant instances sequentially and stops search as soon as a circle discriminates. In computer simulations, the chapter shows that despite its frugality, the social-circle heuristic competes favorably with more complex strategies, especially in environments with a skewed frequency distribution. In two empirical studies, the chapter shows that for predicting people's inferences concerning the relative frequency of real-world events, the heuristic provides a viable alternative to the usual assumption of compensatory instance processing. Finally, the chapter discusses how noncompensatory processing of instance knowledge might account for established regularities in social influence and how the social-circle heuristic relates to norm formation.

Keywords: heuristics, sampling, social network, memory, social circles, norms, attitude, frequency

The probabilistic expectancies we form about others from very limited information are more accurate than we would expect.

Ambady and Rosenthal (1992)

The American presidential race in 1936 was a referendum on President Franklin D. Roosevelt's New Deal economic policies. Roosevelt, running for his second term, would win what was up until then the greatest electoral landslide since the beginning of the United States' two-party system in the 1850s, winning all but eight of the 531 electoral votes, and winning 27,747,636 votes to Alfred Landon's—his Republican rival—16,679,543 (*New York Times*, "Alf Landon," 1987). The presidential race between Roosevelt and Landon was also the beginning of the modern era of political polling—which started with a veritable fiasco. The *Literary Digest*, a venerable magazine and, perhaps, the most prestigious pollster of public opinions in those days, predicted two months before the election: "Governor Alfred M. Landon of Kansas, Republican nominee for President, has better than a two-to-one lead over his Democratic opponent, President Roosevelt, in the first tabulated returns from the *Literary Digest's* 10,000,000 ballot public sentiment test poll on Presidential candidates" ("Landon Is Leading," 1936). The magazine's

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

egregiously botched prediction was based on a “straw poll”: Although it mailed postcards to 10 million people, and received a whopping 2.2 million in return, it had surveyed its own readers, as well as people from two readily available lists of registered automobile owners and telephone users. In 1936, however, millions of Roosevelt supporters were too poor to be able to afford a magazine subscription or a phone, let alone an automobile. The *Literary Digest's* sample was therefore highly unrepresentative.

(p.262) “In the waning days of the 1936 presidential election, a young man from Princeton, New Jersey, ... was becoming increasingly distressed.... He suffered from insomnia, he sucked on his unlit cigarettes, he worried incessantly that he had done something wrong and that his reputation and financial solvency were about to be destroyed” (Moore, 1995, p. 31). The nervous young man was the director of the newly founded *American Institute of Public Opinion*, George Gallup. Gallup understood that the famed *Digest* poll was heading for a disastrous cropper. Gallup brashly announced that the *Digest* would be wrong in the 1936 election. Using the replies of at most 5,000 representatively sampled respondents, 0.23% of the size of the *Digest's* poll, he predicted Roosevelt's landslide victory. The *Literary Digest* went out of business; the “Gallup poll” soon became a household word.

Covering a seemingly boundless variety of topics, contemporary pollsters from Gallup, the Pew Center, and Harris Interactive sample people's opinions, beliefs, preferences, customs, and morals. Probing a small but representative section of the population, they aim to infer what most of us think and feel. The thesis of this chapter is that canvassing samples of people to infer and to predict characteristics of the social world at large is not the prerogative of the pollsters. Like Gallup's interviewers, the human mind can also roam through its personal social spaces to sample instances and garner information, enabling it to make inferences about the social world. There are, however, important differences between professional and the mind's intuitive polling: The mind's samples—drawn from the external social world or from memory—are minute, relative to the

thousands of respondents in Gallup polls. Furthermore, the mind's sample is an unrepresentative one, because it is drawn from a person's social environment, and people tend to know others who are more similar to themselves than to a randomly drawn person (e.g., McPherson, Smith-Lovin, & Cook, 2001). How do people make inferences about the behavior and characteristics of others based on instances sampled from their social environment? And how accurate are the strategies that people might use—given the limitations of the samples available to an individual mind mentioned above?

In what follows, we will describe the rationale and the social rationality of instance-based inference. Moreover, we will propose the *social-circle heuristic* (Pachur, Rieskamp, & Hertwig, 2012), which people could use for making predictions about their social world based on limited search. We will examine the prescriptive and descriptive accuracy of the heuristic, its ecological rationality and boundary conditions, and will explore different domains (such as norm and attitude formation) in which the heuristic might guide our social reasoning.

(p.263) Instance-Based Inference: Polling My Space

Traditionally, public opinion—for instance, prior to elections—is measured by probing voters' intentions with questions such as this: "Assuming that Election Day were to take place this coming Sunday, which party would you vote for?" Although representative polls based on such questions are generally useful, pollsters nevertheless frequently find that their survey results diverge from people's actual behavior in the voting booth ("Will there be an 'Obama effect?'" ; Elder, 2007). But if explicit voting intentions do not predict people's behavior very well, what does? Analyses by Noelle-Neumann offer a surprising answer (e.g., Noelle-Neumann, 1991; Noelle-Neumann & Petersen, 2004). She found that measures that assess the "climate of opinion"—that is, respondents' beliefs about how *others* will vote—are a better predictor of the respondents' behavior than their intentions ("Regardless of your personal opinion, do you think most people ... are holding

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

a favorable opinion of ... party, or don't you think so?"; Noelle-Neumann, 1977, p. 157). Noelle-Neumann's observation foreshadowed the success of prediction markets, with markets not taking a reading of whom people intend to vote for but of who they think will win, and cash wagered indicating the strength of those beliefs (e.g., Wolfers & Zitzewitz, 2004).

The results of Noelle-Neumann suggest that our behavior is strongly influenced by our beliefs and judgments about others. How do people judge the behavior and characteristics of others? One possibility is that, similar to polling institutes, individuals draw a sample from their own social environment and make predictions based on their personal polling results. Such intuitive polling might serve social rationality in several respects. First, probing one's social world is one way of forming a representation of the world's social texture. Second, based on these representations, a person can predict others' beliefs, preferences, and attitudes, and, ultimately, their behavior. Third, judgments about other people's beliefs and behaviors allow individuals to coordinate their own behavior with the behavior of others as well as with social expectations, standards, and norms.

Inferences about social environments can be drawn using either social or nonsocial information. Take the example of a person who attempts to predict which of two soft drinks is more popular. Nonsocial information would be cues such as the prevalence of media ads promoting one or the other, or whether a new variation of the product (e.g., low sugar, new flavor) was introduced recently. To the extent that these indicators are correlated with actual product popularity, they can act as probabilistic cues to infer the relative popularity of the two drinks. There is a large variety of strategies, such as *take-the-best*, *tallying*, or *weighted additive*, that can process nonsocial cues for making inferences (Gigerenzer, Todd, & the ABC Research Group, (p.264) 1999). Alternatively, strategies based on the retrieval of instances assume that a person probes her social environment represented in memory and gauges the relative popularity of Coca-Cola® and Pepsi® from the recallable instances of Coca-Cola and Pepsi aficionados. Of course, due to sampling error, and because this sample is

unlikely to be representative, its characteristics will be imperfect predictors of the characteristics of the population; nevertheless, akin to nonsocial cues, the prevalence of behaviors that people observe in their own samples are correlated with the prevalence of the behaviors in the population. The magnitude of this correlation is one important factor for the success of the mind's polling strategy.

But how do we poll our social world? One possible strategy would be to try to retrieve as many instances as possible from our social environment and aggregate the information in this sample. For example, Hertwig, Pachur, and Kurzenhäuser (2005) obtained evidence that people make judgments about the relative frequency of health risks based on the total number of relevant instances they can recall among their family, friends, and acquaintances. However, people may not always strive for an exhaustive retrieval and instead limit their information search using a simplifying heuristic. Next, we propose such a simple heuristic, the *social-circle heuristic*, which models how people might judge the prevalence of others' beliefs, preferences, characteristics, or behavior in the population.

The Social-Circle Heuristic

The social-circle heuristic can be used to make inferences about our social environment; it can answer questions such as which of two movies, or which of two soft drinks (say, Coca-Cola or Pepsi) is more popular overall, which of two political parties will receive more votes, or which of two behaviors is more frequent (and thus apt) in a given situation. Unlike heuristics such as tit-for-tat, the social-circle heuristic is not a heuristic that navigates a social interaction in which the outcome of a decision also depends on the decision(s) of (an)other person(s). There is no interdependency between the decision or judgment of several persons using the social-circle heuristic. Therefore the social-circle heuristic is a tool for “games against nature” (chapter 1), in which a person needs to make inferences about (social) nature (e.g., the prevalence of behaviors) in order to achieve his ends.

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

Like other heuristics that have been proposed as models of bounded rationality, the social-circle heuristic embodies a sequentially ordered and limited information search. Rather than obeying the classic prescription to include ever more people in the sample as long as the perceived marginal benefits of acquiring additional information exceed the perceived marginal costs (Stigler, 1961), the (p.265) heuristic's search in memory is limited. Importantly, the search follows a systematic order that rests on the well-documented fact that social networks tend to display a hierarchical structure consisting of discrete subgroups (e.g., Hill & Dunbar, 2003; Milardo, 1992; Zhou, Sornette, Hill, & Dunbar, 2005; see also chapter 7). Specifically, the social-circle heuristic assumes that the structure of a person's social network (i.e., the external social world) provides the blueprint for the internal search process for relevant carriers of information—consistent with the increasing evidence of a link between search processes in the external world and in the internal world of memory (e.g., Hills, Todd, & Goldstone, 2008; Pirolli & Card, 1999). Moreover, the heuristic's exploitation of social structures is consistent with Hills and Pachur's (2012) finding that search in social memory is guided by links in a person's social network.

As illustrated in Figure 9-1, it is assumed that an inference about the relative prevalence of two event categories is based on the recognition heuristic (Goldstein & Gigerenzer, 2002) if the name of only one of the categories is recognized. If both are unrecognized, the inference will be made by guessing. The social-circle heuristic applies if the names of both categories are recognized. How does the heuristic's search through the social space proceed? Figure 9-1 represents subgroups in the social network as circles, and the heuristic is assumed to sequentially probe them for critical instances. Based on the well-established phenomenon that people often rely on information about themselves when making inferences about the prevalence of behaviors and characteristics in the population (the so-called *false-consensus effect*; see, e.g., Krueger & Clement, 1994; Ross, Greene, & House, 1977), the social-circle heuristic starts by considering information about

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

the “self.” The self thus represents the starting point for the sampling process (circle 1).

What are the next circles? There are at least two dimensions underlying the structure of a person's social network that could guide further search: altruism and frequency of contact. Altruism typically manifests in kin (Hamilton, 1964a, 1964b) and in reciprocal relationships (Singer, 1981). Adapting altruism as the structuring dimension yields three circles beyond the self-circle: *family* (circle 2); *friends* (circle 3); and *acquaintances* (circle 4), with family relating to kin relationships, and friends and acquaintances to non-kin reciprocal relationships. Alternatively, a person's social network can be structured according to frequency of contact. This structure acknowledges that the people in one's close social proximity are not necessarily family members but may be, for instance, friends or colleagues whom one sees daily. Analyzing the frequency of contact with one's social network members on a daily basis, one finds robust regularities in social environments (chapter 7). When one partitions the social network of a person (i.e., the people a (p.266)

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

person knows) according to frequency of contact, the three circles following the first (“self”) are the people with whom one has contact *at least once a week* (circle 2); *about once a month* (circle 3); and *not more than once in six months* (circle 4).

The social-circle heuristic thus conceptualizes an individual's social network in terms of circles with increasing scope, with smaller circles being more “socially proximal” to the center of the network: the self (see Kahn & Antonucci, 1980; Moreno, 1936).¹ Based on these two alternative structures of the social environment (i.e., altruism and frequency of contact), we distinguish two variants of the social-circle heuristic. Both can be expressed in terms of the following building blocks:

Search rule: Search the social circles for occurrences of the target events (e.g., people who drink Pepsi and Coca-Cola), proceeding sequentially through the circles, starting with the first circle (“self”).

(p.267) *Stopping rule:* If the result of the search within a circle supports one hypothesis concerning the events’ population frequencies (e.g., more people like Pepsi than Coca-Cola), terminate search. If the same number of occurrences (or none) is found for both event categories, proceed to the next circle.

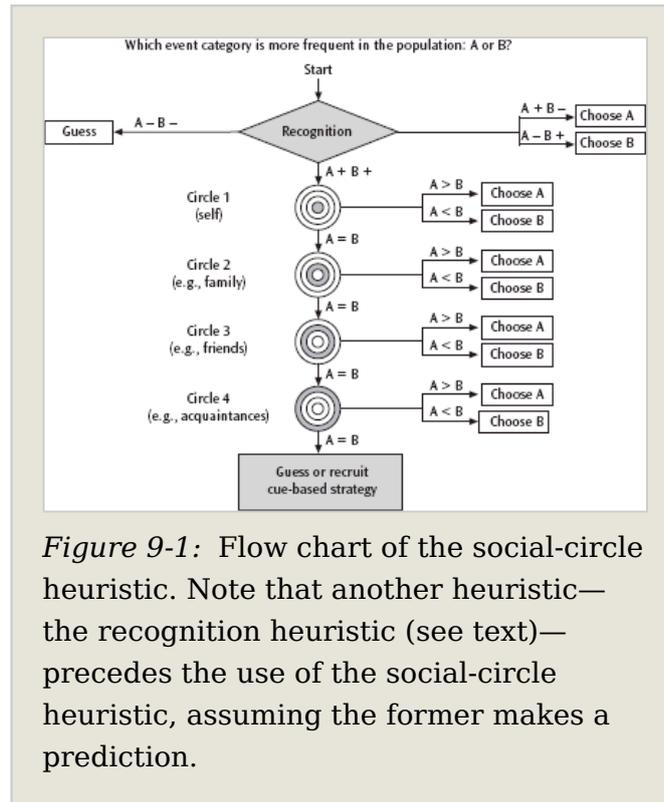


Figure 9-1: Flow chart of the social-circle heuristic. Note that another heuristic—the recognition heuristic (see text)—precedes the use of the social-circle heuristic, assuming the former makes a prediction.

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

Decision rule: Infer that the event category for which more instances are found is the more prevalent one in the population. If the sampled information does not discriminate between the event categories (and no other information is known), either guess or search for semantic cues after the final circle has been searched.

For illustration, consider the Coke versus Pepsi problem. The social-circle heuristic assumes that a person first polls herself (Figure 9-1). If the “self” (i.e., one's own preference concerning Pepsi and Coke) discriminates between the soft drinks (e.g., the “self” fancies Coke), search stops, and the inference is made that more people like Coke than Pepsi. If the self does not discriminate (either because one likes both or neither), search will proceed, and the next circle is examined. Based on the altruism structure, the heuristic will next examine which of the two soft drinks is the preferred brand among one's family members (if they consume any). If this circle does not discriminate between the soft drinks' popularity, the circle of “friends” will be polled, and so on. If the retrieved instances in none of the circles favor one soft drink over the other, the decision maker is assumed to either recruit cue-based strategies or resort to guessing.

The two variants of the social-circle heuristic differ in how search moves through a one's social network. According to the first variant (the altruism structure), search proceeds from the self to people to whom we are genetically related, and then to people outside the family with whom we have reciprocal relationships (cf. Hills & Pachur, 2012). The rationale of this order is that we are likely to have the most reliable, extensive, and easily retrievable knowledge about ourselves, our family members, and close friends with whom we cooperate (Henrich & Henrich, p. 2007, p. 58). According to the second variant (the frequency of contact structure), in contrast, search proceeds from the self to people with whom we have frequent contacts (e.g., Hill & Dunbar, 2003). The rationale of this search order is that frequency of contact has been demonstrated to be a key determinant of the retrieval probability of memory records (Anderson et al., 2004; Anderson & Lebiere, 1998). Henceforth, we refer to the variant of the heuristic exploiting the altruism structure as the

social-circle heuristic, and to the variant using the frequency of contact structure as the *social-circle heuristic_F*.

(p.268) The decision rule of the social-circle heuristic is noncompensatory:² When the search result within one of the inner circles favors one hypothesis, a decision is made. Unlike a compensatory strategy, the heuristic does not look up and does not combine instances across all circles (limited search). As a consequence, instances in more distal circles cannot reverse the judgment, thus avoiding the possibility of cognitive conflict (cf. Hogarth, 1987). As the circles that the heuristic searches are increasingly larger, the heuristic increases its sample size successively until an inference can be made. Thus, the social-circle heuristic uses cognitive resources economically, extending the sample space (and thus investing retrieval efforts) only when necessary (for an investigation of the neural correlates of retrieval effort in the cue-based take-the-best heuristic, see Khader, Pachur, Meier, Bien, Jost, & Rösler, 2011).

The Social-Circle Heuristic: How Frugal and Accurate Is It?

Is the heuristic able to make accurate judgments by using only limited information and based on noncompensatory information processing? To answer this, we evaluated the performance of the social-circle heuristic in a competition that pitted it against two other inference models, both of which also process instances of the target categories (rather than cues) but often use considerably more information than the social-circle heuristic. The competition was implemented in a computer simulation. The task that the models had to solve was to infer which of two event categories has more instances in the population. If no inference could be made based on the number of instances sampled, a random guess was implemented.

The Competitors

Perhaps the most prominent descriptive inference strategy that makes judgments based on recalled instances is the *availability heuristic* (Tversky & Kahneman, 1973). It has been invoked to account for a wide range of frequency judgments

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

concerning social events, such as mortality rates (Lichtenstein, Slovic, Fischhoff, Layman, & Combs, 1978), frequencies of names (Oppenheimer, 2004) or behaviors (Schwarz et al., 1991), and number of sexual partners (Brown & Sinclair, 1999). Tversky and Kahneman's exposition of the heuristic is consistent with two different mechanisms: one that rests on frequency (p.269) (i.e., the number of actually recalled instances) and the other on the ease of recall (see Sedlmeier, Hertwig, & Gigerenzer, 1998). We focus on the first mechanism, and we make use of an implementation that Hertwig et al. (2005) called *availability-by-recall*. Like the social-circle heuristic, *availability-by-recall* assumes that search is limited insofar as it unfolds within the bounds of a person's social network, which Hertwig et al. (2005) defined as consisting of the focal person and his or her family, friends, and acquaintances. Within these bounds, *availability-by-recall* assumes that all recallable instances are considered and no specific order is assumed for their retrieval. As a consequence, *availability-by-recall* is compensatory insofar as, for instance, the frequency of instances in the family circle can be compensated by the frequency of instances in the friends circle. *Availability-by-recall* predicts that the category with the higher number of occurrences in one's social network is more frequent in the target population. In Hertwig et al.'s investigation (2005; Study 2) of how and how well people judge the mortality and incidence rates of various diseases, almost half of respondents were classified as users of *availability-by-recall* (see also Pachur, Hertwig, & Steinmann, 2012).

The third inference strategy in our competition is a sampling model instantiating Wald's sequential analysis (see Wald, 1947). To infer whether event category *A* or *B* occurs more frequently in the population, *Wald's rule* samples "social agents" sequentially, examines whether they belong to one of the event categories, and terminates information search as soon as sufficient support for a particular hypothesis has been accumulated. Specifically, the model considers three hypotheses: (a) that instances of event category *A* occur more frequently than instances of event category *B* (H_{1A}); (b) that *B* occurs more frequently than *A* (H_{1B}); and (c) that both event

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

categories have the same number of instances (H_0). To evaluate these hypotheses, Wald's rule samples pairs of agents sequentially and after each pair compares the accumulated evidence relative to two thresholds representing the two directed hypotheses (i.e., H_{1A} and H_{1B}). As soon as the evidence reaches one of the thresholds, sampling is terminated and the respective hypothesis accepted with the desired level of confidence. In our simulation, the probability α of falsely accepting hypotheses H_{1A} or H_{1B} was set at 0.05, and the probability β of falsely accepting the null hypothesis (i.e., H_0) was set at 0.10.³ In addition to these error probabilities, Wald's rule requires the specification of two (p.270) environmental parameters: first, the expected relative frequency of instances of an event category, π , in the population. We set π equal to the median relative frequency of all event categories in the respective environment (see below). The second was the effect size (Δ) by which the proportions of instances of two event categories A and B differ. We considered two separate levels of Δ , one for the effect size required for H_{1A} to be accepted, Δ_A ; and one for H_{1B} to be accepted, Δ_B . We set Δ_A as the median difference in relative frequency (in the respective environment) among the pairs of event categories that are *more* frequent than π ; Δ_B was set as the median difference among the pairs of event categories that are *less* frequent than π . Further details concerning our implementation of Wald's rule can be found in Box 9-1.

In sum, the purpose of the simulation was to examine how well the social-circle heuristic, using only a subset of the available information, performs relative to availability-by-recall and Wald's rule. Availability-by-recall uses all (recallable) information in a person's social network. In contrast, Wald's rule can also sample instances beyond the bounds of a person's social network, thus representing primarily a statistical benchmark.

The Environments

We tested the strategies in two different environments. Each environment consisted of ten different event categories (e.g., ten infectious diseases), which represented the reference class (see Gigerenzer, Hoffrage, & Kleinbölting, 1991). The event

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

categories differed in their frequency of occurrence. These frequencies were implemented across a population of agents. Each agent was an instance of at most one event category. At the same time, agents made inferences regarding which event category was more frequent in the population of agents in a total of 45 pair comparisons (i.e., the complete pair comparison of the ten event categories). The two environments differed in terms of the frequency distributions across the different event categories. In the first environment, the frequency distribution tracks the distribution of occurrences of the ten most common notifiable infectious diseases in Germany (see Hertwig et al., 2005). As Figure 9-2 shows, this distribution is J-shaped: Only a few event categories are very frequent, whereas most event categories occur infrequently. J-shaped distributions characterize myriad real-world environments (e.g., Hertwig, Hoffrage, & Sparr, 2012; Newman, 2005).⁴ In the uniform (p.271) environment, which we created artificially, the frequency distribution is considerably flatter, with frequency decreasing linearly from the most to the least frequent event category. Each environment was implemented as a toroidal grid (i.e., a matrix where each cell at the border is connected to the cell on the same row or column, respectively, on the opposite side of the matrix), in which each cell represented a social agent that could be an instance of an event category (e.g., a disease; an agent could not be an instance of more than one event category). In both environments, the instances of the event categories were distributed randomly across the agents in the population. The (p.272)

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

density of instances of one of the event categories was relatively high (with 2,394 and 2,350 occurrences among a total of 2,500 agents in the J-shaped and the uniform environment, respectively). The issue of whether the density of instances affects the relative performance of the two models is addressed further below.

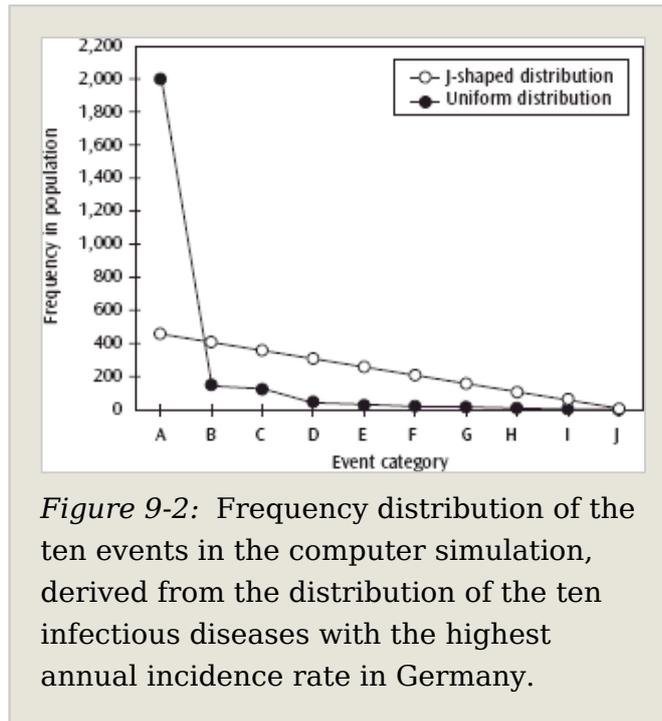


Figure 9-2: Frequency distribution of the ten events in the computer simulation, derived from the distribution of the ten infectious diseases with the highest annual incidence rate in Germany.

Box 9-1: Description of Wald's Rule

Wald's rule infers for each pair comparison whether event category A is more frequent (H_{1A}), category B is more frequent (H_{1B}), or whether both event categories are equally frequent (H_0). At each step during the sampling process, a pair of agents is drawn at random (without replacement) from the population, and the "category value" of each agent in the pair is examined. If the pair discriminates between the event categories—that is, one of the agents (but not the other) is an instance of one of the event categories—the variable n (the number of discriminating cases) will be increased by 1; otherwise another pair of agents is sampled. The variable y (representing the accumulated evidence for the hypotheses) will be increased by 1 if one of the agents (but not the other) is an instance of category A and decreased by one if one of the agents (but not the other) is an

instance of category B. The current value of y is compared to four critical values, y_1 , y_2 , y_3 , and y_4 . These critical values at n are calculated as $y_1(n) = b \times n + a_1$, $y_2(n) = b \times n - a_0$, $y_3(n) = -b \times n - a_0$ and $y_4(n) = -b \times n - a_1$, respectively,

$$\text{with } a_1 = \frac{2 \times \ln \frac{1 - \beta}{\alpha}}{\ln \frac{\pi_{1A}}{1 - \pi_{1A}}}, a_0 = \frac{2 \times \ln \frac{1 - \alpha}{\beta}}{\ln \frac{\pi_{1A}}{1 - \pi_{1A}}}, \text{ and } b = \frac{\ln \frac{1}{4 \times \pi_{1A} \times (1 - \pi_{1A})}}{\ln \frac{\pi_{1A}}{1 - \pi_{1A}}}.$$

Furthermore, π_{1A} is the probability that a pair of agents (sampled randomly) discriminates between the two event categories under H_{1A} and is calculated as $\pi_{1A} = (\pi + \Delta_A) \times (1 - \pi)$. The corresponding probability π_{1B} is calculated as $\pi_{1B} = (\pi - \Delta_B) \times (1 - \pi)$. As $(\pi_{1A} + \pi_{1B}) \neq 1$ in our simulation, π_{1A} and π_{1B} were corrected by subtracting half of the surplus (i.e., $(\pi_{1A} + \pi_{1B})/2$) from each (see Bortz, Lienert, & Böhnke, 2008).

Depending on how the current value of y compares to the four thresholds at n , information search is either continued or stopped. Specifically, if $y(n) > y_1(n)$, H_{1A} will be accepted; if $y(n) < y_4(n)$, H_{1B} will be accepted; and if $y_2(n) > y(n) > y_3(n)$, H_0 will be accepted. Otherwise, information search is continued and the above steps are repeated until one of the conditions is met. If none of the conditions is met after all 2,500 agents in the population have been sampled, then H_0 will be accepted.

As Figure 9-3 shows, in the simulations the social network of each agent was divided into four social circles. The circles were defined in terms of the distance, d , to the focal circle (circle 1), which was made up of the agent itself ($d = 0$). Circle 2 included all four agents with $d = 1$; circle 3 included all eight agents with $d = 2$; and circle 4 included all 28 agents with $d = 3$ or $d = 4$. Altogether, an agent's social network consisted of 40 other agents, and thus the agent could sample up to a maximum of 41 agents (including itself). Consistent with analyses of empirical social networks, the size of the circles grows with increasing distance from the network center (Zhou et al., 2005). Note that in this simulation, the circle structure could represent either the altruism or the contact frequency structures.

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

The social-circle heuristic searches through circles 1 through 4 and terminates search when one circle discriminates. Availability-by-recall, in contrast, searches through all circles of an agent's social network. Wald's rule takes random samples from the entire network until one of the thresholds is reached (or the total environment is sampled), irrespective of the bounds of an agent's social network. In (p.273)

each environment (i.e., J-shaped and uniform) and for each strategy, the simulation was based on 1,000 runs. At each run a randomly picked agent made a total of 45 inferences regarding which of two event categories is more frequent.

The strategies' accuracies were defined as the percentage of correct inferences.

Frugality and Accuracy

How extensive is the poll that the social-circle heuristic takes? In the simulation, the heuristic looks up, on average, 24.7 and 12.3 agents before drawing an inference in the J-shaped and uniform environments, respectively. In both environments, availability-by-recall looks up all 41 agents in an agent's network. On the assumption that rather than looking up all agents in the accessed sample space (non-directed retrieval), only those are sampled that represent an instance of the event categories in question (directed retrieval), the social-circle heuristic looks up, on average, 1.4 and 1.8 agents, compared with 12.2 and 8.5 agents for the availability-by-recall heuristic in the J-shaped and uniform environments, respectively. Put differently, the social-circle heuristic considers only 11.3% (J-

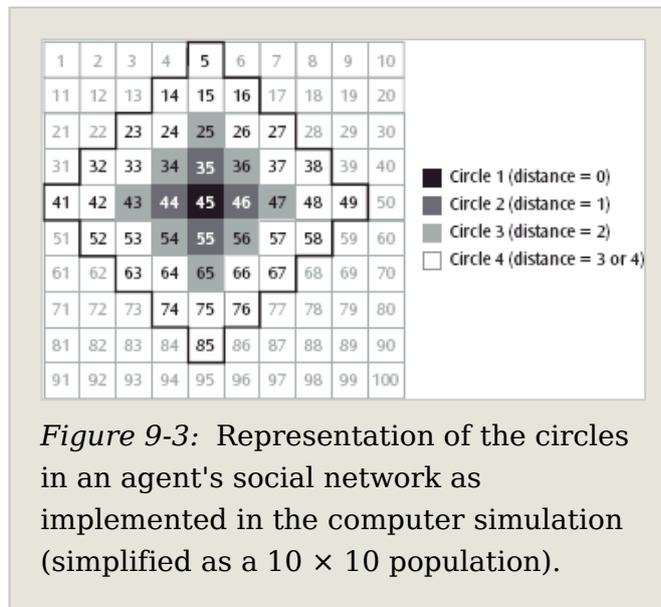


Figure 9-3: Representation of the circles in an agent's social network as implemented in the computer simulation (simplified as a 10 × 10 population).

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

shaped) and 21.1% (uniform) of the information used by availability-by-recall, on average.

How accurate is the social-circle heuristic? As it turns out, the answer depends on the structure of the environment. In the J-shaped environment, the frugal social-circle heuristic, on average, 76.3% correct inferences, nearly indistinguishable from the 77.5% correct (p.274) inferences scored by availability-by-recall. In the uniform environment, however, frugal search exacts a price. With, on average, 83.1% correct inferences, availability-by-recall clearly outperforms the social-circle heuristic's 75.2% correct inferences. How much more accurate is Wald's rule relative to the strategies that are constrained to information sampled in an agent's immediate social network? In the J-shaped environment, Wald's rule achieves an accuracy of 90.6% correct inferences, but requires an average sample size of over 700 agents ($M = 709.4$; $Md = 263.3$), given the standard values of the α and β parameters that we used. In the uniform environment, Wald's rule scored 93.6% correct inferences, requiring nearly 300 agents ($M = 288.1$; $Md = 199.8$).

Taken together, the results suggest that overall there is a tradeoff between accuracy and frugality: The more frugal the strategy (the fewer agents that are looked up), the less accurate the inferences. However, an interesting finding emerges when we distinguish between the two environmental distributions. In the uniform environment, the social-circle heuristic's frugality came at the price of inferior accuracy, whereas in the J-shaped environment no such frugality-accuracy tradeoff occurred (relative to availability-by-recall). That is, the extent to which a fast and frugal polling method can forsake the tradeoff between accuracy and frugality depends on the structure of the environment. Both heuristics were clearly outperformed by Wald's rule. Its predictive superiority, however, comes at the expense of enormous sampling efforts. On average, it sampled between 20 and 30 times as many agents as the social-circle heuristic, thus benefiting from a sample size that lies beyond the bounds of a person operating under constraints of time, knowledge, and memory.

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

To better understand why the social-circle heuristic's performance depends on the structure of the environment, we conducted another set of simulations. Specifically, we determined how the level of accuracy (on the 45 pair comparisons) changes with increasing sample size when samples are drawn randomly from the network. When the sample did not discriminate between a pair of event categories, random choice was implemented. Figure 9-4 shows the results, separately, for the J-shaped and uniform environments. The dotted lines in the graphs indicate the accuracy achieved with a random sample of size 500. Three observations are of relevance: First, in both environments, the information gain as a function of sampling more agents is subject to a diminishing return (cf. Hertwig & Pleskac, 2008). Second, how quickly the return diminishes depends on the environment: In the J-shaped environment, accuracy increases steeply at very small sample sizes but then levels off. Specifically, to match the gain of 14.3 percentage points achieved by increasing the sample size from 0 (p.275) (accuracy: 50%) to 5 (accuracy: 64.3%), one has to boost the sample size from 5 to 50 agents. Relatedly, the largest single gain, 8.7 percentage points, occurs by increasing the sample size from 0, when only random guesses are possible, to the smallest sample size of 1. In the uniform environment, by contrast, the gain of 13.3 percentage points achieved by increasing the sample size from 0 to 5 can be matched by boosting the sample size from 5 to 20. The accuracy gained by increasing the sample size from 0 to 1 amounts to only 3.6 percentage points. Third, we compared the score (black dots in Figure 9-4 Panels A and B) achieved by the social-circle heuristic—which relies on ordered search—with that achieved by random search; that is, when drawing a random sample of the average size required by the social-circle heuristic. In both environments, the social-circle heuristic surpasses that accuracy; the dots lie above the curve, suggesting that the heuristic discards redundant information. In other words, ordered search with conditional search extension (as implemented by the social-circle heuristic) is better than random search.

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

To conclude, our analysis suggests that in an environment with a J-shaped frequency distribution, the benefit of considering additional information diminishes quickly, after some large initial gains. To take advantage of more information in this environment,

(p.276) one has to invest much sampling effort. In the uniform environment, the rate of diminishing returns from additional information is less steep, thus it pays more to keep sampling. In addition, the initial gains are not as large as in the J-shaped environment. These findings suggest that in environments with J-shaped frequency distributions—which abound in the real world (e.g., Hertwig et al., 2012)—inferences can be accurate and frugal at the same time. In uniform environments, however, a frugal strategy is subject to a tradeoff with accuracy to a larger degree.

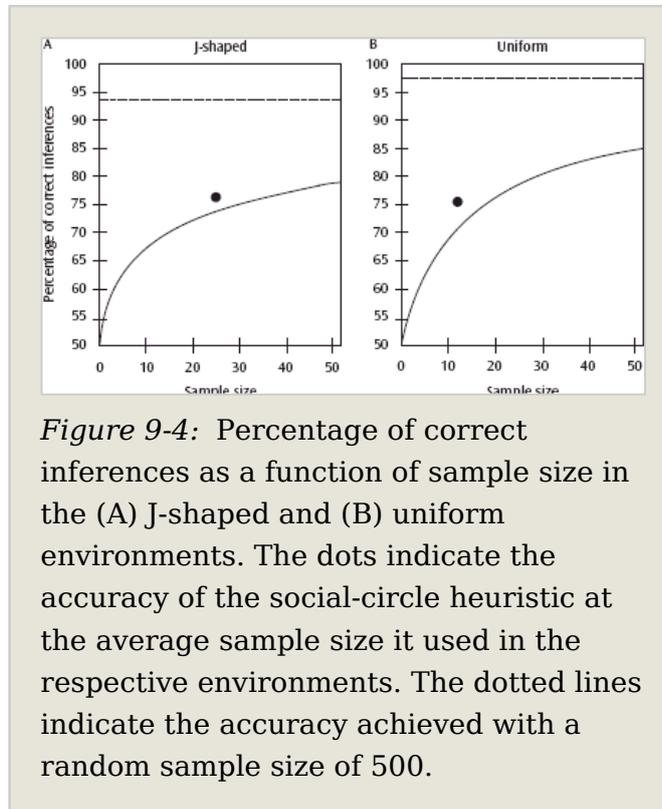


Figure 9-4: Percentage of correct inferences as a function of sample size in the (A) J-shaped and (B) uniform environments. The dots indicate the accuracy of the social-circle heuristic at the average sample size it used in the respective environments. The dotted lines indicate the accuracy achieved with a random sample size of 500.

The Social-Circle Heuristic: How Well Does It Predict People's Inferences?

The computer simulations suggest that, depending on the environment, the social-circle heuristic can be an efficient cognitive tool for judging relative event frequencies. Do people also use it? We next report an empirical study in which we tested how well the social-circle heuristic can predict people's actual inferences in comparison to availability-by-recall. We asked 40 students to judge the relative frequencies of the 24

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

infectious diseases in Germany for which the Robert Koch Institute keeps official records (e.g., Robert Koch Institute, 2001; see Hertwig et al., 2005; Pachur & Hertwig, 2006). For that purpose, we constructed a set of all possible 276 pairs of the diseases, and the task was to infer which one of two infectious diseases has a higher annual incidence rate in Germany. Subsequently, each participant indicated for each disease (a) whether he had heard of it previously (i.e., recognized it) and (b) how many, if any, people in each of his social circles—defined as self, family, friends, and acquaintances, respectively—had been affected by the disease. At first, we focused on the social-circle heuristic based on the altruism structure. Using participants' reported knowledge of instances of the diseases in their social circles, we derived individual-specific and item-specific predictions for both the social-circle heuristic and availability-by-recall. Because Wald's rule may require sample sizes that exceed the size of most people's networks, it is not considered as a descriptive model here.

In addition to examining the descriptive accuracy of the social-circle heuristic, the empirical test allowed us to examine its inferential accuracy under real-world conditions. In the environments used in the simulations, the density of instances (i.e., most agents in the population were instances of one of the event categories) was relatively high. As a consequence, the social-circle heuristic often stopped information search at early circles, allowing us to test the heuristic's accuracy when it considers only very little information. In addition, the instances were distributed randomly across the agents in the population. In real-world environments, however, (p.277) these conditions may not necessarily hold. For instance, the density of occurrences of diseases in a population is lower than in our simulations (because infectious diseases are usually rare events). In addition, infectious diseases are typically not randomly distributed among people, but instead occur in clusters, contrary to the random distribution of instances on the grid in the simulation. How does the heuristic fare as a descriptive model under these conditions?

Which Strategy Predicts People's Inferences Best?

Across the 24 infectious diseases, respondents reported having encountered, on average, 4.2 instances in their social network. Due to this paucity of experience, the social-circle heuristic and availability-by-recall made, on average, predictions in a subset of 11.1% and 10.5% of all 276 pair comparisons, respectively. For 7 of the 40 participants, neither strategy made a single prediction (because these participants did not recall any instances of the diseases in their networks). For each participant, we computed, separately for the social-circle heuristic and availability-by-recall, the percentage of inferences that were in line with the respective strategy (considering all cases in which the strategy did not have to guess and where both diseases were recognized). The social-circle heuristic correctly predicted, on average, 77% of the inferences, slightly less than availability-by-recall, which predicted 78% of the inferences correctly (Kruskal-Wallis test: $z = 1.59$, $p = 0.12$, $r = 0.28$). In other words, both heuristics—one assuming truncated search, the other exhaustive search—predicted actual inferences nearly equally well. One may suspect that the heuristics are almost indistinguishable because—due to the paucity of experience—the social-circle heuristic always had to search up to the final circle to find instances, thus functionally behaving like availability-by-recall. This was not the case. Of all 1,217 inferences (summed across all participants) in which the social-circle heuristic made a prediction, 11%, 29%, 26%, and 34% of the inferences were made on the basis of the first, second, third and fourth circles, respectively. As a consequence, the social-circle heuristic was, on average, more frugal than availability-by-recall, with 1.2 and 1.8 retrieved instances per inference, respectively ($t(32) = 3.54$, $p = 0.01$).⁵

(p.278) Which Strategy Predicts the Environmental Criterion Best?

What price does the social-circle heuristic's frugality exact? To answer this question, we determined the accuracy of the social-circle heuristic and availability-by-recall, separately for each participant. Specifically, we calculated the number of correct inferences of the environmental criterion (i.e., the actual relative event frequencies of the diseases) by each strategy, divided by the number of inferences in which it made

a prediction (again focusing on the cases where both diseases were recognized). The mean accuracy of availability-by-recall was 79%, a level of accuracy that was essentially matched by the social-circle heuristic with 78%. To appreciate this surprisingly high level of accuracy, it is worth reiterating that the knowledge input into the heuristics was participants' very limited sample of, on average, 4.2 instances of infectious diseases. Moreover, participants' sample space (i.e., their social network) was minuscule in comparison to the target population. In the computer simulation, an agent's social network made up about 1.6% of the total population (41 out of 2,500 agents). In the empirical study, a participant's social network could have made up merely about 0.0000017% of the total population of about 82 million Germans.⁶

Admittedly, across *all* pair comparisons, respondents scored only 60.9% correct inferences. However, this less impressive level of accuracy was due to cases in which the heuristics were not applicable. More specifically, when focusing on the cases where the social-circle heuristic or availability-by-recall made a prediction (i.e., when the number of instances discriminated between the two diseases), the participants' actual level of accuracy was, on average, 75.6% and 76.9%, respectively. In contrast, in cases where the heuristics made no prediction because the number of retrieved instances was identical for both diseases, participants' accuracy was only 63.3% and 63.6%, respectively. (The accuracy in the remaining cases—that is, those where at least one disease was not recognized—was 57.3%.)

(p.279) What does this result tell us about people's strategy use? The substantial drop in accuracy between cases where the strategies did and did not make a prediction might indicate that people guess when their instance knowledge does not allow them to derive an inference. Alternatively, given that participants' accuracy (63%) exceeded chance performance (50%; $t[39] = 10.0, p = 0.001$), people might resort to nonsocial cues (see next section) in this case. Based on a different domain (see next section), Pachur et al. (2012) found evidence to support both hypotheses. Half of their respondents guessed, whereas the other half appeared to

resort to cue knowledge when instance knowledge did not enable them to draw an inference.

To conclude, we obtained evidence that people rely on experienced instances sampled from their proximal social world to make inferences about the distal social world (i.e., the population at large). Despite the boundedness of people's experience, availability-by-recall and the social-circle heuristic achieved a high level of accuracy, rendering correct inferences in more than three-fourths of the cases where they made a prediction. The social-circle heuristic fared about as well as availability-by-recall, notwithstanding the fact that the former often relied on even smaller samples than the latter. This finding adds to the small but growing body of evidence that reliance on small samples can give rise to surprisingly competitive inferences (see Fiedler & Kareev, 2006; Hertwig & Pleskac, 2008; Hertwig & Todd, 2003).

Inference About Social Event Frequencies: Based on Instances or Cues?

So far, we have focused on strategies that embody the notion that people behave like intuitive pollsters who probe their social networks for occurrences of the events in question, and conceive of people as carriers of information. As mentioned above, however, there is an alternative approach to judging frequencies of social events: reliance on probabilistic cues (Bergert & Nosofsky, 2007; Gigerenzer et al., 1999; Juslin, Olsson, & Olsson, 2003; Rieskamp & Otto, 2006). For illustration, to judge the frequency of a disease, a person could consider properties of the disease that are indicative of its prevalence: For instance, is it contagious? How long is its incubation period? In light of the strong evidence that people often rely on cue-based inference strategies (Gigerenzer, Hertwig, & Pachur, 2011), they may also employ them to infer event frequencies. Are they? And how does the variant of the social-circle heuristic that structures search according to frequency of contact, the social-circle heuristic_F, fare in predicting people's inferences?

To find out, in the next contest we compare three instance-based strategies to three cue-based strategies (see Pachur et

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

al., 2012). The (p.280) instance-based strategies are: (a) the social-circle heuristic, (b) the social-circle heuristic_F, and (c) availability-by-recall. The cue-based strategies are: (a) *take-the-best* (Gigerenzer & Goldstein, 1996), a lexicographic and noncompensatory heuristic that looks up cues one by one in the order of their validity and terminates search as soon as a cue is found that discriminates between the two event categories; (b) *tallying* (Dawes, 1979), a strategy that looks up all cues, and infers that the event category with the higher number of positive cue values (after subtracting the negative ones) is the more frequent one; and (c) *weighted additive* (WADD), a strategy that has the same policy as tallying, except that the cues are weighted by their validities before they are summed up.

The six strategies were tested against each other in the context of judging the popularity of different sports in Germany. A sport's "popularity" was defined as "the number of registered club members" (Statistisches Bundesamt, 2002). In comparison to infectious diseases, in the sports environment the frequency of instances is relatively high. Therefore, people should be able to retrieve more instances in the process of making an inference. The strategies were tested against each other for all 300 possible pair comparisons of the 25 most popular sports in Germany. Forty participants were presented with the pairs and asked to judge for each pair which of the two sports was more popular. After making these inferences, participants indicated for each sport how many, if any, people in their social network were members of a respective sports club, allowing us to derive the predictions for the instance-based strategies. On average, participants reported 25.6 instances of sport club members (across all sports); substantially more instances than for the infectious diseases studied above. As a consequence, the social-circle heuristic, the social-circle heuristic_F, and availability-by-recall made unambiguous predictions in, on average, 56%, 57%, and 53% of the cases, respectively. In addition, participants reported which of the four network circles (self, family, friends, and acquaintances) these instances belonged to and how frequently they typically had contact with them. Contact information allowed us to assign each recalled instance to one

of the “frequency circles” as defined by the social-circle heuristic_F. Finally, we elicited participants’ knowledge of eight relevant cues—identified in a pilot study—including the sports’ cue values on the different sports, as well as the cue validities. This cue knowledge allowed us to derive the predictions for the cue-based strategies (for more details, see Pachur et al., 2012).

In order to classify participants as users of one of the six strategies, we used a maximum likelihood approach. Accordingly, the extent to which a strategy's predictions matched a participant's inferences was defined by the G^2 measure (e.g., Sokal & Rohlf, 1994), with lower (p.281) values indicating a better fit. We calculated G^2 on those (on average) 39.2% of the items where all six strategies made a prediction. Each participant was assigned to the strategy with the lowest G^2 . If the G^2 of the best strategy was not lower than the G^2 expected under random guessing, the participant was classified as “guessing.” Figure 9-5 shows the resulting classification distribution. Whereas no participant was classified as guessing, about two-thirds were classified as users of an instance-based strategy. Overall, availability-by-recall (23%) and the social-circle heuristic (25%) were the two strategies used by the largest number of participants. The numbers for the other strategies were clearly smaller in comparison, with, for instance, 17% for tallying, 17% for WADD, and 11% for the social-circle heuristic_F.

To summarize, when inferring the frequency of behaviors and characteristics of others, people seem to be more likely to consider instances in their proximal social spaces than to rely on probabilistic cues. In addition, we found that the dimension that appears to structure search and retrieval of instances is not frequency of contact, but kin and reciprocal altruism.

How Ecologically Rational Is Reasoning by Instances?

Inference about event frequencies based on knowledge of concrete instances has in the past often been equated with biased judgments

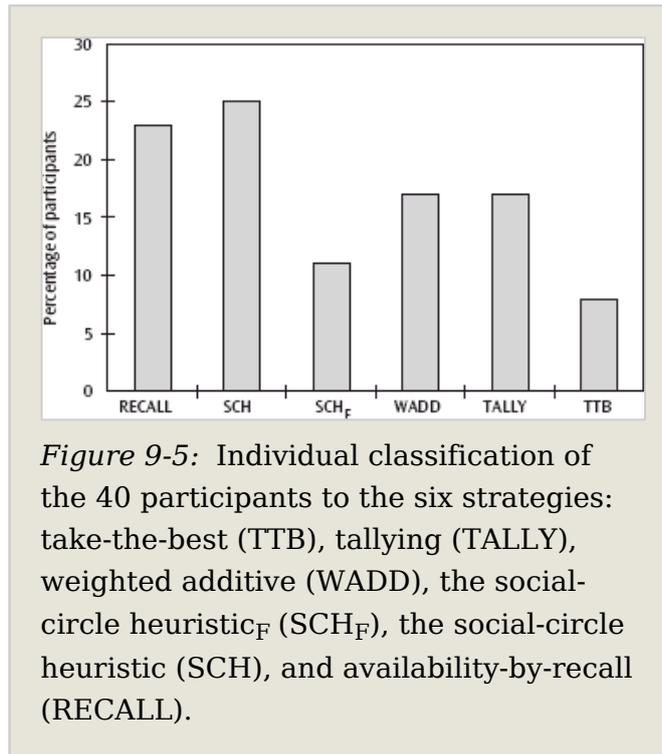
The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

(p.282) (e.g., Lichtenstein et al., 1978; Marks & Miller, 1987; Tversky & Kahneman, 1973). This need not have been so. When proposing the availability heuristic, Tversky and Kahneman initially stressed that “availability is an *ecologically valid* clue for the judgment of frequency

because, in general, frequent events are easier to recall or imagine than infrequent ones” (p. 209; our emphasis). Researchers outside psychology, such as sociodemographers, have recognized the surprising predictive power of knowledge of instances in people's social networks. For instance, Killworth, McCarty, Bernhard, Shelley, and Johnsen (1998) developed a method for estimating the size of hard-to-count subpopulations that is based on individuals' reports of instances in their social networks, and concluded that this method yielded highly reliable estimates. Nevertheless, as emphasized by the notion of ecological rationality (Todd, Gigerenzer, & the ABC Research Group, 2012), the performance of a strategy depends on the structure of the environment.

When Does the Social-Circle Heuristic Do Well?

What characteristics of an environment determine whether the social-circle heuristic succeeds or fails? This question concerns the ecological rationality of the heuristic—that is, the fit between the heuristic and the structure of the environment (cf. chapter 1). As Figure 9-4 Panels A and B demonstrate, the degree of skewness in the frequency distribution is of key importance. In many real-world domains, frequency distributions are such that for very few event categories there



The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

is a very large number of instances, whereas for most event categories there is only a moderate number of instances (e.g., Bak, 1996). Take, for example, the actual frequency distribution of infectious diseases used in our first study (Figure 9-3; see Table 1 in Hertwig et al., 2005): One single disease—gastroenteritis—occurs more frequently than all other diseases combined. By comparison, the domain investigated in the second study, the frequency distribution of people in sports clubs, is considerably less skewed. One way to quantify skewness is to express the distribution of a class of n objects on the target variable y as a power function $y = x^\alpha$, where x refers to the rank of an object (with the n objects ranked according to the target variable in descending order from 1, ..., n). The skewness of a distribution is the value of α that yields the best fit, with a higher α indicating a higher skewness. In Figure 9-6, the distributions of the infectious diseases and sports environments are plotted on a log-log scale (where a power function translates into a straight line), showing that the distribution was considerably more skewed for the diseases ($\alpha = 4.07$; $R^2 = 0.93$) than for the sports ($\alpha = 1.36$; $R^2 = 0.95$).

Why does the skewness of the frequency distribution matter? One possible explanation is that it affects the circle *validities*, defined as (p.283)

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

the proportion of correct inferences when an inference is made based solely on the number of the instances within a given circle. The validity of the “self” circle equals the proportion of the pair comparisons in which a decision maker

is an instance of category A but not of category B, and in which category A is more frequent than B in the population. Analogously, the validities of the other circles equal the proportion of the pair comparisons in which the number of instances, retrieved by the decision maker, of categories A and B is unequal, and in which the more frequent category within the circle is also more frequent in the population. Table 9-1 lists, for all four circles, the circle validities and discrimination rates (the relative frequency with which a circle discriminates between two event categories) for the highly skewed disease environment (for both the computer simulation and the empirical study). In the simulated and the experimental data, the validity of the self circle (based on a sample with $n = 1$) is extremely high; namely, 96% and 91%, respectively. The validity decreases—slightly in the simulated data and substantially in the empirical data—across larger circles, whereas the discrimination rates increase. In the sports environment, in contrast, the circle validities are considerably lower overall and, if anything, increase from the smallest to the largest circle (Table 9-2).

Why does the skewness of the distribution affect the circle validities? In highly skewed environments, the initial circles will discriminate mainly in those cases where one event category is considerably more frequent than the other. The effect sizes (i.e., the objective

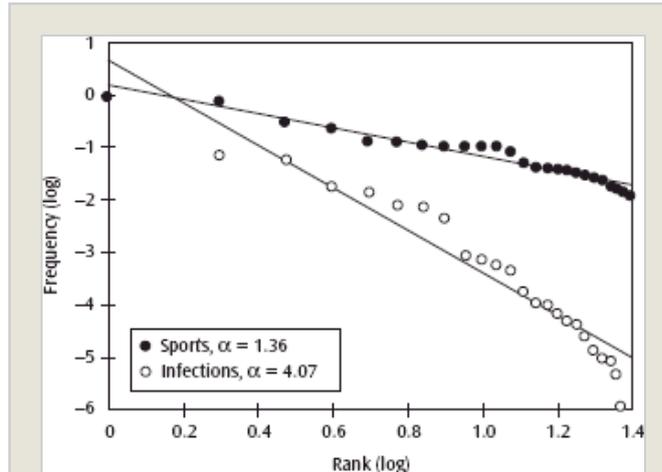


Figure 9-6: The distributions of frequencies in the disease and sport environments, plotted on a log-log scale.

(p.284)

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

Table 9-1: Average Circle Validities and Average Discrimination Rates in the Computer Simulation and the Empirical Study (Involving Infectious Diseases)

Circle	Computer simulation				Empirical study (diseases)		
	d	n of agents in circle	Circle validity (SD)	DR (SD)	Content	Circle validity (SD)	DR (SD)
1	0	1	0.96	0.19 (0.04)	Self	0.91 (0.16)	0.01 (0.02)
2	1	4	0.95	0.30 (0.11)	Family	0.82 (0.25)	0.04 (0.04)
3	2	8	0.93	0.38 (0.12)	Friends	0.81 (0.22)	0.04 (0.04)
4	3 or 4	28	0.90	0.59 (0.11)	Acquaintances	0.73 (0.26)	0.06 (0.06)

Note. d is the distance of the central agent in the network.

The standard deviations (*SD*) of the circle validity and the discrimination rate (*DR*) were calculated across agents (computer simulation) and participants (empirical study), respectively.

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

difference in frequency) are thus very large. Inferences involving comparisons between event categories of medium or low frequencies usually cannot be determined based on the initial circles—because they are too rare to occur in the “self” or “family” circle—and need to be referred to more encompassing circles. That is, in highly skewed environments, the costs of increased sampling error incurred with small sample sizes is offset because only comparisons with very large effect sizes—which are unlikely to be harmed by sampling error—are decided based on initial circles. In less skewed environments,

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

Table 9-2: Average Circle Validities and Average Discrimination Rates in the Empirical Study (Involving Popularity of Sports), Separately for the Social-Circle Heuristic and Its Variant, the Social-Circle Heuristic_F

Circle	Social-circle heuristic			Social-circle heuristic _F		
	Content	Validity (SD)	DR (SD)	Content	Validity (SD)	DR (SD)
1	Self	0.51 (0.19)	0.04 (0.05)	Self	0.51 (0.19)	0.04 (0.05)
2	Family	0.66 (0.24)	0.11 (0.11)	Weekly and more	0.60 (0.20)	0.23 (0.16)
3	Friends	0.60 (0.18)	0.29 (0.18)	Monthly	0.60 (0.22)	0.22 (0.13)
4	Acquaintances	0.55 (0.17)	0.38 (0.21)	6 months and less	0.55 (0.16)	0.40 (0.21)

Note. The standard deviations (SD) of the circle validity and the discrimination rate (DR) were calculated across participants.

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

(p.285) by contrast, the differences between two event categories are usually much smaller, and sampling error will often hamper the ability to correctly decide these comparisons.

Benefits of Reliance on Small Samples

Recently, it has been argued that reliance on small samples has a number of important benefits, such as the early detection of useful binary correlations (e.g., Fiedler & Kareev, 2006; Kareev, 2000, 2005; but see Juslin & Olsson, 2005), or the amplification of differences and therefore reduced difficulty in making a choice (Hertwig & Pleskac, 2008, 2010). The social-circle heuristic and availability-by-recall rely on small samples by merely polling a person's social network or slices of it.⁷ Might this focus on small samples of personally experienced instances be beneficial? It has been argued that distortions in estimates of event frequencies are caused, for instance, by “disproportionate exposure, memorability, or imaginability of various events” (Lichtenstein et al., 1978, p. 551), assuming that the search space in memory extends far beyond a person's social network and includes a virtual circle, populated with incidents conveyed through the mass media (Lichtenstein et al., 1978). Clearly, augmenting the search space in memory by a virtual circle comes at the price of systematic error simply because potential news or entertainment items are selected for their potential to captivate an audience (see, e.g., Combs & Slovic, 1979; Frost, Frank, & Maibach, 1997). In contrast, sampling only within one's social network—although constraining the sample size—guards one against the media's selection of rare, vivid, dramatic, emotional, and sensational events (cf. Pachur, Hertwig, & Steinmann, 2012).

Doubtless, reliance on small samples exacts risks. One such is the risk of being miscalibrated to “clumpiness” in time or space. Illnesses, for instance, often occur in spatial patches or clusters (“hot spots”), such as leukemia near nuclear installations, or increased rates of diseases in underserved areas (e.g., Antunes & Waldman, 2002). Relatedly, many diseases have skewed age distributions. A recent study of the (p.286) total of 229 confirmed human cases of avian influenza type A (H5N1, or “bird flu”) found numerous confirmed cases

among children and young adults, with relatively few cases among older adults (Smallman-Raynor & Cliff, 2007). Clumped spatial or skewed age distributions can compromise the accuracy of instance-based strategies such as the social-circle heuristic. The “young” social network of adolescents, for instance, may under-represent the occurrence of old age diseases such as Alzheimer's. By the same logic, the “old” social networks of elderly people may over-represent the occurrence of these diseases (and under-represent diseases mainly prevalent among the young, such as measles).

Consistently, Benjamin and Dougan (1997; see also Benjamin, Dougan, & Buschena, 2001) found that people's estimates of various mortality risks were more in line with event frequencies in their age cohort than with those in the general population. Interestingly, we observed the same tendency in the sports study. Respondents' inferences about the popularity of various sports were somewhat better tuned to the frequencies in their age cohort (i.e., number of club members aged 27 years and younger) than to the population frequencies. When using the relative frequencies in the cohort rather than those in the entire population as a reference, participants' accuracies was somewhat higher, 64.3% versus 62.9%, $t(39) = 1.7$, $p = 0.05$ (one-tailed). But even this sampling bias may be a blessing in disguise. Despite the common notion that the “world is a village,” people typically do not navigate in all social spheres. Therefore, as Benjamin and Dougan argued, being able to accurately estimate the event frequencies in the population may be less important than being calibrated to the events in one's proximal environment.

When Do People Refrain From Instance-Based Strategies?

The results reported here and in Hertwig et al. (2005) suggest that people often rely on instance-based strategies, and sometimes constrain their information search to very small sample sizes. But any heuristic has boundary conditions. What are those of the social-circle heuristic? First, as with other tools for probabilistic inferences (e.g., recognition heuristic, take-the-best), people are likely not to resort to instance-based strategies when they have direct, conclusive knowledge about the criterion (cf. Gigerenzer et al., 1991; Pachur & Hertwig,

2006). To illustrate, when a person happens to know that soccer is the most popular sport in Germany, he can deduce with certainty that soccer is more popular than any other sport (irrespective of any instance knowledge). Second, people might rely less on instance knowledge sampled in their social environment when they fathom that their social environment represents (p.287) a highly unrepresentative sample. For instance, judging the relative frequency of various professions, a carpenter, working with and perhaps traveling on the job with other carpenters, will hardly make the inference that carpenters are ubiquitous (see also Oppenheimer, 2004; but see Fiedler & Juslin, 2006; Hamill, Wilson, & Nisbett, 1980; Nisbett & Borgida, 1975).

Beyond Event Frequencies: Inferring Norms and Attitudes

Our social networks represent an informationally rich landscape that we can roam when making decisions. How do we use and process this information? One possibility would be to aggregate as much information as possible—as statistical lore would advise us. Alternatively, information processing may be sequential, ordered, and limited—as embodied by the social-circle heuristic. Next, we will show how this notion may help us understand how people learn norms and form attitudes. In both domains, we will assume that the first circle, the self, does not provide answers to the questions posed.

Inferring Social Norms

In social situations it is often desirable to coordinate one's behavior with that of others. Social norms can function as a coordination device. But how do we figure out in a given situation what the social norm is? For instance, suppose a teenager wonders whether or not he should offer his seat to the elderly people who are at the mercy of the shaking bus. The following variant of the social-circle heuristic, omitting the self circle, could provide a quick answer by looking outside, to others, for evidence of how to act:

Search rule: Search the social circle for instances of the behavior in question, proceeding sequentially through the

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

circles involving family, friends, and acquaintances, respectively.

Stopping rule: If, within a circle, there is a majority for one of the two behaviors, stop search. If, within a circle, the same number of instances (or none) for both behaviors is found, proceed to the next circle.

Decision rule: Adopt the behavior of the majority in the circle in which search was stopped. If sampled information does not discriminate, guess which behavior is most appropriate.

Of course, this is not the only heuristic that could be recruited to infer the right social behavior. The teenager may simply imitate the majority behavior of his peers or the behavior of a prestigious (p.288) individual (Henrich & Henrich, 2007; Richerson & Boyd, 2005; chapters 1 and 17). Yet, Epstein (2001) proposed a model for the individual learning of social norms that is closely related to the social-circle heuristic. It consists of a search rule and a stopping rule specifying the sampling procedure within a person's social environment, and a decision rule, according to which the behavior of the majority of persons in the sample is adopted. By way of its stopping rule, the model attempts to reach a decision of how to behave based on as little information as possible. For that purpose, the heuristic first searches for information by checking the distribution F of behaviors within a particular radius, r , (or circle) around the person. The size of the radius indicates the size of the sample taken from the population.⁸ The initial value of r is chosen at random. Let us assume that the majority (say 60%, i.e., $F(r) = 0.60$) of the examined people altruistically offer their seat to an elderly person. The heuristic now extends the sample slightly (i.e., increases r by 1) to check whether this result is robust. If the recommendation from the larger sample is different (i.e., $F(r+1) \neq F(r)$), the person expands the sample until the recommendation from one sample matches that of the next-larger sample. The recommended behavior is then adopted. If, however, the first step—that is, the comparison of the initial sample with the slightly larger sample—yields the same recommendation (i.e., $F(r+1) = F(r)$), the person continues by reducing the sample

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

size and checks whether the smaller sample yields the same recommendation. If it does (i.e., $F(r) = F(r-1)$), the sample is reduced further, until the sample is as small as possible, although still reflecting the recommendation of the next-larger sample size. In other words, Epstein's heuristic will stop information search at a particular sample size if a next-larger sample yields the same recommendation *and* the next-smaller sample yields a different recommendation. Thus, the heuristic makes a decision as soon as the accrued evidence can be confirmed, an idea that has also been proposed for cue-based inference (Karelaia, 2006).

Testing his model in computer simulations, Epstein (2001) observed that over time one behavior becomes dominant, although the other behavior “survived” in small clusters. At the same time, once a particular norm had become entrenched, the model learned to sample only little information (i.e., the resulting radius r was small). The model thus shows how information search becomes increasingly limited when many agents converge on the same behavior; more extensive information search occurs only in regions of the environment in which there is no such convergence. In addition, (p.289) Epstein examined this model under different levels of *noise*, defined as the degree to which a behavior is chosen at random rather than being derived from the distribution of behaviors in the sample. He observed that even with extreme noise, the final group of “circles” on which the agent based a decision was rather small, although no local stability emerged.

Epstein's (2001) model is ecologically rational in the sense that—like the social-circle heuristic—it adjusts the amount of information used to make a decision to the structure of the social environment. The key ecological characteristic here is the diversity of behaviors in the environment. The less diverse the behavior, the smaller the amount of information considered for the final judgment. More extensive information search only occurs when social behavior proves heterogeneous, and when stable norms have not yet formed. In contrast to Epstein's model, the social-circle heuristic has a clearly defined starting point for the sampling process. In

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

addition, because it always starts with a very small sample, the social-circle heuristic does not assume that information, once accessed, is discarded (whereas Epstein's model discards redundant information). One of the key merits of Epstein's model is that it shows how a stable configuration of norms, with a coherent majority of agents following one norm and small niches with agents following a different norm at the fringe of the population, can arise from an arbitrary starting distribution of behaviors. It is an interesting research question for the future to examine whether the social-circle heuristic, using even simpler search and stopping rules than Epstein's model, could generate the same phenomenon.

Forming Attitudes

Studying the 1980 presidential campaigns, MacKuen and Brown (1987) concluded that a “citizen's social circle plays a discernible and important role in the development of political evaluations” (p. 485). Specifically, how people's evaluations of the candidates Reagan and Carter and their respective parties evolved was substantially influenced by the voting intentions of their neighbors, or in Noelle-Neumann's (1977) terms, by the climate of opinions. One interpretation is that, like an intuitive pollster, people sample their immediate social environment to form their political attitudes. This result illustrates a stock-in-trade phenomenon in social psychology: People's views and attitudes of the world are strongly shaped by their social environment (for an overview, see Cialdini & Goldstein, 2004).

Latané (1981) proposed that the impact of other people's opinions is a function of three factors, which are connected multiplicatively: the strength (e.g., authority or expertise) of the person embodying the opinion, the immediacy of that person (in terms of space and (p.290) time), and the number of “opinion carriers” (cf. Stycos, 1952), that is, people holding the opinion. But what are the processes that lead to the influence of these factors on attitude formation? Nowak, Szamrej, and Latané (1990), implementing Latané's social impact theory in a computer simulation, proposed a process according to which a person weights the opinions in her social environment by the number and the mean strength of opinion carriers, divided by their social distance. It is assumed that the

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

resulting “forces” of the different opinions in the environment are compared, and that the opinion that has the greatest force is adopted. Social impact theory shares the common assumption of many theories of human behavior that conflicts are mastered by making tradeoffs, and that weighting and summing are the processes by which such tradeoffs can be made in a rational way (see Brandstätter, Gigerenzer, & Hertwig, 2006). If, for instance, there is a conflict between the opinion expressed by an expert, and the opinion expressed by many laypeople, this conflict will be solved by trading-off the strength of an opinion and the number of opinion carriers.

There is, however, an alternative to the assumption that cognitive processes always compute tradeoffs in terms of weighting and summing of information. In the words of Lopes (1995): “Judgments that are captured algebraically by the idea of weight express themselves in individual choices by the *order* in which operations are carried out” (p. 203, italics in original). According to this view, people might master conflicts by simple sequential heuristics that avoid tradeoffs. Examples of such sequential heuristics are the take-the-best heuristic in probabilistic inference (Gigerenzer & Goldstein, 1996), the priority heuristic in risky choice (Brandstätter et al., 2006), and the social-circle heuristic in inference of event frequencies (Pachur, Rieskamp, & Hertwig, 2012).

Can the social-circle heuristic—which relies on sequential processing—explain the regularities of social impact identified by Latané (1981), which are usually assumed to result from the processes of summing and weighting? Take, for example, the effect of strength (authority) and immediacy of social impact. As it turns out, this effect follows from the architecture of the social-circle heuristic—specifically, from the assumption of limited search. So far we have considered social distance (operationalized as altruism and contact frequency, respectively) as the key dimension guiding search. In domains in which expertise matters, however, search within one's social network could also be ordered according to authority (or a combination of authority and social distance). If so, and if search is terminated once a social circle discriminates, opinion carriers with lower authority or with

greater social distance will simply have a lower probability of being polled. Latané also identified the “number of carriers holding an opinion” as a determinant of social impact. (p.291)

Again, this finding does not necessitate the summing of all opinions in one's social environment. After all—other things being equal—the higher the number of carriers of a particular opinion, the more likely it is that one of them is sampled before information search is stopped, irrespective of when search is stopped. To conclude, there is an alternative to the assumption that attitude formation should be modeled in terms of weighting and summing of all information. The emerging determinants on attitudes may directly follow from the architecture of the simple polling strategies with which people forage for information distributed in their proximal environments.

Conclusion

To infer quantities of the world, we can recruit nonsocial or social information. The members of our social networks afford us one important piece of social information—wittingly and unwittingly, they are carriers of information that we can exploit. Knowing that more people in a network drink Coca-Cola than Pepsi or that more endorse Obama than his Republican opponent can help us predict preferences in the population. Such naturally accessible information in social networks can be reaped using different instance-based heuristics. We considered two: the social-circle heuristic and availability-by-recall. They both appear to describe what a sizeable chunk of respondents do when asked to make inferences about event frequencies. The social-circle heuristic bets on substantially smaller slices of the social environment than availability-by-recall. This frugality works rather well in skewed environments but comes with a substantial price tag in uniform environments. In inferring quantities—including genuinely social quantities such as norms, attitudes, and preferences—the adaptive toolbox may be stacked not just with cue-based but also with instance-based inferential tools. The boundedly rational mind may prove to be also an intuitive pollster of its social environment. (p.292)

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

Notes:

(¹.) Zhou et al. (2005) found evidence that a discretely hierarchical structure with increasingly larger circle sizes is typical for social networks.

(².) Note that the *noncompensatory* character refers to how the sums of occurrences in each circle are processed across circles. Within a circle, an occurrence for one event category can compensate for the occurrence for another event category.

(³.) The simulations showed that performance of the sequential sampling model is relatively robust against changes in these parameters. Interestingly, the model's performance increases while the average required sample size decreases with higher α s and β s (i.e., when the model becomes more liberal). This effect, which occurs in both J-shaped and the uniform environments (see text), is due to a smaller proportion of cases where the model falsely accepts the H_0 (i.e., that both event categories are equally frequent).

(⁴.) J-shaped distributions are often described in terms of a power function (i.e., $y = b \times x^a$). Power-law distributions are among the most prevalent distributions encountered in everyday life. They are related to general growth processes (Gabaix, 1999), and are able to capture phenomena as diverse as the distribution of incomes, sales of books, and the sizes of computer files, moon craters, or cities (Schroeder, 1991).

(⁵.) Note that, in contrast to the computer simulations, this measure of frugality does not include the comparisons in which the number of recalled instances did not discriminate. But given that availability-by-recall and the social-circle heuristic had to guess similarly frequently (both in the simulation and the experiment), this difference is not likely to markedly affect the results.

(⁶.) This estimate is based on the scale-up method (Killworth, McCarty, Bernhard, Shelley, & Johnsen, 1998). According to this method, the social network size of a person, c_i , can be estimated based on the population size t , the number m of recalled instances for event category j and the corresponding

The Mind as an Intuitive Pollster: Frugal Search in Social Spaces

population frequency for the event category, e : Using $t = 82$ million for the population size of Germany, the statistics in Table 9-3, and assuming that the number of recalled instances represents participants' accumulated experiences across a period of ten years (rather than one year, which the incidence rates refer to), we obtain an estimated network size $c = 142$.

(⁷) People differ in terms of the size of their social networks, and larger networks represent, *ceteris paribus*, the population more accurately. Therefore, it is possible that people with larger social networks may make more accurate inferences. There was indeed such a tendency ($r = 0.13$, with network size estimated based on the scale-up methods described in Footnote 5 and using people's reported instance knowledge). This positive relationship disappeared once individual differences in discrimination rates were taken into account. That is, people with larger networks could retrieve more instances. As a consequence, their instances discriminated in more comparisons and required guessing in fewer cases than the sample instances of people with smaller networks, leading to a higher accuracy.

(⁸) In Epstein's (2001) model, agents are assumed to be located on a one-dimensional circle, and r indicates the number of agents to each side that are included in the sample.



Access brought to you by: Max-Planck Society