Cognition 192 (2019) 103994

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Original Articles

The rare preference effect: Statistical information influences social affiliation judgments

Natalia Vélez*, Sophie Bridgers, Hyowon Gweon*

Department of Psychology, Stanford University, United States

ARTICLE INFO

Keywords:
Social cognition
Statistical reasoning
Social categories
Preferences

ABSTRACT

Shared preferences—liking the same things—facilitate and strengthen bonds between individuals. However, not all shared preferences are equally meaningful; sharing a rare preference with someone is often more exciting and meaningful than sharing a common preference. Here we present evidence for the *rare preference effect*: Participants chose to interact with (Experiment 1), and endorsed interactions between (Experiment 2), individuals who shared a rare preference, rather than those who shared a common preference, and this tendency increased with the relative rarity of the preference. While having a preference usually implies knowing *and* liking something, the presence of shared knowledge alone was sufficient to give rise to the rare preference effect (Experiments 3 & 4). Further, we find that social affiliation judgments are modulated by the causal process by which individuals came to have shared knowledge: Participants preferred to interact with someone who acquired a shared preference deliberately rather than accidentally (Experiment 5). In addition to the many cultural and emotional factors that drive mutual attraction, these results suggest that people's decisions about with whom to interact are systematically influenced by the statistics of the social environment.

1. Introduction

Preferences reveal far more than what a person likes or dislikes. Preferences reflect one's personal history and social context, including one's age, gender, socioeconomic class, culture, political affiliation, and personality (Bonneville-Roussy, Rentfrow, Xu, & Potter, 2013; Carney, Jost, Gosling, & Potter, 2008; Rentfrow, Goldberg, & Zilca, 2011; Rozin & Siegal, 2003; Van Eijck, 2001). Perhaps because preferences reveal so much about a person, we are curious to learn about what others like and readily broadcast our own preferences to others. People often bring up their music preferences when becoming acquainted with a stranger (Rentfrow & Gosling, 2006) and form fairly accurate impressions of a stranger's personality traits based solely on their music preferences (Rentfrow & Gosling, 2006, 2007).

Preferences also shape our interactions with others. In particular, we are drawn to people who like what we like: We delight in finding overlaps in preferences and bond with those who enjoy the same hobbies, read the same books, or root for the same sports teams as us. Evidence suggests that people evaluate others who share their music preferences more positively than those who do not (Boer et al., 2011; Lonsdale & North, 2009) and tend to have similar preferences as their friends (Selfhout, Branje, ter Bogt, & Meeus, 2009; Werner & Parmelee, 1979). Preschool-aged children also prefer to approach and learn from people who like the same toys as they do (Fawcett & Markson, 2010a, 2010b), and even infants expect people who have similar preferences to interact positively with one another (Liberman, Kinzler, & Woodward, 2014). Thus, discovering a shared preference is more than just a pleasant coincidence; it provides an opportunity to assess our compatibility with potential social partners and guide our future relationships with them.

However, not all shared preferences are equally meaningful. In the movie *Everyone says I love you*, a love-struck Woody Allen deliberately presents himself as a New Yorker who vacations in Bora Bora, listens to Mahler's 4th symphony, and admires the Italian painter Tintoretto, just to make Julia Roberts (who is also a New Yorker who likes all these things) fall for him. And he succeeds: She not only finds these coincidences fascinating but also finds him attractive. Woody Allen's strategy is not just a clever way to woo someone; it also raises a deeper question about the social significance of discovering someone who likes the same things as you. Among the many things one might like, what makes some shared preferences more meaningful than others?

Indeed, there may be many factors that make a shared preference

* Corresponding authors at: Department of Psychology, Stanford University, 450 Serra Mall, Jordan Hall 280, Stanford, CA 94305, United States.
E-mail addresses: nvelez@stanford.edu (N. Vélez), gweon@stanford.edu (H. Gweon).

https://doi.org/10.1016/j.cognition.2019.06.006
Received 13 September 2018; Received in revised form 30 May 2019; Accepted 4 June 2019

meaningful. Among these factors, here we present evidence for a key factor that systematically influences how people perceive, interpret, and use shared preferences: their *rarity*. In the movie, Julia Roberts may not have fallen for Woody Allen if she routinely ran into people who also share these preferences. You may also have experienced firsthand the thrill of meeting someone who shares a very rare preference with you. We refer to this intuition as the *rare preference effect*: All else being equal, sharing a rare preference with someone may be more meaningful than sharing a common preference.

The current study presents initial empirical evidence for the rare preference effect and investigates its cognitive underpinnings. Below, we motivate the hypothesis that the prevalence of preferences can guide adults' judgments about whom to approach or befriend. We then provide an empirical demonstration of the rare preference effect (Experiments 1–2) and further investigate the representations that give rise to this effect (Experiments 3–5).

1.1. Tracking and inferring the prevalence of preferences

At its core, detecting and making social judgments based on the prevalence of preferences requires a sensitivity to the relative probability of events—that is, distinguishing events that happen with low frequency (rare events) from those that happen with higher frequency (common events) within a given population. Prior research suggests that the roots of this ability may be present early in life. Even infants distinguish differences between common and rare kinds of objects in a population (Xu & Garcia, 2008) and use this information to draw further inferences about the person engaged in the sampling process (Gweon, Tenenbaum, & Schulz, 2010; Kushnir, Xu, & Wellman, 2010).

Yet tracking the statistics of people's *preferences* may pose unique challenges. Unlike the proportion of objects in a box, preferences are not directly observable and must be inferred from others' choices and testimony (Jern, Lucas, & Kemp, 2017; Lucas et al., 2014) or, more recently, from proxy measures such as Billboard charts and streaming counts. Prior work has demonstrated biases and errors in adults' estimates of the prevalence of abstract properties such as beliefs, opinions, attitudes, and habits. Adults systematically distort the prevalence of their own attitudes and behaviors (Monin & Norton, 2003; Nisbett & Kunda, 1985; Ross, Greene, & House, 1977; Suls & Wan, 1987; Suls, Wan, & Sanders, 1988) and ignore base-rate information in favor of descriptions of specific traits when making judgments about others (Kahneman & Tversky, 1973). Even within these studies, however, there is evidence that adults use statistical information appropriately. For example, adults do use base rates when they are the only information provided (Kahneman & Tversky, 1973). Further, while adults tend to inflate the prevalence of their own attitudes, their estimates nonetheless accurately distinguish commonly held attitudes from rarer ones (Nisbett & Kunda, 1985). Numerous studies have shown that adults can also use sparse data to accurately report the distribution of parameters of real-world events, such as their frequency and duration (Griffiths & Tenenbaum, 2006; Hasher & Zacks, 1984; Hertwig & Gigerenzer, 1999; Peterson & Beach, 1967).

Collectively, prior work suggests that, despite some biases, people's statistical intuitions do reflect the overall structure of the world. People track the "social statistics" of their environment, representing not only the prevalence of concrete, observable events but also the prevalence of abstract, unobservable qualities such as people's attitudes and preferences. Although these statistics often manifest as intuitions rather than as exact, formal estimates, they do reflect the relative frequencies and distributions of abstract occurrences.

1.2. Rare preference effect: Influence of statistical information on social decision-making

Building on prior work on intuitive statistical reasoning, the current work focuses on *how* statistical information might affect social judgments. More specifically, we propose that people's beliefs about the prevalence of preferences systematically influence their decisions about with whom they want to interact. Here we outline our three overarching goals.

Our first goal is to address the most basic question: Is the rare preference effect a robust, systematic phenomenon that actually stems from rarity? It is possible that this effect is an illusion that only exists in personal anecdotes or movies. Even if the phenomenon itself is real, the appeal of sharing a rare preference may stem from factors that are confounded with rarity. For example, rare preferences may be considered more socially desirable (see Monin & Norton, 2003; Suls et al., 1988 for related work on attitudes and behaviors) or may be more strongly held than common preferences. Thus, our initial goal is to provide empirical evidence of the rare preference effect. We first demonstrate this effect in a personally relevant and ecologically valid context, where participants report their own preferences and their intuitions about the prevalence of their preferences (Experiment 1, first-person judgments). Here, participants generate statistical information about the real-world prevalence by reconstructing it from their own knowledge and past experience. We then replicate the effect in a minimal, tightly controlled paradigm, where participants see visual displays that explicitly convey statistical information about the prevalence of preferences for novel items in a novel population (Experiment 2, third-person judgments). Together, these experiments provide converging evidence for the rare preference effect and identify a contribution of rarity that is distinct from other attributes that come from participants' prior knowledge of real-world items (e.g., social desirability, strength).

Having demonstrated the presence of the rare preference effect, our second goal is to investigate the scope of this effect: Are people indiscriminately and superficially drawn to rare events, or does the rare preference effect instead reflect a sophisticated use of statistical information? In fact, the rare preference effect may be just one of many ways in which people use information about the prevalence of preferences. For instance, statistical information may also guide inferences about new individuals. Suppose you meet someone whose favorite artists are unknown; in order to maximize the chances of finding a shared preference, you might bring up an artist that is widely liked (e.g., da Vinci) rather than an obscure artist that is less likely to be recognized, let alone liked. Thus, we use third-person judgments (Experiments 2 and 4) to test whether participants use information about the prevalence of preferences flexibly to make a wide range of social judgments. Rather than indiscriminately preferring social partners who have rare preferences, people may prefer someone who has a more common preference or even ignore prevalence information altogether, depending on the context.

Our final goal is to better understand the representations that underlie the rare preference effect. When someone says "My favorite artist is Tintoretto", you learn two things about this person: (1) she knows about the artist and his work, and (2) she enjoys and admires his paintings. That is, preferences provide information about both what people know (henceforth referred to as *knowledge*) and what they like (henceforth referred to as *affinity*[1]). Although it is possible to simply know about Tintoretto without necessarily liking or enjoying his work, or to find his paintings appealing without knowing anything about the artist, knowing and liking often go hand in hand.

Given that expressing a preference for something usually implies both knowledge and affinity, one possibility is that either of these properties is sufficient to give rise to the rare preference effect. Infrequent stimuli are seen as particularly salient and attention-

[1] We operationally define *affinity* as the propensity to like something upon one's first exposure to it, in order to distinguish it from preferences, which are a stable liking for something of which one has prior knowledge and experience.

grabbing (McCarthy, Luby, Gore, & Goldman-Rakic, 1997; Sutton, Braren, Zubin, & John, 1965), and people tend to value rare or scarce items more than common items (Verhallen & Robben, 1994). Thus, all else being equal, people might not only prefer those who share rare knowledge with them, but also those who share rare affinities with them in the absence of prior knowledge. Yet preferences—a stable liking of some activities or items over others—are often the result of a complex chain of events. In order to come to like Tintoretto, one must have had experiences that led one to discover his work in the first place, perhaps by studying Renaissance art, living in Venice, or surrounding oneself with people who frequent art museums. Thus, preferences reflect various aspects of one's cultural knowledge and past experiences. It is possible that the presence of shared knowledge (in the absence of explicit information about affinity) may be sufficient to give rise to the rare preference effect.

In fact, other people's knowledge and affinity can provide qualitatively different information about them. Knowledge is often a reflection of one's current interests, prior background, and social history. This may be especially true for relatively rare or obscure preferences for music, hobbies, or activities; someone who came to learn about Tintoretto presumably sought out particular kinds of artists or was close to people in that niche, whereas someone who came to know about da Vinci could have learned about the artist through many different channels. Thus, shared knowledge between two individuals, when it is rare, can be a good indicator of a broader, meaningful common ground (Clark, Schreuder, & Buttrick, 1983) or even signal the presence of a latent social group to which both individuals belong (Gershman, Pouncy, & Gweon, 2017). On the other hand, affinity—as operationalized here—is separated from a person's prior background or social history. While knowledge in a particular domain implies an active effort to learn about it, affinity reflects a predisposition to find something attractive even without having deliberately sought it out. Although discovering a shared affinity with someone can be just as delightful as discovering shared knowledge, it may not necessarily support further inferences about shared cultural background or social history.

Prior work has found empirical support for the privileged status of shared knowledge over shared affinity in interpersonal relationships. Adults and children prefer people who share personally relevant beliefs with them over those who share arbitrary beliefs (Heiphetz, Spelke, & Banaji, 2014). Preschool-aged children prefer peers who share their knowledge (but not preferences) over peers who share their preferences (but not knowledge) (Soley & Spelke, 2016). Further, children selectively attribute shared knowledge—but not shared preferences—to members of the same cultural group (Soley & Aldan, 2018). These studies suggest that both adults and children readily represent different aspects of preferences that can be shared between individuals (i.e., prior knowledge and spontaneous affinity) and prioritize shared knowledge. These results provide indirect support for our hypothesis that the presence of shared knowledge between individuals may drive the rare preference effect more strongly than shared affinity in the absence of prior knowledge.

The present work explores how adults use statistical information about the prevalence of preferences to choose with whom they would rather interact. We first provide initial evidence for the existence of the rare preference effect (Experiments 1 & 2). We then test the hypothesis that this effect is driven more strongly by the presence of shared knowledge than by shared affinity (Experiments 3 & 4). Finally, Experiment 5 provides further support for the importance of shared background knowledge in people's interpersonal decisions, by asking whether the means by which agents acquired a preference can influence people's decisions, even when rarity is held constant. Our experiments provide converging evidence from two complementary paradigms. In one, we ask participants to provide statistical information about the prevalence of their own preferences (Experiments 1, 3, 5). In the other, we take advantage of a minimal, rigorously controlled context where we provide explicit statistical information about the prevalence of novel preferences in a novel population (Experiments 2, 4).[2]

2. Experiment 1

In Experiment 1, we asked participants to list their favorite bands, books, or movies and to estimate the prevalence of people's preferences for these items. When given the choice between two potential social partners, we predicted that participants would favor an agent who shares a rarer preference with them over someone who shares a more common preference, and that this tendency would increase as the relative rarity of the preference increased.

2.1. Methods

2.1.1. Participants

332 MTurk users participated in an online survey. We report results pooled across two samples: an original experiment (N = 147) and preregistered replication (N = 185; http://osf.io/2tajm/). Results were qualitatively similar across both samples (but see SI 1.2. for separate analyses by sample).[3]

2.1.2. Procedure

First, participants chose one of three activities (watching movies, listening to music, or reading books) and were asked to list their five favorite movies, bands, or books, respectively. Participants were then prompted to estimate the popularity of their favorite items as follows: "If you were to ask 100 strangers about each of these [movies/bands/books], how many people would say that they like them?" Participants indicated their response in two ways, as shown in Fig. 1A: (1) by rank-ordering each item from most to least widely-liked and (2) by providing a numerical estimate of how many of 100 strangers would like each item. An additional group of participants (N = 18) provided inconsistent responses (i.e., by providing a higher numerical estimate for the least widely-liked item than for the most widely-liked item) and were excluded from analysis.

Finally, participants were given a choice between two potential social partners and were asked which agent they would rather talk to. As shown in Fig. 1A, each agent was presented as a chat avatar with a prompt above it saying "I like X!," where X was one of the five items the participant had listed. In order to maximize the difference in the prevalence of the two agents' preferences, one agent (*Common Agent*) always liked the most widely-liked item; the other agent (*Rare Agent*) always liked the least widely-liked item.[4] In this and all subsequent first-person experiments (Experiments 1, 3, & 5), the left/right positioning of Rare and Common Agents was randomized across participants.

2.2. Results

Participants chose the Rare Agent equally often regardless of whether they listed their favorite books, music, or movies (proportion choosing rare by domain: $\chi^2(2) = 3.5$, $p = .17$); we thus collapsed across domains in subsequent analyses.

[2] Materials, data, and analysis scripts for this paper are available at: osf.io/9rb35/.

[3] In the original sample, we aimed for a large sample that is representative of online studies with adults. In subsequent experiments using the first-person paradigm (Experiments 3, 5), we aimed for a consistent sample size of N = 150 per condition. In the preregistered replication, we ran 200 participants (before exclusions) to replicate the effect at 85% power.

[4] A separate group of participants was given a choice between two agents who each liked the first or second most widely-liked items. Because the difference in the prevalence of the two agents' preferences was small, participants in this group did not prefer the agent with the rarer preference; see SI 1.1.

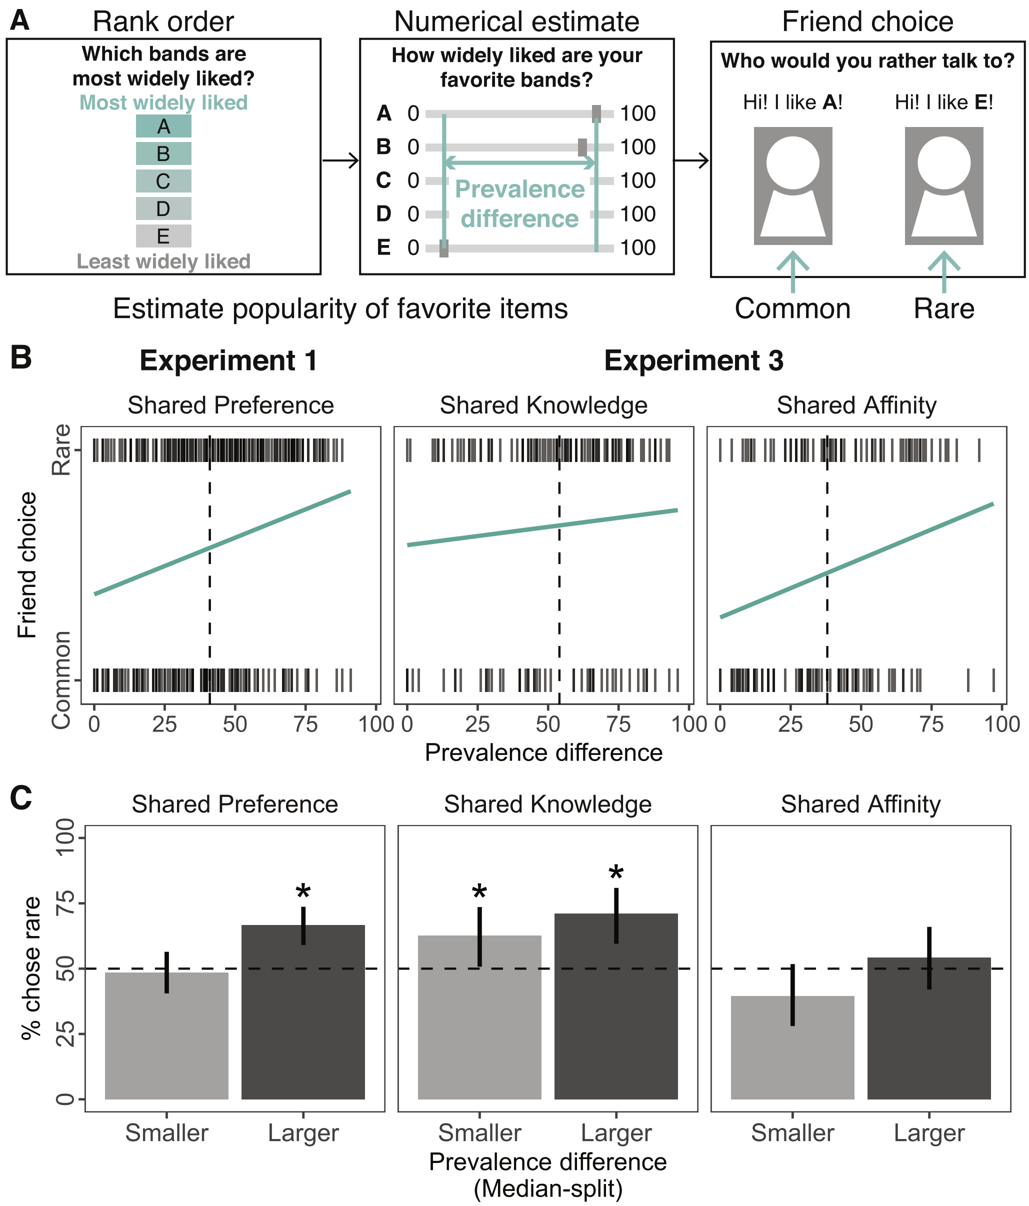

Fig. 1. Experiments 1 & 3 methods and results. (A) Experiment 1 procedure: Participants entered their five favorite bands, books, or movies, then rank ordered items from most to least widely liked and estimated how many of 100 people would like each item. Participants then chose whether they would rather talk to an agent who liked the most (Common Agent) or least (Rare Agent) widely liked item. Experiment 3 used the same procedure with different prompts. (B, C) Relationship between prevalence difference and friend choice. (B) Rug plots show the prevalence difference and friend choice (top = Rare; bottom = Common) of each participant; each tick marks one participant. Solid line shows logistic regression fit; dashed line marks median prevalence difference by condition. (C) Participants were median split by prevalence difference, where "Smaller" includes participants with scores below or at the median, and "Larger" above the median. * show deviations from chance (binomial test, $p < .05$), error bars show 95% CI.

Overall, participants chose the Rare Agent more frequently than the Common Agent (57.8%, two-tailed binomial test: 95% CI = [52.3, 63.2], $p = .004$). We tested the effect of relative rarity using a logistic regression, with the difference in numerical estimates between the most- and least-widely liked items (henceforth "prevalence difference") as a predictor. As expected, participants chose the Rare Agent more often when there was a larger prevalence difference between the two agents' preferences ($\beta = 0.02$, $z = 3.8$, $p = .0001$; see Fig. 1B). We also median-split participants based on prevalence differences. Participants with small prevalence differences were given the choice between two agents whose preferences are similarly prevalent; accordingly, participants in this group did not prefer the Rare Agent (48.4%, 95% CI = [40.5, 56.4], $p = .8$). By contrast, those with larger prevalence differences tended to choose the Rare Agent (66.7%, 95% CI = [59.1, 73.7], $p < .001$; difference between groups: $\chi^2(1) = 10.55$, $p = .001$; Fig. 1C).

In Experiment 1, we explored whether participants' beliefs about the rarity of their own preferences influenced with whom they would rather interact. Our results provide initial evidence for the rare preference effect: Overall, participants preferred the agent who shared a rare preference with them, and this effect was modulated by the relative rarity of the preference. We followed up on this effect in two supplementary experiments. First, we verified that participants' estimates of the prevalence of preferences for their favorite bands align with the bands' popularity on an online music recommendation platform, which provides an independent measure of real-world popularity (SI 2.1.). Second, we replicated the rare preference effect using a continuous measure that allows participants to express a graded preference for, or even indifference between, the two agents (SI 2.2.).

However, because we harnessed people's actual preferences and their own estimates of prevalence, we cannot rule out the possibility that these estimates are confounded with other factors. For instance, people might have chosen the agent who shares their rare preference not because of its rarity but because more socially desirable or strongly held preferences tend to be rare. In Experiment 2, we provide a conceptual replication of the rare preference effect to address this concern and further investigate the scope of this effect.

3. Experiment 2

In Experiment 2, participants made third-party judgments where they introduced potential friends to a target agent, based on the agent's preferences for novel items and on explicitly provided information about how prevalent preferences for these items are among a novel population. We manipulated the prevalence of preferences for each novel item (between subjects) and the preferences of the target agent (within subjects). In this task, participants only had access to information about the prevalence of preferences, without prior knowledge of how socially desirable or strongly held these preferences are. This approach allows us to ask whether the rare preference effect reflects participants' use of statistical information, rather than the effect of other factors that may be confounded with rarity.

Another goal of Experiment 2 was to test whether participants indiscriminately prefer agents who have rare preferences, or whether they instead use statistical information flexibly depending on the context. To this end, we designed four different trial types. As in Experiment 1, we predicted that, when an agent likes both a rare and common game ("Both" trial), participants would introduce that agent to a potential friend who likes the rare game; thus, we aimed to replicate the rare preference effect in a complimentary paradigm. By contrast, in trials where the target agent's preferences are unknown ("Unstated" trial), we expected that participants would attempt to maximize the chances that the target agent would share a preference with the chosen friend and thus choose the friend who has a common preference. Finally, in trials where the target clearly prefers either the common or the rare game ("Common" and "Rare" trials), we predicted that participants would choose the friend who has the same preference, regardless of its prevalence.

3.1. Methods

3.1.1. Participants

Amazon MTurk users ($N = 1142$, average 127 per condition) were randomly assigned to one of nine between-subjects conditions (see Procedure). As in Experiment 1, participants were pooled from two samples: $N = 692$ from the original experiment and $N = 450$ from a replication (preregistration: http://osf.io/2tajm/; see also SI 1.2.). [5]

3.1.2. Procedure

Participants were introduced to novel agents called "Gazorps," who were all polled about whether they liked to play two novel games, "wumbus" and "jibboo". Fig. 2A demonstrates how the results of the poll were presented to participants: Participants first saw an array of 100 grey silhouettes of Gazorps, representing the entire population. On one screen, a random subset of the Gazorps turned purple, to indicate the popularity of wumbus; on the other, a random subset of the Gazorps turned blue, to indicate the popularity of jibboo (see SI 1.4. for additional information regarding this random assignment). We will refer to the more popular of the two games as the "common game", and the least popular as the "rare game"; for example, in Fig. 2A, wumbus is the common game. Participants were randomly assigned to one of nine conditions, which differed in the number of Gazorps who liked each game: 45-45 (baseline), 50-40, 55-35, 60-30, 65-25, 70-20, 75-15, 80-10, or 85-5. The difference between the two numbers was thus analogous to prevalence differences in Experiment 1, representing the difference in popularity between the two games. Table S1 provides sample sizes for each condition.

In the critical test trials, participants were told that they would introduce a shy Gazorp (hence the *target Gazorp*) to one of two potential friends: a Gazorp who likes the common game (henceforth Common Agent), or a Gazorp who likes the rare game (henceforth Rare Agent). Fig. 2B shows the layout of a typical test trial: The target Gazorp (top) was presented alongside the two potential friends (bottom). Each Gazorp's preferences were clearly indicated by displaying a verbal prompt above the Gazorp (e.g, "I like to wumbus and jibboo!") and by overlaying the logo(s) of the Gazorp's preferred game(s) on the image of the Gazorp. The poll results were displayed above each potential friend as a reminder to participants. Participants indicated their response by clicking on a potential friend.

Each participant completed four within-subjects test trials, which differed in the preferences of the target Gazorp (Fig. 2C). In the *Common* and *Rare* trials, the target Gazorp expressed a preference for just one of the games: the commonly liked game (Common) or the rarely liked game (Rare). In the *Both* trial, the target Gazorp expressed a preference for both games. Finally, in the *Unstated* trial, participants were given no information about the target Gazorp's preferences. Trial order was pseudo-randomized, and all superficial aspects of the task were randomized between subjects, including the names of the target Gazorps, the logo and color associated with each game, and the positions of the potential friends in the test trial (see SI 1.2).

The goals of these trial types were twofold: First, to replicate the rare preference effect in the Both trial, and second, to test whether participants' use of statistical information is flexible and context dependent across all trials. The Common and Rare trials represent cases where statistical information should have little effect; we predicted that participants would show a strong tendency to match the target Gazorp with the friend who shares their preference, regardless of its prevalence. The Both trial is the most similar to Experiment 1: The target Gazorp likes both games, just as participants in Experiment 1 shared a preference with both the Rare and Common Agents. Accordingly, we predicted that participants would pair the target with the Rare Agent. Critically, if the results in the Both trial reflect a global bias to favor agents who have rare preferences, participants should choose this agent even in the Unstated trial. However, we hypothesized that participants would make the opposite choice in the Unstated trial (i.e., choose the friend who has the common preference) to maximize the chances that the target Gazorp will have something in common with their potential friend.

3.2. Results

We first examined the 45-45 condition, where both games were equally popular ($N = 135$); to match the terminology used in other conditions, we'll refer to wumbus as the "common" game and jibboo as the "rare" game. In the Common and Rare trials, participants in this condition consistently matched the target to the friend who liked the same game (Common trial: 11.9% chose rare, 95% CI: [6.9, 18.5]; Rare trial: 86.7% chose rare, 95% CI: [79.7, 91.9]). Conversely, and as predicted, participants did not systematically prefer either friend in the Both and Unstated trials (see SI 1.3.); these results suggest that participants' choices were not systematically influenced by superficial aspects of the task. Below, we report results from all conditions except 45-45 ($N = 1007$).

Across all remaining conditions, participants matched the target in the Common and Rare trials with the potential friend who shared the same preference (Common trial: 7.8% chose rare, 95% CI: [6.3, 8.9]; Rare trial: 90.8% chose rare, 95% CI: [88.9, 92.4]). These effects were significant in each individual condition after correcting for multiple comparisons (Fig. S3). Further, the difference in popularity between the two games did not influence participants' tendency to choose the Rare Agent on either of these trials (logistic regression, chose rare ~ prevalence difference; Common: $\beta = -0.005$, $z = -1.1$, $p = .2$; Rare: $\beta = -0.001$, $z = -0.3$, $p = .8$). Thus, when the target liked a single game, participants' judgments were driven by the target's

[5] Our goal in this experiment was to collect 125 participants per condition (75/condition in the original sample, 50/condition in the replication sample). In Experiment 4, we aimed for a consistent sample size of 100 participants/condition.

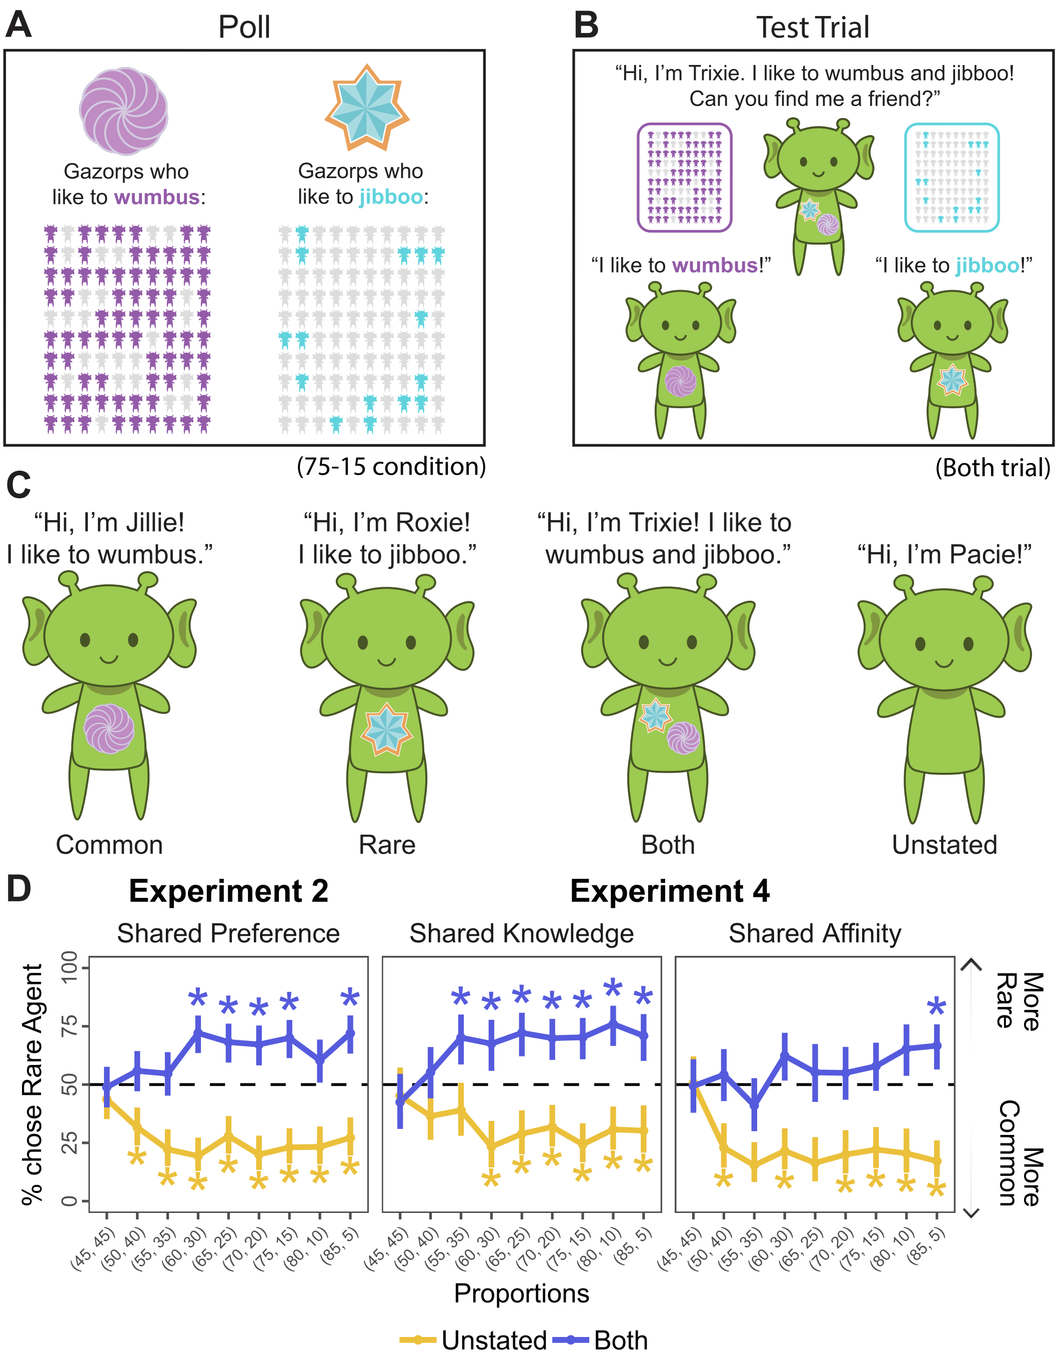

Fig. 2. Experiment 2, 4 methods and results. (A–C) Experiment 2 procedure. (A) Poll: participants saw how many Gazorps like to wumbus and to jibboo. (B) Test trials: Participants matched a target Gazorp (top) with a potential friend who likes the common or the rare game (bottom). (C) Trial types: All participants completed four test trials, which differed in the target Gazorp's preferences. Experiment 4 used the same procedure with different prompts (see Experiment 4 methods). (D) Friend choice based on the relative popularity of the two games (x-axis) and the target's preferences (hue). Error bars deNote 95% CI; asterisks denote significant Bonferroni-corrected contrasts ($p < .05/9$, binomial test).

preferences, while disregarding statistical information.

By contrast, in the Both and Unstated trials, we expected to see participants use prevalence information flexibly to inform their judgments. In the Both trials, we replicated the rare preference effect; most participants paired the target Gazorp who liked both games with the Rare Agent (65.1%, 95%*CI*: [62.1, 68.1], $p < .0001$). Consistent with Experiment 1, this effect was also modulated by the prevalence difference (i.e., condition): participants tended to choose the Rare Agent more often as the prevalence difference increased (logistic regression: $\beta = 0.007$, $z = 2.5$, $p = .01$). By contrast, in the Unstated trial, we found a striking reversal of the rare preference effect: only 24.4% of participants chose the Rare Agent across all conditions (95% CI: [21.8, 27.2], $p < .0001$), and participants' responses did not shift as the prevalence difference increased (logistic regression: $\beta = -0.001$, $z = -0.5$, $p = .6$).

In sum, we replicated the rare preference effect in a complementary task and found that this effect is more than a simple, indiscriminate bias in favor of rare preferences. Given no information about the target's preference, participants used statistical information to match the target with the agent with the most widespread preference; given a clear, single match for the target, participants ignored statistical information altogether. Only when the target liked both games did participants match the target with the agent who held the rare preference, and their tendency to do so increased as the rarity became more extreme. These findings suggest that people's use of statistical information in social affiliation judgments is flexible, selective, and context dependent.

4. Experiment 3

In Experiments 3 and 4, we further explored the nature of the representations that underlie the rare preference effect. We hypothesized that, although preferences provide information about what a person knows (knowledge) and what they like (affinity), the rare preference effect may be driven more strongly by shared knowledge than by shared affinity. In Experiment 3, we adapted the first-person paradigm in Experiment 1 to test whether participants prioritize shared rare knowledge (without an explicitly stated preference) over shared common knowledge, as well as shared rare affinity (without prior knowledge) over shared common affinity.

4.1. Methods

4.1.1. Participants

Amazon MTurk users were randomly assigned to one of two between-subjects conditions: the Shared Knowledge condition (N = 151) and the Shared Affinity condition (N = 143). An additional 8 participants (Shared Knowledge: N = 1; Shared Affinity: N = 7) were excluded from analysis due to inconsistent responses (see Experiment 1 Procedure).

4.1.2. Procedure

The procedure was identical to Experiment 1, except for a minimal change in the prompts. Participants in the Shared Knowledge condition were asked how *widely known* each item is (i.e., how many of 100 people would *know about* each of their five favorite [movies/bands/books]), and were given a choice between two agents who each declared knowledge of one of these items (i.e., "I know about X!"). By contrast, participants in the Shared Affinity condition were asked how *likable* each item is (i.e., how many of 100 people would like each item if given a chance to watch, listen to, or read it for the first time) and chose between two agents who each declared, "I just watched/read/heard X. I like it!"—thus expressing that they like the item without having prior knowledge of it.[6]

4.2. Results

We found no differences in participants' responses by domain (Shared Knowledge: $\chi^2(2) = 2.1$, $p = .4$, Shared Affinity: $\chi^2(2) = 0.13$, $p = .9$). Therefore, we collapsed across domains in the analyses below.

In the Shared Knowledge condition, 66.9% of participants chose the agent who shared their knowledge about the less widely-known item (two-tailed binomial test: 95% CI: [58.8, 74.3], $p < .0001$). Unlike in Experiment 1, participants were not more likely to choose the Rare Agent as the prevalence difference increased (logistic regression, chose rare ~ prevalence difference: $\beta = 0.007$, $z = 1.0$, $p = .3$). A median-split analysis revealed that participants tended to choose the agent who shared rare knowledge with them, both when the prevalence difference was small (below median: 62.7%, 95% CI: [50.7, 73.6], $p = .04$) or large (above median: 71.0%, 95% CI: [59.5, 80.9], $p < .001$). Numerically, more participants with large prevalence differences chose the Rare Agent; however, the difference between the two median-split groups was not significant ($\chi^2(1) = 0.85$, $p = .4$).

By contrast, in the Shared Affinity condition, participants overall did not prefer the agent who shared the rarer affinity (46.8%; two-tailed binomial test: 95% CI: [38.4, 55.3], $p = .5$). A median-split analysis revealed a very different pattern from the Shared Knowledge condition. Participants did not show a preference for either agent when the prevalence difference was small (39.4%, 95% CI = [28.0, 51.7], $p = .09$) or large (54.2%, 95% CI = [42.0, 66.0], $p = .6$, difference between groups: $\chi^2(1) = 2.55$, $p = .11$). Nevertheless, participants were more likely to choose the Rare Agent as the prevalence difference in the two agents' affinity increased (logistic regression: $\beta = 0.02$, $z = 2.7$, $p = .006$). Overall, participants' preference for the Rare Agent was stronger in the Shared Knowledge than in the Shared Affinity condition (Fig. 1B-C; difference between conditions: $\chi^2(1) = 11.2$, $p = .0008$).

These results suggest that the rare preference effect is driven more strongly by shared knowledge than by shared affinity and, in fact, that shared knowledge is sufficient to give rise to this effect. Note that participants' choices were influenced by prevalence differences in the Shared Affinity condition, suggesting that participants did not ignore the prevalence information altogether; rather, they used it differently than participants in the Shared Knowledge condition. By contrast, the influence of the relative magnitude of the prevalence difference was not obvious in the Shared Knowledge condition, mainly because participants preferred the agent who shared relatively rare knowledge with them even when the difference in the prevalence of the two agents' knowledge was small.

5. Experiment 4

As above, we adapted the third-party paradigm in Experiment 2 to test whether participants use statistical information flexibly across contexts, and whether these effects manifest when agents share knowledge (without a stated preference) or affinity (without prior knowledge).

5.1. Methods

5.1.1. Participants

Amazon MTurk users were randomly assigned to two groups: Shared Knowledge (N = 823) and Shared Affinity (N = 752). As in Experiment 2, each group was further subdivided into nine between-subjects conditions (67-113 participants per condition; Table S1). An additional 228 participants (Shared Knowledge: 84; Shared Affinity: 144) were excluded after failing a manipulation check; including these participants does not change the interpretation of our results (see SI 1.7. and Fig. S5).

5.1.2. Procedure

The procedure was identical to that of Experiment 3, except for small changes in the prompts. Participants in the Shared Knowledge group were shown how many Gazorps *knew* how to play each game. In the test trials, they paired a target Gazorp who expressed its knowledge (e.g., "I know how to wumbus and jibboo!") with a Gazorp who knew the common or rare game. By contrast, in the Shared Affinity group, participants were told that all Gazorps were playing both games for the first time, and the poll showed how many Gazorps liked each game after trying it. In the test trials, participants were told about the target Gazorp's affinity (e.g. "I just tried wumbus and jibboo, and I like them!") and were given the choice between two agents who each liked the common or rare game. As in Experiment 3, participants completed four test trials: Common, Rare, Both, and Unstated.

5.2. Results

We first analyzed responses in the 45-45 condition (Shared

[6] A separate group of participants was given a choice between two agents who knew or liked the first and second most widely-known or likable items. As in Experiment 1, participants in this group did not prefer the agent with the rarer knowledge or affinity; see SI 1.5.

Knowledge: N = 73; Shared Affinity: N = 79); as in Experiment 2, we designated jibboo as the "rare" game. Participants reliably matched the target with the friend who knew or liked the same game in both the Common (Shared Knowledge: 21.9% chose rare, 95% CI: [13.1, 33.1]; Shared Affinity: 16.5% chose rare, 95% CI: [9.1, 26.5]) and Rare trials (Shared Knowledge: 74.0% chose rare, 95% CI: [62.4, 83.5]; Shared Affinity: 82.2% chose rare, 95% CI: [72.0, 80.0]; all $ps < 0.0001$). In the Unstated and Both trials, participants did not systematically prefer either friend (see SI 1.6.). Below, we report results in all conditions except 45-45 (Shared Knowledge: N = 750; Shared Affinity: N = 673).

When the target knew or liked just one game, participants ignored statistical information and matched the target agent with the agent who knew or liked the same game, in both the Common (Shared Knowledge: 20.5% chose rare, 95% CI: [17.6, 23.5], Shared Affinity: 15.4% chose rare, 95% CI: [12.8, 18.4]) and Rare trials (Shared Knowledge: 81.3% chose rare, 95% CI: [78.4, 84.1]; Shared Affinity: 85.6% chose rare, 95% CI: [82.7, 88.1], all $ps < .0001$).

In the Both trial, 69.3% of participants in the Shared Knowledge group chose the Rare Agent (95% CI: [65.9, 72.6], $p < .0001$). Participants chose the Rare Agent in almost all conditions except when the prevalence difference was small (i.e., 50-40; Fig. 2D); overall, participants chose the Rare Agent more often as the prevalence difference increased (logistic regression: $\beta = 0.008$, $z = 2.3$, $p = .02$). In the Shared Affinity group, 57.6% of participants chose the Rare Agent (95% CI: [53.8, 61.4], $p < .0001$); however, unlike in the Shared Knowledge condition, this effect was observed *only* in the condition with the highest prevalence difference (i.e., 85-5). Overall, participants in the Shared Affinity condition chose the Rare Agent more often as the prevalence difference increased ($\beta = 0.009$, $z = 2.7$, $p = .006$). Participants chose the Rare Agent more often in the Shared Knowledge group than in the Shared Affinity group (chi-square test: $\chi^2(1) = 20.45$, $p < .0001$).

Finally, in the Unstated trial, we predicted that participants would favor the Common Agent. This effect was observed in both the Shared Knowledge group (30.4% Rare, 95% CI: [27.1, 33.8], $p < .0001$) and in the Shared Affinity group (19.6% Rare, 95% CI: [16.7, 22.8], $p < .0001$), though participants in the Shared Affinity group chose the Common Agent more often ($\chi^2(1) = 21.27$, $p < .0001$). Participants' tendency to choose the Common Agent did not change as the difference between the prevalence of the agents' knowledge or affinity increased (logistic regression, Shared Knowledge: $\beta = -0.004$, $z = -1.23$, $p = .22$; Shared Affinity: $\beta = -0.001$, $z = -0.2$, $p = .8$).

Consistent with the findings in Experiment 3, in Experiment 4 we find further evidence that participants privilege shared knowledge in deciding with whom to interact, and increasingly so as it becomes rarer. Furthermore, we identified two notable patterns of results that could not be gleaned through first-party judgments. First, by manipulating the knowledge and affinity of the target agent, we found that participants did not simply ignore affinity; rather, they used information about both knowledge and affinity flexibly depending on the context. Second, by systematically manipulating the prevalence of agents' knowledge and affinity, we found one edge case where participants do pair agents who share rare affinity—namely, when it is extremely rare.

However, the effects of shared knowledge and affinity are difficult to isolate entirely. In the Shared Affinity condition, we clearly communicated the presence of affinity in the absence of knowledge, as all agents tried each game for the first time. However, in the Shared Knowledge condition, it is more challenging to communicate the presence of knowledge in the absence of affinity, as the absence of affinity could be defined as either indifference or active dislike. Thus, we left the agents' affinity ambiguous by omitting any reference to their enjoyment of the games.

Together, the results from Experiments 3 and 4 suggest that the presence of shared knowledge is sufficient to give rise to the rare preference effect, perhaps because shared knowledge is a better indicator of broader similarities between agents' past experiences that led them to acquire that knowledge. If so, then directly providing information about how people acquired a preference should influence people's judgments, regardless of rarity. In Experiment 5, we test this idea by presenting the choice between two agents who like the same item but differ in how they came to know about it.

6. Experiment 5

In Experiment 5, we adapted the first-person paradigm in Experiments 1 & 3 to test whether participants are sensitive to the causal process by which people arrived at a preference. In this task, participants were given the choice between two agents who shared the same music preference with the participant. Critically, one agent arrived at the preference by deliberately seeking out similar songs; thus, their preference is potentially diagnostic of other, hidden similarities. By contrast, the other agent arrived at their preference accidentally. We predicted that people would be more likely to interact with the agent who deliberately acquired the shared preference for a song over the agent who came to like it accidentally.

6.1. Methods

6.1.1. Participants

Amazon MTurk users were randomly assigned to two between-subject conditions: Common (N = 142) and Rare (N = 158).

6.1.2. Procedure

As in Experiments 1 & 3, participants completed a survey about their own preferences. In this version of the task, participants were always asked about their favorite songs. As in Experiment 1, participants listed their five favorite songs, rank-ordered each song from most to least widely liked, and estimated how many of 100 people would like each song.

Finally, participants chose which of two agents they would rather talk to. Unlike previous experiments, where each agent liked or knew about different items, both agents in this experiment shared the participant's preference for the same song. In two between-subjects conditions, we varied whether both agents liked the most widely-liked item (Common condition) or the least widely-liked item (Rare condition). Both agents indicated their preference with a prompt above their chat avatar that read "I like X!", where X was filled in with the same song for both agents. Critically, each agent had a different prompt below them indicating how they had arrived at their preference. The prompt below the *Deliberate* agent read: "This person found out about it by: searching for similar songs.". The prompt below the *Accidental* agent read: "This person found out about it by: overhearing it at a store.". However, note that participants were never explicitly told that one agent's preference is "deliberate" and the other's is "accidental."

6.2. Results

In this experiment, we asked whether adults prefer to affiliate with people who arrive at a shared preference deliberately, rather than accidentally. Indeed, across both groups, 68.3% of participants (95% CI: [62.7, 73.6]; $p < .0001$, binomial test) preferred the friend who arrived at a shared preference deliberately, rather than the one who arrived at the shared preference accidentally. These results were consistent regardless of the rarity of the preference (Rare: 69.0% Deliberate, 95% CI: [61.1, 76.1]; Common: 67.6% Deliberate, 95% CI: [59.2, 75.2]; difference between conditions: $\chi^2(1) = 0.02$, $p = .9$).

In Experiment 5, we find that people's social affiliation judgments are sensitive not only to the presence of shared knowledge, but also by *how* people acquired that knowledge. Rare preferences may be strong indicators of potential compatibility because they provide information about one's cultural context or imply deliberate attempts to seek out knowledge. In support of this idea, we find that people do prefer those who arrived at a preference deliberately, regardless of its rarity.

However, we acknowledge that this manipulation may imply differences in other aspects of the agents' preferences, such as strength and duration. Together with Experiments 3 and 4, these results suggest that, among the many properties that may be shared between potential social partners, people prioritize those that are more indicative of shared knowledge or past experiences in social affiliation judgments.

7. General discussion

The current work presents support for the hypothesis that people use statistical information about preferences to guide their social decisions, especially when such information signals the presence of a broader shared background. We first experimentally demonstrated the *rare preference effect*, confirming the intuition that discovering a shared preference with someone is more meaningful when it is rare than when it is relatively common. This effect is found in both first-person and third-person judgments, regardless of whether the statistical information was self-generated (Experiments 1 & 3) or explicitly provided (Experiments 2 & 4).

Importantly, we also find that the rare preference effect reflects more than a tendency to indiscriminately favor agents with rare preferences. The current study demonstrates this in three ways. First, participants did not globally favor social partners who have rare preferences; rather, they used information about the prevalence of preferences flexibly, based on the preferences of the target agent (Experiments 2 & 4). Second, although a shared preference implies both knowing and liking the same thing, the rare preference effect was stronger in the presence of shared knowledge than in the presence of shared affinity (Experiments 3 & 4). Finally, even when rarity is held constant, people prefer social partners who arrived at a shared preference deliberately, rather than accidentally (Experiment 5). Taken together, our results suggest that participants represent statistics about what others know and like and that they actively use these representations to guide their interactions with others.

The current results highlight the importance of shared knowledge in forming initial bonds between people (see also Soley & Aldan, 2018; Soley & Spelke, 2016). More specifically, the presence of shared knowledge was sufficient to give rise to the rare preference effect, even in the absence of any explicit information about shared affinity. One possible explanation for these findings is that shared preferences are perceived as an indicator of shared cultural background or prior experiences between individuals. However, this account does not preclude the possibility that shared affinity may also be informative; in fact, we found a similar effect when shared affinity was extremely rare. It may be the case that, as with shared knowledge, a predisposition to like something that is rarely liked by others also indicates some prior experiences that have led the person to like it even upon a single exposure. Another (not mutually exclusive) possibility is that this effect may emerge from the inference that someone who has such an extremely rare tendency is "strange" and might be better matched with someone who shares that tendency.

Our results highlight that humans can harness statistical information about the prevalence of everyday events in flexible, nuanced ways to guide their judgments about social affiliation. Yet we remain cautious about making broader claims about how strongly rarity drives the formation of real-world social relationships. First, in our tasks, participants were either asked to estimate the prevalence of their own preferences or were directly provided information about prevalence. It remains an open question whether people might spontaneously use prevalence to guide their judgments in the absence of explicit prompting. Second, we note that our effect sizes in the first-person paradigm are quite small, particularly in a replication using a continuous measure (see SI 2.2). While the first-person paradigm is more ecologically valid, participants' favorite items likely differ along other dimensions that may also influence their friend choices, such as how strongly they like each item. We thus complement these findings with third person-paradigms that isolate the effect of rarity. Finally, our tasks measured which social partners participants would rather approach (e.g., to whom they would rather talk), rather than who would be best suited for a long-term relationship. We acknowledge that rarity is only one of many factors that may signal initial compatibility, and that its impact on real-world relationships may be difficult to identify.

The current work bridges two fields that have largely developed separately from one another. On one hand, a large body of work in cognitive science has demonstrated both limitations (e.g., Kahneman & Tversky, 1973; Monin & Norton, 2003; Ross et al., 1977; Suls & Wan, 1987; Tversky & Kahneman, 1983) and competencies (e.g., Griffiths & Tenenbaum, 2006; Gweon et al., 2010; Xu & Garcia, 2008) in humans' abilities to estimate, represent, and systematically use statistical information. On the other hand, prior work in social psychology has investigated the role of similarity in behaviors, traits, and cultural backgrounds in social relationships and friendship formation (Byrne & Nelson, 1965; McPherson, Smith-Lovin, & Cook, 2001; Selfhout et al., 2009; Werner & Parmelee, 1979). Our results identify an important, yet unexplored dimension in friendships and mutual attraction: The statistical distribution of these similarities. Given that the personal significance of shared preferences is modulated by statistical information, it is possible that shared personality traits, languages, or cultural backgrounds are also interpreted differently depending on their distributions at the population level. We hope the current work can inspire future studies on how the basic cognitive abilities to represent and reason about the prevalence of abstract events can shape social relationships from early in life to adulthood.

Note that our results reflect people's predictions about hypothetical compatibility, rather than actual compatibility. Indeed, people often become friends despite having markedly different preferences or backgrounds. This naturally invites further questions about *why* discovering a shared preference with someone is so delightful (especially when it is rare), and why shared knowledge might be considered more diagnostic of future friendships. We suggest that the experience of encountering someone who shares a rare preference elicits further inferences about the causal mechanisms that might have generated this coincidence. More specifically, the fact that two individuals share a preference and, in particular, know about the same thing, may be reliable indicators of similar personal histories or cultural backgrounds that allowed both individuals to acquire that knowledge. There may be many possible paths by which a person might come to have a common preference; in order to have a relatively rare preference, one might have to gain deeper knowledge about a particular genre, actively explore new music, and even seek out others who know and like similar music. Thus, it is possible that rare preferences are more meaningful because the set of people who hold a rare preference are perceived as a more coherent social category than people who hold a common preference.

Indeed, past work suggests that people prioritize more coherent categories in inductive inferences (Murphy & Medin, 1985; Nguyen & Chevalier, 2015; Rehder & Hastie, 2004). For example, when given the task of buying a present for someone who likes tennis and skydiving, people tend to choose a gift related to skydiving (Nelson & Miller, 1995). However, this work has largely relied on people's knowledge of real-world categories; skydivers are seen as both a smaller and more coherent group than tennis players (Patalano, Chin-Parker, & Ross, 2006). Here, we find that people prioritize rare preferences even in a minimal context, where they have no prior knowledge of the items or about the agents who like them. Our results suggest that rarity alone may guide social judgments, even in the absence of information about category coherence. Moving forward, our work also opens two potential avenues for future work on how people harness statistical information to learn about social categories. First, statistical properties such as rarity may guide learning about novel social categories; one might consider using the third-person paradigm developed here to test whether groups that share a rare preference are perceived as more coherent than groups that share a common preference, even in the absence of relevant prior knowledge. On the other hand, category coherence may place constraints on the generality of the rare preference effect. In the present work, we focused on domains that have been found to guide impression formation and social relationships in prior work, such as music and games (Rentfrow & Gosling, 2006; Selfhout et al., 2009; Werner & Parmelee, 1979), and we found no differences across these domains. By contrast, it is possible that the effect would be weaker in domains of preferences that are less inductively rich or less strongly held.

Note that the rare preference effect may not be a phenomenon that manifests uniformly across cultures, social contexts, or age groups. First, it is possible that this phenomenon is specific to a particular cultural context that places a premium on distinctiveness and individuality (Kim & Markus, 1999); for instance, in cultures where people place high value on things that are commonly liked, the rare preference effect may be less obvious or even absent. This is not inconsistent with our findings. Our results suggest that people do not indiscriminately favor rare overlaps, but rather make rich, sophisticated inferences based on the nature of the observed similarities. If rare preferences are valuable because they are predictive of other similarities, then they should not be more attractive than common preferences in a context where they are not diagnostic. Thus the tendency to prefer those who share rare preferences may depend on people's beliefs about how various preferences are distributed in their cultural context and the degree to which their own preferences are influenced by others.

Second, the same shared preference may be interpreted differently depending on one's current social environment that provides the basis for estimating its prevalence. Preferences that are common in one context may be rare in another; for example, in Puerto Rico, meeting someone who likes Cultura Profética—a Puerto Rican band—may not necessarily be more surprising or meaningful than meeting someone who likes The Beatles, but it would be a very rare occasion in San Francisco. Thus, the impact of a shared preference may change based on the social environment where it is encountered.

Finally, it is possible that this effect is constrained to a particular point in life. Given that young children are adept statistical reasoners (Gweon et al., 2010; Kushnir et al., 2010; Xu & Garcia, 2008) and are partial to those who share their own preferences (Fawcett & Markson, 2010a, 2010b), it is possible that the rare preference effect arises in childhood. However, even if the cognitive prerequisites for the rare preference effect are present early in life, rare preferences may not be a meaningful indicator of potential compatibility to children, who have relatively less autonomy to explore their preferences. Thus, this effect might instead manifest in adolescence or adulthood, when preferences are treated as distinctive markers of identity (North & Hargreaves, 1999; Rentfrow & Gosling, 2006).

In sum, our study provides a compelling demonstration of the *rare preference effect*: People consider shared rare preferences, as compared to more common preferences, as strong cues for social compatibility. These results show how social statistics—our intuitive beliefs about the prevalence of others' choices and their underlying mental states—can systematically influence our social decisions. By combining intuitions from cognitive science and social psychology, these results provide new insights into the hidden inferential mechanisms that support our everyday intuitions about who we find more interesting, attractive, and compatible.

Acknowledgments

This work was supported in part by a National Science Foundation Graduate Research Fellowship Program grant to NV and a Varieties of Understanding grant from the John Templeton Foundation and James S. McDonnell Scholar Award to HG.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at https://doi.org/10.1016/j.cognition.2019.06.006.

References

Boer, D., Fischer, R., Strack, M., Bond, M. H., Lo, E., & Lam, J. (2011). How shared preferences in music create bonds between people: Values as the missing link. *Personality and Social Psychology Bulletin,* 1–13.

Bonneville-Roussy, A., Rentfrow, P. J., Xu, M. K., & Potter, J. (2013). Music through the ages: Trends in musical engagement and preferences from adolescence through middle adulthood. *Journal of Personality and Social Psychology, 105*(4), 703.

Byrne, D., & Nelson, D. (1965). Attraction as a linear function of proportion of positive reinforcements. *Journal of Personality and Social Psychology, 1*, 659–663.

Carney, D. R., Jost, J. T., Gosling, S. D., & Potter, J. (2008). The secret lives of liberals and conservatives: Personality profiles, interaction styles, and the things they leave behind. *Political Psychology, 29*(6), 807–840.

Clark, H. H., Schreuder, R., & Buttrick, S. (1983). Common ground at the understanding of demonstrative reference. *Journal of Verbal Learning and Verbal Behavior, 22*(2), 245–258.

Fawcett, C. A., & Markson, L. (2010a). Children reason about shared preferences. *Developmental Psychology, 46*, 299–309.

Fawcett, C. A., & Markson, L. (2010b). Similarity predicts liking in 3-year-old children. *Journal of Experimental Psychology, 105*, 345–358.

Gershman, S., Pouncy, H. T., & Gweon, H. (2017). Learning the structure of social influence. *Cognitive Science*.

Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science, 17*(9), 767–773.

Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences, 107*, 9066–9071.

Hasher, L., & Zacks, R. T. (1984). Automatic processing of fundamental information: The case of frequency of occurrence. *American Psychologist, 39*(12), 1372.

Heiphetz, L., Spelke, E. S., & Banaji, M. R. (2014). The formation of belief-based social preferences. *Social Cognition, 32*(1), 22–47.

Hertwig, R., & Gigerenzer, G. (1999). The 'conjunction fallacy' revisited: How intelligent inferences look like reasoning errors. *Journal of Behavioral Decision Making, 12*(4), 275–305.

Jern, A., Lucas, C. G., & Kemp, C. (2017). People learn other people's preferences through inverse decision-making. *Cognition, 168*, 46–64.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review, 80*(4), 237.

Kim, H., & Markus, H. R. (1999). Deviance or uniqueness, harmony or conformity?: A cultural analysis. *Journal of Personality and Social Psychology, 77*, 785–800.

Kushnir, T., Xu, F., & Wellman, H. M. (2010). Young children use statistical sampling to infer the preferences of other people. *Psychological Science, 21*, 1134–1140.

Liberman, Z., Kinzler, K. D., & Woodward, A. L. (2014). Friends or foes: Infants use shared evaluations to infer others' social relationships. *Journal of Experimental Psychology: General, 143*(3), 966.

Lonsdale, A. J., & North, A. C. (2009). Musical taste and ingroup favouritism. *Group Processes & Intergroup Relations, 12*(3), 319–327.

Lucas, C. G., Griffiths, T. L., Xu, F., Fawcett, C., Gopnik, A., Kushnir, T., & Hu, J. (2014). The child as econometrician: A rational model of preference understanding in children. *PLOS One, 9*(3), e92160.

McCarthy, G., Luby, M., Gore, J., & Goldman-Rakic, P. (1997). Infrequent events transiently activate human prefrontal and parietal cortex as measured by functional MRI. *Journal of Neurophysiology, 77*(3), 1630–1634.

McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology, 27*(1), 415–444.

Monin, B., & Norton, M. I. (2003). Perception of a fluid consensus: Uniqueness bias, false consensus, false polarization, and pluralistic ignorance in a water conservation crisis. *Personality and Social Psychology Bulletin, 29*, 559–567.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92*(3), 289.

Nelson, L. J., & Miller, D. T. (1995). The distinctiveness effect in social categorization: You are what makes you unusual. *Psychological Science, 6*(4), 246–249.

Nguyen, S. P., & Chevalier, T. (2015). Category coherence in children's inductive inferences with cross-classified entities. *Cognitive Development, 35*, 137–150.

Nisbett, R. E., & Kunda, Z. (1985). Perception of social distributions. *Journal of Personality and Social Psychology, 48*(2), 297.

North, A. C., & Hargreaves, D. J. (1999). Music and adolescent identity. *Music Education Research, 1*(1), 75–92.

Patalano, A. L., Chin-Parker, S., & Ross, B. H. (2006). The importance of being coherent: Category coherence, cross-classification, and reasoning. *Journal of Memory and Language, 54*(3), 407–424.

Peterson, C. R., & Beach, L. R. (1967). Man as an intuitive statistician. *Psychological Bulletin, 68*(1), 29.

Rehder, B., & Hastie, R. (2004). Category coherence and category-based property induction. *Cognition, 91*(2), 113–153.

Rentfrow, P. J., Goldberg, L. R., & Zilca, R. (2011). Listening, watching, and reading: The structure and correlates of entertainment preferences. *Journal of Personality, 79*(2), 223–258.

Rentfrow, P. J., & Gosling, S. D. (2006). Message in a ballad: The role of music preferences in interpersonal perception. *Psychological Science, 17*, 236–242.

Rentfrow, P. J., & Gosling, S. D. (2007). The content and validity of music-genre stereotypes among college students. *Psychology of Music, 35*, 306–326.

Ross, L., Greene, D., & House, P. (1977). The “false consensus effect”: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology, 13*(3), 279–301.

Rozin, P., & Siegal, M. (2003). Vegemite as a marker of national identity. *Gastronomica, 3*(4), 63–67.

Selfhout, M. H. W., Branje, S. J. T., ter Bogt, T. F. M., & Meeus, W. H. J. (2009). The role of music preferences in early adolescents’ friendship formation and stability. *Journal of Adolescence, 32*, 95–107.

Soley, G., & Aldan, P. (2018). Children and adults selectively attribute shared cultural knowledge to speakers of the same language. *Child Development*.

Soley, G., & Spelke, E. S. (2016). Shared cultural knowledge: Effects of music on young children’s social preferences. *Cognition, 148*, 106–116.

Suls, J., & Wan, C. K. (1987). In search of the false-uniqueness phenomenon: Fear and estimates of social consensus. *Journal of Personality and Social Psychology, 52*(1), 211.

Suls, J., Wan, C. K., & Sanders, G. S. (1988). False consensus and false uniqueness in estimating the prevalence of health-protective behaviors. *Journal of Applied Social Psychology, 18*(1), 66–79.

Sutton, S., Braren, M., Zubin, J., & John, E. (1965). Evoked-potential correlates of stimulus uncertainty. *Science, 150*(3700), 1187–1188.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review, 90*(4), 293–315.

Van Eijck, K. (2001). Social differentiation in musical taste patterns. *Social Forces, 79*(3), 1163–1185.

Verhallen, T. M., & Robben, H. S. (1994). Scarcity and preference: An experiment on unavailability and product evaluation. *Journal of Economic Psychology, 15*(2), 315–331.

Werner, C., & Parmelee, P. (1979). Similarity of activity preferences among friends: Those who play together stay together. *Social Psychology Quarterly,* 62–66.

Xu, F., & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences of the United States of America, 105*, 5012–5015.