Conditional Cooperation under Uncertainty: The Social Description-Experience Gap*

Orestis Kopsacheilis[1], Dennie van Dolder[2] and Ozan Isler[3]

[1] *School of Management, Technical University of Munich*
[2] *Department of Economics, University of Essex*
[3] *School of Economics and Finance, Queensland University of Technology*

May 22, 2024

Abstract

Conditional cooperation is usually investigated in experiments where others' choices are known. In this study, we explore 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 find a description-experience gap: subjects are more likely to cooperate under experience than description when the likelihood of their partner's cooperation is low, while the opposite holds when it is at least 50%. This finding is contrary to expectations from individual choice literature, where rare events typically receive less weight in experiential-based decisions. Our findings indicate that conditional cooperators are less responsive to social information when obtained experientially rather than descriptively, and illustrate how stronger priors under social than under individual uncertainty can account for this disparity.

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

JEL Classifications: C72, C92, D81, D83

*Corresponding author: Orestis Kopsacheilis; orestis.kopsacheilis@tum.de. 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/).

1 Introduction

Human societies flourish through cooperation. A central tendency in human cooperation is the preference for conditional cooperation: many people are willing to cooperate if others do so as well, even if this is not in 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 examined 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 instance, 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 others' behavior (Hayashi et al., 1999; Clark and Sefton, 2001; Van den Assem et al., 2012).

In this paper, we explore conditional cooperation under uncertainty. To this end, 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 the use of deception, using a variation on Bardsley's 'Conditional Information Lottery' (Bardsley, 2000).[1] Furthermore, we vary the modality by which subjects receive this information: subjects 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 instances, relevant descriptive information may be available. For example,

[1] Eliciting subjects' beliefs about others' cooperation in social dilemmas and examining how these beliefs correlate with behavior has 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, the study of conditional cooperation through the correlation of beliefs with actions is inherently limited due to the endogenous formation of beliefs. This limitation is significant, as subjects may project their behavioral tendencies onto others, leading to reverse causality. Moreover, it is impossible to distinguish between a free rider and a cooperator when both believe that others are unlikely to cooperate (Gächter, 2007).

a freelancer who secures a large client through an online labor market will often have access to numerous reviews from other freelancers detailing 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, individuals lack precise descriptive information and must rely on experiences to form expectations. This is the case, for example, with a researcher deliberating whether to engage in a new collaboration or a freelancer who lacks 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 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, stemming from individuals in Experience who sample only a small number of observations and often under-represent the objective frequency of rare events, has been identified as a crucial driver of the underweighting of rare events in the literature on risky choices (Fox and Hadar, 2006; Rakow et al., 2008; Cubitt et al., 2022).[2] However, the gap persists in a reduced form even when accounting for sampling bias (see Wulff et al. 2018, for a meta-analysis). To explore how this statistical bias might influence a potential description-experience gap in social contexts, we introduce two variations of the Experience condition: one where subjects can freely decide how much to sample ('E-Free'), which can lead to 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, and we ensure 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 as responses to social uncertainty—where outcomes depend on the decisions of others—

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

have been found to differ from those to uncertainty caused by 'random acts of nature'.[3] Research on betrayal aversion, for example, 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). Additionally, Costa-Gomes and Weizsäcker (2008), Fetchenhauer and Dunning (2012), and Li et al. (2020) demonstrate that behavior in strategic games is less sensitive to changes in beliefs compared to behavior 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 the likelihood of their partner's cooperation is less than 50%, while the opposite is the case when cooperation is at least 50% likely. Moreover, when comparing the two versions of Experience—with and without sampling bias—we observe that cooperative behavior is practically identical.

In both Description and Experience, we observe behavior that is consistent with conditional cooperation: subjects monotonically increase their cooperation in relation to the likelihood that their partner will cooperate regardless of whether this likelihood is described or learned through experience. Additionally, we explore 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 anticipated, the description-experience gap is driven by conditional cooperators rather than free riders or unconditional cooperators, who should not consider their partner's actions when making their decisions. Moreover, conditional cooperators sample more information than either free-riders or unconditional cooperators.

The pattern of the social description-experience gap we observe deviates from our 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. We find the opposite pattern. Additionally, although sampling bias has been identified as a

[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 support for source-dependent preferences is, among others, provided by Keppe and Weber (1995), Kilka and Weber (2001), Hong Chew et al. (2008), and de Lara Resende and Wu (2010).

crucial driver of the (relative) underweighting of rare events in the literature on individual uncertainty, it has no significant effect on behavior in our study.

To address the disparity with previous literature we conduct the following two analyses. First, we disentangle the gap via the use of two novel indexes: *cooperativeness* and *conditionality*, and find that the gap arises not due to a treatment difference in the overall propensity to cooperate but rather due to 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 towards the types of uncertainty typically used in individual choice experiments can account for 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 hypothesis by observing that 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) explored 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 sampling. They observed a lower proportion of risky decisions under experience compared to description, but only after controlling for sampling bias. They were unable to draw conclusions regarding the moderating effect of probabilities on the gap between description and experienced-based choices, as there was no exogenous variation in rejection probabilities in their experimental design.

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 from repeated and 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 the potential outcomes through repeated experience. Their findings indicated higher cooperation levels when the game was fully described than when learned through experience. Furthermore, 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. 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 given a complete description of the game or only learned about their own payoffs through repeated play. Their findings indicated 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 regarding the structure of the game, our research concentrates on the 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. The ability of our design to experimentally manipulate the likelihood that a subject's partner will take a certain action across the entire probability unit interval was instrumental for identifying this gap between Description and Experience in a social setting. In our experiment, we employed this protocol to study behavior in the Prisoner's Dilemma; however, it can readily be adapted to investigate 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 experience. Section 3 presents our follow-up experiment, in which we test whether people hold stronger priors under social uncertainty compared to abstract individual uncertainty. Last, Section 4 concludes.

2 Described vs. Experienced Social Uncertainty

For our first and main experiment, we implement a novel experimental protocol that allows us to study how people's cooperative behavior responds to changes in the likelihood of being matched to a cooperator and whether the way people learn about this likelihood—through description or experience—matters. In addition, we also explore heterogeneity in behavioral responses across subjects who differ in their cooperative preferences.

Subsection 2.1 discusses three crucial design issues when studying behavior under Description versus behavior under Experience, and explains the choices we make. Subsection 2.2 presents our design, Subsection 2.3 discusses practical procedures for conducting our 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. Even though the overall consensus of these studies is that rare events are less influential in Experience compared to Description, there is substantial variability regarding the size of this gap (see Wulff et al., 2018, for a review). This variability can largely be attributed to differences in the way in which Experience is implemented in the lab.[4]

Here, we consider the three most consequential variations in this experimental design. The first sorts the literature in two main branches according to the payoff repercussions of experiential sampling, in particular: (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).

A second variation regards whether or not the sampling technology allows for sampling bias. In the original paradigm, subjects were given the freedom to decide how much to sample. As a result, subjects' observed relative frequency of an outcome often did not correspond to its objective likelihood. Specifically, subjects who only sample a small number of observations will often under-represent 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 use the so-called matched-sampling paradigm, where subjects have to sample a set amount of times, and the observed frequencies are guaranteed to match the underlying probabilities (Ungemach et al., 2009; Barron and Ursino, 2013; Aydogan and Gao, 2020; Cubitt et al., 2022).

The third variation regards the information that subjects have available regarding potential outcomes. In some studies, subjects are told all possible outcomes and only face uncertainty about their probabilities. In other studies, subjects are also unsure about the set of outcomes and have to learn these through sampling as well.

Each of these variations has been shown to influence the magnitude of the description-experience gap. Camilleri and Newell (2011) compared behavior in sit-

[4]Implementing Description is simply a matter of describing all outcomes and associated probabilities and does not exhibit considerable variation across studies.

uations where sampling entails monetary consequences to behavior in situations where it does not. They found that the gap is more pronounced when sampling has monetary consequences. Furthermore, sampling bias has been identified as the most important 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 diminished 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 about all possible outcomes, as opposed 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 utilize 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). There, however, underweighting in experience only occurs in relative terms: subjects overweight rare events in both Description and in Experience, but this overweighting is less pronounced in Experience.

In our experiment, we implement the paradigm where sampled outcomes entail no monetary consequences. When payoff-consequential sampling has been used in a social context, this has been done by 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 arises endogenously, whereas our goal is to exogeneously manipulate it. Furthermore, such a setup would create a confound between information and prior earnings, as observing cooperation would lead to higher payoffs than observing defection. The sampling paradigm, on the other hand, allows us to manipulate the information received by subjects about their partner's likelihood of cooperation while keeping 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. Finally, we provided subjects with complete information on potential outcomes. Given our focus on social uncertainty regarding others' actions, we aimed to eliminate any un-

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

certainty about the payoff structure of the game by providing a detailed description of the consequences for specific actions taken by the subjects and their partner. We also included comprehension questions to verify that subjects understood the game (see Appendix for the complete instructions). This approach ensured that subjects were fully aware of the rules of the game and that the uncertainty they faced was solely related to the likelihood that their partner would cooperate.

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 depicted in Table 1. Stages 1 and 3 are identical across treatments; Stage 2 contains the experimental manipulation. Subjects are informed at the outset that the study has three stages, but detailed instructions are only provided at the beginning of each stage.

Table 1: Payoff matrix for the Prisoner's Dilemma

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

In Stage 1, subjects are asked for their cooperative decision in a one-shot Prisoner's Dilemma, played with a randomly selected other subject. Subjects do not receive feedback on the decision of their partner. The main purpose of this stage is to elicit decisions that can be used to incentivize the subsequent 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 told that she will play a one-shot Prisoner's Dilemma with subject j, where subject j will play with her first stage decision, whereas subject i herself will be asked to make a new decision.[6] Before making their decision, subjects are provided with the opportunity to acquire information about the cooperation rate of the subpopulation from which their partner will be drawn.

[6] One can 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 some evidence that people are less prone 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 feature, the sequential nature of the game cannot account for any potential treatment differences we observe.

In order 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' ($SPoC$), which is defined as the proportion of Stage 1 cooperative decisions in the subpopulation from which subject j will be drawn and represents the (objective) probability that i will face a cooperative j.

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 notate $SPoC_r$, the r^{th} level of $SPoC$, 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 equal to two, so that the true scenario is always captured by a $SPoC$ of 0, 0.5, or 1. This guarantees that there will always be exactly one level of $SPoC$ from our predetermined seven-item set that is real and six that 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 describes the cooperation rate of the actual subpopulation, it is incentive-compatible for them to treat each task as if it is real. Our implementation is a variant on Bardsley's (2000) 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 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

[7]We thus effectively manipulate the level of cooperation among subjects' potential partners. Therefore, our study bears some resemblance to work on assortative matching in social dilemmas. For example, Gächter and Thöni (2005) investigate repeated Public Goods Games where subjects are either placed in random groups or in 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 provides greater control. By having subjects play against past decisions of their potential partners, we exogeneously fix the cooperation rates that subjects face. 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 bigger 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 increased attrition—especially for the Experience treatments where subjects needed to also go through a sampling stage for each $SPoC$ prior to making a decision.

$SPoC$ 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, for as many times as they like. The sampling process is *with* replacement, so that the observed cooperation probability converges to the objective $SPoC$. In E-Fixed, subjects similarly sample decisions of members of the subpopulation, but they have to do so exactly ten times, and the observed relative frequency of cooperation always matched the objective $SPoC$ level of the scenario presented.To guarantee that this is the case, sampling in E-Fixed is *without* replacement.[10] We do not tell subjects in E-Fixed that the observed frequency matches the objective frequency. The goal of this choice is to avoid sampling bias but keep the exact probability ambiguous for 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 them what they will do if their partner chose 'Keep' and what they will do if their partner chose 'Share'. A key difference between the task in Stage 2 and that in Stage 3 is that unlike most conditional decisions in Stage 2 that 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 with 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). On the basis of their decisions in this game, subjects can be categorized in one of four types: 'conditional cooperators', who match their partner's decision; 'free riders', who always defect; 'unconditional cooperators', who always cooperate; and 'others'.[12]

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

[10] This is the common approach to avoid sampling bias in the literature studying the description-experience gap for individual risky choices (e.g. Ungemach et al. 2009; Barron and Ursino 2013; Glöckner et al. 2016; Aydogan and Gao 2020; Cubitt et al. 2022). Even though the observed relative frequency is identical for every subject facing a given $SPoC$, the order with which events of Cooperation or Defection 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.

[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.

2.3 Experimental procedures

The experiment was conducted online using the Qualtrics survey software. We recruited 1,094 subjects using 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, coupled with the fact that we do not provide subjects with feedback on the outcome of each Prisoner's Dilemma, eliminated the need for real-time interaction between subjects.

Prior to making their decisions, subjects had to correctly answer a comprehension question designed to ensure they understood the payoff consequences of their choice in the Prisoner's Dilemma for both their own and their partner's earnings. Those who failed to answer correctly after three attempts were given a participation fee but were not allowed to proceed with the experiment. This was the case for 43 subjects in total (3.9%, 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, 31 in E-Fixed). This resulted in a final sample size of 990 subjects (mean age=36.0, s.d.=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, but 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 in order to increase statistical power for the analysis of sampling decisions in that treatment. In the end, 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 Appendix B, and the matching protocol for payments is explained in Appendix A.4.

2.4 Results

2.4.1 Cooperation patterns across treatments

Across all treatments, 57.9% of our overall sample cooperated in Stage 1. This cooperation rate did not differ significantly across treatments ($\chi^2(2, 990) = 3.84, p =$

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

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. Figure 1 displays the average cooperation rates in the Prisoner's Dilemma game of Stage 2 across the different $SPoC$ levels for each treatment. In each 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, Figure 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 opts 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 significantly different patterns of cooperation across the three treatments. As shown in Figure 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 below 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 10%).

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.*

[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 Appendix A shows the exact cooperation rates.

Figure 1: Cooperation rates as a function of $SPoC$ across treatments

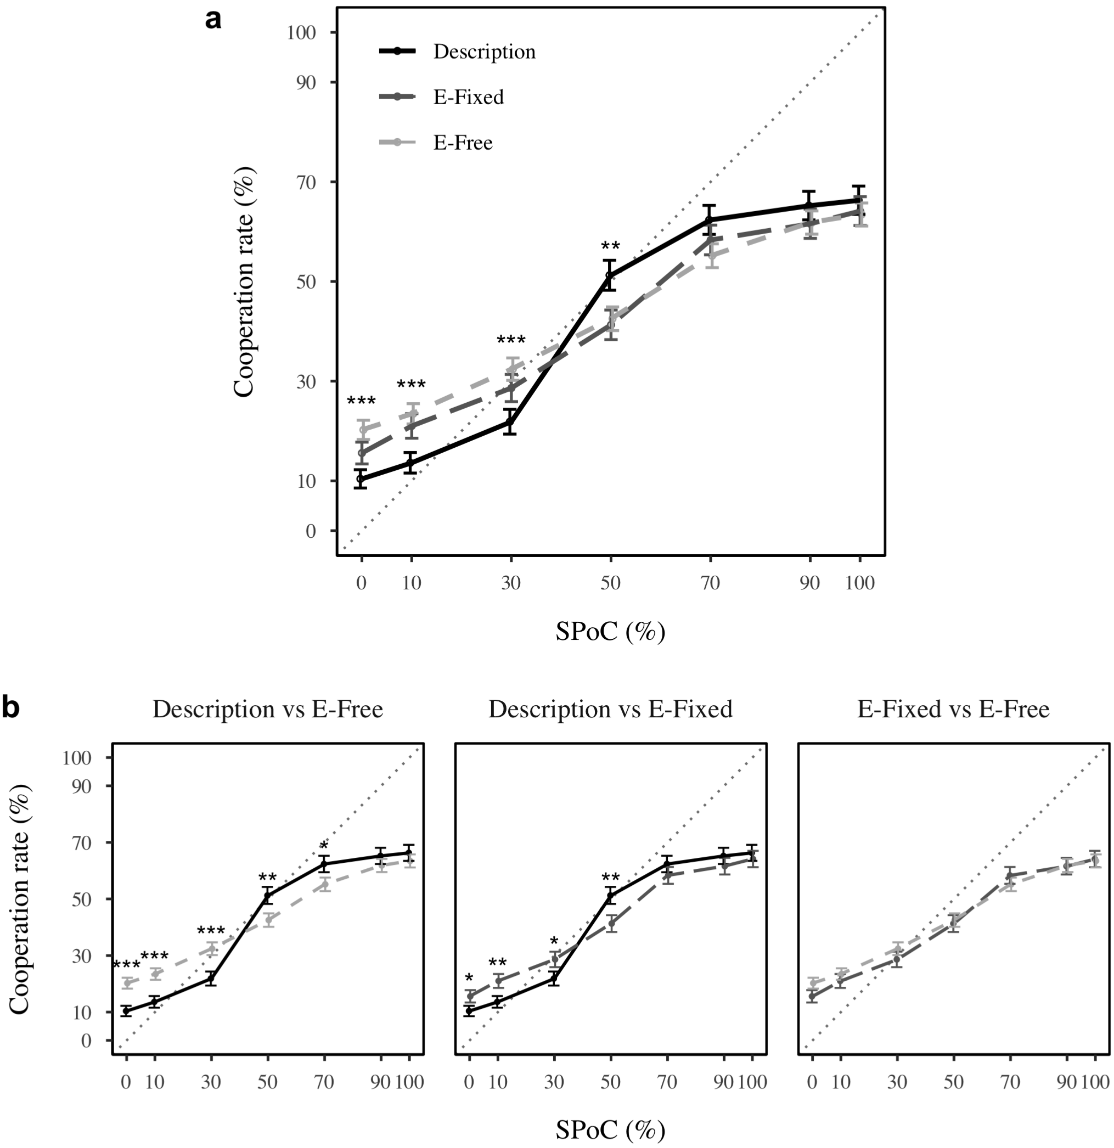

Note. '$SPoC$', the Subpopulation Probability of Cooperation, represents the probability of being matched to a cooperative partner. Pearson's χ^2-tests across all three treatments (a) and for binary comparisons (b). *** $P < 0.001$, ** $P < 0.05$, * $P < 0.1$. Error bars represent standard errors.

We now explore whether the description-experience gap in cooperation can be attributed to sampling bias, the leading driver in 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 only 4 cards 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 was an important factor influencing the description-experience gap in our social context, 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). Instead, 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]

Result 3. *Sampling bias is not a significant driver of the description-experience gap in cooperation.*

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

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

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

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 towards conditional cooperation within a given treatment.[18]

[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 spill-over effects (see Appendix A.2 for more details).

[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 and thus do not materially affect differences between treatments(3.9% in Description, 3.2% in E-Free and 2.5% in E-Fixed; χ^2(2,990)=0.878, p=0.645 test, P = 0.645). 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.

Table 2: *Cooperativeness* and *conditionality* indexes 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 in parentheses. The p-values derive from Kruskal-Wallis tests on individual-level measures of *cooperativeness* and *conditionality* across all three treatments.

Table 2 provides values and statistical comparisons for these two indexes 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, does vary 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 so between Description and E-Fixed, but 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 observed in Figure 1 appear to be driven by a higher degree in *conditionality* in the Description treatment relative to the two experience treatments.

Result 4. *Subjects in the Description treatment react 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 tend to switch from defection to cooperation across treatments offers further insights into this difference in conditionality. As evident in Figure 1, subjects in the Description treatment demonstrate a sharp increase in cooperation when informed that there is a 50% chance that their partner will cooperate. This jump is absent in the two Experience treatments. Figure A2 in Appendix A.3 provides a more detailed analysis, clearly indicating that the increase in cooperation at $SPoC = 0.5$ is the only increase that is markedly different between the Description and Experience treatments. This suggests that many conditional cooperators are prepared 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 Experience treatments.

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 used 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]

Figure 2 depicts subjects' behavior in Stage 2 based on their cooperation types identified in Stage 3, while Table 3 reports the *cooperativeness* and *conditionality* indexes 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 the other types. Unconditional cooperators, on the other hand, consistently score highest on the *cooperativeness* index and score relatively low on *conditionality*. Free riders exhibit the lowest scores on both the *cooperativeness* and*conditionality* indexes.

Statistical tests, reported in Table 3, confirm that treatment differences are driven by conditional cooperators who exhibit different degrees of *conditionality*—but not *cooperativeness*—across treatments. For free riders and unconditional cooperators, no significant treatment differences are observed. This is in line with the

[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 condition were inclined to quickly click through the study, they could be more likely to consistently choose a single strategy—either always defecting or always cooperating. However, the distribution of types did not significantly differ 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.

Figure 2: Cooperation rates as a function of SPoC across treatments by cooperation type

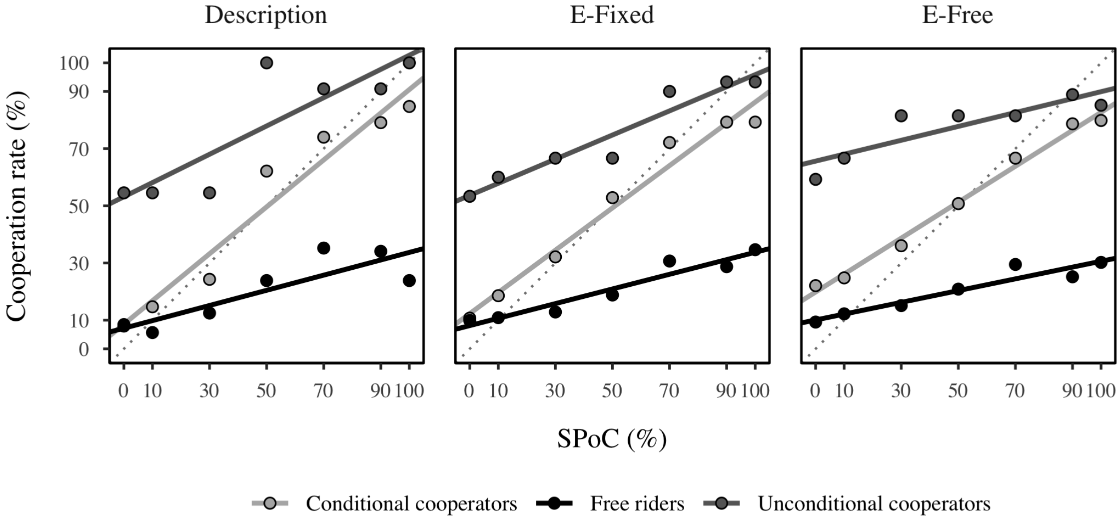

Note. 'CC': Conditional Cooperators; 'FR': Free Riders; 'UC': Unconditional Cooperators. Lines depict linear least squares fits between SPoC and cooperation rates for each type in each treatment.

Table 3: *Cooperativeness* and *conditionality* indexes across treatments 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 in parentheses. The p-values derive from Kruskal-Wallis tests on individual-level measures of *cooperativeness* and *conditionality* across the three treatments.

expectation that the description-experience gap in cooperation can only be caused by people who are sensitive to social information.

Finally, we consider the E-Free treatment in more detail. In this treatment, subjects could decide how much information to sample. Although sampling did not

Figure 3: Cumulative distributions of sampling amount across cooperation types in the E-Free treatment

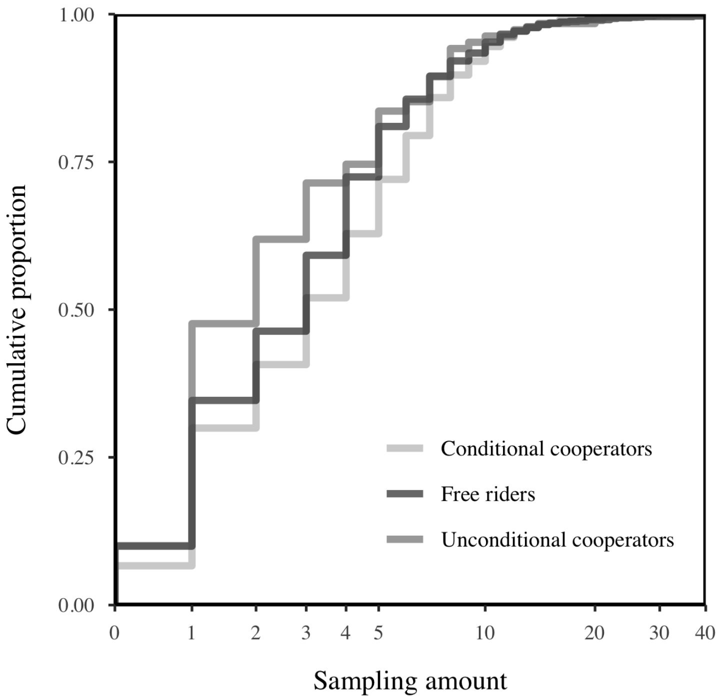

Notes: The figure shows the cumulative distribution of sampling information in the E-free treatment. The x-axis is logarithmic. Standard errors in parentheses. The p-values derive from Kruskal-Wallis tests on individual-level measures of *cooperativeness* and *conditionality* across all three treatments.

entail any monetary cost, it did require additional effort and more time spent on the task. Theoretically, a subject's willingness to incur these non-monetary costs in exchange for information about the cooperativeness of their potential partner 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 either unconditional cooperators or free riders. Hence, we would expect conditional cooperators to collect bigger samples.

Figure 3 displays 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 anticipated, conditional cooperators indeed sampled significantly more than both free riders (clustered Wilcoxon signed rank test; $p = 0.032$) and unconditional cooperators ($p = 0.019$).[20]

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

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 social 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.*

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 common in the individual choice literature—we would expect lower cooperation rates in the former than in the latter when the probability of the partner's cooperation is low. However, we observe the opposite pattern. Second, sampling bias—the leading 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 unaltered even when subjects collected samples whose observed relative frequencies matched the objective probabilities of those events (Result 3).

These differences are unlikely to be caused 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, even though 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 driving 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

Figure 4: Posterior beliefs as a function of observed probability and strength of prior

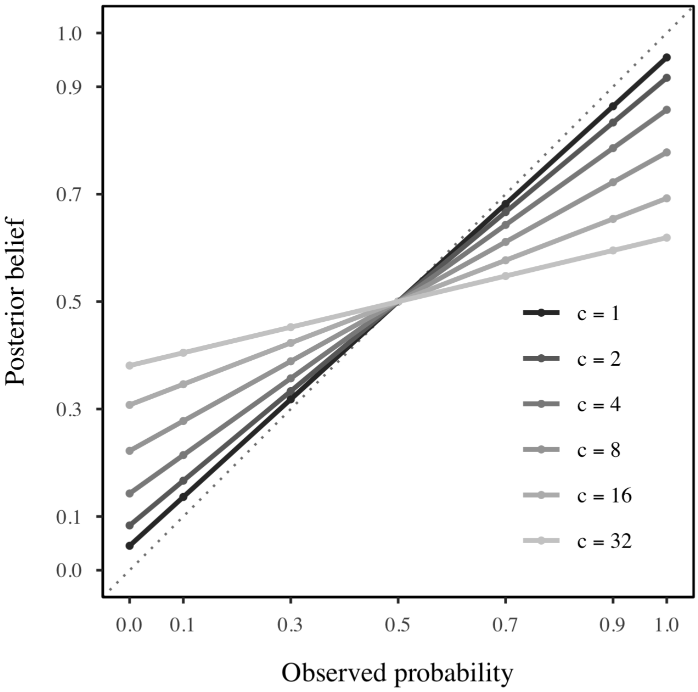

Notes: The figure shows the posterior belief as a function of the observed probability and the strength of prior. Plotted lines model the belief that outcome x will obtain and are based on Carnap (1952)'s tractable equation: $\frac{cp_0+n}{c+N}$, where c is a constant associated with the strength of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

known, there is a degree of ambiguity. Under such ambiguity, the decision-maker's posterior beliefs are some combination of their subjective prior and the information sampled. This difference in the role of priors between description and experience conditions will be reflected in the description-experience gap. Moreover, different contexts—such as individual versus social uncertainty—may invoke different priors, and, therefore, different description-experience gaps (see also Aydogan, 2021).

Specifically, the description-experience gap will be influenced by the strength of prior beliefs. When more weight is placed on prior beliefs, the value of new information is discounted more heavily, leading to posteriors that are less responsive to new information. This phenomenon is visually demonstrated in Figure 4, where we depict posteriors for varying levels of observed probabilities and different strengths of prior belief 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; thus, events with small probabilities seem to be overweighted, while those with high probabilities appear to be underweighted.

We hypothesize that priors regarding the actions of others will generally be stronger than priors regarding winning chances in risky prospects.[22] For instance, when betting on the color of a ball that will be drawn from an urn with an unknown composition, subjects are unlikely to have strong priors about the likelihood of drawing any particular color. Consequently, they will be open to updating their beliefs based on new information. In contrast, when predicting the cooperativeness of others, individuals will likely have 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 under conditions where outcomes are learned through experience—and not when they are objectively described—we would expect subjects in the Experience conditions to be less responsive to new information as compared to those in the Description treatment. This is exactly what we observe. Furthermore, stronger priors in social settings than towards 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 show reduced responsiveness to information under Experience as compared to Description. Additionally, the 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, this reduces the impact of sampling bias on decision-making.

Next, we provide a direct test of whether subjects hold stronger priors in social rather than abstract individual tasks by eliciting subjects' confidence about their beliefs for events that depend on either social or individual uncertainty. We hypothesized that confidence in these beliefs will, on average, be larger for the latter and

[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 previous direct evidence for this hypothesis. However, there is some indirect support. If subjects hold strong priors, they will see relatively little need to sample additional information as they believe that they are already able to predict what will happen. Fleischhut et al. (2022) report that subjects are indeed less likely to sample information when they face social uncertainty than when they face lotteries with similar payoffs.

preregistered our experiment 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 are asked to make an estimate regarding the frequency of an outcome and then state their confidence regarding this estimation.

The key difference between the two treatments lies in the source of uncertainty. In the Social Uncertainty treatment, the decision is made in a social context. Subjects play a one-shot Prisoner's Dilemma, identical to the one in our main experiment (see Table 1). After making their decision, subjects are asked to estimate the prevalence of cooperation decisions 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 with Red and Green cards. Subsequently, they are asked about their belief regarding the prevalence of Red cards in the deck.

The key dependent variable is the confidence regarding their estimate. In both treatments, confidence is elicited in a 7-item Likert scale ranging from '1 - Not at all confident' to '7 - Very confident'.

3.2 Results

Subjects in Individual Uncertainty estimated on average that the percentage of Red cards is 49.5 while subjects in Social Uncertainty estimate that the percentage of subjects who chose to 'Share' was 50.5. The two means are very similar, and their difference is not statistically significant (Mann–Whitney U test: $U = 6651.5, p = 0.4121$). Interestingly, although on average prior beliefs are similar across the two treatments, the variance in Social Uncertainty is larger than in Individual Uncertainty (F-test= $0.14; var_{Ind} = 58.58$, $var_{Soc} = 415.17$, $df = 117$, $p < 0.001$), suggesting that there is much more belief-heterogeneity in this social setting.

Crucially, in line with our hypothesis, we observed that subjects' confidence is statistically significantly higher in estimates conducted in the Social Uncertainty

[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 and 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 comprises 120 observations in the Individual Uncertainty treatment and 118 in the Social Uncertainty treatment.

($\mu = 3.98, sd = 1.42$) than in the Individual Uncertainty ($\mu = 3.38, sd = 1.79$) condition (Mann–Whitney U test: $U = 5500.5, p = 0.002$).

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

4 Conclusion

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

In this paper, we investigate conditional cooperation under conditions of uncertainty. We develop 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 explore how subjects learn about this likelihood. In 'Description', they receive explicit numerical information about this likelihood, while in 'Experience' this information is obtained through a sequential and payoff inconsequential sampling process, where they observe decisions from potential partners.

We identify a significant description-experience gap in cooperation. Subjects in Experience cooperate more than those in Description when the likelihood of their partner's cooperation is less than 50%, while the opposite is the case when cooperation is at least 50%—although this difference was not statistically significant for cooperation likelihood higher than 70%. We disentangle the gap by constructing two novel indexes: *cooperativeness* and *conditionality*, and find that the gap arises not due to treatment differences in the overall propensity to cooperate but rather due to subjects in Description being more sensitive to changes in the likelihood that their partner will cooperate than subjects in Experience. By exploring between-subject heterogeneity, we confirm that conditional cooperators—those who prefer to match their partner's actions—are driving our results. As expected, free-riders or unconditional cooperators do not show significant treatment differences.

The pattern of the description-experience gap that we observe is at odds 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 relative to Description when cooperation is less than 50% rather than the opposite, which is what we observe. Moreover, sampling bias—a key driver of the gap in the individual choice literature—does not 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 play a role 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.

Furthermore, stronger priors in social settings than in settings with 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 show reduced responsiveness to information under Experience as compared to Description. The 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 rather than the typical form of individual uncertainty.

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 we observed in our study. While 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 instance, 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 conditions, 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 individual decision-making literature, Glöckner et al. (2016) observed a gap similar to ours. They attribute this to the greater complexity of their task compared to previous studies. At the same time, Roth et al. (2020) show that in social settings, subjects can sometimes put too little weight on rare events. They examine decisions under experience in a repeated game format and 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 we develop in this paper allows for the systematic manipulation of the likelihood that another subject will take a specific action. The ability to do so 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 used 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 hinge on their beliefs about their partners' actions. This makes it a valuable tool for researchers aiming to explore the description-experience gap in social settings or to study 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 behavior that is typical in a specific setting. This insight is widely utilized to nudge behavior in the 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). Such tactics, however, need to be exercised judiciously: promoting a norm when the socially desirable behavior is infrequent can inadvertently normalize less desirable actions(Cialdini et al., 1990, 1991, 2006). Research that observes such backfiring typically compares situations where negative descriptive norm information is provided with those where no information is given, requiring individuals to rely solely on their subjective priors. Our study contributes to this literature by showing 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 trumps transparency.

References

Abdellaoui, M., L'Haridon, O. and Paraschiv, C. (2011), 'Experienced vs. described uncertainty: Do we need two prospect theory specifications?', *Management Science* **57**(10), 1879–1895.

Aimone, J. A. and Houser, D. (2012), 'What you don't know won't hurt you: A laboratory analysis of betrayal aversion', *Experimental Economics* **15**(4), 571–588.

Aksoy, O. and Weesie, J. (2012), 'Beliefs about the social orientations of others: A parametric test of the triangle, false consensus, and cone hypotheses', *Journal of Experimental Social Psychology* **48**(1), 45–54.

Allcott, H. (2011), 'Social norms and energy conservation', *Journal of Public Economics* **95**(9-10), 1082–1095.

Allcott, H. and Rogers, T. (2014), 'The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation', *American Economic Review* **104**(10), 3003–3037.

Artinger, F., Fleischhut, N., Levati, M. V. and 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', *Management Science* **67**(11), 6934–6945.

Aydogan, I. and Gao, Y. (2020), 'Experience and rationality under risk: Re-examining the impact of sampling experience', *Experimental Economics* **23**, 1100–1128.

Bardsley, N. (2000), 'Control without deception: Individual behaviour in free-riding experiments revisited', *Experimental Economics* **3**(3), 215–240.

Barron, G. and Erev, I. (2003), 'Small feedback-based decisions and their limited correspondence to description-based decisions', *Journal of Behavioral Decision Making* **16**(3), 215–233.

Barron, G. and Ursino, G. (2013), 'Underweighting rare events in experience based decisions: Beyond sample error', *Journal of Economic Psychology* **39**, 278–286.

Bohnet, I., Greig, F., Herrmann, B. and Zeckhauser, R. (2008), 'Betrayal aversion: Evidence from Brazil, China, Oman, Switzerland, Turkey, and the United States', *American Economic Review* **98**(1), 294–310.

Bohnet, I. and Zeckhauser, R. (2004), 'Trust, risk and betrayal', *Journal of Economic Behavior & Organization* **55**(4), 467–484.

Camilleri, A. R. and 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', *Psychonomic Bulletin & Review* **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', *Experimental Economics* **14**(1), 47–83.

Cialdini, R. B., Demaine, L. J., Sagarin, B. J., Barrett, D. W., Rhoads, K. and Winter, P. L. (2006), 'Managing social norms for persuasive impact', *Social Influence* **1**(1), 3–15.

Cialdini, R. B., Kallgren, C. A. and Reno, R. R. (1991), A focus theory of normative conduct: A theoretical refinement and reevaluation of the role of norms in human behavior, *in* 'Advances in experimental social psychology', Vol. 24, Elsevier, pp. 201–234.

Cialdini, R. B., Reno, R. R. and Kallgren, C. A. (1990), 'A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places.', *Journal of Personality and Social Psychology* **58**(6), 1015–1026.

Clark, K. and Sefton, M. (2001), 'The sequential prisoner's dilemma: Evidence on reciprocation', *Economic Journal* **111**(468), 51–68.

Costa-Gomes, M. A. and Weizsäcker, G. (2008), 'Stated beliefs and play in normal-form games', *Review of Economic Studies* **75**(3), 729–762.

Croson, R. T. (2000), 'Thinking like a game theorist: factors affecting the frequency of equilibrium play', *Journal of Economic Behavior & Organization* **41**(3), 299–314.

Cubitt, R. P., Kopsacheilis, O. and Starmer, C. (2022), 'An inquiry into the nature and causes of the description-experience gap', *Journal of Risk and Uncertainty* **65**, 105–137.

de Lara Resende, J. G. and Wu, G. (2010), 'Competence effects for choices involving gains and losses', *Journal of Risk and Uncertainty* **40**(2), 109–132.

Dufwenberg, M., Gächter, S. and Hennig-Schmidt, H. (2011), 'The framing of games and the psychology of play', *Games and Economic Behavior* **73**(2), 459–478.

Erev, I., Glozman, I. and Hertwig, R. (2008), 'What impacts the impact of rare events', *Journal of Risk and Uncertainty* **36**(2), 153–177.

Erev, I. and 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. and Ibanez, L. (2017), 'Social norms and pro-environmental behavior: A review of the evidence', *Ecological Economics* **140**, 1–13.

Fehr, E. and Schurtenberger, I. (2018), 'Normative foundations of human cooperation', *Nature Human Behaviour* **2**(7), 458–468.

Fetchenhauer, D. and Dunning, D. (2012), 'Betrayal aversion versus principled trustfulness—how to explain risk avoidance and risky choices in trust games', *Journal of Economic Behavior & Organization* **81**(2), 534–541.

Fischbacher, U. and Gächter, S. (2010), 'Social preferences, beliefs, and the dynamics of free riding in public goods experiments', *American Economic Review* **100**(1), 541–56.

Fischbacher, U., Gächter, S. and Fehr, E. (2001), 'Are people conditionally cooperative? Evidence from a public goods experiment', *Economics Letters* **71**(3), 397–404.

Fleischhut, N., Artinger, F., Olschewski, S. and Hertwig, R. (2022), 'Not all uncertainty is treated equally: Search under social and nonsocial uncertainty', *Journal of Behavioral Decision Making* **35**(2), e2250.

Fleischhut, N., Artinger, F., Olschweski, S., Volz, K. G. and 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. and Hadar, L. (2006), '"Decisions from experience" = sampling error + prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004)', *Judgment and Decision Making* **1**(2), 159–161.

Fox, C. R. and Tversky, A. (1995), 'Ambiguity aversion and comparative ignorance', *Quarterly Journal of Economics* **110**(3), 585–603.

Frey, B. S. and Meier, S. (2004), 'Social comparisons and pro-social behavior: Testing "conditional cooperation" in a field experiment', *American Economic Review* **94**(5), 1717–1722.

Gächter, S. and Thöni, C. (2005), 'Social learning and voluntary cooperation among like-minded people', *Journal of the European Economic Association* **3**(2-3), 303–314.

Gerber, A. S. and Rogers, T. (2009), 'Descriptive social norms and motivation to vote: Everybody's voting and so should you', *Journal of Politics* **71**(1), 178–191.

Glöckner, A., Hilbig, B. E., Henninger, F. and Fiedler, S. (2016), 'The reversed description-experience gap: Disentangling sources of presentation format effects in risky choice.', *Journal of Experimental Psychology: General* **145**(4), 486–508.

Gächter, S. (2007), Conditional cooperation: Behavioral regularities from the lab and the field and their policy implications, *in* B. S. Frey and A. Stutzer, eds, 'Economics and Psychology: A Promising New Cross-Disciplinary Field', The MIT Press, pp. 19–50.

Hallsworth, M., List, J. A., Metcalfe, R. D. and Vlaev, I. (2017), 'The behavioralist as tax collector: Using natural field experiments to enhance tax compliance', *Journal of Public Economics* **148**, 14–31.

Hayashi, N., Ostrom, E., Walker, J. and Yamagishi, T. (1999), 'Reciprocity, trust, and the sense of control: A cross-societal study', *Rationality and Society* **11**(1), 27–46.

Heath, C. and Tversky, A. (1991), 'Preference and belief: Ambiguity and competence in choice under uncertainty', *Journal of Risk and Uncertainty* **4**(1), 5–28.

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

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

Hong Chew, S., King Li, K., Chark, R. and Zhong, S. (2008), Source preference and ambiguity aversion: Models and evidence from behavioral and neuroimaging experiments, *in* D. Houser and K. McCabe, eds, 'Neuroeconomics', Vol. 20 of *Advances in Health Economics and Health Services Research*, Emerald Group Publishing Limited, pp. 179–201.

Isler, O., Gächter, S., Maule, A. J. and Starmer, C. (2021), 'Contextualised strong reciprocity explains selfless cooperation despite selfish intuitions and weak social heuristics', *Scientific reports* **11**(1), 1–17.

Kelley, H. H. and Stahelski, A. J. (1970), 'Social interaction basis of cooperators' and competitors' beliefs about others.', *Journal of Personality and Social Psychology* **16**(1), 66–91.

Keppe, H.-J. and Weber, M. (1995), 'Judged knowledge and ambiguity aversion', *Theory and Decision* **39**(1), 51–77.

Keser, C. and Van Winden, F. (2000), 'Conditional cooperation and voluntary contributions to public goods', *Scandinavian Journal of Economics* **102**(1), 23–39.

Kilka, M. and Weber, M. (2001), 'What determines the shape of the probability weighting function under uncertainty?', *Management Science* **47**(12), 1712–1726.

Kopsacheilis, O. (2018), 'The role of information search and its influence on risk preferences', *Theory and Decision* **84**(3), 311–339.

Kuhlman, D. M. and Wimberley, D. L. (1976), 'Expectations of choice behavior held by cooperators, competitors, and individualists across four classes of experimental games.', *Journal of Personality and Social Psychology* **34**(1), 69–81.

Li, C., Turmunkh, U. and Wakker, P. P. (2020), 'Social and strategic ambiguity versus betrayal aversion', *Games and Economic Behavior* **123**, 272–287.

Li, X., Molleman, L. and van Dolder, D. (2021), 'Do descriptive social norms drive peer punishment? conditional punishment strategies and their impact on cooperation', *Evolution and Human Behavior* **42**(5), 469–479.

Martin, J. M., Gonzalez, C., Juvina, I. and Lebiere, C. (2014), 'A description–experience gap in social interactions: Information about interdependence and its effects on cooperation', *Journal of Behavioral Decision Making* **27**(4), 349–362.

Mengel, F. (2017), 'Risk and temptation: A meta-study on prisoner's dilemma games', *Economic Journal* **128**(616), 3182–3209.

Neckermann, S., Turmunkh, U., van Dolder, D. and Wang, T. V. (2022), 'Nudging student participation in online evaluations of teaching: Evidence from a field experiment', *European Economic Review* **141**, 104001.

Offerman, T. (1997), *Beliefs and decision rules in public good games: Theory and experiments*, Dordrecht: Kluwer.

Offerman, T., Sonnemans, J. and Schram, A. (1996), 'Value orientations, expectations and voluntary contributions in public goods', *Economic Journal* **106**(437), 817–845.

Peer, E., Brandimarte, L., Samat, S. and Acquisti, A. (2017), 'Beyond the Turk: Alternative platforms for crowdsourcing behavioral research', *Journal of Experimental Social Psychology* **70**, 153–163.

Pletzer, J. L., Balliet, D., Joireman, J., Kuhlman, D. M., Voelpel, S. C. and Lange, P. A. V. (2018), 'Social value orientation, expectations, and cooperation in social dilemmas: A meta–analysis', *European Journal of Personality* **32**(1), 62–83.

Rakow, T., Demes, K. A. and Newell, B. R. (2008), 'Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience-based choice', *Organizational Behavior and Human Decision Processes* **106**(2), 168–179.

Rosner, B., Glynn, R. J. and 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. and Erev, I. (2020), 'On the value of alert systems and gentle rule enforcement in addressing pandemics', *Frontiers in Psychology* **11**, 577743.

Sally, D. (1995), 'Conversation and cooperation in social dilemmas: A meta-analysis of experiments from 1958 to 1992', *Rationality and Society* **7**(1), 58–92.

Selten, R. (1967), Die Strategiemethode zur Erforschung des eingeschräncht rationalen Verhaltens im Rahmen eines Oligopolexperiments, *in* H. Sauermann, ed., 'Beiträge zur experimentellen Wirtschaftsforschung', J.C.B. Mohr, Tübingen, pp. 136–168.

Shafir, E. and Tversky, A. (1992), 'Thinking through uncertainty: Nonconsequential reasoning and choice', *Cognitive Psychology* **24**(4), 449–474.

Starmer, C. and Sugden, R. (1991), 'Does the random-lottery incentive system elicit true preferences? An experimental investigation', *American Economic Review* **81**(4), 971–978.

Thöni, C. and Volk, S. (2018), 'Conditional cooperation: Review and refinement', *Economics Letters* **171**, 37–40.

Tversky, A. and Fox, C. R. (1995), 'Weighing risk and uncertainty', *Psychological Review* **102**(2), 269–283.

Tversky, A. and Kahneman, D. (1992), 'Advances in prospect theory: Cumulative representation of uncertainty', *Journal of Risk and Uncertainty* **5**(4), 297–323.

Ungemach, C., Chater, N. and Stewart, N. (2009), 'Are Probabilities Overweighted or Underweighted When Rare Outcomes Are Experienced (Rarely)?', *Psychological Science* **20**(4), 473–479.

Van den Assem, M. J., Van Dolder, D. and Thaler, R. H. (2012), 'Split or steal? cooperative behavior when the stakes are large', *Management Science* **58**(1), 2–20.

Wulff, D. U., Mergenthaler-Canseco, M. and Hertwig, R. (2018), 'A meta-analytic review of two modes of learning and the description-experience gap.', *Psychological Bulletin* **144**(2), 140–176.