



This article is part of the topic “Computational Approaches to Social Cognition,” Samuel Gershman and Fiery Cushman (Topic Editors). For a full listing of topic papers, see [http://onlinelibrary.wiley.com/journal/10.1111/\(ISSN\)1756-8765/earlyview](http://onlinelibrary.wiley.com/journal/10.1111/(ISSN)1756-8765/earlyview)

Information Sampling, Judgment, and the Environment: Application to the Effect of Popularity on Evaluations

Gaël Le Mens,^{a,b,c} Jerker Denrell,^d Balázs Kovács,^e Hülya Karaman^f

^a*Department of Economics and Business, Universitat Pompeu Fabra*

^b*Barcelona School of Management and Barcelona Graduate School of Economics*

^c*Department of Marketing and Management, University of Southern Denmark*

^d*Warwick Business School, University of Warwick*

^e*Yale School of Management, Yale University*

^f*Marketing Department, Singapore Management University*

Received 24 July 2017; received in revised form 1 June 2018; accepted 13 June 2018

Abstract

If people avoid alternatives they dislike, a negative evaluative bias emerges because errors of under-evaluation are unlikely to be corrected. Prior work that analyzed this mechanism has shown that when the social environment exposes people to avoided alternatives (i.e., it makes them resample them), then evaluations can become systematically more positive. In this paper, we clarify the conditions under which this happens. By analyzing a simple learning model, we show that whether additional exposures induced by the social environment lead to more positive or more negative evaluations depends on how prior evaluations and the social environment interact in driving resampling. We apply these insights to the study of the effect of popularity on evaluations. We show theoretically that increased popularity leads to more positive evaluations when popularity mainly increases the chances of resampling for individuals with low current evaluations. Data on repeat stays at hotels are consistent with this condition: The popularity of a hotel mainly impacts the chances of a repeat stay for individuals with low satisfaction scores. Our results illustrate how a sampling approach can help to explain when and why people tend to like popular alternatives. They also shed new light on the polarization of attitudes across social groups.

Keywords: Information sampling; Learning; Beliefs; Attitudes; Judgments; Biases; Social influence

1. Introduction

How are judgments and attitudes influenced by the individuals and objects with which people get into contact as they navigate their social world? Researchers in psychology have explored this fundamental question using two very different approaches. The first approach focuses on how the social environment affects how people process information about judgment targets. For example, research on priming has demonstrated that being exposed to a particular concept makes some categories more likely to be activated and thus affects the inferences people make about other people (e.g., Macrae & Bodenhausen, 2000). The second approach focuses on how the social environment affects the samples of information to which people are exposed. It has proposed that biases in the samples of information can lead to biased beliefs and judgments (Denrell, 2005; Einhorn & Hogarth, 1978; Fiedler, 2000; Smith & Collins, 2009; see Fiedler & Juslin, 2006, for a review). Several papers in this tradition have focused on what happens when the social environment makes an agent sample a judgment target that she would otherwise have avoided. It has been shown that this tends to lead to more positive attitudes (Denrell, 2007; Denrell & Le Mens, 2007, 2011, 2017; Fazio, Eiser, & Shook, 2004; Le Mens & Denrell, 2011; Le Mens, Kareev, & Avrahami, 2016). In this paper, we revisit this issue.

Why would additional exposure have a systematic effect on attitudes, according to the sampling approach? The key mechanism is the “hot stove effect,” which leads to a negativity bias in evaluations when people learn from experience (Denrell & March, 2001). It works as follows: People are likely to resample options with which they have had positive experiences. This implies that errors of overestimation are likely to be corrected. When people have a negative experience with an option, however, they are unlikely to resample. This implies that errors of underestimation are unlikely to be corrected. Denrell and March (2001) named this asymmetry in error corrections the “hot stove effect” in deference to Mark Twain’s observation about the cat and the hot stove: “We should be careful to get out of an experience only the wisdom that is in it—and stop there; lest we be like the cat that sits down on a hot stove lid. She will never sit down on a hot stove lid again—and that is well; but also she will never sit down on a cold one” (Twain, 1897, p. 124). An important consequence of the hot stove effect is that exposure to an avoided alternative can have a systematic effect on evaluations of this alternative. Such exposure can change an agent’s evaluation from negative to positive—a change that would not have happened otherwise.

In this paper, we demonstrate that additional exposure to avoided alternatives does not always have a positive effect on evaluation. Under some conditions, additional exposure can have a systematic negative effect on evaluations. We also specify sufficient conditions for the emergence of the positive effect of additional exposure on evaluations. This helps delineate the domain of application of the claims about the effect of additional exposure made in earlier work.

We analyze a model in which information sampling is shaped by current evaluations (people are more likely to sample alternatives they like) but also influenced by the social

environment. For example, people may be more likely to be exposed to (and thus sample) alternatives that are popular (i.e., frequently chosen or liked by others). Using a more general theoretical formulation compared to prior work, we show that whether the effect of exposure on evaluations is positive or negative depends how current evaluations and popularity *interact* in driving exposure. For example, if popularity mainly increases the chances of sampling alternatives a decision-maker dislikes, but does not change much the chance of sampling alternatives the decision-maker likes, then higher popularity will be associated with more positive evaluations of an alternative. If popularity mainly increases the chances of sampling alternatives the decision-maker already likes, and does not change much the chance of sampling disliked alternatives, then higher popularity will be associated with more negative evaluations of an alternative.

Using a large dataset of members in a loyalty program in a large hotel chain with more than 4,500 hotels, we estimate how current evaluations and popularity jointly impact the probability of resampling (operationalized as the probability of a repeat stay at a hotel chain). Estimations show that popularity mainly increases the chances of resampling for individuals with low satisfaction scores. Our model implies that in this case, popularity has an indirect and positive effect on quality estimates through its influence on sampling behavior. The sampling mechanism we propose could thus contribute to explaining why people like more popular hotels better than less popular hotels.

At a theoretical level, we also explore the consequences of our mechanism for explaining the polarization of attitudes across social-groups. More generally, our model and data illustrate how a sampling-based approach can help understand how features of the social environment, such as popularity, impact evaluations and judgments. Simon (1955) stressed that judgments are outcomes of cognitive operations on information samples obtained from the environment. While existing explanations focus on how the mind processes the available samples of information (the second stage), our sampling approach emphasizes properties of the information sample on which cognitive processes operate (the first stage).

2. Model

We analyze a simple computational model in which an individual learns about the qualities of two uncertain alternatives from experience. Let $\hat{Q}_{1,t}$ ($\hat{Q}_{2,t}$) denote the quality estimate for Alternative 1 (Alternative 2) at the beginning of period t . The individual updates her quality estimates on the basis of her observations of the payoffs of the alternatives. We also assume that the individual seeks positive experiences: the probability of sampling an alternative is increasing in the decision maker's quality estimate for that alternative (and decreasing in her estimate for the other alternative). To model the influence of the social environment on sampling, we introduce additional parameters. We denote by π_1 (π_2) the environmental factor that pertains to Alternative 1 (Alternative 2). For concreteness, we refer to π_j as the "popularity" of Alternative j .

Unless otherwise noted, we use capital letters to refer to random variables and corresponding non-capital letters to refer to their instantiations and non-random model parameters. Our model is as follows:

Payoffs of the alternatives. In each period, the decision maker samples one of K available alternatives. The payoffs of Alternative k are independent realizations of a random variable with mean μ_k and positive variance σ^2 . The payoff of Alternative k in period t is denoted by $Q_{k,t}$.

Quality estimates. The estimate updating rule has the following form:

$$\hat{Q}_{k,t+1} = \hat{Q}_{k,t} + b_t(Q_{k,t} - \hat{Q}_{k,t}), \tag{1}$$

where b_t is the weight of new sampled information. b_t is allowed to change in every period ($0 < b_t < 1$).

Sampling rule. The likelihood of sampling Alternative k in period t is given by the following Luce choice rule:

$$pS_k(\hat{q}_{1,t}, \dots, \hat{q}_{M,t}, \pi_1, \dots, \pi_M) = \frac{\alpha(\hat{q}_{k,t}, \pi_k)}{\sum_{m=1}^M \alpha(\hat{q}_{m,t}, \pi_m)}, \tag{2}$$

where $\alpha(\cdot, \cdot)$ is a positive function increasing in both of its arguments. $\alpha(\hat{q}_{k,t}, \pi_k)$ can be interpreted as the “attractiveness,” or “utility” of Alternative k .

2.1. Theoretical predictions

Let (\hat{Q}_1, \hat{Q}_2) denote the random variables toward which the quality estimates converge as t becomes large. When the expected value of \hat{Q}_k , denoted $E[\hat{Q}_k]$, increases (decreases) in π_k , we say that popularity has an indirect positive (negative) effect on quality estimates through information sampling.

The following theorem describes the conditions under which the indirect effect of the social environment (popularity) on quality estimates through sampling is positive or negative. The sign of this indirect effect depends on the form of the “utility” function $\alpha(\cdot, \cdot)$.

Theorem 1 Under some mild regularity conditions,¹ we have, for all k :

- (i) If $\frac{\partial^2 \log \alpha(\hat{q}_k, \pi_k)}{\partial \pi_k \partial \hat{q}_k} \leq 0$ then $E[\hat{Q}_k]$ is non-decreasing in π_k .
- (ii) If $\frac{\partial^2 \frac{1}{\alpha(\hat{q}_k, \pi_k)}}{\partial \pi_k \partial \hat{q}_k} \leq 0$, then $E[\hat{Q}_k]$ is non-increasing in π_k .

Proof See the Supporting Information.

Theorem 1(i) states a sufficient condition on $\alpha(\cdot, \cdot)$ for the effect of popularity to be positive. The technical condition on $\alpha(\cdot, \cdot)$ in (i) is known as “log-submodularity” (Karlin & Rinott, 1980). To better understand this condition, imagine increasing π_k from $\pi_{k,1}$ to $\pi_{k,2}$ and

consider how this changes the *ratio* of sampling probabilities, $pS_k(\hat{q}_k, \pi_{k,2})/pS_k(\hat{q}_k, \pi_{k,1})$. According to the sampling rule (Eq. 2), it is equal to $\alpha(\hat{q}_k, \pi_{k,2})/\alpha(\hat{q}_k, \pi_{k,1})$. The “log-submodularity” condition means that this ratio either (a) increases more when \hat{q}_k is *low* than when \hat{q}_k is *high* or (b) the increase is independent of \hat{q}_k . Informally, condition (i) is satisfied when the effect of popularity on sampling is higher when the quality estimates are low than when they are high (or independent of quality).

To illustrate the implications of this condition, consider the special case where sampling follows a logistic choice rule. In this case, the “utility” function is the exponential of a linear function:

$$\alpha(\hat{q}_k, \pi_k) = \text{Exp}(a_0 + a_1\hat{q}_k + a_2\pi_k + a_3\hat{q}_k\pi_k). \tag{3}$$

We assume $a_1 > 0$ and $a_2 > 0$ to ensure that the sampling likelihood is increasing in the quality estimate and in popularity. In this case, we have $\frac{\partial^2 \log \alpha(\hat{q}_k, \pi_k)}{\partial \pi_k \partial \hat{q}_k} = a_3$ and condition (i) holds whenever $a_3 \leq 0$. That is, whenever the value of the interaction term is negative or zero, then condition (i) is satisfied. The theorem implies that popularity has a positive effect on quality evaluations.

More generally, the “log-submodularity” condition can be understood as a weak version of the condition that popularity and quality estimates are “substitutes.” In economics, two goods, x and y , are called “substitutes” if the cross-derivative of the utility function, $\partial^2 u(x, y)/\partial x \partial y$, is negative. If x and y are substitutes this implies that an increase in y increases utility more strongly when x is *small* than when x is large. The “log-submodularity” condition requires that the *logarithm* of $\alpha(\cdot, \cdot)$ has a negative cross-derivative. This condition is weaker than substitutability because substitutability (a negative cross-derivative) implies log-submodularity whereas log-submodularity does not necessarily imply substitutability.²

Condition in (ii) in Theorem 1 states a sufficient condition for the effect of popularity to be negative.

This condition can be understood as a strong version of “complementarity.” Two goods, x and y , are complements if the cross-derivative of the utility function, $\partial^2 u(x, y)/\partial x \partial y$, is positive. If x and y are complements this implies that an increase in y increases utility more if x is *large* than if x is small. The condition in (ii) ($\frac{\partial^2 \frac{1}{\alpha(\hat{q}_k, \pi_k)}}{\partial \pi_k \partial \hat{q}_k} \leq 0$) implies complementarity ($\frac{\partial^2 \alpha(\hat{q}_k, \pi_k)}{\partial \pi_k \partial \hat{q}_k} > 0$), but not the other way around. The condition in (ii) can thus be interpreted as a “strong” complementarity condition. If this condition holds, the indirect effect of popularity on quality estimates via sampling is not positive, but rather negative (or, more generally, non-increasing in π_k).

Informally, Theorem 1 states that whether popularity has a positive or negative effect on quality estimates, via sampling, depends on whether an increase in popularity changes sampling (proportionally) more when (i) the quality estimate is low or when (ii) the quality estimate is high. The intuitive explanation for why this interaction matters is as follows. If popularity increases sampling most for low quality estimates (case i), low quality

estimates have a chance to regress to the mean, that is, upward. If popularity increases sampling most for high quality estimates (case ii), high quality estimates have a chance to regress to the mean, that is, downward.

It is important to note that if the sampling likelihood depends just on the social environment but not on quality estimates, the social environment has no effect on quality estimates. This is because there is no sampling bias in this case: At the beginning of period t , the quality estimate for Alternative k is the weighted average of t or fewer independent and identically distributed observations. Therefore, for all t , $E[\hat{Q}_{k,t}] = \mu_k$.

To illustrate Theorem 1, we simulated the model in a setting with two alternatives with normally distributed payoffs (means $\mu_1 = 2$ and $\mu_2 = 2.5$ and common variance $\sigma^2 = 4$). Thus, Alternative 1 is has lower quality than Alternative 2. We assume that the weight of new evidence in the estimate updating rule is constant and equal to $b = .5$. Consider first the case where $\alpha(\hat{q}_{k,t}, \pi_k) = e^{\hat{q}_{k,t} + \pi_k}$. As we explained above, this is consistent with condition (i) in Theorem 1. In this case, average quality estimates underestimate true qualities, $E[\hat{Q}_{k,t}] < \mu_k$ (Fig. 1). This reflects the “hot stove effect” (Denrell, 2005; Fazio et al., 2004; March, 1996). More important, a higher popularity reduces the extent of underestimation. It follows that when the inferior Alternative 1 is popular (high π_1), it may be estimated to have a higher quality than the superior Alternative 2: $E[\hat{Q}_{1,t}] > E[\hat{Q}_{2,t}]$. Indeed, the probability of mistakenly believing Alternative 1 to be the superior alternative, $P(\hat{Q}_{1,t} > \hat{Q}_{2,t})$, can be shown to increase with π_1 .

Suppose, next, that condition (i) in Theorem 1 does not apply. Then the effect of popularity is not necessarily positive, but it can be negative. To illustrate when this occurs,

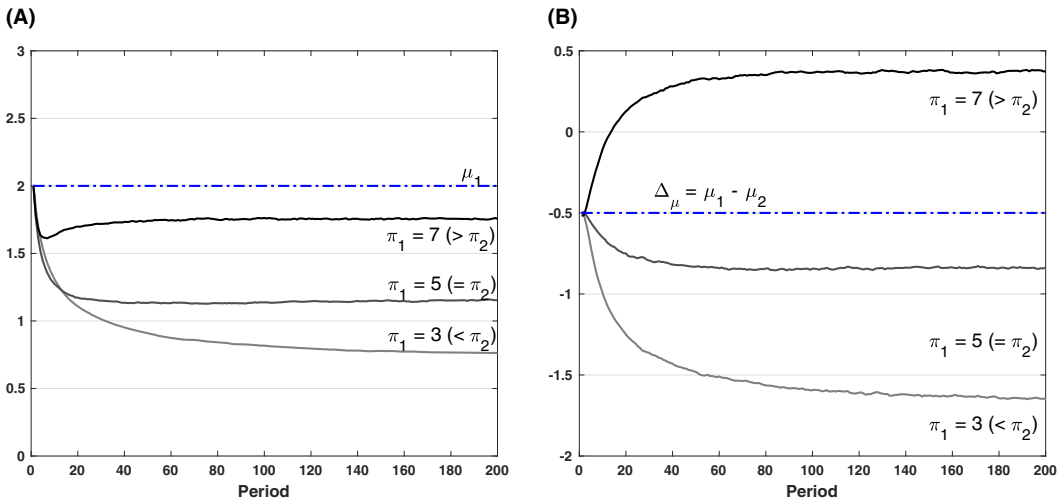


Fig. 1. Simulation in which $\alpha(\hat{q}_{k,t}, \pi_k) = e^{\hat{q}_{k,t} + \pi_k}$ and for which condition (i) in Theorem 1 holds. Panel (A): $E[\hat{Q}_{1,t}]$ as a function of time: Quality estimates decline as a result of the hot-stove effect, but the decline is less when π_1 is large. Panel (B): $E[\hat{Q}_{1,t}] - E[\hat{Q}_{2,t}]$ as a function of time: If the Inferior Alternative 1 is sufficiently popular, learners come to believe it has the higher quality. Based on 10^5 simulations with $\mu_1 = 2$, $\mu_2 = 2.5$, $\pi_2 = 5$, $\sigma = 2$, $b = .5$, and $\alpha(\hat{q}_{k,t}, \pi_k) = e^{\hat{q}_{k,t} + \pi_k}$.

suppose $\alpha(\hat{q}_{k,t}, \pi_k) = e^{\hat{q}_{k,t} + \pi_k + 2\hat{q}_{k,t}\pi_k}$, implying that popularity and quality estimates are complements. The hot-stove effect still operates and qualities are systematically underestimated (see Fig. 2). But, in contrast to the prior setting, the impact of π_1 is no longer positive but negative: The estimated quality of Alternative 1 is *lower* when π_1 is large and higher when π_1 is low.

To illustrate how changes in model parameters affect the size of the effect of the social environment on evaluations, we derived the formula for the expected asymptotic quality estimate for the two alternative model with normally distributed payoffs. We assume the utility function is compensatory, consistent with case *i* in Theorem 1: $\alpha(\hat{q}_{k,t}, \pi_k) = e^{s\hat{q}_{k,t} + \pi_k}$. Here *s* characterizes the sensitivity of the choice to quality estimates. We have the following:

Lemma 1 The expected asymptotic quality estimate for Alternative *k* is given by:

$$E[\hat{Q}_1] = \mu_1 - \frac{sb}{2-b} \sigma^2 \frac{e^{-s\mu_1}}{e^{-s\mu_1} + e^{\pi_1 - \pi_2} e^{-s\mu_2}} < \mu_1. \tag{4}$$

The same holds for Alternative 2.

Proof See the Supporting Information available online.

Just as in the simulations, the expected quality estimate $E[\hat{Q}_1]$ is always lower than quality, $E[\hat{Q}_1] < \mu_1$. The size of the underestimation decreases with the popularity of Alternative 1 π_1 (Fig. 3). At the limit, if $\pi_1 \gg \pi_2$, there is no systematic underestimation

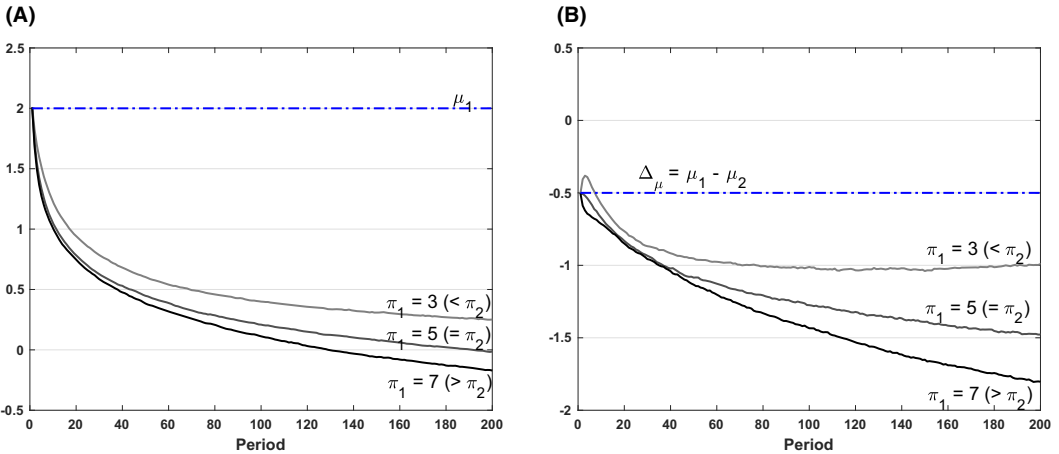


Fig. 2. Simulation in which $\alpha(\hat{q}_{k,t}, \pi_k) = e^{\hat{q}_{k,t} + \pi_k + 2\hat{q}_{k,t}\pi_k}$, for which condition (ii) in Theorem 1 holds, and popularity and quality estimates are “strong complements.” Panel (A): $E[\hat{Q}_{1,t}]$ as a function of time: Quality estimates decline as a result of the hot-stove effect and here the decline is larger when π_1 is large. Panel (B): $E[\hat{Q}_{1,t}] - E[\hat{Q}_{2,t}]$ as a function of time: If the inferior Alternative 1 is popular, learners come to believe it has even lower quality. Based on 10^5 simulations with $\mu_1 = 2$, $\mu_2 = 2.5$, $\pi_2 = 5$, $\sigma = 2$, $b = .5$, and $\alpha(\hat{q}_{k,t}, \pi_k) = e^{\hat{q}_{k,t} + \pi_k + 2\hat{q}_{k,t}\pi_k}$.




Fig. 3. The expected asymptotic quality estimate $E[\hat{Q}_1]$ increases with π_1 . By contrast, $E[\hat{Q}_2]$ decreases with π_1 . Figure obtained by plotting Eq. 4 with $\mu_1 = 2$, $\mu_2 = 2.5$, $\pi_2 = 5$, $\sigma = 2$, $b = .5$ and $s = 1$.

of Alternative 1. Equation 4 indicates that the effect of the social environment on evaluations (the strength of underestimation) is stronger when (a) the variance of the observations (σ^2) is large, (b) the weight of new observations is large (b is close to 1), and (c) the sensitivity of the sampling rule to quality estimates is high (s is large)—see the Supporting Information available online for additional discussion.

3. Application to the effect of popularity on evaluations

In many situations, people have an increased propensity to sample alternatives that are popular (i.e., chosen by many other people) as compared to alternatives that are unpopular (i.e., chosen by few other people). People may decide to go along with the majority and select the more popular alternative to avoid being seen as deviant (Cialdini & Goldstein, 2004; Granovetter, 1978), because of adverse reputation effects from receiving a poor outcome with an unusual alternative (Keynes, 1936), or because they know that those who deviate from the majority opinion tend to be disliked (Gerard & Rotter, 1961). For example, it is difficult for a doctor not to use the “best practice” prevailing in her hospital system in order to treat a given pathology, even if her personal experience with this practice is not positive. It is also safer to choose a popular alternative (Granovetter, 1978) than an unusual one. For example, a researcher will find it easier to get help if he chooses a research method commonly used by his colleagues. The advantage of having help around may motivate him to choose this method even if he does not believe that it is superior or if he got a poor experience with it.

Applied to this kind of setting, our model implies that people will evaluate popular alternatives more positively than unpopular alternatives, even if there are no systematic

differences in quality. The asymmetry in estimates emerges because errors of underestimation are more likely to remain uncorrected for unpopular alternatives than for popular alternatives. The higher the number of periods, the higher the cumulative probability that some errors of underestimation will have emerged and, in turn, the higher the probability that some errors will remain uncorrected. Our explanation is thus most relevant to explaining judgment patterns in empirical settings where people form their attitudes on the basis of repeated experiences, such as when the attitude object is another individual, a service, a sport or leisure activity, a musical artist, a music genre, a restaurant, a hotel, or an investment strategy. Our model is less relevant to settings where people have at most one or two interactions with the attitude object.

Many existing explanations for the positive effect of popularity on quality estimates emphasize conscious popularity-based inference. For example, the “information cascades” and “rational herding” literatures have shown that it is rational to use popularity as a signal of quality (e.g., Banerjee, 1992). Our model deliberately excluded such direct inferences from popularity to quality: We only assumed that popularity affects resampling and hence provides access to additional (unbiased but noisy) payoff information. There is an additional difference: The social learning explanations discussed in the information cascades and rational herding literatures assume that people are aware of the difference in popularity. But our explanation still works when people are *not* aware of such difference. Another explanation for the positive effect of popularity on evaluations focuses on the role of group identity: People might adjust their beliefs and attitudes to conform with the opinion prevailing in a group because they identify with the group (Cialdini & Goldstein, 2004; Turner, 1991). The reason is that similarity of attitudes is an important driver of interpersonal attraction (e.g., Clore, 1976). Our explanation clearly differs from such an identity-based mechanism because the latter consists of a direct effect of popularity on evaluations: The decision maker changes her attitudes when becoming aware of the attitudes of the members of the group. By contrast, our mechanism does not rely on this kind of motivated cognition. The influence of others remains outside the mind: It only affects the information people sample.

Several influential theories, such as cognitive dissonance theory (Festinger, 1957), or self-perception theory (Bem, 1972) have proposed that people adjust their judgments to make them consistent with their behavior. Under our assumption that people are more likely to choose popular alternatives than unpopular alternatives, these theories might also predict that popular alternatives would be evaluated more positively than unpopular alternatives. These theories rely on motivated cognition, whereas our model does not. Therefore, our model is applicable in situations where these theories are unlikely to operate.

While our argument differs from theories based on motivated cognition and popularity-based inferences, we do not challenge the experimental evidence for these mechanisms. Rather, our model suggests a complementary explanation that is likely to be important in naturally occurring environments where popularity affects available information.

3.1. Empirical illustration

Theorem 1 shows that the sign of the indirect effect of popularity on quality estimates through sampling depends on how quality estimate and popularity *interact* in affecting the attractiveness of the alternative. In order to illustrate how this interaction can be measured from field data, we analyzed a large dataset of members in a loyalty program in a large hotel chain with more than 4,500 hotels.

We have data on loyalty members whose first ever experience with the multi-chain hotel group was measured through the satisfaction survey conducted by the hotel group. This first stay took place between 2012 and 2015. The overall satisfaction with the stay is measured on a scale from 1 to 10, where 1 is the worst rating and 10 is the best. We take the satisfaction score given by the customer as her quality estimate of the hotel. Crucial for our purpose is that the subsequent stays of members of the loyalty program are recorded in our data. We aim to predict whether a customer will return to the same hotel (and thus re-sample it) on the basis of her satisfaction with the first stay and the popularity of the hotel.

More precisely, the dependent variable is a binary variable that indicates whether the individual returned for a repeat stay at the first hotel she experienced with the multi-chain hotel group within 180 days of the first stay. We define the *popularity* of a hotel as the number of ratings it received on TripAdvisor in the 365 days preceding the first stay of the customer. The “standardized popularity” is denoted by π_k and refers to the popularity divided by the standard deviation of popularity (which equals 51). Our dataset contains 455,903 individuals and 62,182 re-sampling events.

3.1.1. Analysis

We focus on $pS_{k,2}^i$, the probability that user i visits hotel k within 180 days of the first stay. We assume that this probability can be expressed by a logistic choice rule that is a function of quality estimate and popularity (consistent with Eqs. 2 and 3). Finally, we assume that the attractiveness of all other hotels is 1.³ In other words:

$$pS_{k,2}^i = \frac{e^{a_{0,k} + a_1 \hat{q}_{k,1}^i + a_2 \pi_k + a_3 \hat{q}_{k,1}^i \pi_k}}{1 + e^{a_{0,k} + a_1 \hat{q}_{k,1}^i + a_2 \pi_k + a_3 \hat{q}_{k,1}^i \pi_k}}, \tag{5}$$

where $\hat{q}_{k,1}^i$ is the satisfaction score customer i gave to hotel after her first stay, π_k is the standardized popularity of the hotel, $a_{0,k}$ is a hotel fixed effect, and a_1 , a_2 , and a_3 are parameters to be estimated.

For this specification, in which $\alpha(\hat{q}_k, \pi_k) = \text{Exp}(a_0 + a_1 \hat{q}_k + a_2 \pi_k + a_3 \hat{q}_k \pi_k)$, condition (i) in Theorem 1 holds whenever $a_3 \leq 0$. That is, a sampling approach only predicts positive indirect effect of popularity on quality estimates if the interaction term is negative or zero. Is such an assumption of a non-positive interaction effect plausible?

To find out, we fitted Eq. 5 to the data on hotel visits, using maximum likelihood.⁴ Table 1 shows the estimated values of a_1 , a_2 , and a_3 . As shown, the attractiveness of the

Table 1

Effect of satisfaction with first stay and standardized hotel popularity on likelihood to revisit the hotel based on the estimation of Eq. 5 (with hotel fixed effects)

Satisfaction score: a_1	0.047* (0.003)
Standardized popularity: a_2	0.086* (0.026)
Satisfaction score \times Popularity: a_3	-0.012* (0.002)
Observations	455,903
LR χ^2	264.95

* $p < .001$.

hotel increases with the satisfaction score of the first stay ($a_1 > 0$) and popularity ($a_2 > 0$). Most important, the interaction term is negative ($a_3 < 0$).⁵ This non-positive interaction effect is consistent with the condition (i) in Theorem 1, which implies that in this setting the *indirect* effect of popularity on quality estimates will be positive. Stated differently, the estimates show that popularity can have a positive *indirect* effect on quality estimates via its impact on sampling behavior. The upshot is that the sampling mechanism we propose could explain why people might like more popular hotels better than less popular hotels.

4. Implications for identity signaling and attitude polarization

Our sampling-based mechanism can also contribute to explaining the polarization of attitudes across groups. This happens when people want to *avoid* alternatives that are popular in other groups. This behavior is likely when the activities people choose signal their identities.

By choosing particular type of clothes, hairstyle, or program of education, people signal to others who they are (Berger & Le Mens, 2009; Bourdieu, 1984). The desire to signal their identity can motivate people to engage in activities typically associated with the type of people with whom they want to be identified (McCracken, 1988). Identity signaling also motivates people to avoid activities associated with a group of people from whom one wants to distance oneself (Berger & Heath, 2007). Applied to this setting, our model can provide a novel explanation as to why individuals might shift their attitudes to diverge from the attitudes of groups with which they do not wish to identify (Wood, Pool, Leck, & Purvis, 1996). For example, it has been observed that educated people tend to dislike music they associate with uneducated people (Bryson, 1996).

Suppose there are two groups (A and B, such as teenagers and parents) and two possible activities (1 and 2). Activity 1 is popular in Group A; Activity 2 is popular in Group B. Let π_1^A denote the popularity of Alternative 1 in Group A, π_1^B the popularity of that alternative in Group B, etc. We have $\pi_1^A \gg \pi_1^B$ and $\pi_2^A \ll \pi_2^B$.

If an agent belongs to Group A, she is more likely to adopt a practice that is popular in Group A and unpopular in Group B because she wants to be identified as a member of Group A. More precisely, we assume that if the agent is in Group A, the likelihood that she samples Alternative 1 in period t is:

$$pS_1^A(\hat{q}_{1,t}, \hat{q}_{2,t}) = \frac{e^{s\hat{q}_{1,t} + \pi_1^A - \pi_1^B}}{e^{s\hat{q}_{1,t} + \pi_1^A - \pi_1^B} + e^{s\hat{q}_{2,t} + \pi_2^A - \pi_2^B}}.$$

All the other elements of the model remain the same as in the baseline setting analyzed above. The initial estimates are unbiased; there are no systematic differences between the initial quality estimates of members of the two groups. The following proposition describes the pattern of asymptotic quality estimates:

Proposition 1 Suppose Alternative 1 is much more popular in Group A than in Group B ($\pi_1^A \gg \pi_1^B$) and that Alternative 2 is much more popular in Group B than in Group A ($\pi_2^A \ll \pi_2^B$), we have:

(i) For Group A agents:

$$E[\hat{Q}_1] \sim \mu_1; \quad E[\hat{Q}_2] \sim \mu_2 - \frac{sb}{2-b} \sigma^2.$$

(ii) For Group B agents:

$$E[\hat{Q}_1] \sim \mu_1 - \frac{sb}{2-b} \sigma^2; \quad E[\hat{Q}_2] \sim \mu_2.$$

Proof See Supporting Information available online.

Suppose that the two alternatives have similar qualities ($\mu_1 \sim \mu_2$). Our model implies the emergence of attitude polarization: Members of Group A will tend to evaluate the popular alternative in that group (Alternative 1) more positively than the popular alternative in the other group (Alternative 2). The converse happens for members of Group B.

Prior explanations for the polarization of attitudes across groups have generally invoked some form of motivated cognition: people subconsciously change their preferences to diverge from the attitudes of unwanted groups (e.g., Bryson, 1996) while they strive to adopt attitudes that are similar to the attitudes prevailing in their groups, at least in part because attitude similarity leads to liking (Clore, 1976). Our model does not require that observing the choices of members of the wanted and unwanted groups have such a direct impact on attitudes. It only requires a change in sampling behavior. Our analysis demonstrates that this change in sampling behavior will have an indirect systematic effect on attitudes. For example, a teenager may have a more or less neutral opinion about some music genre. When hearing that his parents like this music, he does not

directly change his opinion. He avoids listening to such music, however, because he feels that if he is seen listening to such music he would appear uncool. Our model implies that such avoidance in behavior, which is not necessarily driven by a personal negative evaluation of the music, will lead to systematic differences in evaluations. More generally, our model shows that the choice of activities influences exposure and learning opportunities, and this creates a systematic evaluative bias against the alternative chosen by most people in the out-group. This, in turn, leads to an evaluative advantage for the alternative popular in the in-group as compared to the alternative that is popular in the out-group.

5. Discussion and conclusion

In this paper, we have shown that when the social environment makes people more likely to sample a particular alternative, people tend to evaluate it more positively without this environmental influence. This occurs when the social environment and the current evaluation are substitutes in the sense that the social environment affects sampling more strongly if the evaluation of the alternative is low rather than high.

To apply this insight into the effect of popularity on evaluations, we noted that in many settings, people are more likely to sample popular alternatives than unpopular alternatives. We also noted that popularity and quality estimates are often substitutes: even if people have a negative evaluation of an alternative, they might select it again if it is sufficiently popular. We found evidence for such an interaction between popularity and evaluations in analyses of the repeat purchase behavior of hotel customers.

Our results do not rely on the fact that people are more likely to *choose* popular alternatives consciously, but rather that they are more likely to *sample* popular alternatives. Our mechanism thus applies to settings where payoff information is more accessible for more popular alternatives even if popularity does not affect choices directly. This can be the case when the decision maker's friends or colleagues are more likely to have experiences with popular alternatives (and share them), or because information about the experiences of others is more easily available through information channels such as the press, online forums or review websites. Our model is thus relevant to settings where the decision maker learns not only from her personal experiences with the alternatives but also from the experiences of others, provided that there is greater access to payoff information for popular alternatives.

Because it focuses on access to information rather than on information processing, our theory does not challenge existing explanations that rely on information processing biases (i.e., motivated cognition) or inferences about quality on the basis of popularity. It provides a complementary perspective to explaining the effect of popularity on evaluations and attitudes. Finally, we note that our theory also applies to settings where the availability of payoff information is influenced by environmental factors other than popularity, such as getting an award (Kovács & Sharkey, 2014), or changes in prices, or non-random ordering of options on websites (Le Mens, Kovács, Kareev, & Avrahami, 2018).

Acknowledgments

This research was funded by the Spanish Ministry of Economics, Industry and Competitiveness Grant #AEI/FEDER UE-PSI2016-75353, Ramon-y-Cajal Fellowship #RYC-2014-15035, the Spanish Ministry of Education Culture and Sports Jose Castillejo Fellowship #CAS15/00225, and BBVA Foundation grant #IN[15]_EFG_ECO_2281 to Gaël Le Mens. Hulya Karaman benefited from financial support by the Goizueta Business School at Emory University.

Notes

1. See footnote 1 in the Supporting Information.
2. Assuming the function and the first derivatives are positive.
3. It is not a problem to assume that the combined attractiveness of all other available hotels is equal to 1 because there is a scaling factor ($e^{a_{0,k}}$ in the expression for the attractiveness of the focal hotel) and what matters is the ratio of the levels of attractiveness.
4. We used the “xtlogit” command in Stata 14.
5. The estimated interaction effect remains the same, at -0.013 , if we remove fixed effects. Ancillary analyses show that other factors affect the sampling likelihood in a similar way, such as the hotel star rating, or the average rating on TripAdvisor.
6. The conditions of positive recurrence are easily verified for the setups where payoffs are normally distributed, the choice rule is the exponential version of Luce choice rule and the weight of new evidence is constant. They also hold for other setups, such as when the payoff distribution follows uniform or Bernoulli distributions and the weight of new evidence remains constant. Problematic settings include configurations where b_t converges quickly to 0 as t becomes large.
7. Consider a modified model in which Alternative k is sampled in *every* period. In that case, the sequence of random variables $(\hat{Q}_{1,t})_{t \geq 1}$ defines a martingale sequence of random variables. We have that if $\sup_{t \geq 1} E[|\hat{Q}_{k,t}|] < \infty$, then $(\hat{Q}_{k,t})_{t \geq 1}$ converges to a random variable with probability 1 (see Billingsley, 1995, p. 468). This limiting random variable has distribution $g_k(\cdot)$.

References

- Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics*, 107, 797–817.
- Berger, J., & Heath, C. (2007). Where consumers diverge from others: Identity signaling and product domains. *Journal of Consumer Research*, 34(2), 121–134.
- Berger, J., & Le Mens, G. (2009). How adoption speed affects the abandonment of cultural tastes. *Proceedings of the National Academy of Sciences USA*, 106(20), 8146–8150.
- Billingsley, P. (1995). *Probability and measure*. New York: Wiley.
- Bem, D. J. (1972). Self-perception theory. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (pp. 1–62). New York: Academic Press.

- Bourdieu, P. (1984). *Distinction: A social critique of the judgment of taste*. Cambridge, MA: Harvard University Press.
- Bryson, B. (1996). “Anything but heavy metal”: Symbolic exclusion and musical dislikes. *American Sociological Review*, 61, 884–899.
- Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, 55, 591–621.
- Clore, G. L. (1976). Interpersonal attraction: An overview. In J. W. Thibaut, J. T. Spence, & R. C. Carson (Eds.), *Contemporary topics in social psychology* (pp. 135–175). Morristown, NJ: General Learning Press.
- Denrell, J. (2005). Why most people disapprove of me: Experience sampling in impression formation. *Psychological Review*, 112, 951–978.
- Denrell, J. (2007). Adaptive learning and risk taking. *Psychological Review*, 114, 177–187.
- Denrell, J., & Le Mens, G. (2007). Interdependent sampling and social influence. *Psychological Review*, 114 (2), 398–422.
- Denrell, J., & Le Mens, G. (2011). Seeking positive experiences can produce illusory correlations. *Cognition*, 119(3), 313–324.
- Denrell, J., & Le Mens, G. (2017). Information sampling, belief synchronization, and collective illusions. *Management Science*, 63(2), 528–547.
- Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, 12(5), 523–538.
- Einhorn, H. J., & Hogarth, R. M. (1978). Confidence in judgment: Persistence of the illusion of validity. *Psychological Review*, 85(5), 395–416.
- Fazio, R. H., Eiser, J. R., & Shook, N. J. (2004). Attitude formation through exploration: Valence asymmetries. *Journal of Personality and Social Psychology*, 87, 293–311.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford University Press.
- Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review*, 107(4), 659–676.
- Fiedler, K., & Juslin, P. (2006). Taking the interface between mind and environment seriously. In K. Fiedler & P. Juslin (Eds.), *Information sampling and adaptive cognition* (pp. 3–29). Cambridge, UK: Cambridge University Press.
- Gerard, H. B., & Rotter, G. S. (1961). Time perspective, consistency of attitude and social influence. *Journal of Abnormal and Social Psychology*, 62, 565–572.
- Granovetter, M. S. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 83, 1420–1443.
- Karlin S., & Rinott Y. (1980). Classes of orderings of measures and related correlation inequalities I. Multivariate totally positive distributions. *Journal of Multivariate Analysis*, 10(4), 467–498.
- Karlin, S., & Rubin, H. (1956). Distributions possessing a monotone likelihood ratio. *Journal of the American Statistical Association*, 51(276), 637–643.
- Keynes, J. M. (1936). *General theory of employment interest and money*. London, UK: Macmillan.
- Kovács, B., & Sharkey, A. J. (2014). The paradox of publicity: How awards can negatively affect the evaluation of quality. *Administrative Science Quarterly*, 59(1), 1–33.
- Le Mens, G., & Denrell, J. (2011). Rational learning and information sampling: On the “naivety” assumption in sampling explanations of judgment biases. *Psychological Review*, 118(2), 379–392.
- Le Mens, G., Kareev, Y., & Avrahami, J. (2016). The evaluative advantage of novel alternatives: An information-sampling account. *Psychological Science*, 27(2), 161–168.
- Le Mens, G., Kovács, B., Kareev, Y., & Avrahami, J. (2018). How endogenous crowd formation undermines the wisdom-of-the-crowd in online ratings. *Psychological Science*, 29(2), 1475–1490.
- Macrae, C. N., & Bodenhausen, G. V. (2000). Social cognition: Thinking categorically about others. *Annual Review of Psychology*, 51, 93–120.
- March, J. G. (1996). Learning to be risk averse. *Psychological Review*, 103, 309–319.
- McCracken, G. (1988). *Culture and consumption: New approaches to the symbolic character of consumer goods and activities*. Bloomington, IN: Indiana University Press.

- Meyn, S., & Tweedie, R. L. (2009). *Markov chains and stochastic stability*. Cambridge, UK: Cambridge University Press.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118.
- Smith, E. R., & Collins, E. C. (2009). Contextualizing person perception: Distributed social cognition. *Psychological Review*, 116(2), 343–364.
- Twain, M. (1897). *Following the equator: A journey around the world*. Hartford, CT: American Publishing Co.
- Wood, W., Pool, G., Leck, K., & Purvis, D. (1996). Self-definition, defensive processing, and influence: The normative impact of majority and minority groups. *Journal of Personality and Social Psychology*, 71, 1181–1193.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Appendix S1. Details of the proofs and sensitivity analyses.

Supplementary Material for ‘Information Sampling, Judgment and the Environment:
Application to the Effect of Popularity on Evaluations’

Gaël Le Mens

Universitat Pompeu Fabra and Graduate School of Economics, Barcelona
Barcelona, Spain

Jerker Denrell

University of Warwick, Coventry, United Kingdom

Balázs Kovács

Yale University, New Haven, United States

Hülya Karaman

Emory University, Atlanta, United States

Supplementary Material for ‘Information Sampling, Judgment and the Environment:
Application to the Effect of Popularity on Evaluations’

Additional Details for the proofs

Proof of the equality in Equation 7 in the proof of Lemma 2.

To explain Equation 7, consider the first term.

$$RHS_1 \hat{=} \int_{r_1} h_{\pi_1, \pi_2}(r_1, \hat{q}_2) pS_1(r_1, \hat{q}_2) \kappa_1(r_1, \hat{q}_1) dr_1. \quad (1)$$

This is the probability that the decision maker has estimates equal to r_1 and \hat{q}_2 , that she samples the first alternative, and that her quality estimate goes from r_1 to \hat{q}_1 which happens with probability $\kappa_1(r_1, \hat{q}_1)$. We integrate over all possible estimates about the first alternative. The second term has a similar structure.

To show that the above joint density in Lemma 2 does satisfy the stability equation (Equation 9 in the Appendix) note first that $pS_1(r_1, \hat{q}_2)$ can be written as follows

$$pS_1(r_1, \hat{q}_2) = \frac{\alpha^{-1}(\hat{q}_2, \pi_2)}{\alpha^{-1}(r_1, \pi_1) + \alpha^{-1}(\hat{q}_2, \pi_2)}. \quad (2)$$

Using this and the formula for the joint density (Equation 7 in the Appendix) to perform the appropriate substitutions in Equation [1](#), we get:

$$\begin{aligned} RHS_1 &= K(\pi_1, \pi_2) \int_{r_1} \alpha^{-1}(\hat{q}_2, \pi_2) g_1(r_1) g_2(\hat{q}_2) \kappa_1(r_1, \hat{q}_1) dr_1, \\ &= K(\pi_1, \pi_2) \alpha^{-1}(\hat{q}_2, \pi_2) g_2(\hat{q}_2) \int_{r_1} g_1(r_1) \kappa_1(r_1, \hat{q}_1) dr_1, \\ &= K(\pi_1, \pi_2) \alpha^{-1}(\hat{q}_2, \pi_2) g_2(\hat{q}_2) g_1(\hat{q}_1). \end{aligned}$$

The last equality is a consequence of the equality in Equation 6 in the Appendix (after a simple change of variable in the integral).

Similarly, for the second term in the RHS of the stability equation, we get:

$$\int_{r_2} h_{\pi_1, \pi_2}(\hat{q}_1, r_2) pS_2(\hat{q}_1, r_2) \kappa_2(r_2, \hat{q}_2) dr_2 = K(\pi_1, \pi_2) \alpha^{-1}(\hat{q}_1, \pi_1) g_1(\hat{q}_1) g_2(\hat{q}_2).$$

Summing the two terms, the RHS of Equation 7 in the Appendix becomes equal to

$$K(\pi_1, \pi_2) \left(\alpha^{-1}(\hat{q}_1, \pi_1) + \alpha^{-1}(\hat{q}_2, \pi_2) \right) g_1(\hat{q}_1) g_2(\hat{q}_2),$$

which is $h_{\pi_1, \pi_2}(\hat{q}_1, \hat{q}_2)$, the stationary distribution.

Details for the proof of Lemma 1.

The double integration of $\hat{q}_1 h_{\pi_1, \pi_2}(\hat{q}_1, \hat{q}_2)$ with respect to \hat{q}_1 and \hat{q}_2 yields:

$$\begin{aligned} E[\hat{Q}_1] &= \frac{e^{-\frac{s^2 \sigma^2 b}{2(2-b)}}}{e^{-s\mu_1 - \pi_1} + e^{-s\mu_2 - \pi_2}} e^{-\pi_1} \int_{\hat{q}_1} \hat{q}_1 e^{-s\hat{q}_1} g_1(\hat{q}_1) d\hat{q}_1 \\ &\quad + \frac{e^{-\frac{s^2 \sigma^2 b}{2(2-b)}}}{e^{-s\mu_1 - \pi_1} + e^{-s\mu_2 - \pi_2}} \mu_1 e^{-\pi_2} \int_{\hat{q}_2} e^{-s\hat{q}_2} g_2(\hat{q}_2) d\hat{q}_2. \end{aligned}$$

Noting that $\int_{\hat{q}_2} e^{-s\hat{q}_2} g_2(\hat{q}_2) d\hat{q}_2$ is the moment generating function of the distribution $g_2(\cdot)$, evaluated at $-s$, we have:

$$\int_{\hat{q}_2} e^{-s\hat{q}_2} g_2(\hat{q}_2) d\hat{q}_2 = e^{-s\mu_2 + \frac{s^2 \sigma^2 b}{2(2-b)}}.$$

Some algebraic manipulations yield:

$$\int_{\hat{q}_1} \hat{q}_1 e^{-s\hat{q}_1} g_1(\hat{q}_1) d\hat{q}_1 = e^{-s\mu_1 + \frac{s^2 \sigma^2 b}{2(2-b)}} \left(\mu_1 - \frac{sb}{2-b} \sigma^2 \right).$$

Summing up the terms, we obtain the desired formula for $E[\hat{Q}_1]$ (Equation 4 in the body of the paper).

Sensitivity of the Predictions to Model Parameters

The amplitude of the sampling bias is moderated by how the environment parameters relate to each other. To see how, it is useful to consider extreme cases. Suppose, for example, that $\pi_1 \gg \pi_2$. In this case, the decision maker will sample Alternative 1 almost no matter what her quality estimates are. There is thus almost no sampling bias for this alternative, the quality estimate is close to the true quality and there is almost no underestimation tendency for this alternative. When $\pi_1 \ll \pi_2$, sampling depends more strongly on quality estimates. The quality estimate is thus subject to the systematic underestimation tendency described above. These additional predictions are formalized in the following corollary to Lemma 1 that deals with what happens when the two alternatives have very different popularities.

Corollary 1. *i) When Alternative 1 is much more popular than Alternative 2, ($\pi_1 \gg \pi_2$),*

$$E[\hat{Q}_1] \sim \mu_1. \quad (3)$$

ii) When Alternative 1 is much less popular than Alternative 2 ($\pi_1 \ll \pi_2$),

$$E[\hat{Q}_1] \sim \mu_1 - \frac{sb}{2-b}\sigma^2 < \mu_1. \quad (4)$$

The same holds for Alternative 2.

Variance of the Observations

When the variance of the observations is low ($\sigma^2 \sim 0$), the environment has little to no effect on quality assessments. Similarly, the decision maker is unlikely to mistakenly believe the inferior alternative to be better than the superior alternative. Like other sampling explanations of judgment biases, our mechanism requires the possibility of making estimation mistakes (e.g., Denrell, 2005; Denrell & Le Mens, 2007, 2011; Fazio et

al., 2004). This is the pattern of error corrections that leads to systematic judgment patterns and biases. If the variance of the observations is low, even few observations will lead to accurate estimates, and no mistake will emerge. This implies that our information sampling mechanism will not operate in such conditions.

Reinforcement learning structure of the sampling process

If the sampling likelihood is not sensitive to the quality of past experiences, systematic errors are unlikely to emerge. This is the case when the sensitivity of the sampling likelihood to quality estimates is low (s is low) or when the weight of new observations is low (b is close to 0). This is because the emergence of the positive effect of the environment on quality assessments relies on the fact that decision makers have a systematic tendency to underestimate the qualities of the alternatives. This occurs because of the increased likelihood to sample an alternative again following positive experiences with that alternative. When s or b is low, such adaptive resampling hardly occurs. This implies, in turn, that systematic underestimation does not emerge and thus that popular and unpopular alternatives are equally likely to be underestimated.

We have assumed so far that the sampling likelihood is increasing both in the quality estimate and in the environment factor π_k . It is worth noting that our formal results remain similar when the sampling likelihood is decreasing in π_k while it is increasing in the quality estimate (e.g., Lieberman, 2000). To see what our model predicts in this case, it is enough to consider the same formulas, but by putting a ‘minus’ sign in front of π_1 and π_2 . For example, Equation 4 in the body of the paper becomes:

$$E[\hat{Q}_1] = \mu_1 - \frac{sb}{2-b}\sigma^2 \frac{e^{-s\mu_1}}{e^{-s\mu_1} + e^{-\pi_1 + \pi_2} e^{-s\mu_2}} < \mu_1. \quad (5)$$

In the section of empirical implications, we build on these insights and analyze a model where the sampling likelihood is increasing in the popularity of the alternative in the group of the decision maker while it is decreasing in its popularity in another group. We

show that such a model leads to a polarization of attitudes across groups.

We could similarly adapt our model to cases where the sampling likelihood is decreasing with the quality estimate. To see what happens in this case, it is enough to replace s by $-s$ in all the formulas. In this case, decision makers have a tendency to overestimate (rather than underestimate) the qualities of the available alternatives. The effect of popularity on quality estimates is negative if the sampling likelihood is increasing with popularity¹, and it is positive if the sampling likelihood is decreasing with popularity.² Although there are fewer settings where people have an elevated tendency to sample alternatives that lead to poor payoffs than settings where people have an elevated tendency to sample again alternatives that lead to positive payoffs, this can happen for example when a journalist tries to uncover stories about unethical or unlawful behavior, or when a wine critique decides to taste again a wine that she found underperforming with respect to her expectations (Laube, 2007).

References

- Denrell, J. (2005) Why most people disapprove of me: Experience sampling in impression formation. *Psychological Review* 112:951-978.
- Denrell, J. (2007). Adaptive learning and risk taking. *Psychological Review*, 114, 177–187.
- Denrell, J., & Le Mens, G. (2007) Interdependent Sampling and Social Influence. *Psychological Review* 114 (2): 398-422.
- Denrell, J., & Le Mens, G. (2011). Seeking positive experiences can produce illusory correlations. *Cognition*, 119(3), 313–324. doi:10.1016/j.cognition.2011.01.007.

¹ In this case,

$$E[\hat{Q}_1] = \mu_1 + \frac{sb}{2-b}\sigma^2 \frac{e^{-s\mu_1}}{e^{-s\mu_1} + e^{\pi_1 - \pi_2} e^{-s\mu_2}}.$$

² In this case,

$$E[\hat{Q}_1] = \mu_1 + \frac{sb}{2-b}\sigma^2 \frac{e^{-s\mu_1}}{e^{-s\mu_1} + e^{-\pi_1 + \pi_2} e^{-s\mu_2}}.$$

Fazio, R. H., Eiser, J. R., & Shook, N.J. (2004). Attitude Formation through Exploration: Valence Asymmetries. *Journal of Personality and Social Psychology*, 87, 293-311.

Laube, J. (2007). When wines merit a second look. *Wine Spectator Blog James Laube's Wine Flights*. URL: <http://www.winespectator.com/blogs/show/id/15264>. Retrieved on Dec. 29, 2013.

Lieberson, S. (2000). *A matter of taste: How names, fashions, and culture change*. Yale University Press.

Le Mens, G., & Denrell, J. (2011). Rational learning and information sampling: On the “naivety” assumption in sampling explanations of judgment biases. *Psychological Review*, 118(2), 379–392. doi:10.1037/a0023010.