Sampling of Social Information: Decisions from Experience in Bargaining

Nadine Fleischhut (nadinefl@mpib-berlin.mpg.de)[1]
Florian Artinger (artinger@mpib-berlin.mpg.de)[1,2]
Sebastian Olschewski (sebastian.olschewski@unibas.ch)[3,4]
Kirsten G. Volz (kirsten.volz@cin.uni-tuebingen.de)[3]
Ralph Hertwig (sekhertwig@mpib-berlin.mpg.de)[1]

[1] Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany
[2] Warwick Business School, University of Warwick, Coventry CV4 7 AL, UK
[3] Centre for Integrative Neuroscience, Otfried-Müller-Str. 25, 72076 Tübingen, Germany
[4] Department of Psychology, University Basel, Missionsstrasse 62a, 4055 Basel, Switzerland

Abstract

Whenever people depend on others, information about how those around them are likely to behave is important for the pursuit and achievement of goals. We investigated how the way such social information is learned, by *description* or *experience*, affects offers in a bargaining situation. Participants learned how often each offer had previously been accepted or rejected, either as probability information or by experiencing others' responses. When participants had to draw a representative sample of responses, the proportion of risky offers decreased under social experience, resulting in a gap to the description condition. When participants could terminate sampling whenever they wanted, however, no description–experience gap was observed. The sampling pattern suggests that participants disregarded probability information and instead relied on the allocation as proxy for risk. Accordingly, a certain amount of social experience seems necessary to overwrite people's initial expectations and change their behavior. Under what conditions people search for social information is crucial for understanding how and when it impacts behavior.

Keywords: decisions from experience; decision–experience gap; bargaining; social influence; learning; sampling; social interaction

Introduction

After the German election in fall 2013, the leadership of the Social Democratic Party put the decision as to whether or not to form a coalition government with the Christian Democrats—the plan favored by the leadership—to a vote among the party's rank-and-file members. Had the majority of votes been against the coalition, the chair of the party, so pundits' opinion, would have had to resign. Yet the popular vote was in favor of a grand coalition, thus giving his agenda the necessary backing. In this and many other social situations, an actor's gains or losses depend on others' behavior. The risks people face are not (or not exclusively) due to chance or natural mechanisms, but lie in uncertainty about how others will respond to their actions. Any information that helps to predict the future behavior of others is therefore instrumental in achieving (or failing to achieve) one's goals.

The relevance of social information (about others' goals, intentions, behavior, etc.) has been investigated in areas such as social comparison, social norms, and social learning (e.g., Carpenter, 2003), with research ranging from field experiments (e.g., Goldstein et al., 2008) to experimental investigations of social interactions in games (e.g., Bicchieri & Xiao, 2009). Here, we were interested in how the way information about others is learned affects social interaction.

When making decisions, people sometimes have access to summary information about others' behavior. For instance, before members of SPD were asked to vote, opinion polls indicated that 78% of SPD voters were in favor of a grand coalition. Often though, people have to rely on observations sampled from their social environment (e.g., Galesic et al., 2012) rather than from the population at large—for instance, when drawing on the experiences of others to decide whether or not it will pay off to accept an unpaid internship with a company in the hope of landing a job there later.

Decisions from experience and from description

Research on risky choice has demonstrated that people's decisions can differ substantially depending on how information is learned (Hertwig & Erev, 2009). In a typical study on risky choice, people are asked to choose between two options, a safe option (with one certain outcome, say $3) and a risky option (e.g., $32 with a probability of 10%; $0 otherwise). Less people choose the risky option when sampling the distribution of outcomes (decisions from *experience*) than when deciding based on a summary statement about probabilities and outcomes (decisions from *description*). If the rare event is undesirable (e.g., –$32 instead of +$32), the pattern reverses. In general, as Hertwig and Erev (2009, p. 518) summarized: "In decisions from experience, people behave as if the rare events have less impact than they deserve according to their objective probabilities, whereas in decisions from *description* people behave as if the rare events have more impact than they deserve (consistent with cumulative prospect theory)."

One explanation for the resulting *description–experience* gap (DE gap) is a sampling error: Because people typically draw only a small number of samples when they are permitted to terminate sampling whenever they want (*free sampling* paradigm), they are likely to undersample rare outcomes. Yet an (attenuated) gap remained even in studies that controlled for sampling error by increasing the size of samples people drew (e.g., Hau et al., 2008; Ungemach et al., 2009). How information about risks is learned is important not only because it affects choices, but because it may trigger different cognitive processes (including recency and evidence aggregation rather than multiplicative calculus; for a discussion, see Hertwig & Erev, 2009; Rakow & Newell, 2010).

To date, studies investigating the impact of social information about group or population behavior have typically provided *descriptive* social information in varying formats: information about the behavior of a small sample of others (Krupka & Weber, 2009); about average behavior (e.g., Bohnet & Zeckhauser, 2004); about majority choices using probabilities (e.g., Bicchieri & Xiao, 2009; Goldstein et al., 2008); or about the distribution of choices typically observed (e.g., Harrison & McCabe, 1996).

In order to understand how and when social information influences behavior, however, it is important to focus not only on the information per se but on how it is learned. In line with this, Martin et al. (2013) reported a DE gap in the social context of an iterated prisoner's dilemma—but this approach concerned information about people's interdependence and not information about the other's behavior.

Predictions for sampling the social environment

To investigate how the way information is learned affects people's behavior in socially uncertain situations, we used a version of the most widely studied social bargaining situation, the ultimatum game (Güth et al., 1982). As in most of the lottery studies used to investigate the DE gap, choices in the mini-ultimatum game (mini-UG; Bolton & Zwick, 1995) are binary: A proposer chooses between two possible allocations of money, which a responder may accept or reject. If the responder accepts, the allocation is implemented; if she rejects, both receive nothing. Because responders tend to reject too low or too unequal offers, the proposer faces a "social" risk due to the responder's option of rejection. We provided proposers with information on how often each allocation had been accepted or rejected in previous experiments, varying the way this information was presented (by *description* or by *experience*). The information allowed proposers to learn how (socially) acceptable and safe or risky each offer was in the population. Extrapolating from research findings on risky choice, we expected to find a DE gap in the same direction:

(1) An offer in which rejection (an undesirable outcome) is rare (< .2) is chosen more often in decisions from *experience* than in decisions from *description*. Conversely, an offer in which acceptance (a desirable outcome) is rare (< .2) is chosen less often in decisions from *experience* than in decisions from *description*.

It is important, though, to recognize the differences between the social and the individual risky choice setting: The risk people face in bargaining in the social context results not from chance, but from responders who also have a stake in the game. On the one hand, the risk of rejection represents a form of punishment if social expectations concerning, for instance, fairness are not met (Fehr & Gächter, 2002). On the other hand, people display a variety of other-regarding behaviors at their own expense even in social contexts without punishment (Fehr & Schmidt, 2006).

Whether learning about rejection rates matters may depend on proposers' motivation: For someone who displays other-regarding motivations (e.g., choosing the option with the smaller difference between herself and the responder), no social information is necessary. For someone who strives for the highest possible outcomes for himself, in contrast, assessing the risk of rejection should matter. It can be a warning sign or even no-go signal. Indeed, a purely self-interested person incurs a higher risk by ignoring rejection rates: Someone who chooses the higher outcome option has a mean risk of 40% of being rejected (MD = 30%) across the empirical choices used in this study, relative to a risk of only 6% (MD = 5%) of someone who choose the more equitable option irrespective of rejection rates. Accordingly, we expected the DE gap to be moderated by people's motivations in terms of social value orientation (Murphy et al., 2011):

(2) The DE gap occurs for participants who pursue their own self-interest and are thus, in theory, concerned with rejection rates (as long as they do not assume everybody else to be a rational and self-interested actor).

Studying how much information people voluntarily acquire opens a further window onto their decision processes. Under the assumption that search behavior depends on people's decision strategy (e.g., Hills & Hertwig, 2010), sample size should be larger for self-interested participants than for those who primarily seek to ensure fair outcomes or what is best for both:

(3) Self-interested participants sample more than those with other-regarding preferences.

Experiment

The goal of the experiment was to investigate whether the way information is learned affects people's choices in a socially uncertain situation in a similar way as it does choices in monetary lotteries. Using a 3 × 2 between-subjects design, we manipulated the way information was learned (*description* vs. *matched sampling* vs. *free sampling*) in two types of decision situations (mini UG vs. lottery).

Before making an offer, proposers in mini-UGs learned how often each allocation had been accepted (rejected) in previous experiments. In the *description* condition, we provided acceptance and rejection rates for both allocations in terms of probability information. In the *experience* conditions, proposers either had to sample 20 times from each option ("*matched sampling*," Ungemach et al., 2009) or stopped sampling whenever they felt ready to make a decision ("*free sampling*," Hertwig et al., 2004). The *matched sampling* condition ensured that participants saw a sample that exactly represented the underlying distribution so that no sampling error was involved. The *free sampling* condition permitted us to study how much information people searched for without any constraints. In all conditions, participants made 12 choices without feedback. In the mini-UG condition, proposers were matched with another person for each choice. To control for other-regarding motivations, we additionally implemented mini-dictator games with allocations identical to the mini-UGs', as well as a validated measure of social value orientation.

To be able to compare choices under social risk with risky choices, we used a binary lottery task as a benchmark. Probabilities were identical to rejection and acceptance rates in the mini-UGs, but no allocation to another person was made.

Methods

Decision situations. We pre-tested 43 mini-UGs to elicit acceptance and rejection rates from participants at the Technical University Berlin. Table 1 lists the 12 choice situations we used in the main experiment. The choice situations presented in the lottery conditions were identical, with the exception that the responder outcome was left out.

To test for the DE gap, we selected only choices in which each allocation was offered by more than 10% of proposers, which as lotteries did not contain a dominant option, and in which the ratio of expected values between the two options ranged between 0.5 and 2. In DE gap studies in risky choice, expected values are typically more similar, because choices can otherwise become trivial and the DE gap disappears. In the mini-UGs, expected values were based on real rejection rates and could not be manipulated freely (see Table 1). Rejection rates were rounded in steps of 5% to permit representation by a sample of 20 draws per option in the *matched sampling* condition.

In all choice situations, one allocation, representing a safe option, was accepted either always or in at least 90% of cases. For the other allocation, the (comparatively) risky option, the probability of being rejected (= receiving a zero outcome) systematically increased across choices: It was either $p < .2$ (Table 1, choice 1–4), $.2 < p < .5$ (choice 5–9), or $.5 < p < .8$ (choice 10–12).

Choices between a sure and a risky option (choice 1–4) are most similar to classical DE gap studies. For choices with two probabilistic options, findings are more ambiguous. Results by Ludvig et al. (2013) suggest that people may focus on the higher reward option—which, in our context, is always the more risky option. In general, more people should choose the risky option in the *experience* than *description* conditions if a rare outcome is undesirable (= rejection) (choice 1–9). Conversely, fewer people should choose the risky option in decisions from *experience* than in decisions from *description* if a rare outcome is desirable outcome (= acceptance) (choice 10–12). Especially for the choice situations that are less common in DE gap studies, the lottery conditions provide a benchmark against which mini-UGs can be compared.

Mini-dictator games used identical allocations yet without probabilities, because in dictator decisions the responder cannot reject the allocation. Social value orientation (SVO) was measured on 16 choices using the SVO Slider measure (Murphy et al., 2011), which is designed to distinguish selfish from altruistic, prosocial, and status-oriented choice patterns. Risk preferences were elicited with a classical lottery task (Holt & Laury, 2002).

Table 1: Choice situation in mini-UGs and lotteries. Columns list outcomes for the risky and safe option for oneself (own) and the responder (other). Reject columns show the probability of each option being rejected. The corresponding zero outcomes are not shown.

Choice	Options						EVratio
	R(isky)		S(afe)		R	S	
	Own	Other	Own	Other	Reject	Reject	
1	55	45	50	50	20	0	0.88
2	80	20	40	40	20	0	1.6
3	70	20	40	40	20	0	1.4
4	80	20	60	40	20	0	1.07
5	80	20	50	30	45	10	0.98
6	40	20	20	80	25	10	1.67
7	70	30	35	35	30	10	1.56
8	65	35	55	45	30	10	1
9	50	15	25	50	45	5	1.16
10	60	5	25	60	75	5	0.63
11	110	5	30	65	75	5	0.96
12	120	0	25	45	80	5	1.01

Information presentation. Before making a decision, proposers received information about the acceptance and rejection rates for each allocation. They were informed that these acceptance and rejection rates were collected in previous experiments at the same lab, from a sample of students drawn from the same population, and could provide information about the present participants' behavior.

In the *description* condition, two options (labelled "X" and "Y") were displayed next to each other on the screen. For each option, the probability of acceptance and rejection was shown (labelled "% accept" and "% reject"), followed by the respective outcomes for proposer and responder ("proposer receives/responder receives").

In the *experience* conditions, two covered card decks represented the two options ("X" and "Y"). Participants were able to learn the allocations and rejection rates by sampling cards from each deck in any sequence they wanted. After clicking on a deck, they would see a card for 800 ms before it was covered again. A card either displayed the outcomes for proposer and responder (meaning that the allocation was accepted), or two zero outcomes (meaning that the allocation was rejected). In the *matched sampling* condition, participants were forced to sample 20 times from each option before they could proceed to the decision screen. The frequencies with which cards occurred were matched exactly to the probabilities provided in the *description* condition, with the order of cards randomized for each participant. In the *free sampling* condition, participants could proceed to the decision whenever they wished, and the sequence of cards was randomly drawn from the distributions used in the *description* condition.

In the lottery conditions, presentation of outcomes and probabilities was identical, except that there was no payoff to a responder and no reference to acceptance or rejection.

Participants and procedure. 268 students (women = 115, mean age = 25.6) at the Technical University Berlin were randomly assigned to one of 12 experimental sessions (*n* = 20–24). All had participated in only few experiments before. After receiving instructions for the first task, participants in mini-UGs had to correctly answer control questions before proceeding, and all conditions included a practice trial to familiarize participants with the task. In the mini-UG conditions, participants made choices in five tasks always presented in the same order: proposer, responder, dictator, SVO, and risk preferences. Responder choices were made second in order not to influence proposer offers, which constituted our main dependent variable. Responders stated for both allocations whether they would accept or reject it without knowing which one was actually offered (strategy method). Participants received no feedback and were informed that they would be randomly matched with another person for each of their decisions. In the lottery conditions, participants made 12 decisions and completed the risk preference task. All choices were randomized in order and screen position, except for the SVO and risk preference task. Finally, participants answered additional questions and provided demographic information. They were paid for one randomly drawn choice from each task (i.e., five in the mini-UG conditions and two in the lottery conditions) in addition to a show-up fee of €3. The experiment lasted approx. 60 minutes for mini-UGs and approx. 30 minutes for lotteries. Participants earned on average €13.10 and €7.40, respectively.

Results

How does the way information is learned affect offers?

Given that proposer choices were identical to choices in the lottery conditions (except for the responder outcomes), we predicted a difference between the *description* and *sampling* conditions in the same direction as observed in the lotteries. However, Figure 1 shows that the effect was different in size and direction.

In the lotteries, the results were as expected. More participants chose the risky option in the *free sampling* condition than in the *description* condition—and vice versa when the rare event was a desirable outcome (Table 1, choice 10–12). In 10 of the 12 choice situations, people chose as if underweighting the rare event of the higher reward option in *free sampling* (binomial test: $p = .039$). The DE gap was significant in 4 of the 12 choice situations ($p < .05$), which was not surprising given that choices were not designed for lotteries and diverged more than usual in expected value. For choices with a ratio of expected values between 0.8 and 1.2, 4 out of 7 gaps were significant. In the *matched sampling* condition, the pattern was qualitatively similar to *free sampling*, with 8 of the 12 choice situations pointing in the predicted direction (binomial test: $p = .39$). Yet the difference to the *description* condition was significant for only one choice situation. The stronger effect in *free sampling* may be due to an additional sampling error that cannot occur in *matched sampling*.

In the mini-UGs, a strong gap emerged between *matched sampling* and *description*—but in the opposite direction as in lotteries. In *matched sampling*, participants chose the risky option less often than in *description* in 11 of the 12 choices (binomial test: $p = .006$), and less often than in *free sampling* in all 12 choices (binomial test: $p < .001$). Although only one choice situation was significant, the pattern across choices was quite consistent. However, no gap was observed between *description* and *free sampling* (binomial test: $p = .38$).

Overall, the gap in mini-UGs was quite strong, but its direction was not in line with our hypotheses derived from the DE gap literature. Instead, fewer participants chose the risky option in *matched sampling* than in *description*, and no gap emerged between *description* and *free sampling*.

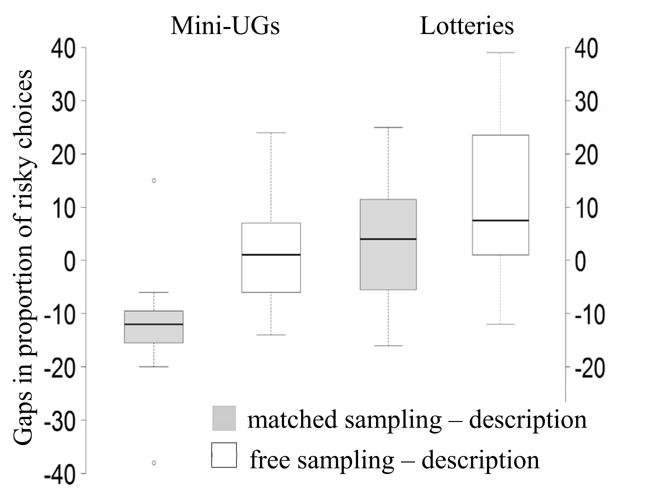

Figure 1: Differences in risky choices between the *sampling conditions* and *description*, separately for mini-UGs and lotteries. Each boxplot displays the distribution of differences in risky choices for all 12 decisions. The dark line shows the median and the whiskers extend to data within 1.5 interquartile range, excluding outliers (dots).

Does the DE gap depend on people's motivation?

Do the results change if we take into account that participants with other-regarding motivations do not require probability information—in other words, does the DE gap emerge for self-interested participants only (hypothesis 2)? When the analysis was restricted to participants motivated by their own self-interest according to the SVO measure, the results were even more pronounced (classifications based on allocations in dictator games correlated with SVO and are not reported). Self-interested participants chose the risky option less often in *matched sampling* than in *description*, again in 11 of the 12 choice situations (binomial test: p = .006), but the effect was now significant for 6 choice situations. At the same time, the difference was absent for people relying on other-regarding motivations (binomial test: p = 1), with no choice situation showing a significant difference.

This indicates that the gap in *matched sampling* was driven by self-interested participants who nevertheless refrained from choosing the risky option with the higher reward for themselves when forced to sample in the *matched sampling* condition. But why did no gap occur in *free sampling*?

Can the DE gap in mini-UGs be explained by sampling behavior?

How people searched for information in the *free sampling* conditions provided a window onto the decision process. Figure 2 shows that the average amount of information sampled differed strongly between the mini-UGs (MD = 7.5) and the lotteries (MD = 24.5), $t(56.25) = -7.20$, $p < .001$. In addition, the median number of samples drawn after a participant saw both allocations in the mini-UGs was 3 (M = 5, SD = 6.5), relative to a median of 12 samples in lotteries (M = 17, SD = 18). In 38% of the cases in which participants saw both allocations, they stopped sampling immediately or one card later, yet in lotteries this happened in only 7% of cases where both outcomes were seen. This pattern suggests that participants in mini-UGs sampled to learn about the allocation only, rather than about the probability. The fact that people strongly undersampled rare events due to small sample sizes, yet no gap emerged between *description* and *free sampling*, further supports the conclusion that probability information was not taken into account. In contrast, a gap occurred in *matched sampling*, where sample size was large and the sample was representative of the underlying distribution.

Consistent with this explanation, social information had only a small impact on proposer offers in *free sampling* and *description*. To test this impact, we subtracted offers of proposers with no information about rejection rates (taken from pilot experiments run to collect rejection rates) from offers in our conditions with social information. The proportion of risky offers without social information was very similar to the proportions in *description*, $t(20.50) = 0.27$, p = .79, and *free sampling*, $t(21.07) = 0.60$, p = .55, although people in these conditions had the chance to use additional social information. Only when people were forced to sample each option 20 times in *matched sampling* did the proportion of risky offers decrease under social experience, $t(22.00) = -1.86$, p = .076, suggesting that a certain amount of experience may be necessary to change behavior. Yet, on their own account, people did not choose to search for probability information in *free sampling*—or appear to use it in *description*.

Why smaller samples in mini-UGs?

Why did people in mini-UGs apparently disregard probability information? Our third hypothesis predicted that people with other-regarding motivations would sample less than people pursuing their self-interest because rejection rates are only central for the decision strategy of the latter.

However, participants with other-regarding motivations did not sample less (MD = 8.7, M = 8.7 SD = 4.2) than self-interested participants (MD = 5.8, M = 9, SD = 6.7), $t(39.98) = 0.207$, p = .84. In addition, a large share of the variance in sample size across participants was left unexplained. Both results raise doubts that other-regarding motivations can explain small samples and why people disregarded probability information.

A different line of explanation starts from the structure of the decision environment. One important question in understanding why participants spurned probability information is how well they can perform without it in a given environment. In the lottery choices used here, the size of outcomes and the risk of getting zero were strongly correlated (r = .56). Yet in the mini-UGs, the size of the responder outcome was an even better proxy for the risk of a zero outcome (r = –.79). This proxy was only available in mini-UGs, which could explain why participants sampled less than in lotteries. Even self-interested people who wanted to assess the risk of rejection did not need to rely on probability information, but could simply use the allocation as a proxy.

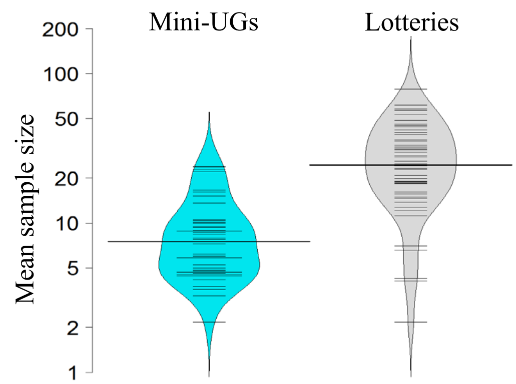

Figure 2: Distribution of participants' mean sample size in the *free sampling* conditions for mini-UGs and lotteries. The long horizontal line represents the median.

Discussion

Information about the likely behavior of others is of crucial importance in many social situations. We used mini-UGs to investigate how the way information is learned (by *description* or *experience*) affects decisions in risky social interaction, relative to decisions in lotteries with identical risks. In lotteries, decision from *description* differed as predicted from decisions in *free sampling*. The gap between *description* and *matched sampling* was much smaller, yet in the same direction.

In the social situation, in contrast, a strong gap emerged between *description* and *matched sampling*—yet in the opposite direction from lottery decisions. Here, fewer people chose the risky option in *matched sampling* than in *description*. The gap was mostly driven by self-interested proposers: When forced to draw a representative sample, they made fewer risky and self-interested choices than in *description* and *free sampling*. At the same time, there was no gap between *description* and *free sampling.* Despite the similarity of the decisions to individual risky choice, neither the direction of the gap nor when it occurred were in line with our predictions derived from lotteries, suggesting different decision processes in the risky social situation.

The second key finding was that people sampled less and stopped earlier in mini-UG than in lotteries, suggesting that proposers in *free sampling* decided based on the allocation rather than on the observed probabilities. Yet, contrary to our prediction, other-regarding motivations were not helpful in explaining the difference in sample size. Instead, a look at the environment revealed that the allocation provided a better proxy for risk than did outcomes in lotteries. As a consequence, even participants who took the risk of rejection into account could rely on the allocation alone instead of on social information. If no information about rejection rates is needed, no DE gap is observed.

Overall, the findings suggest that the DE gap between *description* and *matched sampling* likely has a different explanation in social situations than in risky choice. People's empirical expectations and proxies are likely to differ depending on how they would respond themselves and on what is acceptable in their social environment (Galesic et al., 2012). If self-interested people assume rejection rates to be lower than they in fact are, their expectations and behavior will change only when they are forced to experience a large enough sample to overwrite their initial expectations—as was the case in our *matched sampling* condition. Different expectations may also be responsible for the observed interindividual differences in sample size in *free sampling.* Both claims remain to be tested in future research. Yet the present results already offer insights into how and when social information impacts people's behavior—and highlight the importance of understanding the conditions under which people search for social information in the first place.

Acknowledgments

This research was funded by the Max Planck Institute for Human Development, Berlin.

References

Bicchieri, C., & Xiao, E. (2009). Do the right thing: But only if others do so. *Journal of Behavioral Decision Making*, *22*(2), 191–208.

Bohnet, I., & Zeckhauser, R. (2004). Social comparisons in ultimatum bargaining. *The Scandinavian Journal of Economics*, *106*(3), 495–510.

Bolton, G. E., & Zwick, R. (1995). Anonymity versus punishment in ultimatum bargaining. *Games and Economic Behavior*, *10*(1), 95–121.

Carpenter, J. P. (2003). Bargaining Outcomes as the Result of Coordinated Expectations. *Journal of Conflict Resolution*, *47*(2), 119–139.

Fehr, E., & Gächter, S. (2002). Altruistic punishment in humans. *Nature*, *415*(6868), 137–140.

Fehr, E., & Schmidt, K. M. (2006). The economics of fairness, reciprocity and altruism: Experimental evidence and new theories. In S. Kolm & J. Ythier (Eds.), *Handbook of the economics of giving, altruism, and reciprocity* (pp. 615–691). Amsterdam, The Netherlands: North-Holland.

Galesic, M., Olsson, H., & Rieskamp, J. (2012). Social sampling explains apparent biases in judgments of social environments. *Psychological Science*, *23*, 1515–1523.

Goldstein, N. J., Cialdini, R. B., & Griskevicius, V. (2008). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research*, *35*(3), 472–482.

Güth, W., Schmittberger, R., & Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*, *3*(4), 367–388.

Harrison, G. W., & McCabe, K. A. (1996). Expectations and fairness in a simple bargaining experiment. *International Journal of Game Theory*, *25*(3), 303–327.

Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description–experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, *21*(5), 493–518.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*(8), 534–539.

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, *13*(12), 517–523.

Hills, T. T., & Hertwig, R. (2010). Information search in decisions from experience: Do our patterns of sampling foreshadow our decisions? *Psychological Science*, *21*(12), 1787–1792.

Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, *92*(5), 1644–1655.

Krupka, E., & Weber, R. A. (2009). The focusing and informational effects of norms on pro-social behavior. *Journal of Economic Psychology*, *30*(3), 307–320.

Ludvig, E. A., Madan, C. R., & Spetch, M. L. (2013). Extreme outcomes sway risky decisions from experience. *Journal of Behavioral Decision Making*.

Martin, J. M., Gonzalez, C., Juvina, I., & Lebiere, C. (2013). A description–experience gap in social interactions: Information about interdependence and its effects on cooperation. *Journal of Behavioral Decision Making*.

Murphy, R., Ackermann, K., & Handgraaf, M. (2011). Measuring social value orientation. *Judgment and Decision Making, 6*(8), 771–781.

Rakow, T., & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for future research on experience-based choice. *Journal of Behavioral Decision Making*, *23*(1), 1–14.

Ungemach, C., Chater, N., & Stewart, N. (2009). Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychological Science*, *20*(4), 473–479.