

The Description-Experience Gap in Cooperation*

Ozan Isler[†] Orestis Kopsacheilis[‡] Dennie van Dolder[§]

October 21, 2023

Abstract

Conditional cooperation is usually investigated in experiments where the choices of others are known. In many circumstances, however, there is uncertainty about others' cooperativeness. Using a novel experimental protocol, we manipulate the perceived likelihood of cooperation in a Prisoner's Dilemma, and whether such information is described unambiguously or learned through experience and thus ambiguous. We report on a 'description-experience gap' in which rare events appear to be more influential under experience than under description. This contrasts with earlier results from the individual choice literature. We show how stronger priors under social than individual uncertainty can account for this reversal.

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

JEL codes: C72, C92, D81, D83

*We thank Mohammed Abdellaoui, Robin Cubitt, and Simon Gächter for their valuable comments 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/>).

[†]Centre for Behavioural Economics, Society and Technology, School of Economics and Finance, Queensland University of Technology, Brisbane, Australia.

[‡]Corresponding author. School of Management, Department of Economics and Policy, Technical University of Munich. email:orestis.kopsacheilis@tum.de

[§]Department of Economics, University of Essex.

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 usually investigated in settings where the behavior of others is known with certainty (Fischbacher et al., 2001; Thöni and Volk, 2018). Outside the laboratory, however, the decision to cooperate often has to be made under uncertainty. For example, a researcher starting a project with a new collaborator does not know how cooperative her collaborator will be. Similarly, a freelancer does not always know if her client will pay the agreed amount before completing the job. In such instances, conditional cooperation requires the formation of expectations about others' behavior (Hayashi et al., 1999; Clark and Sefton, 2001; Van den Assem et al., 2012).

Information on the cooperativeness of others can be acquired in different formats. In some cases, relevant descriptive information may be available. For example, a freelancer who finds a big client through an online labor market will often have access to a large number of reviews by other freelance workers detailing their experiences, which would allow for a reasonably accurate as-

assessment of the client's trustworthiness. With sufficiently accurate descriptive information, the decision to cooperate or not can be seen as a decision under risk, where probabilities are objective and known. In many other cases, people lack sufficiently accurate descriptive information and need to rely on personal experience to form expectations. This would, for example, be true for the aforementioned researcher who is deliberating whether or not to enter into a new collaboration and for a freelancer who does not have access to a large corpus of client reviews. Under these conditions, we can think of the decision to cooperate or not as a decision under ambiguity, where probabilities are (at least partially) unknown.

Research on individual decision making suggests that people's choices differ systematically between situations where outcomes and their respective probabilities are objectively described and situations where these are unknown and have to be learned through experience. The most common experimental setup in this literature has subjects make a series of choices between a relatively risky and a relatively safe prospect in two conditions: Description and Experience (Hertwig et al., 2004). In Description, subjects learn about the statistical properties of the prospects through numerical descriptions of the possible outcomes and their associated probabilities. In Experience, on the other hand, subjects learn about these properties by sampling a sequence of independent observations. Early studies in this area found that while rare events appear to be overweighted in description-based choices, they appear to be underweighted in experience-based choices (Hertwig et al., 2004; Hau et al., 2008; Ungemach et al., 2009). More recent studies that employ more diverse choice sets and

structurally estimate weighting functions under Cumulative Prospect Theory (Tversky and Kahneman, 1992), find evidence in the same direction, albeit with a more moderate pattern: rare events appear to also be overweighted in Experience, but to a lesser degree than in Description (Abdellaoui et al., 2011; Kopsacheilis, 2018; Aydogan and Gao, 2020; Cubitt et al., 2021).

The finding that rare events are less influential in experience-based than in description-based choices has become known as the ‘description-experience gap’. Part of this gap can be accounted for by sampling bias (Fox and Hadar, 2006; Rakow et al., 2008). Owing to the skewness of the binomial distribution, people who only sample a small number of observations often under-represent the objective frequency of the rare event and in many cases fail to sample it even once, leaving them unaware of its existence. Nonetheless, most studies that control for sampling bias find that the gap persists (albeit significantly diminished; Hau et al., 2008; Cubitt et al., 2021).¹

Although the description-experience gap has been well-studied in the context of individual decision making (see Wulff et al., 2018, for a review), its implications for decision making under social uncertainty remain largely unexplored. Responses to social uncertainty, in which the optimal choice depends on the decisions of others, have been found to differ from responses to uncer-

¹A notable exception is Glöckner et al. (2016), who find evidence that the direction of the description-experience gap depends on the characteristics of the task employed. In particular, they only observe that subjects appear to give less weight to rare events in Experience than in Description when they restrict their analyses to choice sets in which one prospect is degenerate (i.e., provides an outcome with certainty). When they consider choices between non-degenerate prospects, they find evidence for a reversed description-experience gap, in which rare events appear to receive more weight in Experience compared to Description. Aydogan (2021) provides a theoretical explanation for these findings.

tainty generated by ‘random acts of nature’.² Research on betrayal aversion, for example, suggests that people are less willing to take risk when another person rather than nature determines the outcome (Bohnet and Zeckhauser, 2004; Bohnet et al., 2008; Aimone and Houser, 2012). Furthermore, Costa-Gomes and Weizsäcker (2008), Fetchenhauer and Dunning (2012) and Li et al. (2020) report evidence that behavior in strategic games is less responsive to changes in beliefs than behavior in games against nature. Such differences in responses to social uncertainty and individual uncertainty may also affect a potential description-experience gap.

In addition, research in evolutionary psychology suggests that the effect of experience may be different in social than in individual tasks. In particular, it is argued that human cognitive functions have adapted to detect non-cooperators and that, as a result, people are especially good at remembering past actions of non-cooperators (Cosmides, 1989; Cosmides and Tooby, 1989; Tooby and Cosmides, 2005; Buchner et al., 2009; Bell et al., 2012; Hechler et al., 2016). This feature of memory would have a greater effect when learning from experience than when learning from description—as in the former case people have to rely on their own memory, while in the latter they do not—and could thus lead to a different description-experience gap in social tasks than in individual decision-making tasks.

In this paper, we investigate whether the description-experience gap also

²There is considerable evidence that decisions under uncertainty depend critically on the source generating the uncertainty. The study of the source of uncertainty 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 choices 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).

emerges in a social context. In particular, we study whether and how the format—either descriptive or experiential—through which subjects obtain information regarding the likelihood that their partner in a one-shot Prisoner’s Dilemma will cooperate affects their choice to cooperate.

To this end, we develop a novel experimental protocol that allows for the exogenous and systematic manipulation of the information that subjects receive regarding the probability that their partner will cooperate. Our method is based on the ‘conditional information lottery’ incentivization protocol ([Bardsley, 2000](#)), and works as follows. Each subject plays a one-shot Prisoner’s Dilemma game with their partner, another subject in the session. This partner is randomly drawn from a subpopulation of subjects who have already made their decisions to either cooperate or defect. Before making their decision, subjects receive information regarding the cooperation rate in the subpopulation from which their partner is drawn. In order to observe how subjects condition their cooperative behavior on the likelihood that their partner will cooperate, we present them with seven potential cooperation rates and ask them to make a decision for each situation. The subjects know that only one of the seven situations involves the true cooperation rate of the subpopulation from which their partner is drawn. Because subjects do not know ex-ante which cooperation rate is the actual one, it is incentive compatible for them to treat each situation as if it is payoff relevant. Furthermore, because subjects are fully aware that only one situation involves the true cooperation rate, no deception occurs.

In our first and main experiment, we use the protocol described above to

test for the existence of a description-experience gap in cooperation. We vary the format in which subjects acquire information about the cooperation rate in the subpopulation from which their partner is drawn. In the Description treatment, subjects receive information regarding the frequency of cooperative actions in a numerical format (e.g. '70% chose to cooperate and 30% chose to defect'). In our Experience-Free treatment (E-Free), they instead sample decisions made by members of the subpopulation one at a time and with replacement. To control for sampling bias, we also conduct a modified Experience treatment, Experience-Fixed (E-Fixed), where subjects are required to sample a fixed number of times and the observed relative frequency of cooperation always matches the true probabilities (i.e., sampling without replacement). We investigate the existence of a description-experience gap by comparing the Description treatment with the two Experience treatments. Furthermore, we investigate the impact of sampling bias by comparing the two Experience treatments with each other.

In line with the hypothesis that conditional cooperation is pervasive, we observe that subjects' cooperation rate increases monotonically with the cooperation rate in the subpopulation from which their partner is drawn. However, our treatment comparisons reveal clear evidence for a description-experience gap in cooperation. Interestingly, this gap differs from the prevalent findings in individual risky decisions in two important ways. First, the direction of the gap is 'reversed': rare events appear more influential under experience than under description. Second, we find that sampling bias does not markedly affect this gap in cooperation: aggregate patterns in cooperation observed in

both variations of Experience are almost identical. This is in sharp contrast to the literature on risky choice, where sampling bias has been found to be the most important driver of the description-experience gap (Fox and Hadar, 2006; Rakow et al., 2008; Cubitt et al., 2021).

In addition, we elicit conditional cooperation preferences using the strategy method, the standard approach in the literature (Selten, 1967; Fischbacher et al., 2001). In particular, we measure preferences for conditional cooperation by allowing subjects to condition their choices on the actual behavior of their partners, removing any uncertainty. Our findings indicate that preferences elicited under information certainty predict behavior under uncertainty to a striking degree of accuracy. Free riders and unconditional cooperators should not care about information regarding others' intentions or actions. Consistent with such preferences, we find that the observed description-experience gap is primarily driven by conditional cooperators. Furthermore, conditional cooperators sample more information than free-riders or unconditional cooperators in the E-Free treatment, in which subjects can decide how much social information to accumulate.

Having observed that the description-experience gap under social uncertainty differs from what we would expect given the common finding in individual choice experiments, the question of what causes this difference arises. We argue that stronger priors under social uncertainty than towards the types of uncertainty used in the typical individual choice experiment can account for at least part of this difference. If people hold stronger priors, their posterior beliefs will be less responsive to new information. This will give rise

to a ‘regression to the mean’ effect leading them to 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 under social uncertainty than under uncertainty in individual choice.

To the best of our knowledge, only two previous studies have investigated the potential description-experience gap in a social context. [Artinger et al. \(2012\)](#) studied the description-experience gap in a Public Goods Game, but in their study the uncertainty concerned the relative benefit of contributing to the public good, not the actions of others. They find no indication of a description-experience gap in cooperation. [Fleischhut et al. \(2014\)](#) studied decisions from description and experience in an ultimatum bargaining setting where subjects learned how often particular offers had been rejected in a previous experiment before making their offer. They find some indication that the proportion of risky decisions is lower under experience than under description, but only when sampling bias is eliminated. In the presence of sampling bias they find no difference between description and experience-based choices. They are unable to draw conclusions regarding the moderating effect of probabilities on the gap between description and experienced based choices, as there is no exogenous variation in rejection probabilities in their design. Our study adds to this previous literature by focusing on uncertainty in the social domain of cooperation, by introducing controlled variation in the probability that others will cooperate, and by investigating heterogeneity in the degree to which subjects seek out and respond to social information on the basis of their cooperative preferences.

Our study provides the first clear evidence for a description-experience gap in a social context. We develop and validate a flexible and deception-free experimental protocol that allows for the exogenous and systematic manipulation of subjects' expectations regarding the likelihood that their partner will cooperate. Consequently, we identify an important mechanism underlying the difference between social and individual uncertainty: people have significantly stronger priors when the resolution of uncertainty is arbitrated by the action of a fellow human rather than a random chance mechanism.

2 Described vs. Experienced Social Uncertainty

Our first and main experiment is designed to investigate whether people's cooperative choices differ systematically between situations where they learn about the likelihood that their partner will cooperate through description or through experience. To investigate this question, we develop a novel experimental protocol that allows for the systematic manipulation of the information that subjects receive regarding the probability that their partner will cooperate. In addition, we also explore heterogeneity in behavioral responses across subjects who differ in their cooperative preferences.

2.1 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.³

³One 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 normative 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). 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.

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.

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.

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 subjects 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).

The way i obtains information about the $SPoC$ 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’.⁴ 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 matches the objective $SPoC$ level of the present scenario.

Finally, in Stage 3 we elicit cooperation preferences using the strategy method. Again, we ask subjects for their decision in a Prisoner’s Dilemma, but now we allow them to condition their decision directly on that of their partner. 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

⁴Technically, 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 A.3.

decisions in Stage 2 that take place under (social) uncertainty, the conditional decision 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 Good 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'.⁵

2.2 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).⁶ 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, alleviate the necessity of real-time interaction between subjects.

Prior to making their decisions, subjects had to correctly answer a comprehension question designed to test whether they understood the payoff consequences of their choices in the Prisoner's Dilemma for their own and their partner's individual earnings. Subjects who failed to provide correct answers

⁵The '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.

⁶Our selection criteria were that subjects were UK residents and had an approval rating of 90 and above.

to the comprehension question in three attempts were given a participation fee, but were not allowed to proceed to the experiment. This was the case for 43 subjects in total (3.9%, 17 in Description, 16 in E-Free, 10 in E-Fixed). Furthermore, 61 other subjects failed to complete the entire task (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 for approximately 20 minutes on average. Subjects received a fixed participation fee of £1.25, and a variable payment that could vary between £0.00 and £1.50. Only one choice per subject, selected at random, determined the variable payment. The average total earning was approximately £2.00 per subject. The allocation of subjects to treatment 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 increase statistical power for the analyses related to sampling decisions in that treatment. In the end, we observe the behavior of 279 subjects in the Description treatment, 276 in the E-Fixed treatment, and 435 in the E-Free treatment. Appendix [A.3](#) presents the instructions used in the experiment, and Appendix [A.2](#) explains the matching protocol used to determine payments.

2.3 Results

Across all treatments, 57.9% of our overall sample cooperated in Stage 1. This percentage did not differ significantly across treatments ($\chi^2(2, 990) = 3.84, p = 0.146$) and is similar to other experiments employing one-shot Pris-

oner’s Dilemmas (Sally, 1995; Mengel, 2017).

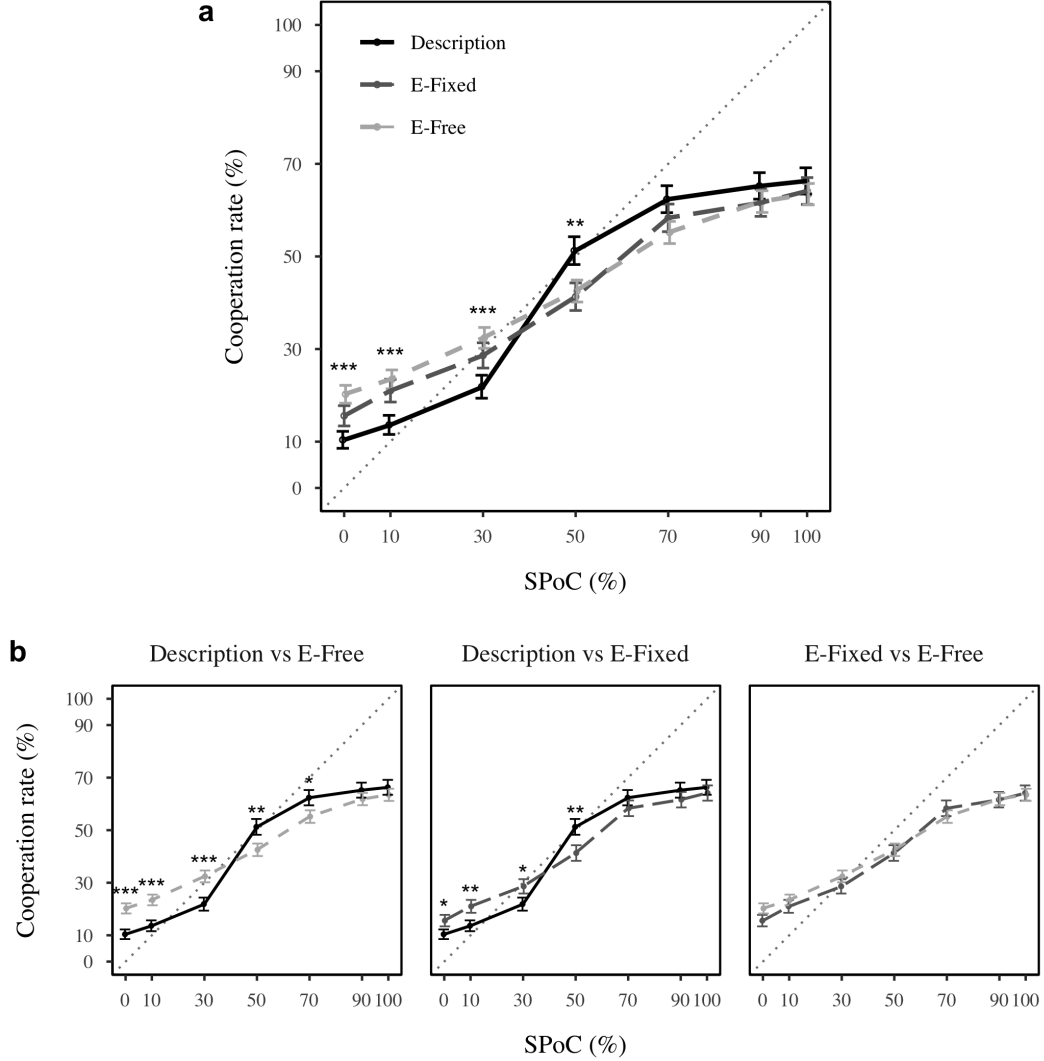
Next, we turn to the main task of Experiment 1. Figure 1 presents the average cooperation rates in the Prisoner’s Dilemma game of Stage 2 across the different $SPoC$ levels for each treatment. In each treatment, the cooperation rate increases with the $SPoC$ level: subjects are more likely to cooperate when the probability of being matched with a cooperator is higher. This pattern suggests that there is a considerable tendency towards conditional cooperation. Nevertheless, Figure 1 also suggests that there is a tendency for unconditional behavior. Specifically, a sizable fraction of subjects cooperate even when the probability of being matched with a cooperator is zero ($SPoC = 0$) or defect when the probability of being matched with a cooperator is one ($SPoC = 1$).⁷

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

We now move on to the investigation of the description-experience gap in cooperation. Our statistical analysis suggests that there are significantly different patterns of cooperation across the three treatments (see Figure 1, and Table A1 in Appendix A). As can be seen in Figure 1, cooperation rates in Experience are statistically significantly higher than those in Description when cooperation is relatively infrequent (i.e., for $SPoC < 0.5$). At the same time, cooperation rates in the Experience treatments are below those in the Description treatment when cooperation is relatively frequent (for $SPoC \geq 0.5$), but except for $SPoC = 0.5$ this difference is not statistically significant.

⁷Only subjects in the description treatment know this probability for sure. Subjects in the experience treatments do not know whether or not the cooperation rate that they observe reflects the objective $SPoC$.

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 agent. 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.

Therefore, although we observe a description-experience gap in cooperation, this pattern is the opposite of the canonical finding in individual deci-

sions under risk and uncertainty: rare events appear to have a larger effect on cooperation when information is acquired through experience than when it is acquired through description. This difference is particularly pronounced when cooperation is rare (see the leftmost region of Figure 1).

Result 2 *There is a significant description-experience gap in cooperation. This gap is particularly pronounced when the likelihood of cooperation is low, where people in Experience tend to cooperate more than those in Description.*

Whereas the description-experience gap in individual risky decisions is observed for rare or infrequent events, in our social contest we also find significant differences between Description and Experience at $SPoC = 0.5$. Although cooperation is significantly less likely in Description compared to Experience at $SPoC = 0.3$ ($\chi^2(2, 990) = 9.32, p = 0.009$), the opposite is the case at $SPoC = 0.5$ ($\chi^2(2, 990) = 6.97, p = 0.031$). We will get back to this sharp change when we analyze the changes in cooperation rates across $SPoC$ levels.

We now consider whether the description-experience gap in cooperation is due to sampling bias—the leading factor of a description-experience gap in individual risky decisions. As expected, E-Free exhibits significant sampling bias: subjects in E-Free sampled relatively little, with a median sample of 4 cards per round. As a result, in 63% of all cases where a sample was obtained, the relative observed frequency misrepresented $SPoC$ by 10 percentage points or more.

If sampling bias were an important driver of the description-experience gap in our social setting, we would have observed significant differences between

E-Free and Description (where sampling bias is present) and between E-Fixed and Description (where sampling bias is exogenously eliminated). Instead, the two description-experience gaps are very similar. Furthermore, χ^2 -tests do not reject the null hypothesis of equal cooperation rates between the two Experience treatments for any level of *SPoC* (for all seven tests: $\chi^2(2, 990) < 2.43, p > 0.119$).

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*. Formally, these indexes are calculated as follows:

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

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

C_{ir} takes the value 1 (0) if subject i decides to cooperate (defect) at $SPoC_r$, where n stands for the total number of subjects in a given treatment. The *cooperativeness* index in Equation 1 represents average cooperation in a treatment across all levels of *SPoC*. The *conditionality* index in Equation 2 on the other hand, captures the overall change in cooperation between $SPoC = 0$ to $SPoC = 100$.⁸ Intuitively, values of *conditionality* that are closer to 1 are sug-

⁸A sizeable proportion of people (approximately a third of the total sample) switched their action more than once when considering all *SPoC* scenarios. The *conditionality* index is unaffected by these inconsistencies as it ignores such intermediate switches.

Table 2: *Cooperativeness* and *conditionality* indexes across treatments

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

gestive of a stronger tendency for conditional cooperation in that treatment.⁹

Table 2 reports values and statistical comparisons for these two indexes across treatments. Cooperativeness does not differ across the three treatments (Kruskal-Wallis test, $H(2) = 0.190$, $p = 0.909$) but *conditionality* does (Kruskal-Wallis test, $H(2) = 10.6$, $p = 0.005$). Moving to binary comparison between treatments, we find that *conditionality* differs statistically significantly between Description and E-Free and marginally significantly 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

⁹Behavior corresponding to ‘reverse conditional cooperation’ ($C_{i1} = 1$ and $C_{i7} = 0$) is rare and symmetrically distributed across treatments. Therefore, it cannot drive any treatment differences (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, $P = 0.645$). Excluding these cases from the analysis does not change our results. We are agnostic about the interpretation of reverse conditional cooperation behavior: it may be due to error or misunderstanding, or it may represent a rare type of cooperation preference.

vs E-fixed $U = 57,109, p = 0.213$). Thus, the differences observed between treatments 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 *People in Description react more strongly to social information than those in Experience.*


The *conditionality* index captures the change in cooperation between $SPoC = 0$ and $SPoC = 1$, but it does not tell us at which SPoC levels the increase in cooperation takes place. To investigate this, Figure 2 plots the slope of the response function at each transition from $SPoC_r$ to $SPoC_{r+1}$.¹⁰ Formally, this slope is calculated through Equation 3. Intuitively, the steeper the slope, the stronger the reaction in a treatment to the transition.

$$slope(r, r + 1) = \frac{1}{n} \sum_{i=1}^n \frac{C_{i(r+1)} - C_{ir}}{SPoC_{r+1} - SPoC_r} \quad (3)$$

Figure 2 highlights the sharp increase in cooperation in Description at the transition to $SPoC = 0.5$, i.e. where there is a 50% chance that the match would cooperate. The increase in cooperation between $SPoC = 0.3$

¹⁰An alternative to looking at the slope would be to look at the SPoC transition at which people switch from defection to cooperation. We consider the slope analysis to be more appropriate for two reasons. First, the slope analyses does not require us to exclude subjects who switched more than once. This is especially important because instances of multiple switching are not equal across treatments; subjects in the E-Free treatment are more likely to switch multiple times than those in the other two treatments (30.8% in Description, 31.9% in E-Fixed and 39.3% in E-Free; $\chi^2(2,990)=6.87, p=0.032$). This is understandable, as subjects in the E-Free treatment often do not accurately sample the true SPoC value. Second, by calculating the slopes we control for the fact that $SPoC$ is not equidistantly distributed across the unit interval (for example, there are 20 percentage points between $SPoC = 70$ and $SPoC = 90$ but only 10 percentage points between $SPoC = 90$ and $SPoC = 100$).

Figure 2: Change in the cooperation rate over the change in SPoC



Note. ‘ $X \rightarrow Y$ ’: transition from $SPoC = X$ to $SPoC = Y$. The horizontal dotted line gives the point where the slope is equal to unity. Error bars represent standard errors.

and $SPoC = 0.5$ in the Description treatment is remarkable: it is the only transition in any of the treatments where the level of cooperation increases more sharply than the $SPoC$ value itself ($slope = 1.52$), in all other transitions the increase in the level of cooperation is substantially lower than the increase in $SPoC$ (all $slopes < 1.0$).

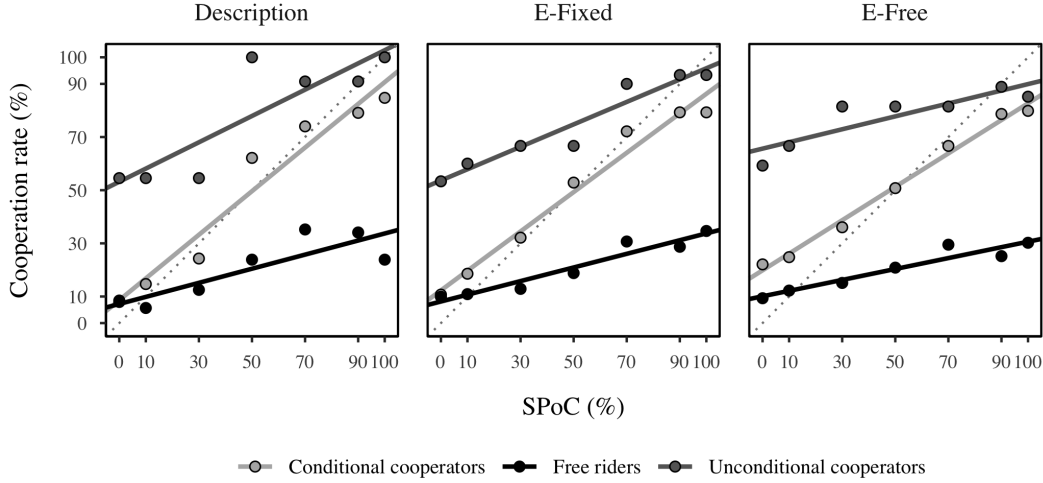
Result 5 *The only significant difference in reactions to transitions of cooperation, occurs towards $SPoC = 0.5$. When subjects in Description learn that at least half of their sub-group also cooperates, they switch from defecting to cooperating in disproportionately high frequencies.*

Cooperative behavior diverges significantly between Description and Experience at $SPoC = 0.5$, despite the fact that neither cooperation nor defection can be characterized as rare in this case. The slope analysis shows that this difference occurs due to a strong reaction by subjects in Description, relative to those in Experience, to the information that half of their potential partners chose to cooperate. It is at this point that the cooperation pattern between Description and experience flips: although subjects in Experience were more cooperative than those in Description for $SPoC < 0.5$, the opposite is the case for $SPoC \geq 0.5$.

We now turn to the analysis of actions in Stage 3, where we elicited conditionally cooperative preferences in the one-shot Prisoner’s Dilemma under certainty. Decisions for the game were elicited under certainty regarding the other player’s actions, allowing us to categorize subjects into cooperation types in a way that has become standard in the literature (Fischbacher et al., 2001). Overall, the majority of subjects was categorized as conditional cooperators (Description: 63.4%; E-Free: 59.3%; E-Fixed: 50.7%). The second most frequent category was that of free riders (Description: 31.5%; E-Free: 32.0%; E-Fixed: 36.6%) while unconditional cooperators were in the minority (3.9%; 6.2%; 10.9%). Those who did not fit into any of these three categories were rare (Description: 1.1%; E-Free: 2.5%; E-Fixed 1.8%).¹¹

¹¹The relative frequency of the different cooperation types was not equal across the three treatments ($\chi^2(6, 990) = 17.54, p = 0.007$), with a lower prevalence of conditional cooperators in E-Fixed (50.7%) than in 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 we used to test for sampling bias, may be due to its enforcement of a long sampling process. If subjects in this condition were motivated to quickly click through the study, then they would be more likely to act unconditionally (e.g., always defecting or always cooperating). Nevertheless, the type distributions did not differ between the E-Free

Figure 3: 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.

Figure 3 depicts behavior in Stage 2 according to cooperation type identified in Stage 3, whereas Table 3 reports cooperation indexes for each type across the treatments. Results clearly indicate that type categorization, elicited under information certainty, successfully predicts Stage 2 behavior in all treatments. In particular, those who are categorized as conditional cooperators in Stage 3 are substantially more sensitive to changes in the likelihood that their partner will cooperate in Stage 2 as compared to the other types. Likewise, those categorized as unconditional cooperators score highest on the *cooperativeness* index and score low on *conditionality*, while those who are categorized as free riders have the lowest *cooperativeness* score and also the lowest *condi-*

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.

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)
<i>p</i>	0.768	0.877	0.789	0.0002	0.529	0.403

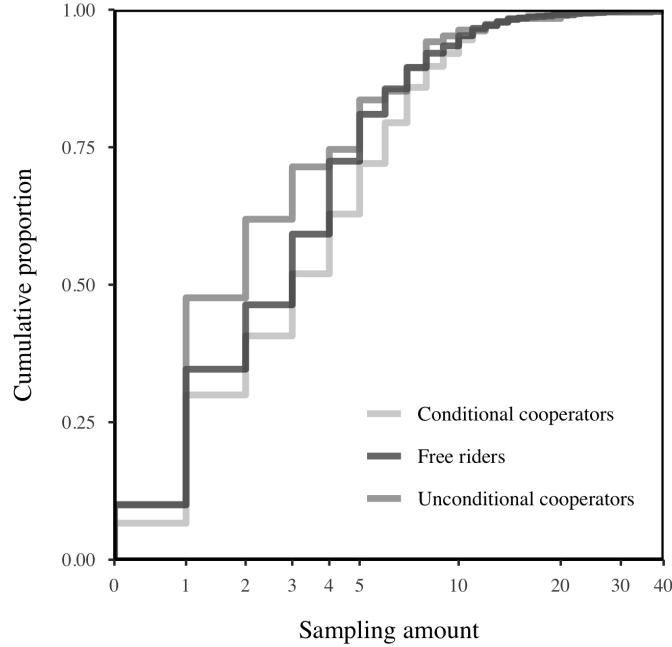
Notes: ‘CC’: Conditional Cooperators; ‘FR’: Free Riders; ‘UC’: Unconditional Cooperators. Standard errors in parentheses. The *p*-values derive from Kruskal-Wallis tests on individual-level measures of *cooperativeness* and *conditionality* across the three treatments.

tionality score.

Statistical tests, reported on Table 3, verify that treatment differences are driven by conditional cooperators who exhibit different degrees of *conditionality* (but not *cooperativeness*) across treatments. There are no significant treatment differences among free riders or conditional cooperators. This is in line with the expectation that the description-experience gap in cooperation can only be driven by people who care about social information.

Finally, we consider the E-Free treatment in more detail, where subjects could decide how much information to sample. Although sampling did not entail any monetary cost, it does require exerting more effort and spending more time on the task. Theoretically, the willingness to incur costs in return for information about the *cooperativeness* of the environment should depend

Figure 4: 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.

on social preferences. In particular, conditional cooperators should be more interested in others' behavior than either unconditional cooperators or free riders. Hence, conditional cooperators should collect bigger samples.

Figure 4 plots the cumulative distribution of sampling amounts for each type in E-Free. Average sampling amount was 4.1 draws per round for conditional cooperators, 3.7 for free riders and 3.1 for unconditional cooperators. As predicted, conditional cooperators sampled significantly more than both free

riders (clustered Wilcoxon signed rank test¹²; $p = 0.032$) and unconditional cooperators ($p = 0.019$).

Result 6 *The cooperation preferences elicited in Stage 3 under information certainty are highly predictive of Stage 2 behavior under social information uncertainty. Specifically:*

6.1 *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.*

6.2 *The description-experience gap in cooperation is driven by conditional cooperators.*

6.3 *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 probabilistic information was obtained affected cooperative behavior (Result 2). However, the description-experience gap in this social context differed in two important ways from the canonical finding in the individual choice literature. First, its pattern is reversed: people cooperate more when cooperation is rare in the Experience treatments compared to the Description treatment, suggesting that rare events are


¹²See [Rosner et al. \(2006\)](#) for more details on this test.

more influential in Experience rather than in Description. Second, sampling bias—the leading contributor of the description-experience in individual choice literature—was not a significant driving factor. Rather, the social description-experience gap remained unaltered even when subjects collected samples whose relative frequencies matched the objective probabilities of those events (Result 3).

Here, we propose a parsimonious explanation for the description-experience gap that we observed that can account for the discrepancy with the individual choice literature. Our argument relies on the role of priors in belief updating. If decision makers have complete knowledge of objective outcomes and probabilities, as they do under Description, then prior beliefs are irrelevant. In Experience, however, probabilities are not fully known. In such cases, the posterior belief of the decision maker is some combination of the decision maker’s subjective prior and their observed sampling information. Naturally, the description-experience gap will reflect this asymmetry in the updating process. Furthermore, different settings—in this case that of individual or social uncertainty—may invoke different priors, and, hence, differences in the resulting description-experience gap (see also [Aydogan, 2021](#)).

According to this perspective, the extent of the gap will depend on the strength of priors. The more the weight placed on the prior belief the more the value of new information is discounted, and thus the less responsive posteriors are to new information. Figure 5 provides a visual demonstration of this argument by simulating posteriors for different levels of observed probabilities

Figure 5: 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.

and different strengths of prior belief, using [Carnap's \(1952\)](#) rule of updating.¹³

It is easy to see that as the strength of the prior increases, the posterior belief becomes flatter. Such flat posteriors give rise to 'regression to the mean' effects, where small probability events appear to be overweighted and high probability

¹³[Carnap's \(1952\)](#) rule offers a tractable way of tracking belief updating in decisions from experience (see also [Aydogan, 2021](#)).

events appear to be underweighted.

Our explanation rests on the assumption that priors regarding social actions and outcomes will generally be stronger than priors in abstract, non-social settings. When betting on the color of a ball that will be drawn from an urn with an unknown composition of balls, subjects are unlikely to have strong priors regarding the likelihood of any particular color and therefore willing to update their beliefs in light of new information. In contrast, when predicting others' cooperativeness, subjects will have stronger priors based on experience and perhaps even moral views which would make them more reluctant towards updating prior beliefs. For example, research has shown that there is considerably heterogeneity in the beliefs about the cooperativeness of others in Prisoner's Dilemmas, and that these beliefs strongly correlate with a subject's own social value orientation (Aksoy and Weesie, 2012; Pletzer et al., 2018).

If people hold strong priors, then they should be less responsive to the information that they obtain in Experience compared to that in Description. This is exactly what we observe. Furthermore, if people give relatively little weight to new social information this would also explain why sampling bias had less of an effect in the social rather than in the individual choice setting.

Evidence on sampling behavior provides some tentative indirect support for our hypothesis that people hold stronger priors under social than under individual uncertainty. 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. Indeed, the median sampling amount

of 4 draws per deck we observe in E-Fixed is considerably smaller than that reported by most previous studies.¹⁴ Corroborating this hypothesis, [Fleischhut et al. \(2014, 2018\)](#) report a similar tendency for reduced exploration when subject face social uncertainty rather than when facing lotteries with similar payoffs.

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 be on average larger for the latter. We preregistered our methods and hypothesis.¹⁵

3.1 Experimental design and procedures

We recruited 241 subjects through Prolific, who were randomly assigned to one of two treatments: Individual Uncertainty or Social Uncertainty.¹⁶ 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 difference between the two treatments relates to the source of uncertainty. In the Social Uncertainty condition, the decision is made in a social context. Subjects play a one-shot binary Prisoner's Dilemma, identical to the

¹⁴[Hills and Hertwig \(2010\)](#) report that the median subject sampled each option 9 times which seems to be close to the modal search effort in this literature (see also [Wulff et al., 2018](#)).

¹⁵The preregistration can be accessed at: https://aspredicted.org/5H5_CH7

¹⁶Similar to our first experiment, our selection criteria were that subjects were UK residents and had an approval rating of 90 and above. Three subjects were not allowed to resume in the study as they repeatedly failed to answer the comprehension question. All 3 discarded observations were in the Social treatment. Our analysis comprises of 120 observations in Individual Uncertainty and 118 in Social Uncertainty.

one in our main experiment (see Figure 1). After making their decision, they 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. In both treatments, subjects are asked about their confidence regarding their estimate.

The key dependent variable is the confidence regarding their prior. 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 50.5 of subjects chose to 'Share'. The two means are very similar and their difference 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 (enter statistics), suggesting that there is much more belief-heterogeneity in this social setting.

Crucially, in line with our hypothesis, we observed that subjects' confidence is higher in estimates conducted in the Social Uncertainty ($\mu = 3.98, sd = 1.42$) than in the Individual Uncertainty ($\mu = 3.38, sd = 1.79$) condition

(Mann–Whitney U test: $U = 5500.5, p = 0.002$).

Result 7 *People 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 occasions there is uncertainty about the cooperativeness of others. In this paper, we investigate how the likelihood of cooperation as well as how the format in which this information is acquired affects conditional cooperation. We exogenously manipulated the likelihood of cooperation and communicated to subjects either through explicit and numerical information (Description) or by asking them to discover this information through a sequential sampling procedure (Experience).

We contribute to the relevant literature by identifying—for the first time—a significant gap between descriptive and experiential learning in social settings where cooperation can emerge. When cooperation is rare, people are significantly more likely to cooperate if they learn about others’ cooperative intentions through an experiential learning process that requires sequential sampling rather than through explicit and numerical descriptions. The pattern of the social description-experience gap we captured is of the opposite direction than that of the canonical finding in individual risky choices. (e.g., [Hertwig et al. 2004](#); [Hertwig and Erev 2009](#); [Wulff et al. 2018](#), but see [Glöckner](#)

[et al. 2016](#) for a notable exception).

We propose an explanation for this “reverse” gap and provided evidence for it in a follow-up experiment. Specifically, we argued that the strength of prior beliefs regulate the influence of rare events in domains of ambiguity. We demonstrate in a simulation how sticky priors lead to less updating and flatter posterior beliefs that over-emphasize the likelihood of rare cooperative events. Furthermore, we show empirically that such sticky priors are more prominent under social uncertainty rather than individual uncertainty. One reason for this discrepancy between the two domains could be the abstractness of the task features in individual uncertainty. In the canonical experimental setup, uncertainty is represented by instruments such as urns of balls or decks of cards that are unlikely to be associated with strong priors. Conversely, the features of the decision context in our social uncertainty treatments have a social connotation and are therefore less abstract and more likely to relate to people’s past experiences and social value orientations.

To those researchers who wish to further investigate the cognitive, preferential, emotional, informational or otherwise asymmetries between individual and social uncertainty, our study offers the following important methodological tool. The novel experimental protocol that we develop in this paper allows for the systematic manipulation of subjects’ expectations regarding the likelihood that their partner will take a particular action. We validate our method by comparing observed behavior to established measures and find a remarkable degree of consistency. In our experiment, we use this protocol to investigate behavior in the Prisoner’s Dilemma, but it can easily be adjusted to study

behavior in different games where players' choices are contingent on beliefs of their partners action.

One implication of our findings with policy relevance is in the domain of nudging behavior through the provision of descriptive norms. Policy makers have often used descriptive statistics to mobilize collective action. Such tactics, however, need to be exercised judiciously: when the socially desirable behavior is rare, using descriptive norms can back-fire and the socially harmful but individually beneficial action can—inadvertently—become normalized (Cialdini et al., 1990, 1991, 2006). Research that observes such backfiring typically compares the situation in which people receive the (negative) descriptive norm information with the situation in which people receive no information at all and thus have to rely purely on their own prior. Our study adds to this literature by showing that communicating a negative descriptive norm is harmful even if people are otherwise able to infer the descriptive norm through experiential sampling. In short: when many act badly, 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):3003–37.
- 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*, Articles in Advance.
- 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.
- Bell, R., Buchner, A., Erdfelder, E., Giang, T., Schain, C., and Riether, N. (2012). How specific is source memory for faces of cheaters? Evidence for categorical emotional tagging. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38(2):457–472.

- 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.
- Buchner, A., Bell, R., Mehl, B., and Musch, J. (2009). No enhanced recognition memory, but better source memory for faces of cheaters. *Evolution and Human Behavior*, 30(3):212–224.
- Carnap, R. (1952). *The continuum of inductive methods*, volume 198. University of Chicago Press Chicago.
- 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*, volume 24, pages 201–234. Elsevier.
- 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.
- Clark, K. and Sefton, M. (2001). The sequential prisoner's dilemma: Evidence on reciprocation. *Economic Journal*, 111(468):51–68.
- Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with the wason selection task. *Cognition*, 31(3):187–276.
- Cosmides, L. and Tooby, J. (1989). Evolutionary psychology and the generation of culture, part II. case study: A computational theory of social exchange. *Ethology and sociobiology*, 10(1-3):51–97.
- Costa-Gomes, M. A. and Weizsäcker, G. (2008). Stated beliefs and play in normal-form games. *Review of Economic Studies*, 75(3):729–762.
- Cubitt, R. P., Kopsacheilis, O., and Starmer, C. (2021). An inquiry into the nature and causes of the description-experience gap. *Journal of Risk and Uncertainty (forthcoming)*.
- 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.
- 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. (2018). Not all uncertainty is treated equally: Search under social and nonsocial uncertainty. Working paper.
- Fleischhut, N., Artinger, F., Olschewski, 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*, pages 1048–1053. Cognitive Science Society.
- 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.
- 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.
- 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.
- Hau, R., Pleskac, T. J., Kiefer, J., and Hertwig, R. (2008). The description–experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, 21(5):493–518.
- 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.

- Hechler, S., Neyer, F. J., and Kessler, T. (2016). The infamous among us: Enhanced reputational memory for uncooperative ingroup members. *Cognition*, 157:1–13.
- 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.
- Hills, T. T. and Hertwig, R. (2010). Information Search in Decisions From Experience: Do Our Patterns of Sampling Foreshadow Our Decisions? *Psychological Science*, 21(12):1787–1792.
- 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 Houser, D. and McCabe, K., editors, *Neuroeconomics*, volume 20 of *Advances in Health Economics and Health Services Research*, pages 179–201. Emerald Group Publishing Limited.
- 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.
- 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.
- Li, C., Turmunkh, U., and Wakker, P. P. (2020). Social and strategic ambiguity versus betrayal aversion. *Games and Economic Behavior*, 123:272–287.
- Mengel, F. (2017). Risk and temptation: A meta-study on prisoner’s dilemma games. *Economic Journal*, 128(616):3182–3209.
- 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.
- 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 eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments. In Sauermann, H., editor, *Beiträge zur experimentellen Wirtschaftsforschung*, pages 136–168. J.C.B. Mohr, Tübingen.
- 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.
- Tooby, J. and Cosmides, L. (2005). Conceptual foundations of evolutionary psychology. In Buss, M., editor, *The Handbook of Evolutionary Psychology*, pages 5–67. Wiley Online Library.

- 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.

A Appendix

A.1 Cooperation rate

Table A1: Cooperation rates across treatments

SPoC	Description	E-Free	E-Fixed	p -value
0	10.4	20.2	15.6	0.002
10	13.6	23.4	21.0	0.005
30	21.9	32.4	28.6	0.009
50	51.6	42.5	41.3	0.031
70	62.4	55.1	58.3	0.163
90	65.2	61.8	61.6	0.592
100	66.3	63.4	64.1	0.732

Note: The Sub-population Probability of Cooperation (SPoC) is the probability of being matched to a cooperative agent in a given scenario. The p -values are for Pearson's χ^2 -tests across all three treatments ($df = 2, N = 990$).

A.2 Matching protocol and payment

Subjects were randomly assigned to a treatment and given an index number within this treatment (e.g. 1,2,3...). They were then assigned to subpopulations of 3 members. To illustrate this process assume that there are 6 subjects in a treatment and that they are assigned to two subpopulations - A and B - such that $A = \{1, 2, 3\}$ and $B = \{4, 5, 6\}$. The matching algorithm consists of the following steps:

- 1 and 2 are matched for Stage 2
- 4 and 5 are matched for Stage 3

- 3 and 6 get matched across subpopulations and get paid for Stage 1
- if there is a number that is indivisible by 3:
 - if two players are left then they are matched for Stage 1
 - otherwise, if one player is left unmatched, she gets the maximum payoff (£3.75)

Matchings for Stage 2 and Stage 3 require that one of the two players acts according to their Stage 2 action while the other according to their Stage 1 action. This is randomly decided for each pair.

A.3 Instructions and experimental interface

Unless specified otherwise, all screens were encountered by subjects across all three treatments.

Figure B1: Welcome screen

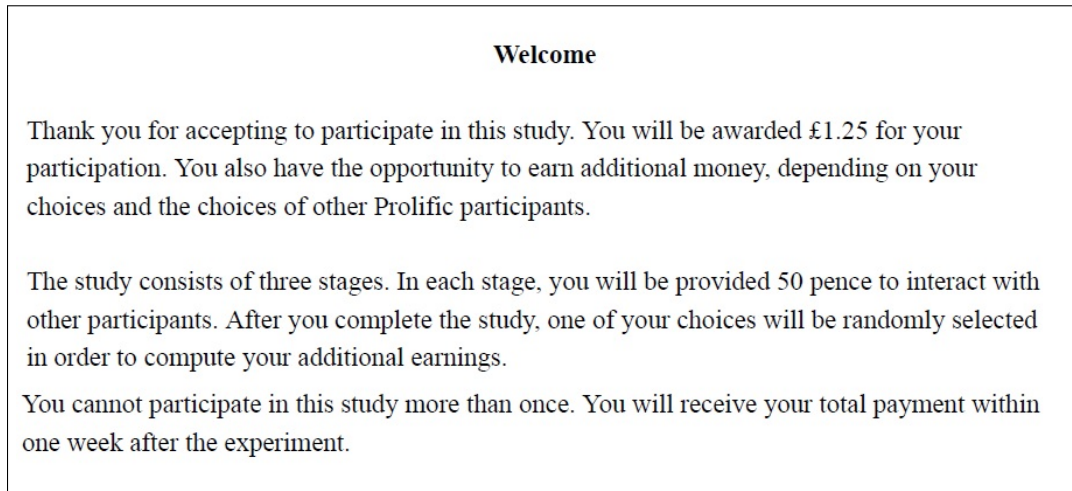


Figure B2: Stage 1/A

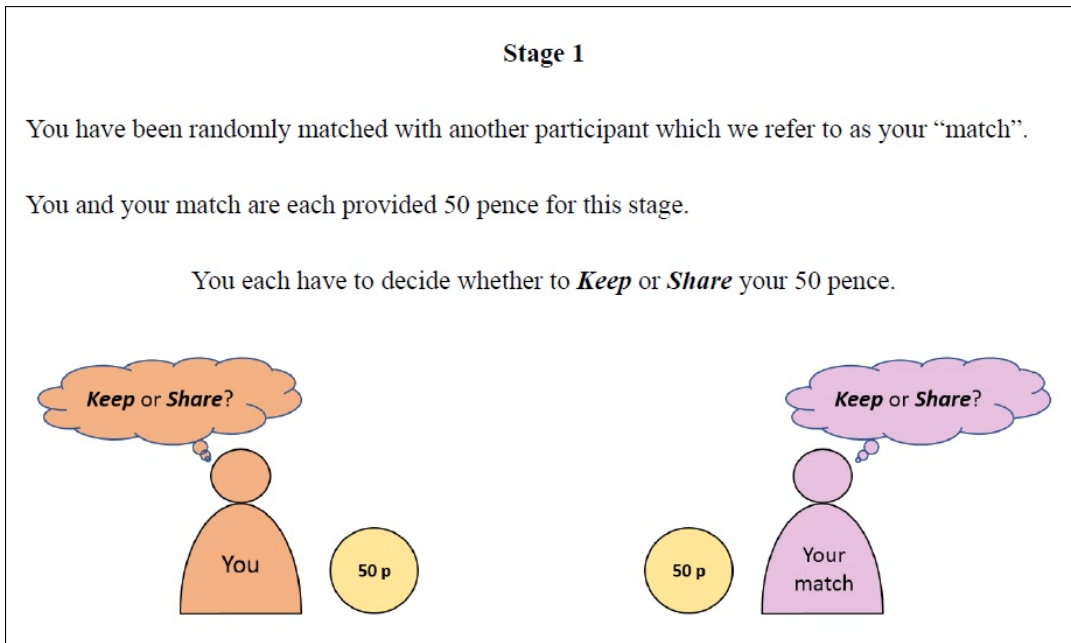


Figure B3: Stage 1/B

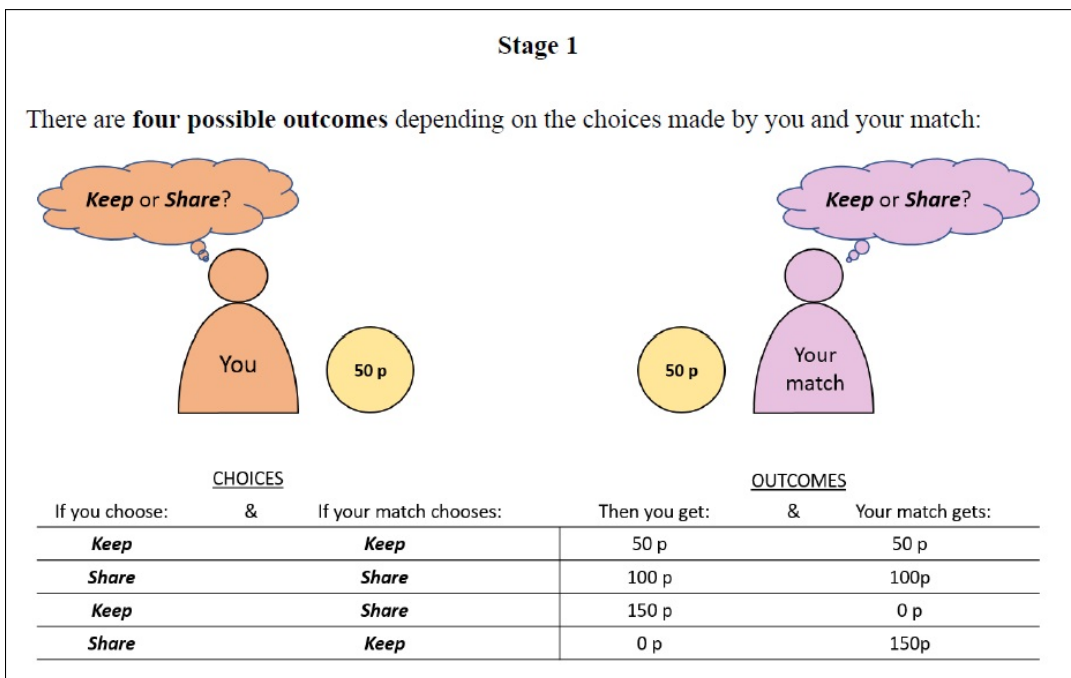


Figure B4: Stage 1: Decision

Your Decision for Stage 1

Now please make your decision:

KEEP

SHARE

If you choose:	&	If your match chooses:	Then you get:	&	Your match gets:
<i>Keep</i>		<i>Keep</i>	50 p		50 p
<i>Share</i>		<i>Share</i>	100 p		100p
<i>Keep</i>		<i>Share</i>	150 p		0 p
<i>Share</i>		<i>Keep</i>	0 p		150p

Note. This is the decision interface for Stage 1.

Figure B5: Stage 2/A

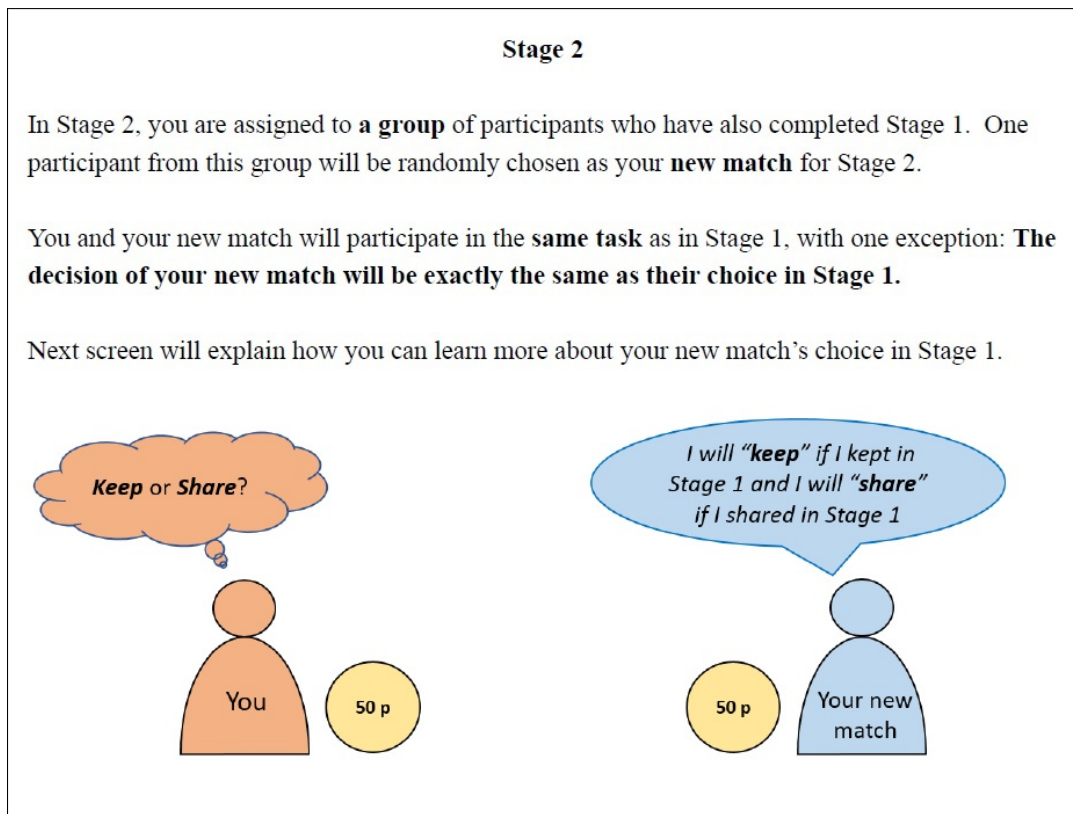


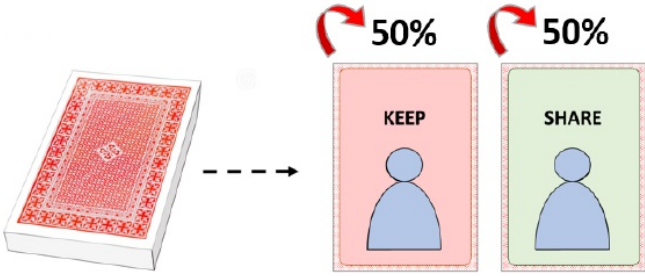
Figure B6: Stage 2/B: Description only

Stage 2

Before making your decision whether to keep or share your 50 pence with your new match, you will be given the opportunity to learn about the choices of your potential group members.

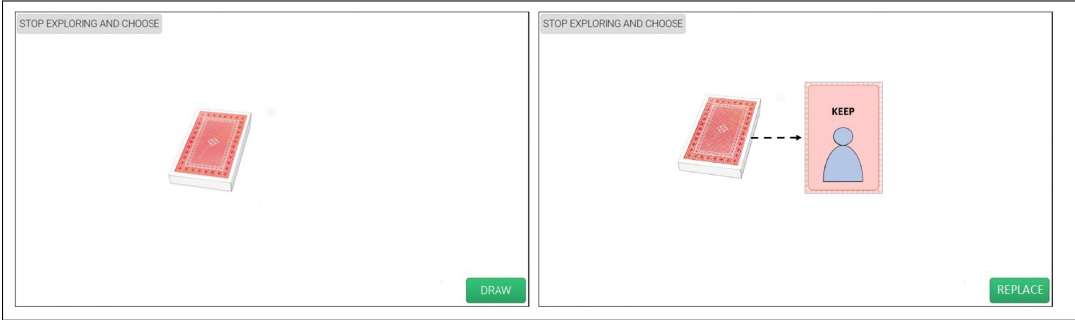
Suppose the deck below represents your actual group. Then the cards in the deck represent the Stage 1 choices of **each participant in your group including your new match** (but excluding your own choice).

The example below shows that 50% of the cards say "Keep" and 50% of the cards say "Share". Therefore, 50% of your group chose to "Keep" and 50% of your group chose to "Share". After you see the distribution of the two types of cards you will be asked to submit your decision.



Note. This screen was encountered only by subjects in Description. subjects in Experience saw instead the screen in the next Figure.

Figure B7: Stage 2/B: Experience only




Note. These two screens were encountered only by subjects in E-Free. subjects in E-Fixed saw a similar demonstration but there was no ‘STOP EXPLORING AND CHOOSE’ button on the top-left of the screen. Moreover, the ‘REPLACE’ button was replaced with one that read ‘NEXT CARD’ as in E-Fixed sampling was without replacement.

Figure B8: Stage 2/C

Stage 2

You will see **seven different decks of cards**.

Only **one** of the seven decks corresponds to the **actual group** you have been assigned to. The other six decks present hypothetical situations.



You do not know which deck describes your actual group. Therefore, you will make seven independent decisions to keep or share, one for each deck.

Only one of your seven decisions will be used to determine the outcome of Stage 2, which will be your decision for the deck that represents **the actual group** you are assigned to.

Hence, you should consider each deck independently of the other decks and make your decision assuming that the deck that you learn about in fact describes your actual group.

Figure B9: Stage 2/D

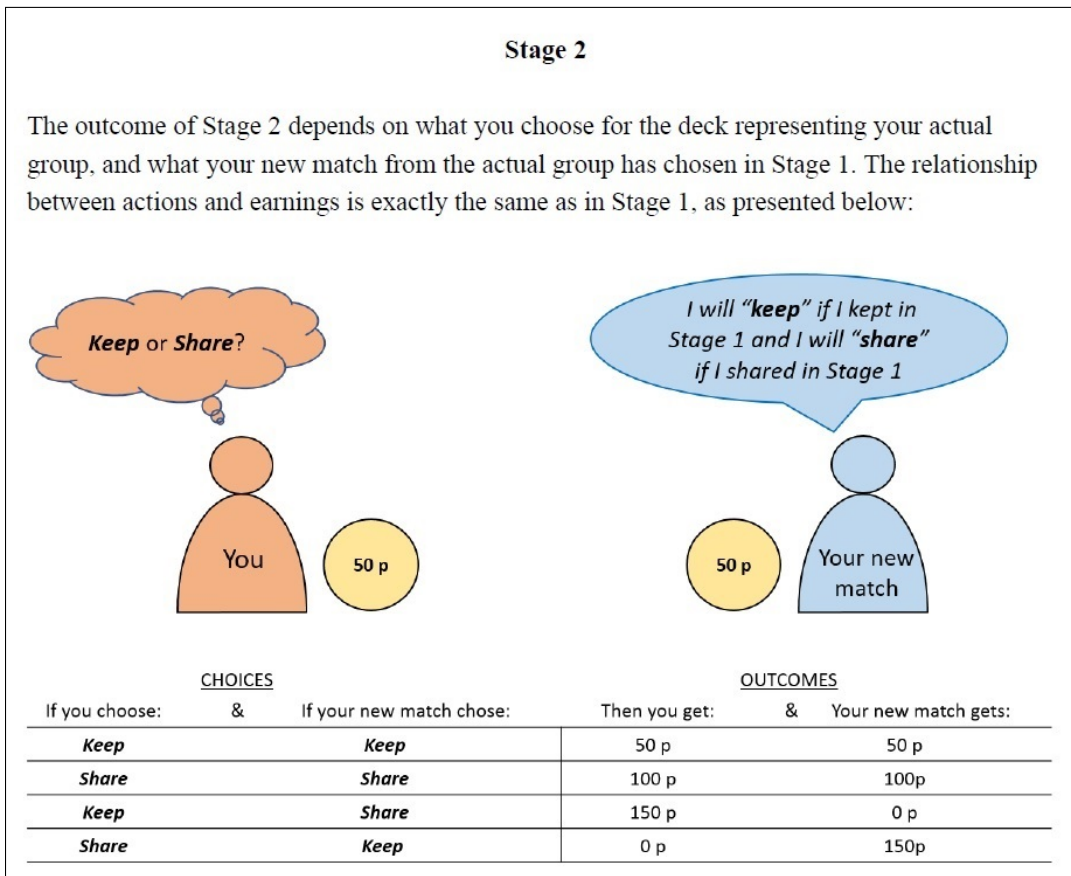


Figure B10: Stage 2: Decision

Your Decision

Please decide. Your decision will determine your earnings if your assigned group is actually represented by the deck that you just explored.

SHARE
 KEEP

If you choose:	&	If your new match chose:	Then you get:	&	Your new match gets:
<i>Keep</i>		<i>Keep</i>	50 p		50 p
<i>Share</i>		<i>Share</i>	100 p		100p
<i>Keep</i>		<i>Share</i>	150 p		0 p
<i>Share</i>		<i>Keep</i>	0 p		150p

Note. This screen follows the screen where subjects learn about the distribution of each scenario. Examples of how this information is obtained for each scenario can be seen in Figure B6 for Description and Figure B7 for Experience.

Figure B11: Stage 3/A

Stage 3

You have been randomly matched with another participant which we refer to as your “Stage 3 match”.

You and your Stage 3 match are each provided 50 pence for this stage.

You and your Stage 3 match will participate in the **same task** as in the previous two stages. Similar to Stage 2, **the decision of your Stage 3 match will be exactly the same as their choice in Stage 1.**

You will make a new choice whether to *Keep* or *Share* your 50 pence for **each possible decision that your Stage 3 match could have made in Stage 1.**

The diagram illustrates the decision-making process in Stage 3/A. On the left, an orange figure labeled 'You' is shown with a thought bubble containing the text 'Keep or Share?'. Next to the figure is a yellow circle labeled '50 p'. On the right, a green figure labeled 'Your Stage 3 match' is shown with a speech bubble containing the text 'I will "keep" if I kept in Stage 1 and I will "share" if I shared in Stage 1'. Next to the figure is a yellow circle labeled '50 p'.

Figure B12: Stage 3/B

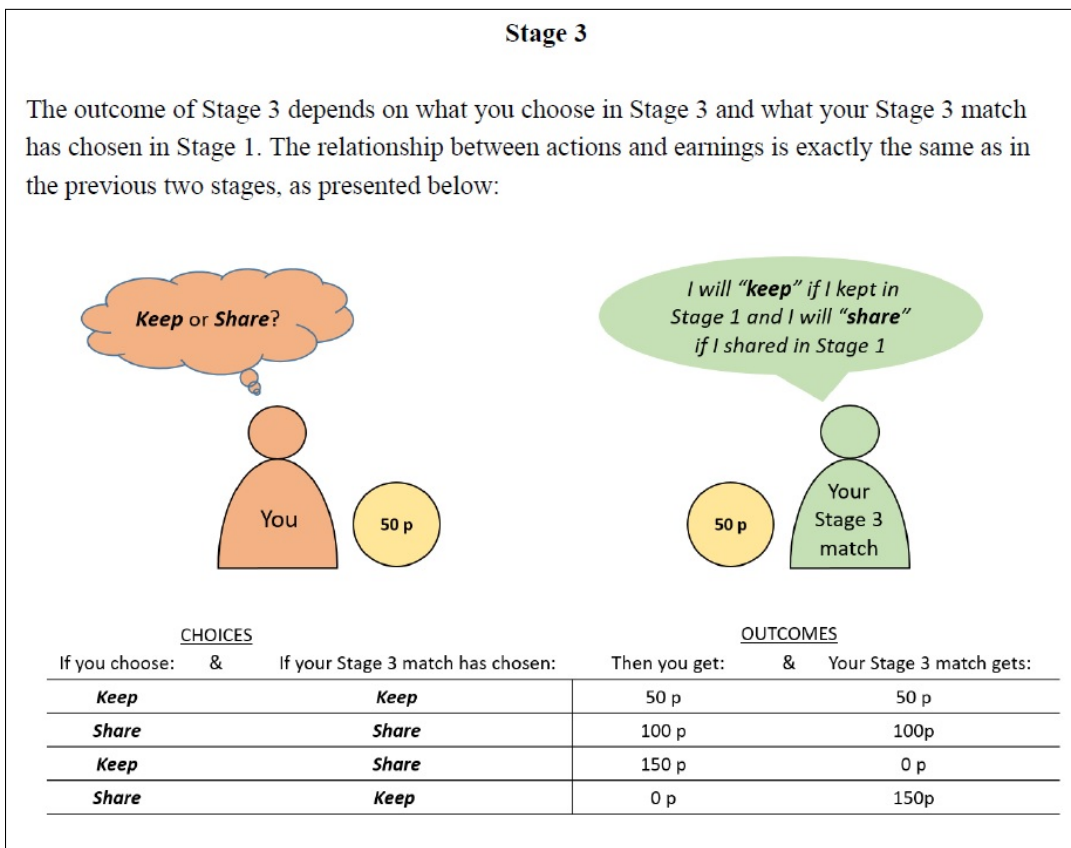


Figure B13: Stage 3/C

