Games and Economic Behavior 152 (2025) 241–256

Contents lists available at ScienceDirect

Games and Economic Behavior

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

Conditional cooperation under uncertainty: The social description-experience gap ☆

Orestis Kopsacheilis [a,b,*], Dennie van Dolder [c], Ozan Isler [d]

[a] TUMCS for Biotechnology and Sustainability, Technical University of Munich, Germany
[b] School of Management, Technical University of Munich, Germany
[c] Department of Economics, University of Essex, United Kingdom of Great Britain and Northern Ireland
[d] School of Economics, The University of Queensland, Australia

ARTICLE INFO

JEL classification:
C72
C92
D81
D83

Keywords:
Decisions from description
Decisions from experience
Prisoner's dilemma
Cooperation
Social uncertainty
Ambiguity

ABSTRACT

Conditional cooperation is typically studied in experimental settings where the behavior of others is known to subjects. In this study, we examine conditional cooperation under uncertainty. Using a novel experimental design, we exogenously manipulate the likelihood that a subject's partner in a Prisoner's Dilemma will cooperate. Information about the partner's cooperation is either presented descriptively or learned through experiential sampling. We observe a description-experience gap: subjects are more likely to cooperate under experience than under description when their partner's probability of cooperation is low, while the opposite holds when it is 50% or higher. This result contrasts with expectations deriving from the individual choice literature, where rare events are typically underweighted in experience-based decisions. We find that the gap we observe is driven by conditional cooperators being less responsive to social information acquired experientially than to that acquired descriptively. Furthermore, we show that stronger priors held by subjects under social uncertainty compared to individual uncertainty can account for the disparity with the individual choice literature.

1. Introduction

Human societies flourish through cooperation. A central aspect of human cooperation is the preference for conditional cooperation: many people are willing to cooperate if others do so as well, even when it conflicts with their material self-interest (Chaudhuri, 2011; Fehr and Schurtenberger, 2018). People are more likely to contribute to public goods (Keser and Van Winden, 2000; Fischbacher et al., 2001; Fischbacher and Gächter, 2010; Thöni and Volk, 2018; Isler et al., 2021), vote (Gerber and Rogers, 2009), donate to charity (Frey and Meier, 2004), pay taxes (Hallsworth et al., 2017), and conserve energy (Allcott, 2011; Allcott and Rogers, 2014) to the extent that others do the same.

Conditional cooperation is typically studied in settings where the behavior of others is known with certainty (Fischbacher et al., 2001; Thöni and Volk, 2018). Outside the laboratory, however, decisions to cooperate are often made under uncertainty. For example,

☆ An earlier version of this paper was circulated under the title "The Description-Experience Gap in Cooperation". We thank Mohammed Abdellaoui, Robin Cubitt, Chris Starmer, and Simon Gächter, as well as participants in various conferences and workshops, for useful comments. The research was supported by the Economic and Social Research Council [grant numbers ES/K002201/1, ES/P008976/1] under the Network for Integrated Behavioural Science. Ozan Isler also acknowledges funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie fellowship number 658186 (https://ec.europa.eu/).

* Corresponding author.
E-mail address: orestis.kopsacheilis@tum.de (O. Kopsacheilis).

https://doi.org/10.1016/j.geb.2025.04.012
Received 25 October 2022

Available online 4 May 2025
0899-8256/

a researcher embarking on a project with a new collaborator may not know how cooperative her collaborator will be. Similarly, a freelancer may be unsure whether a client will fulfill payment agreements. In such scenarios, conditional cooperation requires the formation of expectations about the behavior of others (Hayashi et al., 1999; Clark and Sefton, 2001; Van den Assem et al., 2012).

In this paper, we investigate conditional cooperation under uncertainty. We develop a novel experimental protocol that enables us to exogenously manipulate the likelihood of a subject's partner cooperating in a Prisoner's Dilemma, without using deception, employing a variation of Bardsley's "Conditional Information Lottery" (Bardsley, 2000).[1] Furthermore, we also vary the modality by which subjects receive this information: they are either provided with a description of the cooperation rate in the population from which their partner will be drawn ('Description') or can learn about this cooperation rate by sampling decisions of potential partners ('Experience'). In Experience, sampling does not directly affect payoffs; it only provides information.

Information about the cooperativeness of others can be acquired in various forms. In some cases, relevant descriptive information may be readily available. For example, a freelancer who lands a large client through an online labor market will often have access to numerous reviews from other freelancers that describe their experiences. This access allows for a reasonably accurate assessment of the client's trustworthiness. With sufficiently accurate descriptive information, the decision to cooperate or not can be viewed as a decision under risk, a special case of uncertainty where probabilities are objective and known. However, in many cases, people lack precise descriptive information and must rely on personal experiences or observations of others' experiences to form expectations. This is the case, for example, when a researcher deliberating whether to engage in a new collaboration or a freelancer who does not have access to a substantial corpus of client reviews. Under such conditions, the decision to cooperate can be considered a decision under ambiguity, where uncertainty extends beyond risk, and where probabilities are at least partially unknown.

Research on individual decision-making suggests that people's choices systematically differ between situations where outcomes and their probabilities are described and those where they are learned through experiential sampling. In its most common manifestation, the "description-experience gap" is consistent with the hypothesis that rare events are less influential in experience-based than in description-based choices (Barron and Erev, 2003; Hertwig et al., 2004; Hertwig and Erev, 2009; Wulff et al., 2018). Sampling bias, which arises when individuals in Experience sample only a small number of observations and, therefore, causing their sample to frequently underrepresent the objective frequency of rare events, has been identified as a critical factor driving the underweighting of rare events in risky choices (Fox and Hadar, 2006; Rakow et al., 2008; Cubitt et al., 2022).[2] However, the gap persists in a reduced form even after accounting for sampling bias (see Wulff et al. 2018, for a meta-analysis). To examine how this statistical bias might influence a potential description-experience gap in social contexts, we investigate two variations of the Experience condition: one in which subjects can freely decide how much to sample ('E-Free'), allowing for biased samples, particularly if they sample sparingly, and one where sampling bias is eliminated by design ('E-Fixed'). In E-Fixed, subjects sample a fixed, predetermined number of times, ensuring that the observed frequency always matches the true likelihood.

Although the description-experience gap has been extensively studied in the context of individual uncertainty, its implications for decision-making under social uncertainty have received relatively little attention. This oversight is notable because responses to social uncertainty—where outcomes depend on the decisions of others—have been found to differ from those to uncertainty caused by "random acts of nature".[3] For instance, research on betrayal aversion indicates that individuals are less inclined to take risks when outcomes are determined by another person rather than by nature (Bohnet and Zeckhauser, 2004; Bohnet et al., 2008; Aimone and Houser, 2012). Furthermore, studies by Costa-Gomes and Weizsäcker (2008), Fetchenhauer and Dunning (2012), and Li et al. (2020) find that behavior in strategic games is less sensitive to changes in beliefs than in games against nature.

In our main experiment, we observe a statistically significant description-experience gap: subjects are more likely to cooperate in treatments where they learn through experiential sampling than in the treatment where probabilities are objectively described when their partner's likelihood of cooperation is below 50%, while the opposite is the case when this likelihood is at least 50%. Additionally, behavior is nearly identical across the two Experience treatments—whether sampling bias is present or eliminated.

In both Description and Experience treatments, we observe behavior consistent with conditional cooperation: subjects' cooperation increases monotonically with the likelihood that their partner will cooperate, regardless of whether this likelihood is described or learned through experience. We also examine between-subject heterogeneity. Specifically, we measure preferences for conditional cooperation using the strategy method, allowing subjects to base their choices on the actual behavior of their partner, thereby removing any uncertainty (Selten, 1967; Fischbacher et al., 2001). Our findings show that preferences elicited under information certainty accurately predict behavior under conditions of risk and ambiguity. As expected, the description-experience gap is driven by conditional cooperators rather than by free riders or unconditional cooperators, who should not consider their partner's actions when making their decisions. Furthermore, conditional cooperators sample more information than either free riders or unconditional cooperators.

[1] The elicitation of subjects' beliefs about others' cooperation in social dilemmas and the study of how these beliefs correlate with behavior have a long history in both psychology (Kelley and Stahelski, 1970; Kuhlman and Wimberley, 1976; Aksoy and Weesie, 2012; Pletzer et al., 2018) and economics (Offerman et al., 1996; Offerman, 1997; Croson, 2000; Fischbacher and Gächter, 2010; Dufwenberg et al., 2011). However, studying conditional cooperation through the correlation of beliefs with actions is inherently limited due to the endogenous nature of belief formation. This limitation is significant, as subjects may project their behavioral tendencies onto others, leading to reverse causality. Additionally, it is impossible to distinguish between a free rider and a cooperator if both believe that others are unlikely to cooperate (Gächter, 2007).

[2] We discuss the literature on the description-experience in individual uncertainty in greater detail in Subsection 2.1.

[3] There is considerable evidence that decisions under uncertainty depend on the source generating the uncertainty. The study of source-dependent uncertainty preferences was advanced by Tversky and colleagues in the 1990s (Heath and Tversky, 1991; Tversky and Kahneman, 1992; Tversky and Fox, 1995; Fox and Tversky, 1995). Additional empirical evidence for source-dependent preferences is provided by Keppe and Weber (1995), Kilka and Weber (2001), Hong Chew et al. (2008), and de Lara Resende and Wu (2010).

The pattern of the social description-experience gap we observe deviates from expectations based on the canonical finding in the literature on individual uncertainty. If subjects underweight rare events in experiential settings relative to descriptive ones, we would expect lower cooperation rates in Experience than in Description when the probability of the partner's cooperation is low. Instead, we observe the opposite pattern. Furthermore, while sampling bias has been identified as a key driver of the (relative) underweighting of rare events in studies on individual uncertainty, it has no significant effect on behavior in our experiment.

To address the disparity with the previous literature, we conduct two analyses. First, we disentangle the gap by constructing two novel indices: *cooperativeness* and *conditionality*. We find that the gap does not result from a treatment difference in the overall propensity to cooperate, but rather from subjects in Description being more sensitive to changes in the likelihood of cooperation than subjects in Experience. Second, using the Bayesian updating model by Carnap (1952), we show that stronger priors under social uncertainty than those under the types of uncertainty typically used in individual choice experiments can explain the difference. When individuals hold stronger priors, their posterior beliefs will be less responsive to new information. In a second experiment, we find confirmatory evidence for this sticky prior hypothesis: subjects are more confident in their prior about others' likelihood to cooperate in a Prisoner's Dilemma than in their prior about winning chances in a risky prospect.

To the best of our knowledge, only a few studies have examined the description-experience gap in social contexts. Artinger et al. (2012) examined this gap in a stochastic Public Goods Game, focusing on the uncertainty about the benefits of contributing rather than the uncertainty about the actions of others. Subjects were either informed about the potential benefits and their probabilities or learned these through sampling. Their findings revealed no significant description-experience gap in cooperation. Fleischhut et al. (2014) investigated the gap in an ultimatum bargaining scenario, where subjects learned about the rejection rates of offers from a previous experiment, either through description or through sampling. They observed a lower proportion of risky decisions under experience compared to description, but only after controlling for sampling bias. However, due to the lack of exogenous variation in rejection probabilities, their design did not allow them to draw conclusions about the moderating effect of probabilities on the gap between description- and experience-based choices.

While our study—and the two aforementioned studies—investigate the social description-experience gap by comparing learning from description to learning through experiential sampling without payoff consequences, two other studies have explored this gap by comparing learning from description to learning through repeated, payoff-consequential play. Martin et al. (2014) had subjects engage in an iterated Prisoner's Dilemma, where they either received a complete description of the game or learned about potential outcomes through repeated experience. Their findings showed higher cooperation levels when the game was fully described than when learned through experience. Moreover, cooperation was more prevalent when players could learn both their own and others' payoffs through experience, as opposed to only receiving feedback on their own outcomes. Similarly, Erev and Greiner (2015) conducted experiments on a repeated asymmetric stag-hunt game with two equilibria: one efficient with equal payoffs and one inefficient with unequal payoffs. Subjects were either provided with a complete description of the game or only learned about their own payoffs through repeated play. Their findings showed that subjects were more likely to play the efficient and equal equilibrium under description than under experience.

Unlike the studies by Artinger et al. (2012), Martin et al. (2014) and Erev and Greiner (2015), which focus on uncertainty about the structure of the game, our research focuses on uncertainty regarding others' actions. Furthermore, in contrast to Fleischhut et al. (2014), we exogenously manipulate the likelihood of cooperation using a novel experimental protocol that we introduce and validate. The key observation underlying our main findings is that people are less responsive to changes in the likelihood of cooperation when learning through social experience rather than through descriptions. Our design's ability to experimentally manipulate the likelihood of a subject's partner taking a specific action across the entire probability unit interval was instrumental for identifying this gap between Description and Experience in a social setting. While we employed this protocol to study behavior in the Prisoner's Dilemma, it can readily be adapted to study behavior in other games where players' decisions hinge on their beliefs about their partners' actions.

The paper proceeds as follows. Section 2 presents our main experiment, in which we study social behavior under description and under experience. Section 3 describes our follow-up experiment, in which we test whether people hold stronger priors under social uncertainty compared to abstract individual uncertainty. Finally, Section 4 discusses our findings and concludes.

2. Described vs. experienced social uncertainty

In our primary experiment, we introduce a novel experimental protocol to investigate how individuals' cooperative behavior responds to changes in the likelihood of being matched with a cooperator, and whether the mode of learning about this likelihood—through description or experience—affects these responses. In addition, we examine heterogeneity in behavioral responses across subjects with differing cooperative preferences.

Subsection 2.1 outlines three key design considerations when studying behavior under Description versus Experience, and explains the rationale behind our methodological choices. Subsection 2.2 details our experimental design, Subsection 2.3 describes the practical procedures for conducting the experiment, and Subsection 2.4 presents the results.

2.1. Implementing "Experience" in the lab

The description-experience gap has been extensively studied in the context of individual decision making. Although the general consensus in these studies is that rare events are less influential in Experience than in Description, there is considerable variability in the size of this gap (see Wulff et al., 2018, for a review). This variability can largely be attributed to differences in the way Experience is implemented.[4]

Here, we consider the three most consequential variations in this experimental design. The first categorizes the literature into two main branches based on the payoff implications of experiential sampling: (i) studies in which subjects sample information without monetary consequences (often referred to as "sampling paradigm", Hertwig et al. 2004), and (ii) studies in which sampling does entail monetary consequences (also known as the "clicking paradigm", Barron and Erev 2003).

The second variation concerns whether the sampling technology allows for sampling bias. In the original paradigm, subjects were free to decide how much to sample. As a result, the relative frequencies of the outcomes observed by the subjects often did not correspond to their objective probabilities. Specifically, subjects who sample only a small number of observations will often underrepresent the true likelihood of rare events and, in many cases, fail to sample such events even once, leaving them unaware of their existence. To study the description-experience gap in the absence of sampling bias, other studies employ the so-called matched-sampling paradigm, in which subjects have to sample a fixed number of times, and it is ensured that the observed relative frequencies match the underlying probabilities (Ungemach et al., 2009; Barron and Ursino, 2013; Glöckner et al., 2016; Aydogan and Gao, 2020; Cubitt et al., 2022).

The third variation concerns the information available to subjects about the potential outcomes. In the original sampling paradigm, subjects face uncertainty about both outcome probabilities and set of possible outcomes. To focus on the effect of probability weighting, later studies inform subjects about the set of possible outcomes and focus on the uncertainty surrounding their probabilities (e.g. Abdellaoui et al. 2011; Kopsacheilis 2018; Aydogan and Gao 2020).

Each of these variations has been shown to influence the magnitude of the description-experience gap. Camilleri and Newell (2011) compared behavior in situations where sampling entails monetary consequences to behavior in situations where it does not, and found that the gap is more pronounced when sampling has monetary consequences. Furthermore, sampling bias has been identified as a key contributor to this gap (see Wulff et al. 2018 for a meta-analysis). Cubitt et al. (2022) compared behavior between treatments with and without sampling bias and report that the gap is substantially reduced when the relative frequencies in the observed sample match the objective probabilities. Lastly, Erev et al. (2008) demonstrate that decisions made under experience more closely resemble those made under description when subjects in the experience condition are informed of all possible outcomes, compared to when this information is not provided—a phenomenon they refer to as "the mere presentation effect".

Although these variations affect the magnitude of the description-experience gap, they do not alter its direction.[5] Even studies that use the sampling paradigm, control for sampling bias, and inform subjects about all possible outcomes report a significant gap (e.g. Aydogan and Gao 2020; Cubitt et al. 2022). However, in these studies, underweighting in experience occurs only in relative terms: subjects overweight rare events in both Description and Experience, but this overweighting is less pronounced in Experience.

In our experiment, we implement the paradigm in which sampled outcomes have no monetary consequences. When payoff-consequential sampling has been used in a social context, it typically involves having subjects play a repeated game with another subject (Martin et al., 2014; Erev and Greiner, 2015). Such a setup is not suitable for our objective, as the cooperation rate emerges endogenously, whereas our goal is to manipulate it exogenously. Additionally, this setup would create a confound between information and prior earnings, as observing cooperation would result in higher payoffs than observing defection. By contrast, the sampling paradigm allows us to manipulate the information subjects receive about their partner's likelihood of cooperation while holding other variables constant. This paradigm is also simpler to implement, as it does not require synchronous interaction.

Regarding sampling bias, we designed two treatments: one in which sampling bias is likely to occur and one in which it is eliminated by design. This enables us to examine the effect of sampling bias on behavior under social uncertainty. In addition, we provide subjects with complete information about potential outcomes. Given our focus on social uncertainty regarding others' actions, we aimed to eliminate any uncertainty about the payoff structure of the game by providing a detailed description of the consequences of specific actions taken by the subject and their partner. We included comprehension questions to verify that subjects fully understood the game (see Online Appendix B for the complete instructions). This approach ensured that subjects were fully aware of the rules of the game and that the uncertainty they faced pertained solely to the likelihood of their partner cooperating.

2.2. Experimental design

The experiment consists of three treatments and employs a between-subjects design. Each treatment has three stages. In each stage, subjects play one-shot Prisoner's Dilemmas with the payoff structure shown in Table 1. Stages 1 and 3 are identical across treatments, while Stage 2 contains the experimental manipulation. Subjects are informed at the outset that the study comprises three stages, but detailed instructions are only provided at the beginning of each stage.

In Stage 1, subjects are asked to make their cooperative decision in a one-shot Prisoner's Dilemma, played with a randomly selected other subject. Subjects do not receive feedback on their partner's decision. The primary purpose of this stage is to elicit decisions that will be used to incentivize choices in the following two stages.

In Stage 2, subjects are asked to make decisions for seven independent one-shot Prisoner's Dilemmas. Each subject i is informed that she will be re-matched with another subject, j, who will be randomly selected from a subpopulation of subjects. Subject i is

[4] By contrast, implementing Description simply involves describing all outcomes and their associated probabilities, and does not vary considerably across studies.

[5] For the robustness of the gap with payoff-inconsequential sampling, see the meta-analysis by Wulff et al. (2018); for robustness when controlling for sampling bias, see, for example, Ungemach et al. (2009) and Barron and Ursino (2013); for robustness to the mere presentation effect, see Abdellaoui et al. (2011) and Kopsacheilis (2018).

Table 1
Payoff matrix for the Prisoner's Dilemma.

	Keep	Share
Keep	50,50	150,0
Share	0,150	100,100

told that she will play a one-shot Prisoner's Dilemma with subject j, where subject j will act according to her first stage decision, whereas subject i will be asked to make a new decision.[6] Before making their decision, subjects are given the opportunity to acquire information about the cooperation rate of the subpopulation from which their partner will be drawn.

To observe how subjects condition their level of cooperation on the likelihood that their partner will cooperate, we ask subjects to make choices for seven different potential subpopulations—one real and six hypothetical. Across the seven scenarios, we systematically vary the "Subpopulation Probability of Cooperation" ($SPoC$), defined as the proportion of Stage 1 cooperative decisions in the subpopulation from which subject j will be drawn. This represents the (objective) probability that i will face a cooperative j. The order in which scenarios are encountered is randomized for each subject.

We consider seven levels of $SPoC$ that span the probability spectrum: $\{0, 0.1, 0.3, 0.5, 0.7, 0.9, 1\}$. Let r be an index that runs through the different levels of $SPoC$ in ascending order. We denote the r^{th} level of $SPoC$ as $SPoC_r$, with $SPoC_1 = 0$ and $SPoC_7 = 1$. Unbeknownst to subjects, the actual size of the subpopulation from which their partner will be drawn is set to two, ensuring that the true scenario is always captured by a $SPoC$ of 0, 0.5, or 1. Thus, the design ensures that exactly one level of $SPoC$ from the seven predefined levels is real, while the remaining six are hypothetical.[7]

Subjects are asked to make a decision for each potential subpopulation, without feedback, with the understanding that only the decision for the real scenario will be payoff-relevant. Importantly, as subjects do not know which scenario corresponds to the actual subpopulation's cooperation rate, it is incentive-compatible for them to treat each task as if it were real. Our implementation is a variant of Bardsley's (2000)'s conditional information system, which combines elements of the strategy method (Selten, 1967) and the random incentive system (see for example Starmer and Sugden, 1991).[8]

The way subject i obtains information about $SPoC$ varies across treatments. Subjects are randomly assigned to one of three treatments: decisions from description (Description), decisions from experience with free sampling (E-Free), and decisions from experience with fixed sampling (E-Fixed). In Description, subjects learn about $SPoC$ through numerical descriptions. For example, when $SPoC = 0.7$, the screen displays "*70% of your group chose to Keep and 30% of your group chose to Share*".[9]

In E-Free, subjects can sample decisions of members of the subpopulation one at a time, as many times as they like. The sampling process is conducted *with* replacement, ensuring that the observed cooperation probability converges to the objective $SPoC$. In E-Fixed, subjects similarly sample decisions of members of the subpopulation, but they must do so exactly ten times, and the observed relative frequency of cooperation always matched the objective $SPoC$ level of the presented scenario. To guarantee this, sampling in E-Fixed is conducted *without* replacement.[10] We do not inform subjects in E-Fixed that the observed frequency matches the objective frequency. The goal of this design choice is to eliminate sampling bias but keep the exact probability ambiguous to the subject, as ambiguity is considered an integral component of learning from experience (Abdellaoui et al., 2011).[11]

Finally, in Stage 3, we elicit cooperation preferences using the strategy method in a sequential Prisoner's Dilemma. Specifically, we ask subjects what they would do if their partner chose *Keep* and what they would do if their partner chose *Share*. A key difference between the tasks in Stage 2 and Stage 3 is that, unlike most conditional decisions in Stage 2, which take place under social uncertainty, the conditional decisions in Stage 3 are made under certainty. As in Stage 2, subjects' partners are randomly determined and play using their Stage 1 decision. This procedure is based on the standard method for eliciting conditionally cooperative preferences in Public Goods Games (Fischbacher et al., 2001). Based on their decisions in this game, subjects can be categorized into one of four

[6] One could argue that subjects in Stage 2 are effectively playing a sequential one-shot Prisoner's Dilemma with imperfect information about j's action. From a theoretical perspective, the fact that the game is sequential rather than simultaneous is irrelevant. However, there is evidence that people are less likely to cooperate if they know that their match has already made their decision (Shafir and Tversky, 1992; Bardsley, 2000). Given that all games played in Stage 2 share this characteristic, the sequential nature of the game cannot account for any potential treatment differences we observe.

[7] Our manipulation of cooperation levels among subjects' potential partners bears some resemblance to research on assortative matching in social dilemmas. For example, Gächter and Thöni (2005) investigate repeated Public Goods Games in which subjects are either assigned to random groups or to groups composed of individuals with similar cooperative preferences. They find that groups of like-minded cooperators are able to sustain higher levels of cooperation than the most cooperative randomly formed groups. Compared to such studies, our experimental design offers greater control. By having subjects play against past decisions of their potential partners, we exogenously fix the cooperation rates that subjects encounter. Additionally, by employing a variant of Bardsley's "Conditional Information Lottery" (Bardsley, 2000), we maintain full control over these cooperation rates.

[8] We opted against using larger subpopulations (e.g. with 10 subjects) due to time considerations. Eliciting cooperation for the four additional $SPoC$ values that are possible at this group size $(0.2, 0.4, 0.6, 0.8)$ would have markedly increased the duration of the experiment and likely led to higher attrition—especially for the Experience treatments, where subjects also had to complete a sampling stage for each $SPoC$ before making a decision.

[9] Technically, the term "subpopulation" is more accurate than the term "group", as the latter is commonly used to imply an interaction among all group members (such as in Public Goods Games). Nonetheless, we chose to use the term "group" in the instructions, as we deemed it more intuitive for the subjects. Details of the instructions can be found in Online Appendix B.

[10] Even though the observed relative frequency is identical for every subject facing a given $SPoC$, the order in which cooperation or defection events are sampled is randomized for each subject and each $SPoC$.

[11] Cubitt et al. (2022), however, report that whether or not subjects are informed about the fact that they have observed the objective frequency does not materially affect their choices.

types: *conditional cooperators*, who match their partner's decision; *free riders*, who always defect; *unconditional cooperators*, who always cooperate; and *others*.[12]

2.3. Experimental procedures

The experiment was conducted online using the Qualtrics survey software. We recruited 1,094 subjects through Prolific (www.prolific.co; Peer et al. 2017).[13] The one-shot nature of the Prisoner's Dilemma in all three stages of our experiment, combined with the fact that subjects did not receive feedback on the outcome of each Prisoner's Dilemma, eliminated the need for real-time interaction between subjects.

Before making their decisions, subjects were required to correctly answer a comprehension question designed to ensure they understood the payoff consequences of both their own and their partner's choice in the Prisoner's Dilemma for both subjects' earnings. Those who failed to answer correctly after three attempts were given a participation fee but were not allowed to proceed with the experiment. In total, 43 subjects (3.9%) failed the comprehension check (17 in Description, 16 in E-Free, 10 in E-Fixed). Additionally, 61 subjects did not complete the experiment (5.6%; 9 in Description, 21 in E-Free, and 31 in E-Fixed). This resulted in a final sample size of 990 subjects (mean age = 36.0, SD = 12.1; 61.6% female).

The experiment lasted approximately 20 minutes on average. Subjects were paid a fixed participation fee of £1.25 and could earn an additional variable amount between £0.00 and £1.50. Only one choice per subject, selected at random, determined their variable payment. The average total earnings were approximately £2.00 per subject. The allocation of subjects to treatments was random; however, the likelihoods were not uniform—subjects were more likely to be assigned to the E-Free treatment than to the other two treatments. This was done to increase statistical power for the analysis of sampling decisions in that treatment. In total, we observed the behavior of 279 subjects in the Description treatment, 276 in the E-Fixed treatment, and 435 in the E-Free treatment. The instructions used in the experiment are detailed in Online Appendix B, and the matching protocol for payments is explained in Online Appendix A.3.

2.4. Results

2.4.1. Cooperation patterns across treatments

Across all treatments, 57.9% of the overall sample cooperated in Stage 1. This cooperation rate did not differ significantly across treatments ($\chi^2(2, 990) = 3.84, p = 0.146$) and is similar to that observed in other experiments studying one-shot Prisoner's Dilemmas (Sally, 1995; Mengel, 2017).

Next, we turn to the main task of Experiment 1. Fig. 1 presents the average cooperation rates in the Stage 2 Prisoner's Dilemma across the different $SPoC$ levels for each treatment. In every treatment, the cooperation rate increases monotonically with the $SPoC$ level, indicating that subjects are more likely to cooperate when the probability of being matched with a cooperator is higher. This pattern highlights a strong tendency towards conditional cooperation.

At the same time, Fig. 1 also reveals signs of unconditional behavior. Notably, a sizable fraction of subjects choose to cooperate even when the probability of being matched with a cooperator is zero ($SPoC = 0$), and an even larger fraction choose to defect even when this probability is one ($SPoC = 1$).[14]

Result 1. *Subjects' willingness to cooperate increases monotonically with the probability that their partner will cooperate.*

There are significant differences in cooperation patterns across the three treatments. As shown in Fig. 1, cooperation rates in the Experience treatments are statistically significantly higher than those in the Description treatment when cooperation is relatively infrequent ($SPoC < 0.5$).[15] Conversely, cooperation rates in the Experience treatments are lower than those in the Description treatment when cooperation is more frequent ($SPoC \geq 0.5$); however, this difference is not statistically significant, except at $SPoC =$ 0.5 and 0.7 (the latter only when comparing Description with E-Free, and only at the 10% level).

Therefore, although we observe a description-experience gap in cooperation, the pattern we identify diverges from the canonical finding in individual decision-making under risk and uncertainty. If rare events were less influential in our Experience treatments relative to our Description treatment, we would expect lower cooperation rates in the former than in the latter when the probability of the partner's cooperation is low. Contrary to this expectation, we observe the opposite pattern.

Result 2. *There is a significant description-experience gap in cooperation. This gap is particularly pronounced when the likelihood that one's partner will cooperate is low. In these scenarios, subjects in Experience treatments tend to cooperate more than those in the Description treatment.*

[12] The *other* category consists of subjects who defect when their partner cooperates and cooperate when their partner defects. They are sometimes referred to as "reverse conditional cooperators" and usually represent a tiny minority.

[13] Our selection criteria required subjects to be UK residents with an approval rating of 90 or above.

[14] Only subjects in the Description treatment are certain of this probability. Subjects in the Experience treatments do not know whether the cooperation rate that they observe accurately reflects the objective $SPoC$.

[15] Table A1 in the Online Appendix shows the exact cooperation rates.

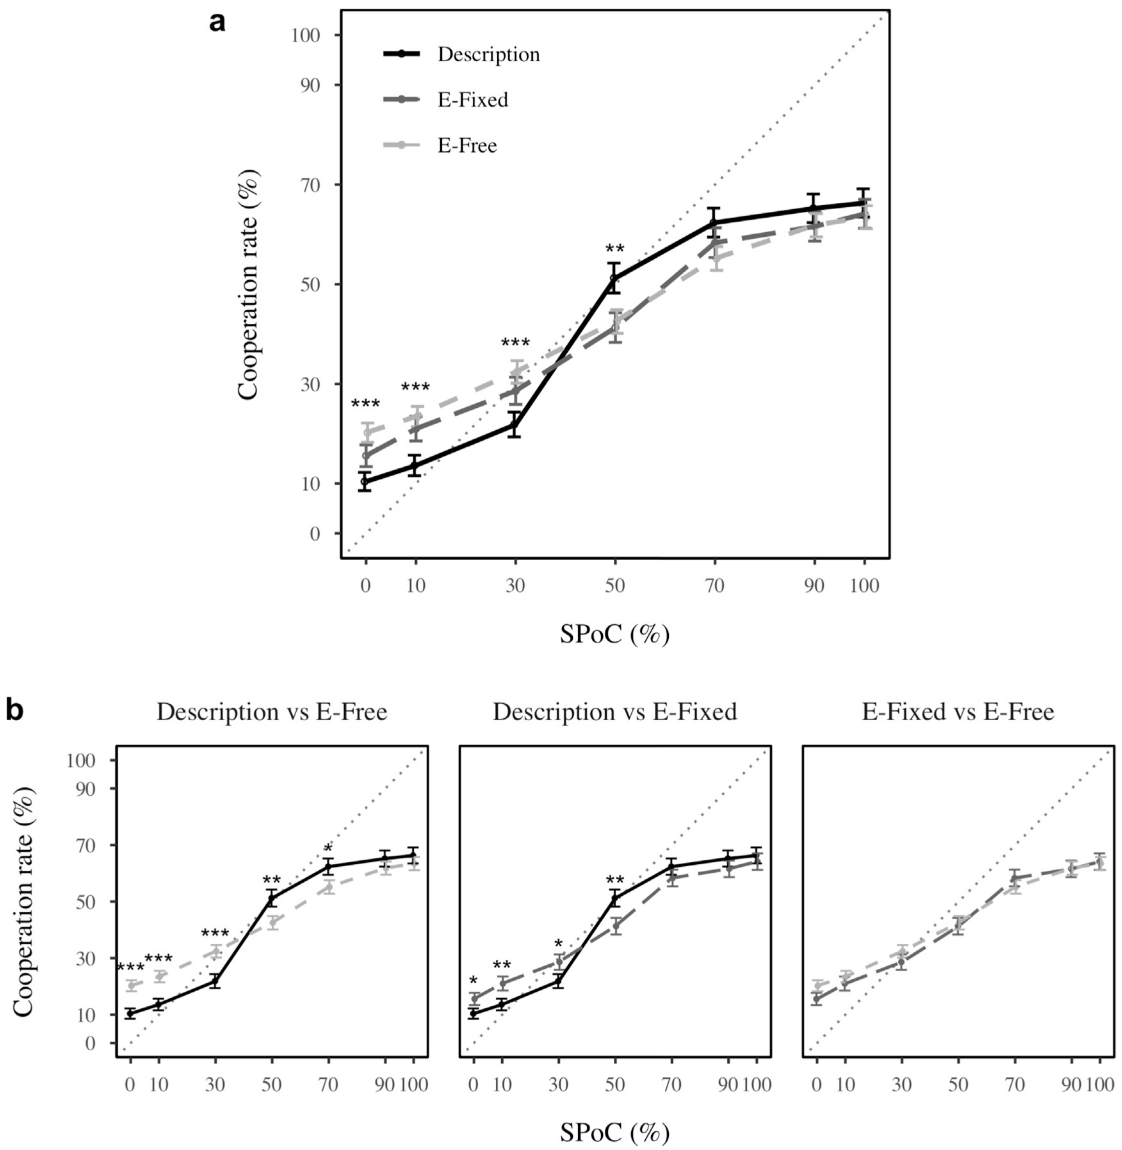

Notes: $SPoC$, the Subpopulation Probability of Cooperation, represents the probability of being matched with a cooperative partner. Pearson's χ^2-tests are conducted across all three treatments (a) and for binary comparisons (b). *** $P < 0.001$, ** $P < 0.05$, * $P < 0.1$. Error bars represent standard errors.

Fig. 1. Cooperation rates as a function of $SPoC$ across treatments.

We now explore whether the description-experience gap in cooperation can be attributed to sampling bias, which is the primary driver of the description-experience gap observed in individual risky decisions. As expected, there was significant sampling bias in the E-Free treatment: subjects in E-Free typically sampled only a few times, with a median of just 4 samples per round. Consequently, in 63% of cases where a sample was obtained, the observed frequency of cooperation deviated from the actual $SPoC$ by 10 percentage points or more.

If sampling bias were an important factor influencing the description-experience gap in our social setting, we would expect significantly different gaps between Description and E-Free (where sampling bias is present) than between Description and E-Fixed (where sampling bias is eliminated). However, the two description-experience gaps are very similar. Moreover, χ^2-tests do not reject the null hypothesis of equal cooperation rates between the two Experience treatments for any level of $SPoC$ (across all seven tests: $\chi^2(2, 990) < 2.43, p > 0.119$).[16]

[16] Our randomization protocol enables us to qualitatively replicate all the above findings even when analyzing only the initial $SPoC$ scenario encountered by each subject. This shows that our findings are not driven by spillover effects (see Online Appendix A.1 for more details).

Table 2
Cooperativeness and *conditionality* indices across treatments.

	Cooperativeness	Conditionality
Description	0.416	0.559
	(0.016)	(0.034)
E-Free	0.427	0.432
	(0.015)	(0.027)
E-Fixed	0.415	0.486
	(0.019)	(0.033)
p	0.909	0.005

Notes: Standard errors are in parentheses. The p-values derive from Kruskal-Wallis tests on individual-level measures of *cooperativeness* and *conditionality* across all three treatments.

Result 3. *The description-experience gap in cooperation is not statistically significantly influenced by sampling bias.*

To shed more light on the behavioral aspects of the description-experience gap in cooperation, we introduce two indices: *cooperativeness* and *conditionality*. These indices are defined as follows:

$$cooperativeness = \frac{1}{n}\frac{1}{7}\sum_{i=1}^{n}\sum_{r=1}^{7} C_{ir} \quad (1)$$

$$conditionality = \frac{1}{n}\sum_{i=1}^{n}(C_{i7} - C_{i1}) \quad (2)$$

where C_{ir} represents whether subject i decides to cooperate (1) or defect (0) at $SPoC_r$, and n is the total number of subjects in a given treatment. The *cooperativeness* index, given by Equation (1), measures the average level of cooperation across all $SPoC$ levels within a treatment. The *conditionality* index, given by Equation (2), captures the overall change in cooperation from the scenario where one's partner will not cooperate ($SPoC = 0$) to the situation where one's partner will definitely cooperate ($SPoC = 1$), within a given treatment.[17] Intuitively, values of *conditionality* that are closer to 1 indicate a stronger tendency toward conditional cooperation within a given treatment.[18]

Table 2 presents the values and statistical comparisons for these two indices across treatments. The *cooperativeness* index does not differ significantly across the three treatments (Kruskal-Wallis test, $H(2) = 0.190$, $p = 0.909$), indicating that overall levels of cooperation are similar. The *conditionality* index, on the other hand, varies significantly across treatments (Kruskal-Wallis test, $H(2) = 10.6$, $p = 0.005$). Further binary comparisons between treatments reveal that *conditionality* differs significantly between the Description and E-Free treatments, and marginally between Description and E-Fixed, but it does not differ significantly between the two Experience treatments (Mann–Whitney U test: Description vs E-Free, $U = 68,328$, $p = 0.001$; Description vs E-Fixed, $U = 41,528$, $p = 0.066$; E-free vs E-fixed $U = 57,109$, $p = 0.213$). Thus, the differences observed between treatments in Fig. 1 appear to be driven by a higher degree of *conditionality* in the Description treatment relative to the two Experience treatments.

Result 4. *Subjects in the Description treatment respond more strongly to changes in the probability that their partner will cooperate than those in the Experience treatments.*

Investigating the specific $SPoC$ values at which subjects switch from defection to cooperation across treatments provides further insights into this difference in conditionality. As evident in Fig. 1, subjects in the Description treatment exhibit a sharp increase in cooperation when informed that there is a 50% chance their partner will cooperate. This jump is absent in the two Experience treatments. Fig. A2 in the Online Appendix provides a more detailed analysis, clearly indicating that the increase in cooperation at $SPoC = 0.5$ is the only increase that stands out between the Description and Experience treatments. This suggests that many conditional cooperators are willing to cooperate if they know there is at least a 50% probability their partner will cooperate—a threshold that is more difficult to identify in the Experience treatments.

[17] Approximately one-third of the total sample switched their action more than once across the different $SPoC$ scenarios. The *conditionality* index is not influenced by such inconsistencies, as it disregards intermediate changes, focusing solely on the endpoints.

[18] Instances of “reverse conditional cooperation” (where $C_{i1} = 1$ and $C_{i7} = 0$) are rare and approximately evenly distributed across treatments (3.9% in Description, 3.2% in E-Free, and 2.5% in E-Fixed; $\chi^2(2{,}990) = 0.878$, p = 0.645 test, P = 0.645). Therefore, they do not materially affect differences between treatments, and excluding these cases does not alter our findings. The reasons behind reverse conditional cooperation may include error, misunderstanding, or a rare type of cooperative preference.

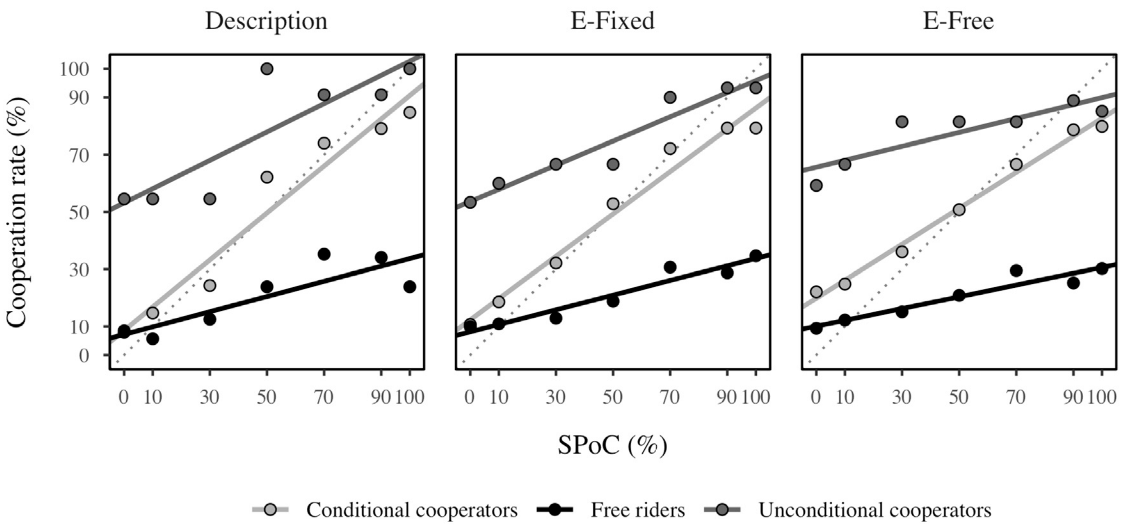

Notes: Lines depict linear least squares fits between $SPoC$ and the cooperation rate for each cooperation type in each treatment.

Fig. 2. Cooperation rates as a function of $SPoC$ across treatments, grouped by cooperation type.

2.4.2. Heterogeneity based on cooperation types

We now turn to heterogeneity analyses based on the cooperative preferences elicited in Stage 3, where we employed the strategy method in a one-shot sequential Prisoner's Dilemma to categorize subjects into distinct cooperation types. Subjects were asked to make their decision to cooperate conditional on their partner's action—thus, under conditions of certainty.

Overall, the majority of subjects were categorized as conditional cooperators (Description: 63.4%; E-Free: 59.3%; E-Fixed: 50.7%). The second most frequent category was free riders (Description: 31.5%; E-Free: 32.0%; E-Fixed: 36.6%). Unconditional cooperators were a minority in all treatments (Description: 3.9%; E-Free: 6.2%; E-Fixed: 10.9%). Subjects who did not fit into any of these three categories were rare (Description: 1.1%; E-Free: 2.5%; E-Fixed: 1.8%).[19]

Fig. 2 depicts subjects' behavior in Stage 2 based on their cooperation types identified in Stage 3, while Table 3 reports the *cooperativeness* and *conditionality* indices for each type across the treatments. The results clearly demonstrate that the type categorization, elicited under information certainty, predicts behavior in Stage 2 across all treatments. Specifically, subjects categorized as conditional cooperators in Stage 3 are substantially more responsive to changes in the likelihood of their partner's cooperation in Stage 2 compared to other types. Unconditional cooperators, on the other hand, consistently score highest on the *cooperativeness* index and relatively low on *conditionality*. Free riders exhibit the lowest scores on both the *cooperativeness* and*conditionality* indices.

Statistical tests, reported in Table 3, confirm that treatment differences are driven by conditional cooperators, who exhibit varying degrees of *conditionality*—but not *cooperativeness*—across treatments. No significant treatment differences are observed for free riders and unconditional cooperators. These findings are in line with the expectation that the description-experience gap in cooperation can only arise among individuals who are sensitive to social information.

Finally, we examine the E-Free treatment in greater detail. In this treatment, subjects could decide how much information to sample. Although sampling did not entail any monetary cost, it did require additional effort and time spent on the task. Theoretically, a subject's willingness to incur these non-monetary costs in exchange for information about their potential partner's cooperativeness should depend on their cooperation type. Specifically, conditional cooperators—who base their cooperation on the behavior of others—should exhibit greater interest in obtaining this information than unconditional cooperators or free riders. Therefore, we would expect conditional cooperators to collect larger samples.

Fig. 3 presents the cumulative distribution of sampling amounts for each cooperation type in the E-Free treatment. On average, conditional cooperators sampled 4.1 draws per round, compared to 3.7 for free riders and 3.1 for unconditional cooperators. As expected, conditional cooperators sampled significantly more than both free riders ($p = 0.032$, clustered Wilcoxon signed rank test) and unconditional cooperators ($p = 0.019$).[20]

[19] Statistical analyses revealed significant differences in the frequency of cooperation types across treatments ($\chi^2(6, 990) = 17.54$, $p = 0.007$). Notably, E-Fixed had a lower prevalence of conditional cooperators (50.7%) compared to E-Free (59.3%; $\chi^2(1, 711) = 5.05, p = 0.025$) and Description (63.4%; $\chi^2(1, 555) = 9.16, p = 0.002$). The lower prevalence of conditional cooperators in E-Fixed, which was designed to test for sampling bias, may be due to its enforcement of a lengthy sampling process. If subjects in this treatment were inclined to quickly click through the study, they might have been more likely to consistently choose a single strategy—either always defecting or always cooperating. At the same time, the distribution of types did not differ significantly between the E-Free and Description treatments ($\chi^2(3, 714) = 3.95, p = 0.266$), and our findings in this section hold when restricted to this pairwise comparison.

[20] See Rosner et al. (2006) for more details on this test.

Table 3
Cooperativeness and *conditionality* indices across treatments, grouped by cooperation type.

	Cooperativeness			Conditionality		
	CC	FR	UC	CC	FR	UC
Description	49.6	20.5	77.9	76.3	15.9	45.5
	(1.7)	(2.2)	(8.2)	(3.8)	(5.4)	(16.5)
E-Free	51.3	20.3	77.8	57.8	20.9	25.9
	(1.8)	(2.1)	(4.9)	(3.4)	(4.5)	(8.8)
E-Fixed	49.3	20.9	74.8	68.6	24.8	40.0
	(2.2)	(2.6)	(3.8)	(4.3)	(4.6)	(10.5)
p	0.768	0.877	0.789	0.0002	0.529	0.403

Notes: CC: Conditional Cooperators; FR: Free Riders; UC: Unconditional Cooperators. Standard errors are in parentheses. The *p*-values derive from Kruskal-Wallis tests on individual-level measures of *cooperativeness* and *conditionality* across the three treatments.

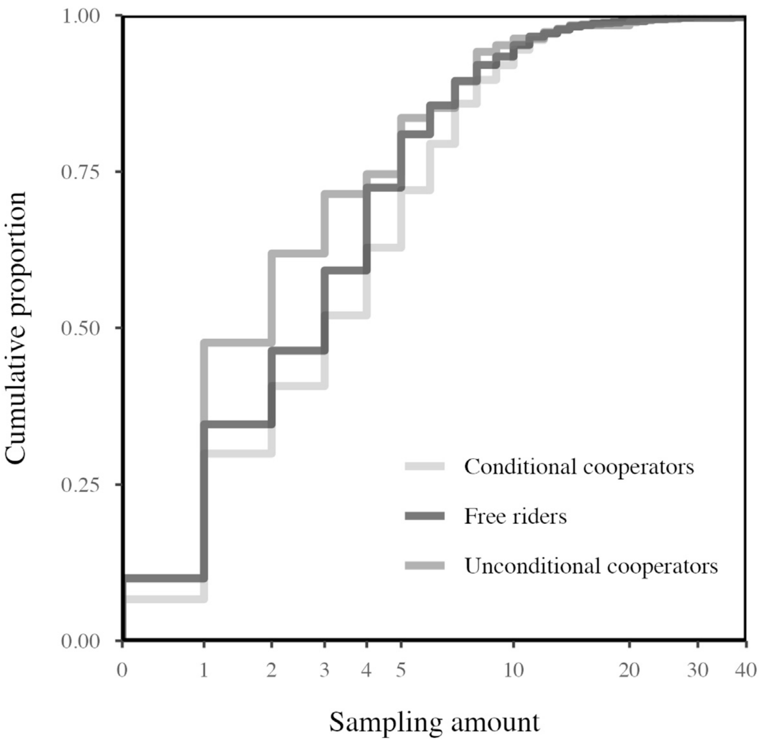

Notes: The figure presents the cumulative distribution of sampled information in the E-free treatment across the three main cooperation types. The x-axis is displayed on a logarithmic scale.

Fig. 3. Cumulative distributions of the sampling amount across cooperation types in the E-Free treatment.

Result 5. *The cooperation preferences elicited in Stage 3 under information certainty are highly predictive of Stage 2 behavior under social uncertainty. Specifically:*

5.1 *Cooperativeness (highest for unconditional cooperators and lowest for free riders) and conditionality (high for conditional cooperators and low for other types) scores are consistent with elicited cooperation preferences.*
5.2 *The description-experience gap in cooperation is driven by conditional cooperators.*
5.3 *Conditional cooperators sample more social information in E-Free than free riders and unconditional cooperators.*

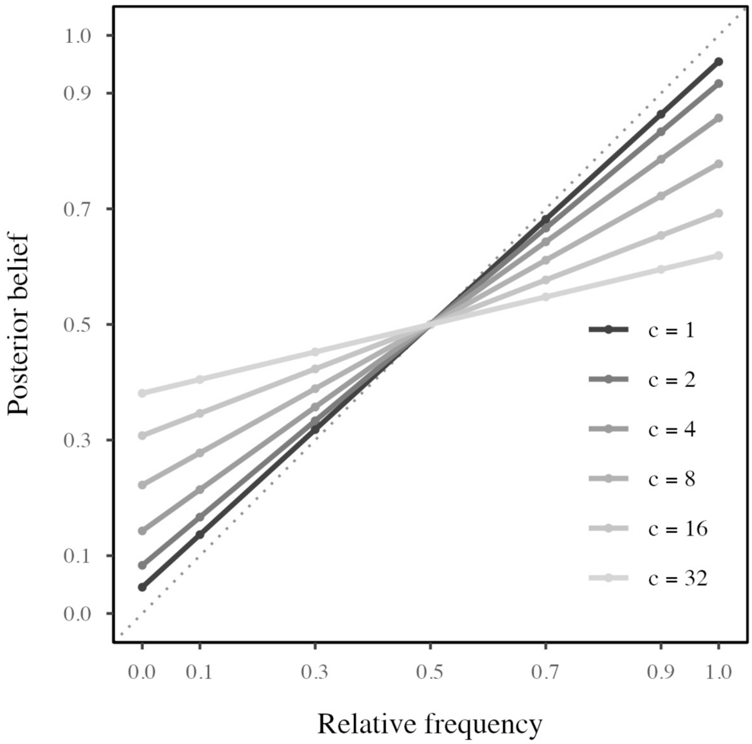

Notes: The figure shows the posterior belief as a function of the observed relative frequency and the strength of the prior. Plotted lines represent the belief that outcome x will occur and are based on Carnap (1952)'s tractable equation: $\frac{cp_0+n}{c+N}$, where c is a constant associated with the strength of the prior, ranging from low ($c = 1$) to high ($c = 32$), N is the total number of observations, which we set to 10, n is the number of occurrences of x, and p_0 is the prior belief of x, fixed at $p_0 = 0.5$ for this example. Lighter colors correspond to stronger prior beliefs. The dotted line represents the diagonal where the strength parameter c is equal to 0.

Fig. 4. Posterior beliefs as a function of the observed relative frequency and the strength of the prior.

3. Beliefs in social vs. individual uncertainty

Our first and main experiment reveals a significant description-experience gap in cooperation: the way in which subjects learned about the likelihood that their partner would cooperate affected their cooperative behavior (Result 2). This gap in a social context deviates from the canonical findings in the individual choice literature in two ways. First, its pattern is reversed: if rare events were less influential in our Experience treatments compared to our Description treatment—as is commonly observed in the individual choice literature—we would expect lower cooperation rates in the former than in the latter when the probability of the partner cooperating is low. However, we observe the opposite pattern. Second, sampling bias—the primary contributor to the description-experience in the individual choice literature—was not a significant driving factor in our social context. The social description-experience gap remained unchanged when subjects collected samples in which the observed relative frequencies matched the objective probabilities of those events (Result 3).

These differences are unlikely to be driven by features of our experimental design, such as "sampling without monetary consequences" or providing "full transparency of outcomes" to subjects. As detailed in Subsection 2.1, although these features have been shown to reduce the size of the gap in previous studies on individual choice, they do not reverse it. This suggests that other differences between our social context and the typical individual choice setup are responsible for the divergent results.

Here, we propose a parsimonious explanation for our key findings and the discrepancies with the individual choice literature. Our explanation centers on the role of priors in belief updating. In the Description treatment, where decision-makers have complete knowledge of objective outcomes and probabilities, prior beliefs are irrelevant. In contrast, in the Experience treatments, where probabilities are not fully known, some degree of ambiguity persists. Under such ambiguity, the decision-maker's posterior beliefs are shaped by both their subjective prior and the information sampled. This difference in the role of priors between Description and Experience treatments will be reflected in the description-experience gap. Moreover, different contexts—such as individual versus social uncertainty—may invoke different priors and, consequently, different description-experience gaps (see also Aydogan, 2021).

Specifically, the description-experience gap will be influenced by the strength of prior beliefs. When greater weight is placed on prior beliefs, the value of new information is discounted more heavily, resulting in posteriors that are less responsive to new information. Fig. 4 illustrates how posterior beliefs vary with observed relative frequencies and prior strength, using Carnap's (1952) rule of updating.[21] As the strength of the prior increases, the resulting posterior belief curve becomes flatter, leading to a "regression to the mean" effect. Consequently, events with small probabilities appear to be overweighted, while those with high probabilities appear to be underweighted.

We hypothesize that priors regarding the actions of others are generally stronger than priors regarding winning chances in risky prospects.[22] For instance, when betting on the color of a ball drawn from an urn with an unknown composition, subjects are unlikely to have strong priors about the likelihood of drawing any particular color. Consequently, their beliefs will be highly responsive to new information. In contrast, when predicting the cooperativeness of others, individuals likely hold stronger priors, informed by their own experience and, perhaps, even moral views, making them more reluctant to update these beliefs.

This belief updating framework can account for the description-experience gap observed in our study and the deviations from past research. Given that prior beliefs should influence behavior only when outcomes are learned through experience—and not when they are objectively described—we would expect subjects in the Experience treatments to be less responsive to new information compared to those in the Description treatment. This is exactly what we observe. Furthermore, stronger priors in social settings, compared to those in settings involving abstract individual uncertainty, can explain the difference with previous studies. If subjects hold weak priors regarding their winning chances in risky prospects, they are unlikely to exhibit a significant reduction in responsiveness to information under Experience as compared to Description. Moreover, adherence to prior beliefs in social interactions can also explain why sampling bias—a key driver of the gap in individual decision scenarios—has minimal impact in our study. If strong priors render additional information less influential, the impact of sampling bias on decision-making is reduced.

Next, we provide a direct test of whether subjects hold stronger priors regarding social uncertainty compared to abstract individual uncertainty by eliciting their confidence in their beliefs about events that depend on either social or individual uncertainty. We hypothesize that confidence in these beliefs will, on average, be higher for the former, and preregistered our experimental design and hypothesis.[23]

3.1. Experimental design and procedures

We recruited 241 subjects through Prolific and randomly assigned them to one of two treatments: "Individual Uncertainty" or "Social Uncertainty".[24] In both treatments, subjects were asked to estimate the likelihood of an outcome and then state their confidence in this estimate.

The key difference between the two treatments lies in the source of uncertainty. In the Social Uncertainty treatment, decisions are made in a social context. Subjects play a one-shot Prisoner's Dilemma identical to the one used in our main experiment (see Table 1). After making their decision, they are asked to estimate the prevalence of cooperation among other subjects.

In the Individual Uncertainty treatment, the social context is removed. Subjects first guess the color of a randomly drawn card from a deck containing red and green cards. Subsequently, they are asked to estimate the prevalence of red cards in the deck.

The key dependent variable is subjects' confidence in their estimate. In both treatments, confidence is measured using a 7-point Likert scale ranging from "1 - Not at all confident" to "7 - Very confident".

3.2. Results

Subjects in the Individual Uncertainty treatment, estimated, on average, that the percentage of red cards was 49.5, while subjects in the Social Uncertainty treatment estimated that the percentage of subjects who chose to *Share* was 50.5. This difference is not statistically significant (Mann–Whitney U test: $U = 6651.5, p = 0.412$).

Although priors are centered around 50% in both treatments, their distribution is far less dispersed for Individual Uncertainty than for Social Uncertainty. Specifically, the variance in the Individual Uncertainty treatment is smaller than in the Social Uncertainty treatment (F-test= 0.14; $var_{Ind} = 58.58$, $var_{Soc} = 415.17$, $df = 117$, $p < 0.001$), suggesting greater belief heterogeneity in the social setting.[25]

Crucially, in line with our hypothesis, subjects' confidence is statistically significantly higher in estimates made under Social Uncertainty ($\mu = 3.98, SD = 1.42$) than under Individual Uncertainty ($\mu = 3.38, SD = 1.79$; Mann–Whitney U test: $U = 5500.5, p = 0.002$).

Result 6. *Subjects are more confident in their prior belief under Social Uncertainty compared to Individual Uncertainty.*

[21] Carnap's (1952) rule provides a tractable method for modeling belief updating in decisions from experience (see also Aydogan, 2021).

[22] To the best of our knowledge, there is no direct evidence for this hypothesis. However, there is some indirect support. If subjects hold strong priors, they will perceive relatively little need to sample additional information, as they believe they can already predict what will happen. Fleischhut et al. (2022) report that subjects are indeed less likely to sample information when facing social uncertainty than when facing lotteries with similar payoffs.

[23] The preregistration can be accessed at: https://aspredicted.org/5H5_CH7.

[24] As in our first experiment, our selection criteria required subjects to be UK residents with an approval rating of 90 or above. Three subjects were excluded from the study as they repeatedly failed to answer the comprehension question correctly. All three excluded subjects were from the Social Uncertainty treatment. Our analysis includes 120 observations in the Individual Uncertainty treatment and 118 in the Social Uncertainty treatment.

[25] See Fig. A.3 in the Online Appendix for the corresponding histograms of these distributions.

4. Conclusion and discussion

Many people are conditionally cooperative: they cooperate if others do so as well. Conditional cooperation has predominantly been studied under conditions of information certainty. However, in many real-world scenarios, there is uncertainty about the cooperativeness of others.

In this paper, we investigate conditional cooperation under uncertainty. We introduce a novel experimental protocol that allows us to exogenously manipulate the likelihood of a subject's partner cooperating in a Prisoner's Dilemma. Drawing on insights from studies on the "description-experience gap" observed in decisions under individual uncertainty, we vary how subjects learn about this likelihood: they either receive explicit numerical information about the cooperation rate in the population from which their partner will be drawn ("Description") or they learn about this rate through a sequential, payoff-inconsequential sampling process in which they observe decisions made by potential partners ("Experience").

We identify a significant description-experience gap in cooperation. Subjects in the Experience treatments cooperate more than those in the Description treatment when the likelihood of their partner's cooperation is below 50%, whereas the opposite is observed when the likelihood is at least 50%—although this difference is not statistically significant for cooperation likelihood above 70%.

To disentangle the gap into its distinct behavioral components, we construct two novel indices: *cooperativeness* and *conditionality*. Our findings reveal that the gap does not stem from treatment differences in the overall propensity to cooperate but rather from subjects in Description being more sensitive to changes in their partner's likelihood of cooperation compared to those in Experience.

We examine between-subject heterogeneity by measuring subjects' cooperative preferences, following Fischbacher et al. (2001). We find that these preferences, elicited under information certainty, strongly predict subjects' behavior under uncertainty. Crucially, conditional cooperators—those who prefer to match their partner's actions—are driving the description-experience gap that we observe. As expected, free riders and unconditional cooperators do not exhibit significant treatment differences.

The pattern of the description-experience gap observed in our study contrasts with the canonical finding in the individual choice literature, where subjects act as if they underweight rare events in Experience relative to Description. If this were the case in our setting, we would expect lower cooperation rates in Experience compared to Description when cooperation is below 50%. However, we observe the opposite pattern. Moreover, sampling bias—a key driver of the gap in the individual choice literature—does not materially affect the gap in our setting.

We propose a parsimonious explanation for our results that can account for this disparity and test it in a follow-up experiment. Specifically, we demonstrate that when decisions are made under conditions of ambiguity—where outcome probabilities are not fully known—stronger priors render posterior beliefs and, consequently, behavior less sensitive to new information. Decisions under Experience are made under ambiguity, whereas decisions under Description are not, meaning that priors should influence behavior in the former but not in the latter. This differential influence of priors can account for the direction of the gap observed in our experiment.[26]

Furthermore, stronger priors in social settings—compared to those in settings involving abstract individual uncertainty—can help explain the difference with previous studies. If subjects hold weak priors regarding their winning chances in risky prospects, they are unlikely to show reduced responsiveness to information under Experience compared to Description. Adherence to prior beliefs in social interactions can also explain why sampling bias has minimal impact in our study. If strong priors render additional information less influential, this reduces the impact of sampling bias on decision-making. In a follow-up experiment, we confirm that subjects indeed hold stronger priors under social uncertainty than under the typical form of individual uncertainty.

Although our sticky-prior argument explains why the flattening of the response curve in Experience compared to Description is more pronounced in our social setting than in individual choice settings, it does not account for why the pattern in individual choice experiments often moves in the opposite direction. To understand why rare events receive less weight in Experience relative to Description in settings of individual uncertainty, additional factors must be considered—factors that are either absent in our social setting or overshadowed by the influence of strong priors. For a comprehensive discussion of the factors driving the description-experience gap in individual choice contexts, we refer readers to Wulff et al. (2018) and Cubitt et al. (2022).

An implication of our explanation is that any factor leading to likelihood insensitivity under ambiguity could generate a description-experience gap similar to the one observed in our study. Although our data show that subjects hold stronger priors when predicting others' cooperation in a Prisoner's Dilemma compared to when predicting their winning odds in an abstract lottery, this does not imply that priors are always stronger in social settings than in individual choice scenarios. For example, individuals may hold strong priors about less abstract, non-social uncertainties, such as tomorrow's weather or stock market movements. If this is the case, the description-experience gaps observed when people predict such outcomes would be similar to the one we observe.

To deepen our understanding of the divergent behavior observed under Description and Experience, future research should investigate the factors that influence people's sensitivity to new information, along with other cognitive and preferential factors that drive this gap. For instance, in a notable exception within the individual decision-making literature, Glöckner et al. (2016) observed a gap similar to ours, which they attribute to the greater complexity of their task compared to previous studies. In contrast, Roth et al. (2020) show that in social settings, subjects sometimes place too little weight on rare events. Examining decisions under experience

[26] The sticky-prior argument may also help explain why the gap we observe is more pronounced when the likelihood of the partner cooperating is low than when it is high. In the Description treatment of Study 1, we observe that subjects' cooperation rate is relatively stable when the likelihood that their partner will cooperate is above 50% but drops sharply when it falls below 50%. Given that most subjects' priors are at or above 50% (as shown in Study 2, Fig. A3 in the Online Appendix), there is greater potential for new information to decrease rather than increase the cooperation rate. Thus, there is more scope for the effect of sticky priors to manifest when the actual likelihood of the partner's cooperation is low rather than when it is high.

in a repeated-game format, they find behavior consistent with subjects severely "undervaluing" rare events. These findings point to a rich area for further investigation, suggesting that subtle aspects of the decision-making environment can significantly influence both the direction and magnitude of the description-experience gap.

The novel experimental protocol developed in this paper allows for the systematic manipulation of the likelihood that another subject will take a specific action—without the use of deception. This capability was crucial for identifying the gap between description and experience within our social context. We validated our protocol by demonstrating that established measures of subjects' cooperation types, elicited using the strategy method, align remarkably well with behavior observed in our experiment. While we employed our protocol to study behavior in the Prisoner's Dilemma, it is highly flexible and can be easily adapted to investigate behavior in other games where players' decisions depend on their beliefs about other players' actions. Therefore, this protocol serves as a valuable tool for researchers exploring the description-experience gap in social settings or studying behavior under social uncertainty more broadly.

One policy implication of our findings pertains to the use of social norm interventions. People tend to conform to descriptive social norms, defined as the behaviors that are typical in a given setting. This insight is widely used to nudge behavior in a desired direction, for example, by informing people that most others file their taxes on time or engage in environmentally friendly actions (Gerber and Rogers, 2009; Farrow et al., 2017; Hallsworth et al., 2017; Li et al., 2021; Neckermann et al., 2022). However, such tactics must be applied judiciously: promoting a norm when socially desirable behavior is infrequent can inadvertently normalize undesirable actions (Cialdini et al., 1990, 1991, 2006). Research on these backfiring effects typically compares situations in which negative descriptive norm information is provided with those in which no information is given, requiring individuals to rely solely on their subjective priors. Our study contributes to this literature by demonstrating that communicating a negative descriptive norm is harmful, even when individuals could otherwise infer the norm through experiential sampling. In short, when many act poorly, ambiguity is preferable to transparency.

Declaration of competing interest

The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

They acknowledge financial support by the Economic and Social Research Council [grant numbers ES/K002201/1, ES/P008976/1] under the Network for Integrated Behavioural Science, which was used for research that resulted in this article. This paper contains data from experiments run with human subjects. Ethics approval was obtained by the Nottingham School of Economics Research Ethics Committee.

Ozan Isler also acknowledges funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie fellowship number 658186 (https://ec.europa.eu/).

Appendix. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.geb.2025.04.012.

Data availability

Data will be made available on request.

References

Abdellaoui, M., L'Haridon, O., Paraschiv, C., 2011. Experienced vs. Described uncertainty: do we need two prospect theory specifications? Manag. Sci. 57 (10), 1879–1895.
Aimone, J.A., Houser, D., 2012. What you don't know won't hurt you: a laboratory analysis of betrayal aversion. Exp. Econ. 15 (4), 571–588.
Aksoy, O., Weesie, J., 2012. Beliefs about the social orientations of others: a parametric test of the triangle, false consensus, and cone hypotheses. J. Exp. Soc. Psychol. 48 (1), 45–54.
Allcott, H., 2011. Social norms and energy conservation. J. Public Econ. 95 (9–10), 1082–1095.
Allcott, H., Rogers, T., 2014. The short-run and long-run effects of behavioral interventions: experimental evidence from energy conservation. Am. Econ. Rev. 104 (10), 3003–3037.
Artinger, F., Fleischhut, N., Levati, M.V., Stevens, J.R., 2012. Cooperation in a risky environment: Decisions from experience in a stochastic social dilemma. Technical report. Jena Economic Research Papers.
Aydogan, I., 2021. Prior beliefs and ambiguity attitudes in decision from experience. Manag. Sci. 67 (11), 6934–6945.
Aydogan, I., Gao, Y., 2020. Experience and rationality under risk: re-examining the impact of sampling experience. Exp. Econ. 23, 1100–1128.
Bardsley, N., 2000. Control without deception: individual behaviour in free-riding experiments revisited. Exp. Econ. 3 (3), 215–240.
Barron, G., Erev, I., 2003. Small feedback-based decisions and their limited correspondence to description-based decisions. J. Behav. Decis. Mak. 16 (3), 215–233.
Barron, G., Ursino, G., 2013. Underweighting rare events in experience based decisions: beyond sample error. J. Econ. Psychol. 39, 278–286.
Bohnet, I., Greig, F., Herrmann, B., Zeckhauser, R., 2008. Betrayal aversion: evidence from Brazil, China, Oman, Switzerland, Turkey, and the United States. Am. Econ. Rev. 98 (1), 294–310.
Bohnet, I., Zeckhauser, R., 2004. Trust, risk and betrayal. J. Econ. Behav. Organ. 55 (4), 467–484.
Camilleri, A.R., Newell, B.R., 2011. When and why rare events are underweighted: a direct comparison of the sampling, partial feedback, full feedback and description choice paradigms. Psychon. Bull. Rev. 18 (2), 377–384.
Carnap, R., 1952. The Continuum of Inductive Methods. University of Chicago Press.
Chaudhuri, A., 2011. Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature. Exp. Econ. 14 (1), 47–83.
Cialdini, R.B., Demaine, L.J., Sagarin, B.J., Barrett, D.W., Rhoads, K., Winter, P.L., 2006. Managing social norms for persuasive impact. Soc. Influen. 1 (1), 3–15.

Cialdini, R.B., Kallgren, C.A., Reno, R.R., 1991. A Focus Theory of Normative Conduct: A Theoretical Refinement and Reevaluation of the Role of Norms in Human Behavior. Advances in Experimental Social Psychology, vol. 24. Elsevier, pp. 201–234.
Cialdini, R.B., Reno, R.R., Kallgren, C.A., 1990. A focus theory of normative conduct: recycling the concept of norms to reduce littering in public places. J. Pers. Soc. Psychol. 58 (6), 1015–1026.
Clark, K., Sefton, M., 2001. The sequential prisoner's dilemma: evidence on reciprocation. Econ. J. 111 (468), 51–68.
Costa-Gomes, M.A., Weizsäcker, G., 2008. Stated beliefs and play in normal-form games. Rev. Econ. Stud. 75 (3), 729–762.
Croson, R.T., 2000. Thinking like a game theorist: factors affecting the frequency of equilibrium play. J. Econ. Behav. Organ. 41 (3), 299–314.
Cubitt, R.P., Kopsacheilis, O., Starmer, C., 2022. An inquiry into the nature and causes of the description-experience gap. J. Risk Uncertain. 65, 105–137.
de Lara Resende, J.G., Wu, G., 2010. Competence effects for choices involving gains and losses. J. Risk Uncertain. 40 (2), 109–132.
Dufwenberg, M., Gächter, S., Hennig-Schmidt, H., 2011. The framing of games and the psychology of play. Games Econ. Behav. 73 (2), 459–478.
Erev, I., Glozman, I., Hertwig, R., 2008. What impacts the impact of rare events. J. Risk Uncertain. 36 (2), 153–177.
Erev, I., Greiner, B., 2015. The 1-800 critique, counter-examples, and the future of behavioral economics. In: Handbook of Experimental Economic Methodology. Oxford University Press, New York, pp. 151–165.
Farrow, K., Grolleau, G., Ibanez, L., 2017. Social norms and pro-environmental behavior: a review of the evidence. Ecol. Econ. 140, 1–13.
Fehr, E., Schurtenberger, I., 2018. Normative foundations of human cooperation. Nat. Hum. Behav. 2 (7), 458–468.
Fetchenhauer, D., Dunning, D., 2012. Betrayal aversion versus principled trustfulness—how to explain risk avoidance and risky choices in trust games. J. Econ. Behav. Organ. 81 (2), 534–541.
Fischbacher, U., Gächter, S., 2010. Social preferences, beliefs, and the dynamics of free riding in public goods experiments. Am. Econ. Rev. 100 (1), 541–556.
Fischbacher, U., Gächter, S., Fehr, E., 2001. Are people conditionally cooperative? Evidence from a public goods experiment. Econ. Lett. 71 (3), 397–404.
Fleischhut, N., Artinger, F., Olschewski, S., Hertwig, R., 2022. Not all uncertainty is treated equally: search under social and nonsocial uncertainty. J. Behav. Decis. Mak. 35 (2), e2250.
Fleischhut, N., Artinger, F., Olschweski, S., Volz, K.G., Hertwig, R., 2014. Sampling of social information: decisions from experience in bargaining. In: 36th Annual Meeting of the Cognitive Science Society. Cognitive Science Society, pp. 1048–1053.
Fox, C.R., Hadar, L., 2006. "Decisions from experience" = sampling error + prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004). Judgm. Decis. Mak. 1 (2), 159–161.
Fox, C.R., Tversky, A., 1995. Ambiguity aversion and comparative ignorance. Q. J. Econ. 110 (3), 585–603.
Frey, B.S., Meier, S., 2004. Social comparisons and pro-social behavior: testing "conditional cooperation" in a field experiment. Am. Econ. Rev. 94 (5), 1717–1722.
Gächter, S., Thöni, C., 2005. Social learning and voluntary cooperation among like-minded people. J. Eur. Econ. Assoc. 3 (2–3), 303–314.
Gerber, A.S., Rogers, T., 2009. Descriptive social norms and motivation to vote: everybody's voting and so should you. J. Polit. 71 (1), 178–191.
Glöckner, A., Hilbig, B.E., Henninger, F., Fiedler, S., 2016. The reversed description-experience gap: disentangling sources of presentation format effects in risky choice. J. Exp. Psychol. Gen. 145 (4), 486–508.
Gächter, S., 2007. Conditional cooperation: behavioral regularities from the lab and the field and their policy implications. In: Frey, B.S., Stutzer, A. (Eds.), Economics and Psychology: A Promising New Cross-Disciplinary Field. The MIT Press, pp. 19–50.
Hallsworth, M., List, J.A., Metcalfe, R.D., Vlaev, I., 2017. The behavioralist as tax collector: using natural field experiments to enhance tax compliance. J. Public Econ. 148, 14–31.
Hayashi, N., Ostrom, E., Walker, J., Yamagishi, T., 1999. Reciprocity, trust, and the sense of control: a cross-societal study. Ration. Soc. 11 (1), 27–46.
Heath, C., Tversky, A., 1991. Preference and belief: ambiguity and competence in choice under uncertainty. J. Risk Uncertain. 4 (1), 5–28.
Hertwig, R., Barron, G., Weber, E.U., Erev, I., 2004. Decisions from experience and the effect of rare events in risky choice. Psychol. Sci. 15 (8), 534–539.
Hertwig, R., Erev, I., 2009. The description-experience gap in risky choice. Trends Cogn. Sci. 13 (12), 517–523.
Hong Chew, S., Li King, K., Chark, R., Zhong, S., 2008. Source preference and ambiguity aversion: models and evidence from behavioral and neuroimaging experiments. In: Houser, D., McCabe, K. (Eds.), Neuroeconomics. In: Advances in Health Economics and Health Services Research, vol. 20. Emerald Group Publishing Limited, pp. 179–201.
Isler, O., Gächter, S., Maule, A.J., Starmer, C., 2021. Contextualised strong reciprocity explains selfless cooperation despite selfish intuitions and weak social heuristics. Sci. Rep. 11 (1), 1–17.
Kelley, H.H., Stahelski, A.J., 1970. Social interaction basis of cooperators' and competitors' beliefs about others. J. Pers. Soc. Psychol. 16 (1), 66–91.
Keppe, H.-J., Weber, M., 1995. Judged knowledge and ambiguity aversion. Theory Decis. 39 (1), 51–77.
Keser, C., Van Winden, F., 2000. Conditional cooperation and voluntary contributions to public goods. Scand. J. Econ. 102 (1), 23–39.
Kilka, M., Weber, M., 2001. What determines the shape of the probability weighting function under uncertainty? Manag. Sci. 47 (12), 1712–1726.
Kopsacheilis, O., 2018. The role of information search and its influence on risk preferences. Theory Decis. 84 (3), 311–339.
Kuhlman, D.M., Wimberley, D.L., 1976. Expectations of choice behavior held by cooperators, competitors, and individualists across four classes of experimental games. J. Pers. Soc. Psychol. 34 (1), 69–81.
Li, C., Turmunkh, U., Wakker, P.P., 2020. Social and strategic ambiguity versus betrayal aversion. Games Econ. Behav. 123, 272–287.
Li, X., Molleman, L., van Dolder, D., 2021. Do descriptive social norms drive peer punishment? Conditional punishment strategies and their impact on cooperation. Evol. Hum. Behav. 42 (5), 469–479.
Martin, J.M., Gonzalez, C., Juvina, I., Lebiere, C., 2014. A description-experience gap in social interactions: information about interdependence and its effects on cooperation. J. Behav. Decis. Mak. 27 (4), 349–362.
Mengel, F., 2017. Risk and temptation: a meta-study on prisoner's dilemma games. Econ. J. 128 (616), 3182–3209.
Neckermann, S., Turmunkh, U., van Dolder, D., Wang, T.V., 2022. Nudging student participation in online evaluations of teaching: evidence from a field experiment. Eur. Econ. Rev. 141, 104001.
Offerman, T., 1997. Beliefs and Decision Rules in Public Good Games: Theory and Experiments. Kluwer, Dordrecht.
Offerman, T., Sonnemans, J., Schram, A., 1996. Value orientations, expectations and voluntary contributions in public goods. Econ. J. 106 (437), 817–845.
Peer, E., Brandimarte, L., Samat, S., Acquisti, A., 2017. Beyond the turk: alternative platforms for crowdsourcing behavioral research. J. Exp. Soc. Psychol. 70, 153–163.
Pletzer, J.L., Balliet, D., Joireman, J., Kuhlman, D.M., Voelpel, S.C., Lange, P.A.V., 2018. Social value orientation, expectations, and cooperation in social dilemmas: a meta-analysis. Eur. J. Pers. 32 (1), 62–83.
Rakow, T., Demes, K.A., Newell, B.R., 2008. Biased samples not mode of presentation: re-examining the apparent underweighting of rare events in experience-based choice. Organ. Behav. Hum. Decis. Process. 106 (2), 168–179.
Rosner, B., Glynn, R.J., Lee, M.-L.T., 2006. The Wilcoxon signed rank test for paired comparisons of clustered data. Biometrics 62 (1), 185–192.
Roth, Y., Plonsky, O., Shalev, E., Erev, I., 2020. On the value of alert systems and gentle rule enforcement in addressing pandemics. Front. Psychol. 11, 577743.
Sally, D., 1995. Conversation and cooperation in social dilemmas: a meta-analysis of experiments from 1958 to 1992. Ration. Soc. 7 (1), 58–92.
Selten, R., 1967. Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments. In: Sauermann, H. (Ed.), Beiträge zur Experimentellen Wirtschaftsforschung. J.C.B. Mohr, Tübingen, pp. 136–168.
Shafir, E., Tversky, A., 1992. Thinking through uncertainty: nonconsequential reasoning and choice. Cogn. Psychol. 24 (4), 449–474.
Starmer, C., Sugden, R., 1991. Does the random-lottery incentive system elicit true preferences? An experimental investigation. Am. Econ. Rev. 81 (4), 971–978.
Thöni, C., Volk, S., 2018. Conditional cooperation: review and refinement. Econ. Lett. 171, 37–40.

Tversky, A., Fox, C.R., 1995. Weighing risk and uncertainty. Psychol. Rev. 102 (2), 269–283.
Tversky, A., Kahneman, D., 1992. Advances in prospect theory: cumulative representation of uncertainty. J. Risk Uncertain. 5 (4), 297–323.
Ungemach, C., Chater, N., Stewart, N., 2009. Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? Psychol. Sci. 20 (4), 473–479.
Van den Assem, M.J., Van Dolder, D., Thaler, R.H., 2012. Split or steal? Cooperative behavior when the stakes are large. Manag. Sci. 58 (1), 2–20.
Wulff, D.U., Mergenthaler-Canseco, M., Hertwig, R., 2018. A meta-analytic review of two modes of learning and the description-experience gap. Psychol. Bull. 144 (2), 140–176.