

Social Judgments from Adaptive Samples

Jerker Denrell,

University of Oxford

Oxford, United Kingdom

Email: jerker.denrell@sbs.ox.ac.uk.

Gaël Le Mens

Universitat Pompeu Fabra and Barcelona GSE

Barcelona, Spain

Email: gael.le-mens@upf.edu.

Social Judgments from Adaptive Samples

“Why is it that people sometimes believe ‘what is not so’?”

Thomas Gilovich, 1991.

A vast amount of research has documented the systematic judgment errors that people make. They use inaccurate stereotypes about those belonging to other groups, they develop superstitious beliefs, they are overconfident about their predictive abilities... Much of the existing social psychological literature suggests that such biases emerge because people process information inaccurately; they use flawed hypotheses tests and motivated reasoning, and rely on biased heuristics (e.g., Fiske & Taylor, 2007; Tversky & Kahneman, 1974).

Recently, however, a number of scholars have challenged the perspective that fallible judgment is driven by flawed processing of information. Noting that the environment often produces unrepresentative samples of information, they have demonstrated that even if individuals process available information correctly, their judgments will be subject to systematic error patterns akin to those previously explained by invoking information-processing deficiencies (see Fiedler & Juslin, 2006 for a review). This sampling approach has, for example, produced alternative explanations for in-group bias, illusory correlation (Fiedler, 2000) and overconfidence (Juslin, Winman & Olsson, 2007).

In this chapter, we argue that an important source of such information bias is ‘adaptive sampling,’ defined here as the tendency of decision makers to select again, and thus continue to sample, activities that led to positive experiences but to avoid activities that led to poor experiences. Such a tendency is basic to most learning mechanisms.¹ It is *adaptive* because it ensures that decision makers avoid activities with consistently poor outcomes. Nevertheless, it generates a sample bias: the likelihood that an individual will take another sample and get more information about the outcome of an activity depends on past experiences with that alternative. In particular, the decision maker will stop learning about alternatives she (possibly mistakenly) evaluates negatively.

We argue that adaptive sampling can cast new light on several well-known judgment biases in social psychology, such as in-group bias in impression formation (Denrell, 2005), risk-aversion (Denrell & March, 2001, Denrell, 2007), more positive assessments of popular alternatives (Denrell & Le Mens, 2011b; Le Mens & Denrell, 2011), illusory correlations (Denrell & Le Mens, 2011a) and social influence (Denrell & Le Mens, 2007).

¹ It is the foundational principle of reinforcement learning algorithms (Sutton and Barto, 1998). See also the ‘law of effect’ (Thorndike, 1911).

Our perspective emphasizes an understudied interaction between two important aspects of belief formation. The first aspect pertains to the goal of the decision maker. People sometime select alternatives or interact with other individuals just to learn about those. But, they often care about the immediate outcomes of their choices, and not just about the informational content of their experiences: their goal is often to ultimately have positive, enjoyable, experiences. This desire for positive experiences is important because it provides the rationale for the adaptive selection of alternatives, based on past experiences used as predictors for future experiences. The second aspect is the role of the environment as a factor that constrains and enables access to information and thus indirectly influences decision-making and judgment. Constrained access to information, combined with adaptive sampling, can explain what seem to be ‘irrational’ judgments without assuming flawed information processing.

This chapter is organized as follows. We start by delineating what we mean by adaptive sampling and discussing its most basic implication: the tendency to underestimate the value of uncertain alternatives. Then, we illustrate how this tendency emerges using computer simulations of a simple learning model. We then discuss the role of access to information and describe the various biases that can be explained by our approach. Finally, we conclude the chapter by a discussion of the rationality of adaptive sampling.

Adaptive sampling in decision making and judgment

What is adaptive sampling?

To explain what we mean by adaptive sampling, consider the following restaurant review:

“Not a happy dining experience. The service was absent minded, and the curry really wasn’t up to very much. I won’t be going here again.” Tim (San Jose Mercury News, 12/15/2005).

Clearly, Tim’s experience was poor. As a result, Tim’s impression of the restaurant is negative and he has decided to avoid the restaurant in the future.

This is adaptive sampling at work – the probability that Tim will go to the restaurant again is low because his experience was poor and his impression negative. To avoid another poor experience, Tim avoids the restaurant in the future. If Tim’s experience had been positive, however, Tim would probably have gone to the restaurant again and, doing so, would have had access to information about it. Thus, the probability that Tim will sample the restaurant again, go there and experience the food again, depends on Tim’s previous experiences – positive experiences lead to further sampling while negative experiences lead to avoidance.

More generally, we refer to the strategies that increase the probability of sampling alternatives with favorable past outcomes and reduce the probability of sampling alternatives with poor past outcomes as ‘adaptive sampling’ schemes. These strategies are adaptive because they use information obtained from experience to adapt their behavior and to ultimately improve the average outcome obtained by those who adopt it. By changing the probability of sampling in response to past experiences, decision makers ensure that alternatives with consistently poor outcomes are avoided and that alternatives with consistently good outcomes are pursued. (To be sure, such approach-avoidance strategies are adaptive only to the extent that past experiences are valid predictors for the qualities of future experiences, that is, to the extent that the environment is relatively stable).

Because of this beneficial feature, adaptive sampling is basic to almost any experiential learning process and is thus very often observed. People tend to continue to interact with others with whom they have had good experiences and they tend to engage again in activities they have enjoyed. Conversely, people tend to avoid individuals with whom they did not get along well, and they tend to avoid activities they did not find enjoyable.

Adaptive sampling leads to a negativity bias

Despite the fact that adaptive sampling is both a sensible and a common sampling strategy, it generates a subtle sample bias that has systematic consequences for belief formation (Denrell & March, 2001; Gilovich, 1991; March, 1996).

The basic problem resides in the very nature of adaptive sampling: the likelihood that an individual will take another sample and get more information about the quality of an activity depends on the outcome of past experiences with that activity. As a result, the probability of sampling is not fixed, as in randomized experiments, but is contingent on the history of experiences with the alternative. Because the probability of sampling is higher for alternatives with good past outcomes, more information will be gained about such alternatives than about alternatives with poor past outcomes. It can be shown that this information asymmetry leads to a negativity bias, or systematic tendency to underestimate the value of uncertain alternatives (Denrell, 2005; Denrell & March, 2001).

To explain this, let’s reflect on the consequences of Tim’s initial experience at the restaurant. Following his poor experience, he will avoid the restaurant. And unless he obtains additional information about it in some other way, his negative impression will persist.

Suppose, by contrast, that Tim’s initial experience with the restaurant was positive and that he leaves the restaurant with a positive impression. If he follows an adaptive sampling strategy, he is likely to go to that same restaurant again in the future. In

information terms, Tim is likely to sample that alternative again. By doing so, Tim will gain additional information about the restaurant, and will be able update his impression. If this second diner is a really poor experience, Tim might develop an overall negative impression of the restaurant, despite his positive first experience.

The crucial aspect of this story is that a poor second experience can overcome a positive initial experience. But the opposite correction cannot occur. This is because when the first experience is negative, there is simply no second experience. And thus, there is no possibility for upward correction. This, in turn, implies that negative impressions are more stable than positive impressions. And overall, this process leads to a general tendency to underestimate the value of the restaurant.

'Underestimation' versus 'Overestimation'

Another way to explain this outcome is to note that, due to adaptive sampling, errors or 'overestimation' are likely to be corrected, while errors of 'underestimation' are likely to persist.

To explain this, let us make Tim's restaurant story more specific. Suppose that 50% of the time the restaurant serves good meals and 50% of the time it serves bad meals. Tim gives good meals a 4 and bad meals a 2 on a 1 to 5 scale. Tim initially does not know anything about the quality of the restaurant.

When Tim experiences a good meal, and gives it a 4, he overestimates the quality of the restaurant. Because he is likely to revisit the restaurant, however, he can experience a generally more representative set of outcomes and correct his error of overestimation. After a few meals, he might develop a rather accurate assessment of the quality of the restaurant. Formulated differently, his assessment of the quality of the restaurant is likely to "regress to the mean" (see Fiedler and Krueger, This Volume, for an overview of regression effects).

Contrast this to what happens if Tim initially experiences a bad meal. He gives the restaurant a 2, which corresponds to an underestimation of the quality of the restaurant. Because Tim is likely to avoid the restaurant in the future, he cannot experience a more representative set of outcomes and thus cannot correct his error of underestimation. As a result, his assessment will not regress to the mean after a negative experience.

Overall, this asymmetry in the probability of correcting errors of over and underestimation implies that errors of underestimation will be more likely than errors of overestimation.

Errors in Social Perception and Adaptive Sampling

As noted by Jussim, Stevens, and Salib (this Volume) social perception is not terribly accurate. Why do such errors occur and are they not eventually corrected as social perceivers gain further information? Much research in social psychology is motivated by this question and several mechanisms have been discussed in the literature that could explain why erroneous perceptions persist (e.g., confirmation biases, self-fulfilling prophecies). We argue that adaptive sampling is a less explored but potentially important mechanisms for why some errors are not corrected. Erroneous and negative perceptions might imply that social perceivers avoid further sampling.

There are obviously many exceptions to the adaptive sampling assumptions. We discuss them at a latter stage in this chapter but now turn to an illustration of the emergence of the negativity bias using computer simulations of a simple learning model.

A simulation model of adaptive sampling

Model

Consider an individual, T, who has to decide, repeatedly, whether to engage in an activity. For example, T may have to decide whether to go to a restaurant.

We assume the payoff from the activity is uncertain. For example, the quality of the meals served at a restaurant may vary from day to day, or with variations of the quality of the ingredients. To model this variability, we assume that T's payoff from selecting the activity in period t , $X(t)$, is drawn from a normal distribution with mean 0 and standard deviation 1.

Every time T chooses the activity, T learns more about it and updates his impression of the activity. We use a simple model to capture this type of learning. We assume that the updated impression is a weighted average of the old impression and the experienced payoff (Anderson, 1981; Hogarth & Einhorn, 1992; Kashima & Kerekes, 1994). That is,

$$I(t) = (1-b) I(t-1) + bX(t),$$

where $I(t)$ is the impression of A at the end of period t and b is a parameter regulating the weight of the new experience. For most of the simulations below we will assume that $b = 0.5$.² We assume that T can only learn about the activity in periods when he chooses it. In periods when T does not choose the activity, we assume that her impression remains the same:

² The exact value of b does not qualitatively affect the results we discuss in this chapter, provided that b is higher than 0 and lower than 1. The value of b , however, affects the size of the effects and how quickly they unfold over time.

$$I(t) = I(t-1).$$

We assume that the initial impression is equal to 0. That is, the initial impression is unbiased – it is equal to the expected payoff of the activity.

The adaptive sampling assumption is implemented by assuming that the probability, $P(t)$, that T will choose the activity in period t is an increasing function of the impression. More precisely, we assume that $P(t)$ is a logistic function of the impression:

$$P(t) = \frac{1}{1 + \exp(-s * I(t - 1))}.$$

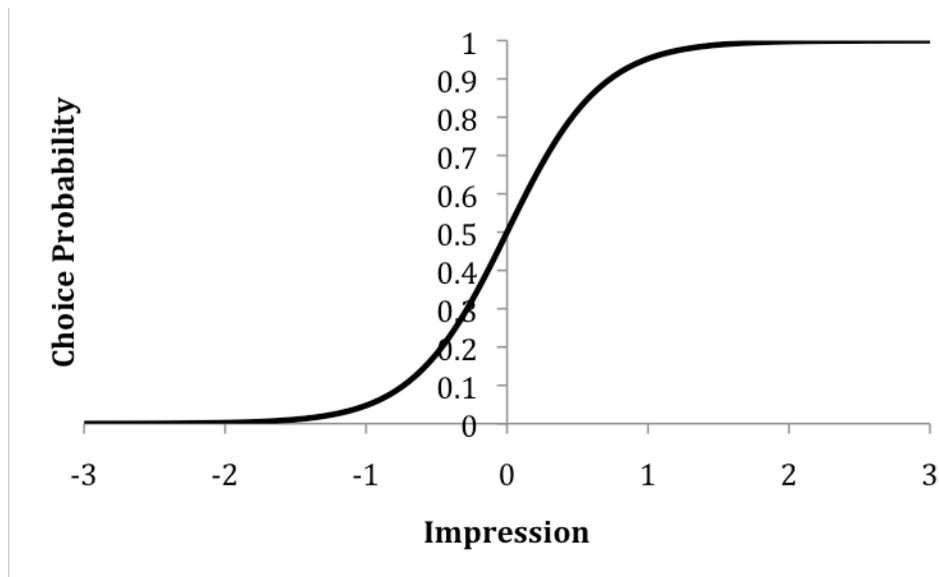


Figure 1: Plot of how the choice probability, $P(t)$, varies with the impression, $I(t-1)$.

This logistic choice rule has often been used to model choices under uncertainty (e.g. Luce, 1959). Here $s > 0$ is a parameter regulating how sensitive the choice probability is to the impression. If s is large then $P(t)$ is close to one whenever $I(t-1) > 0$ and close to zero whenever $I(t-1) < 0$. If s is close to zero then the choice probability is close to one half even for positive impressions. In the simulations below, we will assume that s equals 3 (see Denrell, 2005, for a discussion of the effect of varying s).

Figure 1 plots how the choice probability, $P(t)$, varies with the impression, $I(t-1)$, when $s = 3$.

Negativity bias

The above model implies that T's impression is a weighted average of all of T's experiences. Nevertheless, due to adaptive sampling, T's impression will be negatively biased, as illustrated on the graphs of Figure 2. This figure plots the distribution of T's impression at the end of period 10 together with the normal distribution with mean zero and variance one. The distribution of the impressions is negatively skewed like in the restaurant example discussed in the previous section: most impressions are negative. The average impression is also negative; it is -0.31. Moreover, in period 10, T chooses the activity in only 35% of the simulation runs.

The reason for this negativity bias is the asymmetry in error corrections discussed above. If T's impression is positive, and thus 'overestimates' the expected payoff (0 in that case), then he is likely to choose the activity again and his updated impression will be a weighted average of the previous impression and of the latest payoff. Because the expected payoff is zero, T's new impression will tend to decline towards zero (this is the well-known phenomenon of regression to the mean). Moreover, there is a chance that $X(t)$ is so negative that T's new impression will become negative. If T's impression is negative, T 'underestimates' the expected payoff. In that case, T is unlikely to select the activity again and T's impression will not be updated, that is, it will remain negative. Stated differently, negative impressions tend to be more 'sticky.'

The above reasoning suggests that an important condition for the negativity bias to emerge is that there is a possibility for errors in estimation of the quality of an alternative. If it were enough to select the alternative just once to fully know its value, the negativity bias would not emerge. Conversely, the negativity bias will be stronger the higher the likelihood and the amplitude of estimation errors. This line of reasoning leads us toward an important implication of adaptive sampling: the emergence of seemingly risk-averse behavior.

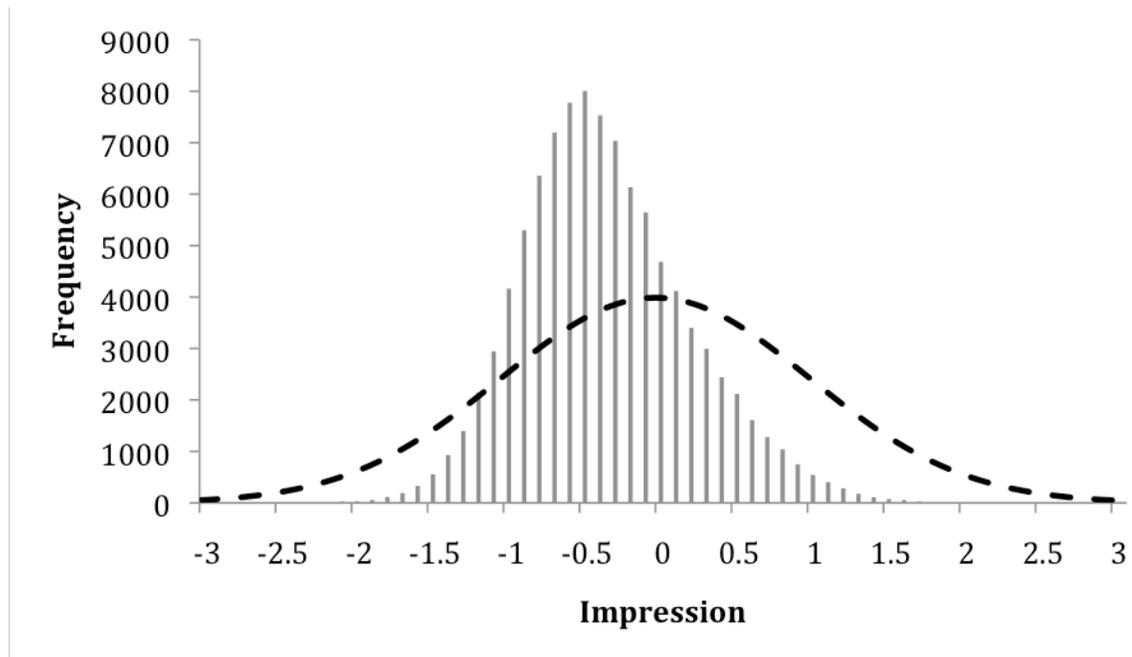


Figure 2: The distribution of T's impression after period 10 compared to the normal distribution of the payoff (based on 100 000 simulations).

Risk-Averse Behavior

The model of adaptive sampling produces seemingly risk-averse behavior (Denrell, 2007; Denrell & March, 2001; March, 1996). That is, T is less likely to select the activity whenever its payoff distribution is more variable. To illustrate this, we simulated the learning model for different values of the standard deviation of the payoff distribution. Figure 3 shows that the probability that T will choose the activity in period 10 decreases when the standard deviation of the payoff distribution increases. In other words, T behaves as if he were 'risk-averse.'

This effect occurs because a more variable payoff distribution tends to produce more extreme negative payoffs that lead to premature avoidance of the alternative. What matters is that the stronger the error of underestimation, the lower the likelihood that it will be corrected. Large errors of overestimation are less consequential because they will generally be quickly corrected (they lead to further selections of the alternative). Overall, this implies that both the choice probabilities and impressions tend to be lower for more variable alternatives than for less variable alternatives.

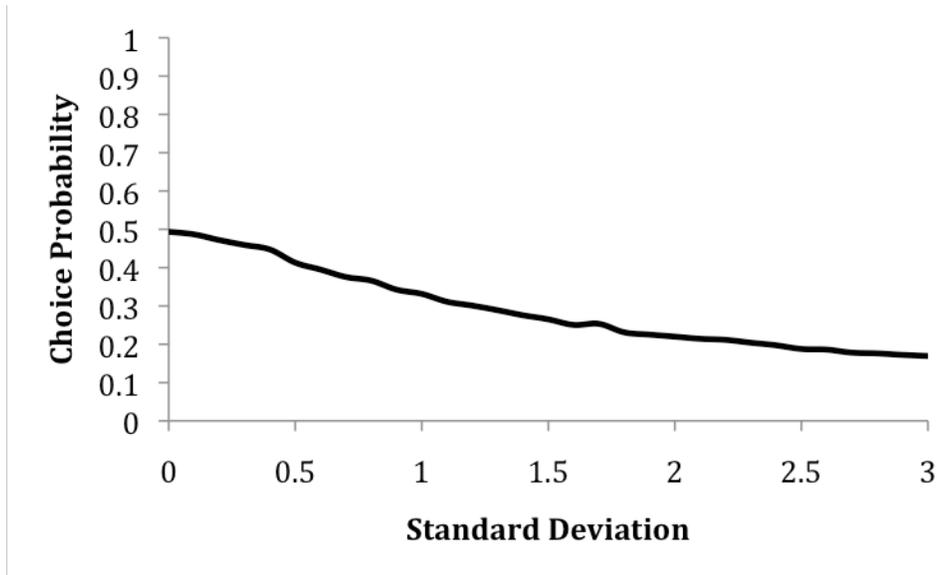


Figure 3: The probability that T will choose the object in the tenth period is decreasing with the standard deviation of the payoff distribution. Each point is the average of 10000 simulations.

The above model shows how a tendency to choose a less variable alternative can be the result of learning. Thus, the model offers an explanation of seemingly “risk-averse” behavior that differs from explanations that attribute risk-aversion to a concave utility function (Arrow, 1971) or to loss-aversion (Kahneman and Tversky, 1979). In the above model, seemingly risk-averse behavior can emerge as a result of learning, even if the decision maker is formally ‘risk neutral,’ that is, has a linear utility function and does not directly care about the variability of the outcomes. What happens in the above model is that decision makers’ impressions of the more variable alternatives will be lower. That is, decision makers learn to avoid a risky alternative because they have a poor impression of these alternatives. Denrell (2007) formally demonstrated that such risk-averse behavior emerges for a broad class of learning models.

Information Access and Judgment

The most important implication of adaptive sampling is that it can explain why access to information can have a systematic effect on judgment and choice even if this information is unbiased. In this section, we discuss how this can be the case. And, in the next section, we will show that this can cast a new light on a number of judgment biases reported in social psychology and the decision-making and judgment literatures.

Why access to information has a systematic effect on judgment and choice

The explanation of the negativity bias through adaptive sampling relied on the assumption that people obtain information about the quality of an alternative only if they actively select it. If they do not select the alternative, they do not get additional information and thus do not update their beliefs about it. But this condition does not always hold: sometimes people have access to information about an alternative even if they do not select it. For example, people might not control which individuals or activities they get exposed to. For example, people may have to continue to work with colleagues they find disagreeable or incompetent. In some other settings, people can get access to information about others without interacting with them. For example, you might learn about the achievement of one of your colleagues even if you do not interact with him or her. In this latter setting, even if you select or avoid activities and interactions based on past experience, access to information might not be fully determined by your adaptive sampling strategy. When this is the case, adaptive sampling does not lead to a negativity bias.

This simple observation has the following important implication: if people follow an adaptive sampling strategy, then estimates of the quality of an alternative are systematically and positively influenced by access to information. In particular, having access to one additional observations of the payoff distribution they are learning about tends to increase the tendency to evaluate this alternative positively. The reason is that such access to unbiased information eliminates the negativity bias that adaptive sampling otherwise would have led to. This is what is commonly known in the probability literature as regression to the mean (see also the chapter by Fiedler & Krueger, this volume, for an overview of regression effects in social psychology).

Simulation model of the impact of information access

To demonstrate formally this systematic effect of access to unbiased information, suppose we change the above simulation model in the following way: in each period there is a probability r that T will be able to observe the payoff of the uncertain alternative even if she does not choose it (Denrell, 2005). The above discussion of the effect of one more observation suggests that the higher the probability that T will be able to observe such ‘foregone’ payoffs, the higher the probability that T will have a positive impression of the alternative. This intuition is confirmed by the simulation results reported on Figure 4, which plots the probability that T’s impression is positive in period ten as a function of r . When $r = 1$, T observes information about the uncertain alternative in every period, and his impression is thus an unbiased estimate of the quality of the alternative. For lower values of r , there is some negativity bias, but with a lower magnitude than what happens when sampling is strictly adaptive as in the previous section. For example, if $r = 0.3$ the probability of a positive impression in period 10 is 37%, while it is 49% when $r = 0.8$.

A similar effect would also have emerged if T had been ‘forced’ to revisit the restaurant, when he would have avoided it on the basis of a negative impression. This type of situation might happen if other restaurants are closed and there is no other alternative for a meal. Alternatively, an individual might go to the restaurant with his or her friends. If your friends really want to go to that restaurant and the quality of your company matters more than the quality of your meal, you might still go. Regardless of the reason, such ‘forced’ choice would lead to an information sample that is less affected by prior experiences than in the base case. If, at the limit, forced choice occurs whenever the decision maker would not have selected the alternative on the basis of his impression, the impression will be an unbiased estimate of the quality of the alternative.

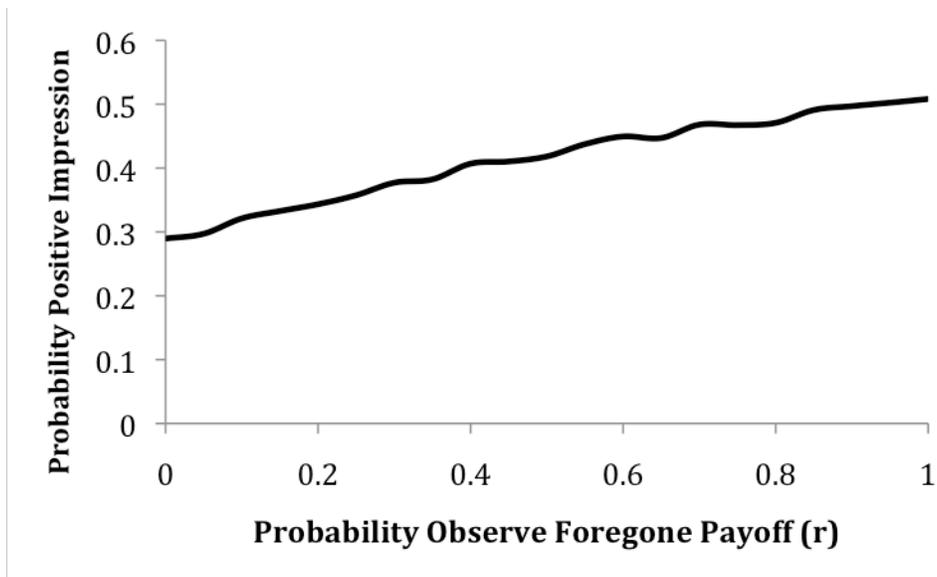


Figure 4: The probability that T’s impression is positive in period 10 is increasing with r , the probability that T will be able to observe the payoff of the uncertain alternative even in period when T does not choose the uncertain alternative. Each point is the average of 10 000 simulations.

Experimental evidence

An experiment on attitude formation by Fazio, Eiser and Shook (2004) illustrates how adaptive sampling can lead to a negativity bias and how access to information can eliminate it. Participants participated in a survival game, in which they had to eat beans with positive energy levels and avoid beans with negative energy levels in order to survive. Initially they did not know what type of beans generated positive energy levels, but they could learn from experience. They only learned from their own experience, that is, by choosing beans and experiencing their outcome. They did not learn about the energy value of the beans they avoided.

At the end of the game, participants were presented with beans of different types and asked to estimate if these beans were of the positive or negative type. The basic result was that a negativity bias emerged. Specifically, participants made more errors of ‘underestimation’ than ‘overestimation’ - they were more likely to mistake a positive bean for a negative than vice versa. The reason was that if they suspected, perhaps falsely, that a bean of a particular type was negative, they avoided these types of beans in the future. In contrast, if they suspected that a bean was positive, they continued to select this type of bean. In other words, they followed a strategy of adaptive sampling.

In one manipulation, Fazio and his colleagues changed the information structure so that participants learned about the energy values also about the beans they avoided. In this case, there was no negativity bias. Instead, participants were equally likely to make an error of under- or overestimation. We also obtained similar results in experimental investigations of variations of the multi-armed bandit setting (Le Mens & Denrell, 2010).

Adaptive Sampling and Judgment Biases

Now that we have explained the nature of adaptive sampling and its most basic implication for the role of information on belief formation and choices, we can illustrate how it can help explain several phenomena that have been viewed as puzzling and irrational, including in-group bias, social influence or illusory correlations. The usual explanations of these phenomena attribute them to mental ‘flaws’, such as inaccurate perception and biased information processing (Fiedler, 2000). In a series of articles (Denrell, 2005, 2007; Denrell & Le Mens, 2007, 2011; Le Mens & Denrell, 2011), we have suggested that adaptive sampling may provide an alternative explanation.

In putting forward this alternative explanation, we do not want to suggest that existing accounts are incorrect. There is substantial experimental evidence that heuristic processing can give rise to some of the effects we describe below. Nevertheless, our work suggests an alternative, complementary explanation, that may be important in settings where information is not provided to people but has to be actively sampled.

In-group bias in impression formation

Why do people develop more positive opinions of those close to them? College students have more positive opinions of their roommates than of other students (Festinger, Schachter & Back, 1950; Segal, 1974) and members of ethnic groups have more positive opinions of their own groups (Dasgupta, 2004; Hewstone, Rubin, & Willis, 2002). Explanations for this tendency have focused on how flawed hypothesis

testing, confirmation biases, motivated reasoning, or prior expectations can distort impressions (DiDonato, Ulrich, & Krueger, 2011; Wood, 2000).

Denrell (2005) suggested that adaptive sampling could provide an alternative explanation. Key to this explanation is the observation that more information is usually available about in-group members that one does not personally interact with, simply because one tends to be more connected to the in-group. In addition, people are more likely to continue to interact with in-group members they have a negative impression of than with out-group members they have a negative impression of. For example, Levin, van Laar and Sidanius (2003) show that while individuals often avoid members of other groups if initial encounters are negative, they tend to continue to interact with in-group members. In some cases, it may also be difficult to avoid interacting with in-group members one dislikes, such as members of the same family or department.

This implies that the negativity bias generated by adaptive sampling will be eliminated or at least attenuated for in-group members. As shown above, this leads to more positive evaluations of in-group members as compared to out-group members.

Denrell (2005) also showed that this explanation implies that evaluations of out-groups will be more positive for larger out-groups, simply because it is more likely that one will come into contact with them, at work or in schools etc. As shown by the recent meta-analysis of contact and prejudice by Pettigrew and Tropp (2006), there is substantial evidence for the underlying idea that contact leads to reduction in prejudice. Moreover, several studies have shown that larger out-groups, because of the increased probability of contact, lead to reductions in prejudice (see Denrell, 2005, for a review).

The explanation offered by adaptive sampling is clearly distinct from other perspectives. Rather than emphasizing flawed perception or preconceived notions, it stresses how access to information can give rise to negative stereotypes about out-groups. From a normative point of view this shift in perspective is important. It demonstrates why de-biasing individuals may not be enough. Rather, more even access to information, perhaps through more even access to formal and informal contacts, may be essential to the elimination of biased stereotypes.

Evidence from a recent natural experiment on the effect of additional interactions on racial is consistent with this suggestion. Shook and Fazio (2008) analyzed the evolution of automatically activated racial attitudes toward African Americans of White freshmen that were randomly assigned an African American roommate or a White roommate. The racial attitudes of those with an African American roommate became more positive after one quarter but the racial attitudes of those with a White

roommate did not change. While several interpretations are possible, this study suggests that additional exposure can help correct the negativity bias against out-group members.

Social Influence

Why do proximate others tend to develop similar attitudes? Previous explanations of such social influence have focused on why an individual would be motivated to agree with the opinions or beliefs of others (e.g. Cialdini & Goldstein, 2004; Wood, 2000).

In Denrell and Le Mens (2007), we have shown that motivated reasoning and imperfect information processing are not necessary to explain social influence. Rather, a social influence effect can also emerge because of adaptive sampling. Our explanation focuses on how an individual A can indirectly influence the attitude of another individual B by affecting the activities that B samples and get exposed to. Surprisingly, this effect emerges even if the outcomes experienced by B, when sampling an activity, are *independent* of A's experiences.

To explain how this mechanism works, consider two friends, A and B. Suppose that A likes a restaurant, while B does not. Usually, B would then avoid the restaurant. However, if B and A are friends, B may sometimes join A at the restaurant, if B cares more about her friendship with A than about the food. By sampling the restaurant again, B gets new information, which might change her attitude from negative to positive, that is, closer to A's attitude. This would not have happened, however, if A also had a negative impression. In that case, they would have both avoided the restaurant.

As this example shows, influence over sampling can *indirectly* lead to influence over attitudes. In the above example, B did not change her negative attitude to the restaurant simply because A had a positive attitude. Thus, A had no direct influence over B's attitude. But A had an indirect influence over B's attitude by changing B's sampling behavior. More generally, this social influence through interdependent sampling provides a novel mechanism for why public conformity in behavior might lead to private acceptance.

Denrell and Le Mens (2007) showed that this mechanism provides a simple account of existing findings in the literature. For example, it explains why beliefs tend to be influenced more in the direction of those of powerful people than in the direction of people with less power. The explanation, according to our mechanism, is that powerful people have more influence over what activities others get exposed to. And, in turn, they have more influence on their attitudes.

The explanation of social influence through interdependent sampling leads to novel empirical predictions. In particular, it implies an asymmetry in social influence.

Consider, again, two individuals, A and B. The model of interdependent sampling implies that A's attitude is more influential if B's attitude is negative than if B's attitude is positive. The reason is that if B's attitude is negative, B is likely to avoid sampling and might only sample if A has a positive attitude. If B instead has a positive attitude, B might sample anyway, whether A's attitude is positive or not. We found that this asymmetric pattern of social influence is present in Newcomb's (1961) longitudinal data on students' attitudes towards each other.

Preference for popular alternatives

The systematic positive effect of additional information on impression formation also has interesting population-level implications. It suggests that there will be a sample bias that favors established and popular alternatives over novel and potentially superior alternatives.

The general idea is that the social environment tends to provide additional information about popular alternatives, even if the decision maker does not personally choose them. For example, popular restaurants get reviewed, and thus one can learn about those even if one does not attend them. But information about new, or unpopular venues is harder to access, and one often has to go there to learn about the venue. If the decision maker avoids a popular restaurant following poor experiences, she might still learn about it by reading reviews, and learn that it is not that bad. This might not happen for the unpopular restaurant. This asymmetry in terms of access to information can help explain why more popular alternatives are often more positively evaluated.

Popularity can also influence impression formation in a different way, through its effect on sampling. In some cases, people might feel compelled to try out popular alternatives, even if they do not believe they are of high quality. As a result of this additional incentive to sample, people may get more positive evaluations of the quality of popular alternatives.

One reason why people might want to sample popular alternatives, even if they do not believe that their quality is the highest, is that the payoff from adopting an alternative may increase with the number of others adopting the same alternative. For example, in evaluating an operating system for a personal computer people may care about both its reliability (i.e. quality) and the number of others who have chosen the same operating system. Ideally, they would like to choose an operating system with the highest quality, but such an operating system might be less useful if few others have adopted it because sharing programs with others is also important. Alternatively, people may decide to go along with the majority and select the most popular alternative to avoid being seen as deviant (Cialdini & Goldstein, 2004; Granovetter, 1978), or because of adverse reputation effects to receiving a poor outcome with an unusual alternative (Keynes, 1936).

Denrell and Le Mens (2011b) show that in the presence of such external influences on sampling, quality assessments of the alternatives will also be biased towards the popular alternative when people do not know the qualities of the alternatives but can only learn about them from their own experiences. When there are only two alternatives, most people will come to believe that the most popular alternative is also of superior quality, even when it is not.

The reason is that here popularity affects opportunities for error corrections: If an agent mistakenly believes that the most popular alternative is the worst, she is likely to discover her mistake. But if she mistakenly believe that the least popular alternative is the worst, she is unlikely to discover her mistake. To see why, suppose that alternative 1 is the best. In addition, suppose that, by chance, most people have come to select alternative 2. If an agent incorrectly believes that alternative 1 is the worst, she is likely to avoid it, because it is also unpopular. As a result, her negative estimate of the quality of alternative 1 remains unchallenged and therefore persists. Now, suppose that alternative 1, rather than alternative 2, is the most popular. If the agent incorrectly believes that alternative 1 is the worst, she might still want to select it again, in order to gain the benefit of coordination. Because she obtains some additional information about the quality of alternative 1 when selecting it again, she might discover that it is not that bad, and even superior to alternative 2. This asymmetry in error corrections leads to an overall tendency to underestimate the quality of unpopular alternatives.

Illusory correlations

Adaptive sampling also suggests a novel explanation for the emergence of illusory correlations in person perception (Denrell & Le Mens, 2011b). In standard studies on illusory correlations, experimental participants observe a set of items each characterized by a pair of attribute values (X, Y). Existing theories explain illusory correlation by proposing that some observations receive more weight in the computation of the correlation than others (e.g. Allan, 1993; Hamilton & Gifford, 1976). This assumption of differential weighing is not necessary to our explanation based on adaptive sampling.

To see how adaptive sampling can explain illusory correlations, consider the following example. Suppose you learn about two traits of individuals you meet at a Swing dance venue. Suppose you learn, from experience, about whether the people you meet are good dancers and agreeable individuals.

Suppose that dancing skills and agreeableness are uncorrelated in the population. That is, an individual who is a good dancer is not more or less likely to be more agreeable than an individual who is not a good dancer. If you want to interact with people that are good on at least one of the two dimensions (i.e good dancer and/or agreeable) you will end up perceiving the two attributes to be *positively* correlated.

The key to the emergence of this illusory correlation is that you may stop interacting with somebody depending on your assessment of her dancing skills and agreeableness. Suppose you believe a given individual i to be a poor dancer. If you find i disagreeable, you are unlikely to interact with her again, and thus your belief about her poor dancing skill will tend to persist even if she is in fact a good dancer. If, on the contrary, you believe i to be agreeable, you are likely to dance with her again. Doing so, you might discover that she is in fact a skilled dancer. Overall, this sequential process of belief formation and information sampling implies that the distribution of estimates will diverge from the distribution of attributes in the population of swing dancers. Because combinations of estimates that lead to avoidance (low perceived dancing skills, low perceived agreeableness) will be more stable than combinations of estimates that lead to further sampling (e.g. low perceived dancing skills, high perceived agreeableness), combinations of estimates that lead to avoidance will be over-represented.

Now, suppose that you want to interact with people that are good on the two dimensions. That is, you only want to interact with those you perceive to be good dancers *and* agreeable. In this case, adaptive sampling will lead you to believe that the two attributes are *negatively* correlated in the population. More generally, the sign of the illusory correlation depends on how the decision maker combines estimates in making her sampling decisions (for further explanation and boundary conditions, see Denrell & Le Mens, 2011b).

This model provides an alternative account of phenomena such as the ‘halo’ effect in person perception or the documented tendency for people to like proximate others better than distant others. Again, these phenomena are shown to emerge from adaptive sampling – the key to the above results is how the attributes of others are sampled and when you stop sampling those.

Conclusion: The Rationality of Biased Judgments

Our argument so far has been that decision makers who follow well-known learning processes will end up making seemingly biased judgments because of the sample bias generated by adaptive sequential sampling. Would not a rational person, who understands the effect of sample bias, be able to correct for it and thus avoid the above judgment biases?

Interestingly, the answer is no. It can be demonstrated that several of the above judgment biases continue to hold even if it is assumed that the decision-maker is rational, follows Bayes’ rule in updating beliefs, and is aware of the sample bias in the available data (Denrell, 2007; Le Mens and Denrell, 2011).

Consider first the basic negativity bias: most decision makers end up underestimating the uncertain alternative. One might suspect that this result emerges only when decision makers do not follow Bayes' rule in updating their estimates of the value of the uncertain alternative. As Denrell (2007) showed, however, the basic result continues to hold even if decision makers are rational and follow the optimal (expected payoff-maximizing) learning strategy. That is, decision makers follow Bayes' rule, they have an accurate prior, and they are able to compute the optimal amount of experimentation. Even under these conditions most decision makers will end up underestimating the uncertain alternative.

This sounds paradoxical - how can rational decision makers have a tendency to underestimate the value of the uncertain alternative? The key to resolving this paradox is first to realize that when outcomes matter, rational decision makers will not prioritize accuracy at the cost of obtaining poor outcomes. In particular, if they have tried the uncertain alternative several times and payoffs have been negative, it is optimal to avoid that alternative and choose the known alternative instead. By doing so they will stop getting further information about the uncertain alternative and will not be able to correct errors of underestimation. As a result, errors of underestimation will be more likely than errors of overestimation as in the case when heuristic choice rules are used (like in our simulations). Nevertheless, when decision makers are rational estimates are correct on average (the expected estimates of the alternatives, across decision makers, are unbiased). Thus, even if a rational decision maker is aware that she is more likely to underestimate than overestimate the value of the uncertain alternative, she would have no incentive to change her estimate. There is nothing paradoxical about this; it simply reflects the fact that the distribution of estimates is unbiased as well as skewed (with most estimates being negative).

Consider, next, the effect of information access on beliefs and preferences. Le Mens and Denrell (2011) show that this effect also continues to hold even if decision makers are rational in the sense that they update their beliefs following Bayes' rule, they are aware of the possible sample bias, and they follow an optimal policy of experimentation (an optimal learning policy). Thus, even if decision makers were rational they would end up being more likely to believe that an alternative for which information is more accessible is superior.

These results illustrate how adaptive sampling and constraints on information access can, without the further assumption of biased information processing, lead to biased judgments. In line with recent work on 'rational analysis,' this suggests that what appears to be irrational behavior could possibly be a rational solution to a problem different from the one that the researcher had in mind (Anderson, 1990; Dawes &

Mulford, 1996; Klayman & Ha, 1987; Oaksford & Chater, 1994). In making this suggestion we are by no means claiming that cognitive biases are unimportant - there is conclusive evidence that they have substantial effects. Rather, our approach suggests alternative, complementary, explanations that may be important in settings where information has to be sequentially sampled and people care about outcomes as well as accuracy. Our approach also has important normative implications. For example, to eliminate the in-group bias, it may not be enough to de-bias how people process social information. Rather, information about out-group members needs to be provided.

References

- Allan, L. G. (1993). Human contingency judgments: Rule based or associative? *Psychological Bulletin*, 114, 435-48.
- Anderson, N. H. (1981). *Foundations of information integration theory*. New York: Academic Press
- Anderson, J. (1990). *The adaptive character of thought*. Lawrence Erlbaum.
- Arrow, K. J. (1971). *Essays in the theory of risk bearing*. Chicago, IL: Markham.
- Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, 55, 591– 621.
- Dasgupta, N. (2004). Implicit ingroup favoritism, outgroup favoritism, and their behavioral manifestations. *Social Justice Research*, 17(2):143– 69
- Dawes, R., & Mulford, M., (1996). The false consensus effect and overconfidence: Flaws in judgment or flaws in how we study judgment? *Organizational Behavior and Human Decision Processes*, 65(3): 201–211.
- 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. (2008). Indirect Social Influence. *Science*, 321(5885): 47 - 48
- Denrell, J., & Le Mens, G. (2007). Interdependent Sampling and Social Influence. *Psychological Review*, 114 (2): 398-422.
- Denrell, J., & Le Mens, G. (2011a). Seeking positive experiences can produce illusory correlations, *Cognition*, forthcoming.
- Denrell, J., & Le Mens, G. (2011b). Learning To Be Satisfied with the Status Quo, *working paper*.
- Denrell, J. & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, 12, 523-538.

- DiDonato, T.E., Ullrich, J., & Krueger, J.I. (2011). Social Perception as Induction and Inference: An Integrative Model of Intergroup Differentiation, Ingroup Favoritism, and Differential Accuracy. *Journal of Personality and Social Psychology*, 100, 66-83.
- Einhorn, H. J., & Hogarth, R. M. (1978). Confidence in judgment: Persistence of the illusion of validity. *Psychological Review*, 85, 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., Schachter, S., & Back, K. W. (1950). Social pressures in informal groups: A study of human factors in housing. Stanford, CA: Stanford University Press.
- Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review*, 107, 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, England: Cambridge University Press.
- Fiske, S. T., & Taylor, S. E. (2007). *Social Cognition: from Brains to Culture* (3rd ed.). New York: McGraw-Hill.
- Gilovich, T. (1991). *How we know what isn't so: The fallibility of human reason in everyday life*. New York: Free Press.
- Granovetter, M.,S. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 83, 1420-1143.
- Keynes, J., M. (1936). *General Theory of Employment Interest and Money*. Macmillan, London, UK.
- Hamilton, D. L., & Gifford, R. K. (1976). Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology*, 12, 392-407.
- Hewstone, M., Rubin, M., & Willis, H. (2002). Intergroup Bias. *Annual Review of Psychology*, 53 (1): 575.
- Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24, 1–55.
- Juslin, P., Winman, A., & Olsson, H. (2000). Naive empiricism and dogmatism in confidence research: A critical examination of the hard-easy effect. *Psychological Review*, 107, 384-396.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decisions under risk. *Econometrica*, 47, 263-291.
- Klayman, J., & Ha, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, 94, 211–228.

- Levin, S., van Laar, C., & Sidanius, J. (2003). The Effects of Ingroup and Outgroup Friend- ships on Ethnic Attitudes in College: A Longitudinal Study. *Group Processes & Intergroup Relations*, 6 (1): 76-92.
- Le Mens, G., & Denrell, J., (2010). A Systematic Effect of Access to Information on Impression Formation, *unpublished manuscript*.
- Le Mens, G., & Denrell, J. (2011). Rational Learning and Information Sampling: On the 'Naivety' Assumption in Sampling Explanations of Judgment Biases, *Psychological Review*, forthcoming.
- Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. New York: Wiley.
- March, J. G. (1996). Learning to be risk averse. *Psychological Review*, 103: 309-19.
- Newcomb, T. M. (1961). *The acquaintance process*. New York: Holt, Rinehart & Winston.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101(4): 608–630.
- Oaksford, M. and Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford University Press, USA.
- Pettrigrew, T. F., and Tropp, L.R. (2006), A meta-analytic test of intergroup contact theory. *Journal of Personality and Social Psychology*, 90(5): 751-783.
- Segal, Mady W. 1974. Alphabet and Attraction: An Unobtrusive Measure of the Effect of Proximity in a Field Setting. *Journal of Personality and Social Psychology*, 30, 654-657.
- Smith E.R., & Collins E.C. (2009). Contextualizing person perception: Distributed social cognition. *Psychological Review*, 116, 343–364.
- Sutton, R., & Barto, A. G. (1998). *Reinforcement learning*. Cambridge, MA: The MIT Press.
- Thorndike, E. L. (1911). *Animal intelligence: Experimental studies*. Lewiston, NY: Macmillan.
- Tversky, A. & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, 27: 1124-1131
- Wood, W. (2000). Attitude change: Persuasion and social influence. *Annual Review of Psychology*, 51, 539–570.