CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The University of Nottingham
UNITED KINGDOM · CHINA · MALAYSIA

Discussion Paper No. 2018-10

Felix Kölle and
Thomas Lauer

October 2018

Cooperation, Discounting, and the Effects of Delayed Costs and Benefits

CeDEx Discussion Paper Series
ISSN 1749 - 3293

CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit http://www.nottingham.ac.uk/cedex for more information about the Centre or contact

Suzanne Robey
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0)115 95 14763
suzanne.robey@nottingham.ac.uk

The full list of CeDEx Discussion Papers is available at

http://www.nottingham.ac.uk/cedex/publications/discussion-papers/index.aspx

Cooperation, Discounting, and the Effects of Delayed Costs and Benefits

Felix Kölle* and Thomas Lauer

University of Cologne

This version: October 2018

Abstract

Numerous studies have investigated how people resolve intertemporal trade-offs in individual decision making, but little is known about how the timing of costs and benefits affects behavior in strategic decision situations. Here, we experimentally study how delayed costs and/or benefits affect cooperation in a social dilemma situation. We find that cooperation is substantially reduced (increased) when only the benefits (costs) of cooperation are shifted towards the future. We show that the change in contributions can be explained by (*i*) a shift in the beliefs about others' cooperativeness, (*ii*) a shift in subjects' willingness to conditionally cooperate, and (*iii*) a subject's degree of impatience. We further demonstrate that the amount of economic incentives needed to close the cooperation gap are substantial, indicating discount rates in our strategic context of about 50 percent, much higher than the ones typically observed in individual decision contexts. Finally, when both costs and benefits are delayed to the same extent, contribution levels do not change, indicating that cooperation is time-consistent.

Keywords: Public goods; cooperation; discounting; time preferences; delay

JEL Classification Numbers: H41; D63; C92.

*Corresponding author: Kölle: University of Cologne, Albertus Magnus Platz, 50923 Cologne, Germany, Phone: +49 (0)221 470 5317, E-mail: *felix.koelle@uni-koeln.de*. We would like to thank Loukas Balafoutas, Stefania Bortolotti, Michalis Drouvelis, Guillaume Frechette, Simon Gächter, Axel Ockenfels, Theo Offerman, Bettina Rockenbach, Matthias Sutter, Roel van Veldhuizen, Jeroen van der Ven, Lukas Wenner, and participants from various seminars and conferences for helpful comments and suggestions. Financial support from the Center for Social and Economic Behavior (C-SEB) at the University of Cologne is gratefully acknowledged.

1 Introduction

Many economic decisions involve a time dimension forcing people to make intertemporal trade-offs between costs and benefits that occur at different points in time. Numerous studies, both theoretically and empirically, have investigated these trade-offs in the context of *individual decision making* in which choices only affect oneself, such as saving, human capital accumulation, dieting, or exercising (see Frederick et al., 2002, for an overview). The timing of when costs and benefits realize, however, is also important in *strategic decision situations* in which outcomes also depend on the choices of others. Yet, so far only very little is known about how intertemporal trade-offs unfold in such situations, and it is far from obvious whether people discount payoffs similarly across individual and strategic decisions as the latter are much more complex, and, hence, subjects might rely on very different cues or strategies. Furthermore, in strategic contexts other-regarding concerns may affect behavior, too, and these concerns might interact with individuals' time preferences in non-trivial ways. As such, it is important to study both of these dimensions in conjunction rather than in isolation as most previous literature did. Here, we provide a systematic and comprehensive analysis of how the timing of costs and benefits affects game-play in a strategic decision situation, and highlight the channels through which the factor time affects behavior. Our results reveal that delaying the consequences of actions affects behavior through multiple channels, and that discounting in strategic contexts may exceed the degree of discounting typically observed in individual decision contexts.

We study the question of discounting in strategic contexts in a collective action problem. Studying collective action problems is interesting because they are ubiquitous in nature and exist at all levels in human society, at the global scale and in many local communities and organizations. Furthermore, some of the most pressing challenges of mankind, such as limiting the warming of the planet or reducing the over-exploitation of natural resources, are important examples of collective action problems in which the costs and benefits of cooperation do not realize immediately and at the same time. In particular, in both of these examples there is a need of immediate action (reducing CO_2 emissions, refraining from over-exploitation), but the rewards from these costly efforst can only be claimed in the future.[1] Intuitively, because most people exhibit at least some degree of impatience (i.e., they discount future outcomes relative to immediate ones), incentives for cooperation are weakened when the benefits of collective action can only be reaped in the future while free-riding pays off

[1]In extreme cases these benefits might not even be claimed by oneself but only by subsequent generations (see Schelling (1995) for a discussion on intergenerational discounting and Hauser et al. (2014) for a recent experimental test along these lines).

immediately. On the contrary, if costs rather than the benefits of cooperation are delayed, this should ease cooperation as free riding becomes less attractive. Note, however, that in these cases people might find it optimal to change own behavior even in the absence of any discounting, as beliefs about others' cooperativeness and others' degree of impatience matter for behavior, too.

We investigate the effects of delayed costs and benefits on cooperation with the help of a laboratory experiment. Using an experimental approach has the major advantage of allowing us to tightly manipulate the timing of decision and the timing when consequences realize, something that is hard to achieve in the field. We base our analysis on a standard linear public goods game, an experimental version of a collective action problem. We employ a 2×2 between-subjects design, in which we vary when the costs and benefits of cooperation realize, either "immediately" (in 1 day) or delayed (in 12 months). Compared to the situation when both costs and benefits are realized immediately, we find that cooperation levels significantly drop by about 50% when benefits of cooperation are delayed, and that cooperation significantly increases by about 70% when costs of cooperation are delayed. When both costs and benefits are delayed to the same extent, in contrast, we find no significant differences in cooperation compared to the baseline without any delay, indicating that cooperation is time-consistent.

We identify three channels through which the timing of costs and benefits affects cooperative behavior. First, in line with our theoretical predictions (see Section 2), we find that cooperation differences are particularly pronounced for subjects who display a high degree of impatience. We follow previous literature and measure subjects' degree of impatience by asking them to make a series of binary decisions between smaller-sooner and larger-later rewards using a multiple price list design (compare Coller and Williams, 1999; Harrison et al., 2002; Andersen et al., 2008). Second, we find that subjects significantly shift their beliefs about other people's cooperativeness; they become more pessimistic (optimistic) about others' contributions if benefits (costs) of cooperation are delayed. Third, on top of beliefs, we also observe a significant shift in subjects' cooperation preferences, which, using a variant of the strategy method, we elicit by asking each participant to indicate how much he or she wants to contribute to the public good for each rounded average contribution level of other group members (compare e.g., Fischbacher et al., 2001; Fischbacher and Gächter, 2010; Gächter et al., 2017). In particular, as indicated by a significant level shift in the average conditional contribtion schedule, we find that people become less (more) willing to match others' contribution when the benefits (costs) of cooperation are delayed. At the same time, however, we observe that the general principle of reciprocity, i.e., the willingness to contribute more the more others contribute, is equally strong in all treatments as indicated by a similar slope of the conditional contribution schedule.

In a second study, we then quantify the amount of monetary incentives needed in order to close the cooperation gaps observed in our two asymmetric treatments. That is, in the case of delayed benefits we ask by how much economic incentives have to be increased in order to raise contributions up to the level without any delay, while in the case of delayed costs we study by how much economic incentives can be decreased such that contributions are similar to the level in the no delay case. This allows us to calculate a discount factor from our strategic context in which payoffs are interdependent, and compare it to the discount factor from our individual decision task. We find that in the case of delayed benefits, economic incentives (in terms of the marginal per capita return (MPCR)) need to be increased by 92% to raise contributions up to case without any delay. When costs rather than benefits are delayed, in contrast, economic incentives can be decreased by 52%. In both cases, the shift in economic incentives correspond to an implicit yearly discount rate of about 50%, which is much higher than the ones observed in individual choice tasks, where discount rates typically takes values of around 30% (compare e.g. Harrison et al. (2002); Dohmen et al. (2010)). To the best of our knowledge, this is the first paper that demonstrates that delaying the consequences of actions in strategic environments can lead to exacerbated discounting. At the individual level, we further find some moderate correlation between discounting across our individual and strategic decision context, indicating that other factors than ones own degree of impatience matter for discounting behavior in strategic contexts, too.

Our paper contributes to two so far largely unrelated strands of the literature. The first concerns the large literature on collective action problems (Olson, 1965; Hardin, 1968; Andreoni, 1990; Ledyard, 1995; Cornes and Sandler, 1996; Ostrom, 2000; Zelmer, 2003; Gächter and Herrmann, 2009; Chaudhuri, 2011). Yet, despite its obvious relevance for many real-world collective action problems (see above), the question of how the timing of costs and benefits affects cooperation in these type of situations has been largely ignored so far. The second strand of literature focuses on the question of how agents resolve intertemporal trade-offs between costs and benefits. This literature, however, has almost exclusively studied behavior in the context of individual decision making, in the absence of any externalities on others (see Frederick et al., 2002, for an overview).[2]

We are aware of only very few attempts that directly try to combine these two types of decision contexts. The paper that is most closely related to ours is Jacquet et al. (2013) who, using a repeated threshold public goods game, vary whether the benefits of cooperation

[2]A few studies have investigated discounting in social but non-strategic settings, such as dictator games (Kovarik, 2009; Dreber et al., 2016; Kölle and Wenner, 2018) and charitable giving (Breman, 2011; Andreoni and Serra-Garcia, 2017).

are realized the next day or in 7 weeks. They find that in the latter case cooperation drops by 23%. Similarly, Deck and Jahedi (2015a) find that delaying the benefits in a Prisoner's Dilemma Game by two weeks reduces cooperation rates from 63% to 45%.[3] We extend the insights from these studies in several dimensions. First, we not only consider the case of delayed benefits but also study the cases when either costs or both costs and benefits are delayed. This allows us to test whether people discount costs and benefits of cooperation to the same extent, or whether there are asymmetries between these two. Second, in contrast to both of these studies, we not only describe aggregate effects at the cooperation level, but also uncover the behavioral mechanisms behind this effect by investigating the interplay between subjects' degree of impatience, their beliefs about others' cooperativeness, and their preferences towards cooperation. As such, we provide novel insights into the micro-foundation of discounting in strategic situations. Finally, by determining how much economic incentives are needed to offset the cooperation gap, we can obtain a measure of discounting in a strategic context which we can link to discounting in individual, non-strategic contexts.[4]

We organize the rest of the paper as follows. In Section 2, we provide an overview of the general setup as well as derive theoretical predictions of how the timing of costs and benefits of cooperation affect contribution behavior. Section 3 describes the experimental design and procedures. Section 4 contains the empirical analysis. In Section 5, we describe the design and the results of our second study. Section 6 concludes.

2 The general setup & theoretical considerations

As a workhorse for studying cooperation, we use a standard linear public goods setting. Consider a group of n individuals. Each individual is given a budget of m tokens which she can either contribute to a public good, c_i, or use for her private consumption, x_i, subject to the budget constraint $x_i + c_i = m$. Using this simple setup, we experimentally vary when the payoffs from the private consumption and the public good are realized: either

[3] Apart from the Prisoner's Dilemma game, they further study the effects of delayed payments on Stag-Hunt game choices, as well as bidding behavior in a first and second price auction. Furthermore, Deck and Jahedi (2015b) study behavior in an intertemporal contest game, but find little evidence for any time discounting.

[4] Two further related papers are Fehr and Leibbrandt (2011) and Davis et al. (2016). Both studies separately measure individuals' time preferences and cooperation behavior in a static social dilemma game without any delayed payments. Consistent with the findings from our baseline treatment in which there is no delay, they find that a subject's degree of impatience and her degree of cooperativeness are largely unrelated. Interestingly, however, Fehr and Leibbrandt (2011) find that time preferences measured in the lab predict cooperation behavior outside the lab in a situation that entails an intertemporal component. In particular, they find that more patient fisherman use more sustainable fishing instruments that are less likely to exploit the collectively used fishing grounds.

"immediately" (in 1 day) or delayed (in 12 months) (see Section 3.4 below for details on how payments were implemented). The present value of individual i's income stream from her private consumption and the consumption of the public good can then be written as

$$\pi_i = I_x x_i + I_c \theta \cdot \sum_{j=1}^{n} c_j \tag{1}$$

where the parameter θ denotes the marginal per capita return (MPCR). If $\frac{1}{n} < \theta < 1$, we have a social dilemma situation in which there is a trade-off between individual and social payoff maximization. To capture the timing of when payoffs from the private and the public good are realized, I_x and I_c are indicator functions defined by

$$I_x = \begin{cases} 1 & \text{if } x_i \text{ is realized immediately} \\ \delta_i & \text{otherwise} \end{cases}$$

and

$$I_c = \begin{cases} 1 & \text{if } c_i \text{ is realized immediately} \\ \delta_i & \text{otherwise} \end{cases}$$

and where δ_i denotes individual i's discount factor. We assume $0 \leq \delta_i \leq 1$.

We study this basic decision situation in a 2×2 between-subjects design, leading to a total of four treatments (see Table 1 for an overview). In our baseline treatment, *NowNow*, we consider a situation in which both the benefits from the private and the public good are realized immediately ($I_x = I_c = 1$). This treatment serves as our first benchmark as it reflects the scenario studied by most previous literature. As a second benchmark, in *LaterLater* ($I_x = I_c = \delta_i$) we consider a situation in which the benefits from the private and the public good are both realized with a delay of 12 months. This resembles, for example, a situation in which one has to make a commitment about whether or not to contribute to a common good in the future, such as committing to volunteer in future local community events or to donate part of the end-of-year bonus to charity.

Our remaining two treatments entail an asymmetry with regard to at which points in time the payments from the private and the public good are realized. In particular, in *NowLater* the costs of cooperation, i.e., the forgone private consumption, are realized immediately, while the benefits of cooperation, i.e., the payments from the public good, are delayed ($I_x = 1, I_c = \delta_i$). In treatment *LaterNow*, in contrast, the benefits of cooperation are realized immediately, while the costs of cooperation (in form of foregone private consumption) have to be borne only in the future ($I_x = \delta_i, I_c = 1$). The former treatment resembles many real-world situations in which there is a need for immediate action, but where the fruits of these costly efforts can only be reaped in the future. Think, for example, about the problem of climate change where costly reductions of CO_2 emissions are needed today, while the results of these efforts will only realize in the long-run. Or, think of a team of workers or researchers which have to decide to collaborate or not in order to receive a joint end-of-year bonus or a joint future publication, respectively. Situations in which costs but not benefits are delayed, in contrast, occur for example when collaborative action yields an immediate benefit for the team and the time not dedicated to other individual projects have detrimental effects for, e.g., future promotions, or when governments provide public goods on credit that has to be repaid only in the future.

		Benefits from public good	
		"Immediate" *(in 1 day)*	Delayed *(in 12 months)*
Benefits from private consumption (costs of cooperation)	"Immediate" *(in 1 day)*	*NowNow*	*NowLater*
	Delayed *(in 12 months)*	*LaterNow*	*LaterLater*

Table 1: Experimental treatments.

In the following, we analyze how the timing of costs and benefits of cooperation affects predicted contribution levels to the public good. Under the standard assumption that individuals are fully rational and only interested in maximizing their own discounted monetary payoffs, inserting the budget constraint into (1) and taking the first derivative with respect to c_i yields the following first-order conditions:

$$\frac{\partial \pi_i}{\partial c_i} = \begin{cases} -1 + \theta & \text{if } x_i \text{ and } c_i \text{ are realized immediately} \\ -1 + \delta_i \theta & \text{if } x_i \text{ is realized immediately and } c_i \text{ delayed} \\ -\delta_i + \theta & \text{if } x_i \text{ is realized delayed and } c_i \text{ immediately} \\ -\delta_i + \delta_i \theta & \text{if } x_i \text{ and } c_i \text{ are realized delayed} \end{cases} \tag{2}$$

In each row, the first term denotes the (discounted) marginal costs of contributing to the public good, while the second term denotes the (discounted) marginal benefits of contributing to the public good. Given the range of values for θ and δ_i, it is straightforward to see that the dominant strategy for individual i is to free ride and contribute nothing to the public good ($c_i = 0$) in all cases, except when x_i is realized delayed, c_i is realized immediately, and δ_i is sufficiently low. In this case, if $\delta_i < \theta$, the benefits of cooperation outweigh the discounted costs, and full cooperation ($c_i = m$) becomes the dominant strategy. Hence, under the assumption of pure self-interest we do not expect any differences across the *NowNow*, *NowLater*, and *LaterLater* treatments, but contributions are expected to be higher for those individuals in *LaterNow* who are sufficiently impatient.

There is, however, plenty of evidence indicating that many people are not solely motivated by their own monetary payoffs, but that they also care about fairness and the well-being of others. For example, even when nobody is predicted to contribute anything, evidence from previous public good experiments suggest that there is some positive amount of voluntary cooperation, even in anonymous one-shot interactions (see e.g., Gächter and Herrmann, 2009; Chaudhuri, 2011, for an overview). A variety of models of other-regarding preferences (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Charness and Rabin, 2002; Dufwenberg and Kirchsteiger, 2004; Falk and Fischbacher, 2006) have been established that can explain behavior observed in the laboratory and in the field (see e.g., Sobel, 2005; Fehr and Schmidt, 2006, for an overview).

In the following, we analyze to what extent the predictions above change if we incorporate some form of social preferences. We use a modified version of the Fehr and Schmidt (1999) model that, in addition to full free riding, can capture two behavioral regularities observed in previous literature: (i) positive rates of cooperation, i.e., $c_i > 0$, and (ii) conditional cooperation, i.e., contributions that are increasing in (the beliefs about) others' contributions.[5] In particular, we allow for the possibility that players dislike inequality, i.e., they experience a loss in utility if they receive a (discounted) material payoff that is either higher or lower than the (discounted) average payoff of the other group members, $\bar{\pi}_j = \frac{1}{n-1} \sum_{j \neq i}^{n} \pi_j$. Note that different to the original specification by Fehr and Schmidt (1999), we use the average rather than individual payoffs of others as the reference point. We chose this modification for two reasons. First, in our experiment (see below) subjects were only provided with information about average but not individual behavior of others, which renders individual payoff comparisons difficult. Second, as shown by Hartig et al. (2015), even if individuals are provided with information about individual contribution behavior, the majority of people who do react positively to others' contributions are indeed mainly guided by the average contribution of others (rather than, e.g., by the highest or the lowest contribution). A player i's objective function can then be written as

[5]Note, however, that our predictions are not bound to the Fehr and Schmidt (1999) model but hold for a larger class of social preferences models such as, e.g., a model of altruism or warm glow of giving (Andreoni, 1989; Anderson et al., 1998).

$$U_i(\pi_i, \bar{\pi}_j) = \pi_i - \alpha_i \cdot max\{\bar{\pi}_j - \pi_i, 0\} - \beta_i \cdot max\{\pi_i - \bar{\pi}_j, 0\} \quad (3)$$

where α_i measures individual i's marginal psychological costs of disadvantageous inequality and β_i that of advantageous inequality. Following Fehr and Schmidt (1999), we assume $\alpha_i \geq \beta_i \geq 0$ and $\beta_i < 1$. Given the linearity of the utility function, depending on an individual's marginal disutility from inequality as well as her degree of impatience, she will either free-ride completely, exactly match others' average contribution $\bar{c}_j = \frac{1}{n-1} \sum_{j \neq i}^{n} c_j$, or contribute her whole endowment m. In particular, maximizing (3) with respect to c_i, the best response function yields:

$$c_i^* = \begin{cases} m & \text{if } I_c\theta > I_x(1 + \alpha_i) \\ \bar{c}_j & \text{if } I_x(1 - \beta_i) < I_c\theta \leq I_x(1 + \alpha_i) \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

The basic intuition for this result is as follows. For each token contributed to the public good, a player earns a (discounted) monetary payoff of $I_c\theta$. Every dollar used for private consumption, instead, yields a (discounted) monetary benefit of I_x minus a (discounted) non-pecuniary loss of at most $I_x\beta_i$ from increasing inequality (in case others contribute a positive amount to the public good). Therefore, if $I_x(1-\beta_i) > I_c\theta$, it is a dominant strategy for player i to contribute nothing. If, however, her marginal dis-utility from advantageous inequality becomes sufficiently strong such that $I_x(1 - \beta_i) \leq I_c\theta$, she finds it optimal to exactly match others' contributions in order to reduce inequality. In that case, as shown by Fehr and Schmidt (1999), multiple equilibria with positive contributions $c_i = \bar{c}_j \in [0, m]$ exist. In some cases, if $I_c\theta \geq I_x(1 + \alpha_i)$, a player even finds it optimal to contribute her full endowment irrespective of others' contributions. Note, however, that this condition can only be fulfilled if the benefits from the private consumption are delayed, the benefits from the public good are immediate, discounting is sufficiently high, and player i does not suffer too much from disadvantageous inequality (i.e., α_i is sufficiently small). In all other cases, player i at most contributes the amount (she believes) others contribute.

To better compare how the timing of the costs and benefits of cooperation affects predicted contribution levels, from (4) we can derive threshold values, $\tilde{\beta}_i$, the degree of advantageous inequality concerns has to exceed in order to make (conditional) cooperation feasible. The threshold values are given by

$$\begin{aligned} \tilde{\beta}_i^{NowNow} &= 1-\theta \\ \tilde{\beta}_i^{NowLater} &= 1-\delta\theta \\ \tilde{\beta}_i^{LaterNow} &= \tfrac{\delta-\theta}{\delta} \\ \tilde{\beta}_i^{LaterLater} &= 1-\theta \end{aligned}$$

From this it follows that $\tilde{\beta}_i^{NowLater} \geq \tilde{\beta}_i^{NowNow} = \tilde{\beta}_i^{LaterLater} \geq \tilde{\beta}_i^{LaterNow}$ as $1-\delta\theta \geq 1-\theta \geq \frac{\delta-\theta}{\delta}$. The intuition for this result is straightforward. First, compared to the benchmark case without any delay (*NowNow*), if players discount future payoffs, delaying the benefits from the public good (*NowLater*) makes cooperation more costly and, hence, a higher degree of inequality aversion is needed in order to make an individual willing to match others' contributions. Second, because delaying the benefits from the private good (*LaterNow*) makes cooperation relatively cheaper, in that case a lower degree of inequality aversion is sufficient to induce cooperation. Last, if both the benefits from the private and the public good are delayed, the degree of inequality aversion needed to sustain cooperation is the same as in the no delay case. The reason is that delaying all payments does not affect the marginal trade-off between cooperating or not because δ enters both sides of the equation and, hence, cancels out. Note that while these strict cut-off predictions, which are due to the linearity of the Fehr and Schmidt model (Fehr and Schmidt, 1999), certainly represent only a simplification of the expected behavior, we believe that they are nevertheless useful in providing an intuition of how the timing of the costs and benefits of cooperation affects the comparative statics of contributions.

Based on these considerations, and under the assumption that α_i, β_i, and δ_i are equally distributed across treatments, we expect $c_i^{LaterNow} > c_i^{NowNow} = c_i^{LaterLater} > c_i^{NowLater}$ (Hypothesis 1). Furthermore, since in *NowLater* (*LaterNow*) the threshold value $\tilde{\beta}_i$ has to exceed in order to make player i willing to conditionally cooperate is decreasing (increasing) in δ_i, we expect the differences in cooperation between *NowNow*, *NowLater*, and *LaterNow* to be larger the higher subjects' degree of impatience, i.e., the lower δ (Hypothesis 2).

3 The Experiment

Our experiment consists of three parts. In Part 1 and 2 of the experiment, subjects played two versions of a one-shot linear public goods game based on the decision situation described above. Subjects were matched into groups of $n = 4$, and each group member was given an endowment of $m = 20$ tokens. In both parts, group members then had to decide how many (if any) of their 20 tokens they want to contribute to a group project (Voluntary Contribution

Mechanism (VCM)). Tokens contributed were doubled and shared equally among all four group members, irrespective of their individual contribution decisions, i.e., $\theta = 0.5$. Hence, while each token contributed to the public good was worth 0.5 monetary units (MU) for each group member, each token not contributed was worth 1 MU to the individual. For Part 1 and 2 of the experiment, subjects were randomly assigned to one of the four experimental treatments described above, *NowNow*, *NowLater*, *LaterNow*, or *LaterLater* (compare Table 1). At the beginning of Part 3, which was identical across all treatments, groups were resolved. We then elicited a proxy for subjects' time preferences by asking them to make a series of decisions involving a trade-off between smaller-sooner and larger-later rewards. In the following, we describe each of the three parts in more detail.

3.1 Part 1: Elicitation of cooperation attitudes

Previous literature has highlighted the importance of both cooperation attitudes (a_i) and beliefs (b_i) for explaining cooperation outcomes (c_i) (Fischbacher and Gächter, 2010; Fischbacher et al., 2012; Gächter et al., 2017). The reason is that people who differ in their ex-ante attitudes towards cooperation can ex-post make the same cooperation decision. To see why, consider a 'conditional cooperator' who is willing to cooperate only if she believes others do so too. But there may also be 'free riders', who never want to contribute to the public good irrespective of their beliefs how much others contribute. A conditional cooperator who believes that others do not contribute and a person with a free rider attitude both contribute nothing: their ex-post behavior is observationally equivalent, despite different ex-ante attitudes. Thus, if cooperation is a function of attitudes and beliefs, the challenge is to separate them empirically. This is particularly interesting in our setting, as it allows us to investigate how the timing of costs and benefits affects the underlying mechanisms behind cooperation. That is, we can disentangle to what extent differences in contributions are due to differences in attitudes, differences in beliefs, or a combination of both factors.

Therefore, in Part 1 of our experiment, in each treatment we elicited subjects' attitudes towards cooperation using a variant of the strategy method (Selten, 1967; Fischbacher et al., 2001). This design allows eliciting an individual's willingness to cooperate as a function of the other group members' contributions. In particular, subjects are asked to make an unconditional and a conditional contribution to the public good. In the unconditional contribution, subjects are simply asked how much they want to contribute to the public good. In the conditional contribution, subjects are asked to fill a contribution table in which they have to specify a vector a_i of contributions as a function of all possible average contributions of the other group members. To guarantee incentive compatibility, in each group a random mechanism selects three members for whom the unconditional contribution is payoff-relevant, and one member for whom the conditional contribution is payoff-relevant. For this subject, the conditional contribution is calculated according to the (rounded) average unconditional contribution of the three other group members and her contribution table. To get a clean measure of an individual's attitude towards cooperation, the game is only played once which avoids any confounds due to repeated game incentives. Previous studies have shown that on average people are conditionally cooperative, but that there is pronounced heterogeneity in cooperation attitudes across individuals.

3.2 Part 2: Elicitation of beliefs and cooperation outcomes

In Part 2 of our experiment, subjects played a one-shot direct-response public goods game in which all group members had to simultaneously decide on their contributions, c_i. To avoid any spillover effects between parts, at the beginning of Part 2 groups were randomly re-matched and no information about Part 1 behavior was given until the very end of the experiment.[6]

After subjects decided on their contribution, they were asked for their beliefs, b_i, about the average contribution of the other three group members. Beliefs were elicited as "most likely intervals" using the Most Likely Interval elicitation rule (MLI) method proposed by Schlag and van der Weele (2015). In particular, subjects were asked—in an incentive-compatible way—to state an interval in which the others' average contribution was believed to belong. The interval was stated by specifying an upper bound U and a lower bound L. The monetary incentives were such that subjects were paid whenever the true realization lay in the stated interval, but payments were strictly decreasing in the width of the reported interval, i.e., subjects faced a trade-off between the chance of their guess being correct and receiving higher earnings in case the guess was correct. Specifically, payments were calculated according to the following function, where $\bar{c_j}$ denotes the realization of the others' average contribution[7]

$$\pi_i^{belief}(L, U, \bar{c_j}) = \begin{cases} 0.25 \times (20 - (U - L)) & \text{if} \quad \bar{c_j} \in [L, U] \\ 0 & \text{if} \quad \bar{c_j} \notin [L, U] \end{cases} \tag{5}$$

[6]For the re-matching we used a perfect stranger matching protocol, i.e., subjects in Part 2 of the experiment interacted only with people with whom they have not interacted before. We did not control for order-effects between Part 1 and Part 2 as Fischbacher et al. (2012) show that playing either the strategy-method or the direct-response game first has no effect on behavior.

[7]The parameter 0.25 was chosen to calibrate the earnings from the belief elicitation and the decisions in the public goods game. $20 - (U - L)$ ensures that if subjects specified the whole action space as their interval (and hence their guess would be always correct), earnings were 0. See Schlag and van der Weele (2015) for more details and a discussion of the theoretical properties of the MLI method.

The advantage of this method compared to eliciting beliefs as point estimates (as usually done in these type of games) is that it not only allows to collect beliefs about general tendencies, i.e., whether cooperation is believed to be high or low, but also to get a measure about the underlying precision/uncertainty of people's beliefs. Narrow intervals correspond to high precision (low uncertainty), while wide intervals are an expression of low precision (high uncertainty). This is interesting because it allows us to investigate whether changing the timing of costs and benefits in our public goods setting not only shifts beliefs about others' cooperation levels, but also affects the perceived strategic uncertainty about the other group members' behavior. This could be the case because adding the time dimension introduces an additional layer of complexity to the decision problem - contributions now not only depend on a person's cooperativeness but also on her degree of patience - which, in turn, could increase the uncertainty about others' behavior.

3.3 Part 3: Elicitation of time preferences

In Part 3 of the experiment, we elicited a proxy for individuals' time preferences. As shown in Section 2, we expect the size of our treatment effects to vary depending on subjects' degree of impatience - they should be large for very impatient subjects and small for very patient subjects - which is why they serve as an important control variable.

Eliciting time preferences has become an important component of experimental studies, yet there is still an ongoing debate about how to best measure discounting (see Andreoni et al. (2015) and Cheung (2016) for recent discussions of different methods). While some researchers rely on a multiple price list (MPL) design (cf. Coller and Williams, 1999; Harrison et al., 2002; Andersen et al., 2008), others have proposed a different method based on convex time budgets (CTB) (Andreoni and Sprenger, 2012a;b). Given that the focus of our paper is on understanding the effects of delay on cooperation, rather than precisely estimating individuals' time preference parameters, we decided to use a MPL design, which is slightly shorter and simpler than the CTB technique.[8] In our choice list, we ask subjects to choose between twenty smaller-sooner and larger-later rewards. All twenty choices were presented in a table similar to Table A1 (see Appendix A). The smaller-sooner reward was a fixed

[8]One possible disadvantage of the multiple price list method is that it implies inconsistent time preferences if subjects switch more than once between the smaller-sooner and larger-later rewards. To minimize the number of subjects switching back and forth between the two options, we provided a pop-up message in case of multiple switching points asking subjects whether they are sure that they want to switch again. This procedure seemed to work. We observe a very high rate of consistency with 98% of subjects exhibiting a unique switching point from the smaller-sooner to the larger-later reward. For the remaining four subjects, we use the midpoint of the first and the last row in which they switch to the larger-later option as the switching point. All results are robust to alternative specifications or exclusion of these subjects.

payment of 50 Euro, while the larger-later reward increased from 51.25 Euro in the first row to 75 Euro in the last row (in increments of 1.25 Euro). As in the case of our public goods game, sooner rewards were paid with a small delay of one day, while later rewards were paid with a delay of 12 months. Hence, the larger-later reward in row one gave an annual interest rate of 2.5 percent, while in the last row it yielded a rate of return of 50 percent.

From the row in which a subject switches from the smaller-sooner to the larger-later reward, we obtain an index of impatience. In particular, the more often an individual chooses the larger-later reward, i.e., the less money is necessary to induce her to wait 12 months, the more patient she is. Moreover, under the assumption of locally linear utility and annual compounding, an individual's switching point provides a bound on her discount rate. If, for example, a risk-neutral subject prefers the immediate payment for all amounts up to 60 Euro and then switches to the later option for all following choices, we can infer that her annual discount rate lies between 20% and 22.5%, or, alternatively, $0.8 \leq \delta \leq 0.83$.

3.4 Procedures

At the beginning of the experiment, subjects were informed about the three-part nature of the experiment. We did not, however, reveal any specific information about the details of each of the three parts. Instead, subjects were first introduced to the basic decision situation of the public goods game. They then had to successfully complete a comprehension test consisting of several questions about the comparative statics of the game as well as the timing of their outcomes (see online Appendix B for an English version of the instructions including screenshots of the experiment as well as a a copy of the control questions). After that, the instructions for Part 1 were distributed and read aloud. Once all subjects finished Part 1, we distributed and read aloud instructions for Part 2. At the end of Part 2, groups were resolved and subjects received instructions for Part 3. As before, no information about behavior of others was revealed. At the end of the experiment, subjects then received detailed information about the results from each part. Subjects' earnings were determined by the sum of their payoffs from Part 1 and Part 2. In Part 3, we randomly selected one subject for payment, and for this subject the earnings from Part 3 were added to the remaining earnings from Part 1 and 2. In addition, each subject received a show-up fee of €4. On average, subjects earned about €18.30, and sessions lasted about one hour.

We recruited a total of $n = 248$ (*NowNow*: $n = 64$, *NowLater*: $n = 64$, *LaterNow*: $n = 56$, *LaterLater*: $n = 64$) students from various disciplines using the online recruiting software ORSEE (Greiner, 2015) and computerized the experimental sessions using z-Tree (Fischbacher, 2007). In the recruitment email, it was explicitly stated that subjects would receive parts of their payment via bank transfer and that they therefore need to be willing to share their bank information with us.[9] Nothing, however, was mentioned about the possibility of delayed payments. This was done to avoid any effects due to self-selection into our experiment (apart from being willing to receive payments via bank transfer).

One major challenge when studying inter-temporal choices is to make all choices equivalent except for their timing. In particular, transaction costs associated with experimental payments have to be held constant across all time periods. We undertook several precautionary measures to equate transaction costs as much as possible. Most importantly, all sooner and later payments were made via bank transfer at the same day of the experiment.[10] That is, irrespective of the timing of the payment, all payments were transmitted to the bank at the same time by one of the authors. While sooner payments were initiated "immediately", later payments were arranged as *scheduled payments* in which we specified the exact date in the future when the actual transfer had to be carried out. Subjects were fully informed about all these procedures and that they would receive an email with a confirmation of both transfers within two days after the experiment. Subjects were further told that bank transfers typically take at least one working day. These procedures not only ensure similarity in physical transaction costs, but also makes sooner and later payments equally credible. In addition, at the end of the experiment all subjects received a written "payment agreement form" which listed all payments at every point in time a subject would receive. It also entailed a subject's bank information and email-address as well as the full contact information and the signature of one of the authors.[11] Hence, subjects had the possibility to contact us at any time in case of any problem with regard to their payments.

By introducing a small front-end delay of 1 day, we further aimed at limiting the scope for immediacy effects and present bias. Indeed, recent evidence suggests that present bias only occurs when payments are truly immediate, and that adding a small delay to immediate payments is sufficient to prevent present bias to occur (see Andreoni and Sprenger (2012a); Balakrishnan et al. (2017)). Hence, we believe that differences in behavior across treatments are mainly driven by standard exponential discounting rather than quasi-hyperbolic

[9]It was also made clear, however, that these information will solely be used for transaction purposes and that after all payments have been made, this data will be deleted. All subjects that showed up to any of our sessions were fine with these procedures.

[10]The show-up fee of €4 was paid in cash at the end of the experiment. This was required by the internal rules of the laboratory.

[11]See online Appendix C for a copy of the payment agreement form. Note that compared to standard laboratory experiments, one might argue that our payment procedures are somewhat less anonymous. However, it was made very clear to the subjects that at no point in time we would link their personal information to their behavior. Furthermore, given that the procedure was exactly the same across all treatments, it should not affect any of our treatment comparisons.

discounting (Laibson, 1997; O'Donoghue and Rabin, 1999; Frederick et al., 2002), although we cannot rule out the effect of present bias completely. In any case, since the payment procedures are exactly the same across all treatments, this should not affect our treatment comparisons. For the later payment we chose a delay of 12 months (*i*) to make the delay meaningful as in many of the real-world examples discussed above, and (*ii*) to avoid any seasonal effects that could arise if subjects have preferences for receiving money in a certain month, e.g., before Christmas or holidays.

4 Results

We start our analysis by comparing cooperation levels in the one-shot direct-response game in Part 2 of our experiment, which gives us a clean measure of subjects' cooperativeness in the absence of any strategic incentives. We then use the subjects' responses from the remaining parts of the experiment to investigate to what extent the observed differences can be explained by individuals' degree of impatience (Part 3), their attitudes towards cooperation (Part 1), and their beliefs about others' cooperativeness (Part 2).

4.1 The effects of delayed costs and benefits on cooperation

Figure 1 and Table 2 display the cooperation levels for each of our four treatments. In our baseline treatment, *NowNow*, average contributions amount to 9.6, which is very consistent with cooperation rates observed in previous studies that typically lie between 40% and 50% (Zelmer, 2003).[12] We observe very similar cooperation rates in *LaterLater*, where both the costs and the benefits of cooperation are delayed; the mean contributions (amounting to 10.2) as well as the overall distribution of contributions are both statistically indistinguishable from the amounts observed in *NowNow* (t-test, $p = 0.693$; Mann-Whitney U test, $p = 0.801$).[13] In the two asymmetric treatments, in contrast, contributions change substantially. When the benefits of cooperation are delayed but the costs of cooperation (in form of foregone private consumption) are immediate (*NowLater*), we observe a significant and substantial decrease in contributions to only 4.6, 52% less than in *NowNow* (t-test, $p < 0.001$; Mann-Whitney U test, $p < 0.001$). In the case in which costs delayed but benefits are immediate (*LaterNow*), in contrast, contributions significantly increase by 71% to 16.4 (t-test, $p < 0.001$; Mann-Whitney U test, $p < 0.001$). While the change in contributions

[12] This is reassuring as it gives us further confidence that our recruitment and payment procedures did not introduce any systematic bias or selection effect with regard to general cooperativeness.

[13] If not stated otherwise, all reported p-values correspond to two-sided tests.

	Mean (sd) contributions	Full free-riding ($c_i = 0$)	Full cooperation ($c_i = 20$)
NowNow	9.59 (7.62)	21.88%	25.00%
NowLater	4.61 (5.44)	45.31%	4.69%
LaterNow	16.45 (6.56)	8.93%	71.43%
LaterLater	10.16 (8.45)	31.25%	32.26%
Pairwise comparisons			
NowNow vs. NowLater	$p < 0.001$	$p = 0.005$	$p = 0.001$
NowNow vs. LaterNow	$p < 0.001$	$p = 0.053$	$p < 0.001$
NowNow vs. LaterLater	$p = 0.693$	$p = 0.230$	$p = 0.330$

Notes: Column 1 displays mean contributions in each of the four treatments with the standard deviation in parentheses. p-values according to t-tests (column 1) or χ^2-tests (columns 2 and 3).

Table 2: Contributions in the one-shot game in Part 2.

in the latter seems (in absolute terms) somewhat larger than in *NowLater*, the difference between the two is not statistically significant (F-test, $p = 0.398$).[14]

To provide some more detail, Table 2 also reports the fraction of subjects contributing either nothing or their whole endowment to the public good. It reveals a clear shift in the distribution of contributions across treatments. Compared to *NowNow*, the fraction of zero (full) contributors significantly increases (decreases) when the benefits of cooperation are delayed. The reversed picture is observed in the case of delayed costs. If, instead, both costs and benefits are delayed, the fraction of people contributing fully or nothing is very similar compared to our baseline treatment. We summarize these findings in our first result:

> **Result 1:** *Compared to the situation in which both costs and benefits of co-operation are realized immediately, delaying the benefits (costs) of cooperation significantly decreases (increases) contribution levels by 52% (71%). Delaying both costs and benefits, in contrast, has no significant impact on the level of co-operation.*

Our aggregate results are very consistent with our theoretical predictions derived in Section 2 (Hypothesis 1). The theory, however, also predicts heterogeneous effects across agents depending on their degree of time discounting, δ_i (Hypothesis 2). In particular, it predicts that our treatment differences should be largest for those individuals who discount

[14]Very similar results are obtained when analyzing the unconditional contributions from the attitude elicitation task in Part 1 of the experiment.

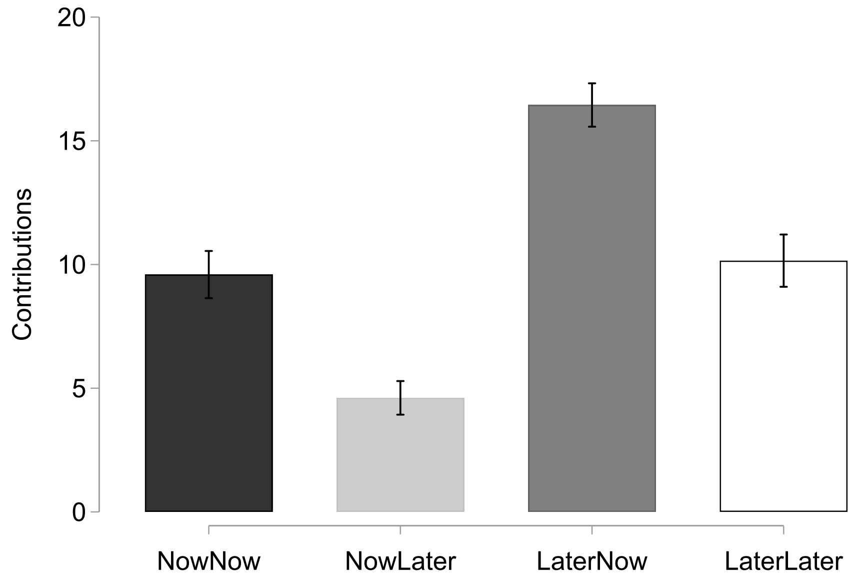

Figure 1: Average contributions in the one-shot direct-response game in Part 2 ($\pm$ 1 s.e.m.).

the future very strongly, i.e., subjects with a very small δ_i, whereas these differences are expected to be small (close to zero) for individuals who discount future payments only very little, i.e., who have a $\delta_i \approx 1$.

Recall that in Part 3 of our experiment, each individual had to choose between twenty smaller-sooner and larger-later rewards (compare Table A1 in Appendix A). In the following, we use the number of late choices as a proxy for an individual's time preference; the more often she chooses the larger-later reward (i.e., the less money she demands as a compensation for delaying the €50 payment by one year), the more patient she is.[15] Importantly, we find no significant differences in our time preference measure across the four treatments (Kruskal-Wallis test, $p = 0.301$). We take this as evidence that the randomization into treatments worked, and that having experienced different versions of the public goods game in Part 1 and 2 of the experiment had no effect on behavior in Part 3 (see Figure A1 in Appendix A for a distribution of choices in our time preference elicitation task).

To examine the predictions of Hypothesis 2, we analyze whether individual time preferences are predictive for cooperation behavior within our four treatments. As expected, when the costs and benefits of cooperation are realized at the same time, we find no signifi-

[15] As discussed in Section 3.3., we can also use the row in which a subject switches from the smaller-sooner to the larger-later reward to calculate bounds on her discount rate. The median discount rate falls in the range of 30 to 32.5 percent, which is very much in line with findings from previous studies using a similar design. For example, in a representative sample of 1000 German adults, Dohmen et al. (2010) find a median annual discount rate between 27.5 and 30 percent. Similarly, Harrison et al. (2002) report an average annual discount rate of 28.1 percent.

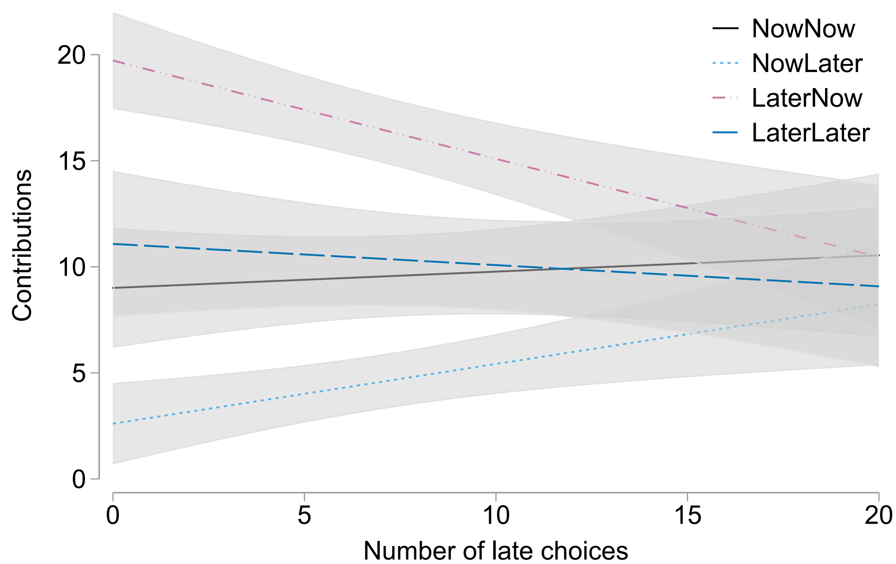

Figure 2: Fitted lines from linear regressions including a 95% confidence interval.

cant correlation between an individual's degree of impatience and her contribution behavior, neither in *NowNow* (Spearman's rank correlation: $\rho = 0.04, p = 0.746$) nor in *LaterLater* ($\rho = -0.11, p = 0.400$). This indicates that there is no direct link between cooperativeness and (im)patience per se, which is in line with previous evidence (Fehr and Leibbrandt, 2011; Davis et al., 2016). Time preferences matter, however, in our two asymmetric treatments. where either costs or benefits of cooperation are delayed. In particular, when benefits of cooperation are delayed but costs are immediate (*NowLater*) we find a significant positive relationship between subjects' degree of patience and their level of contributions ($\rho = 0.39, p = 0.002$). That is, the more patient a subject the less she suffers from delaying the benefits and, thus, the lower her incentives to decrease her contributions. The opposite effect is observed when costs rather than benefits are delayed (*LaterNow*). Here we find that the more patient subjects are, the less they contribute ($\rho = -0.55, p < 0.001$).

These results are illustrated by Figure 2 depicting fitted lines (including 95% confidence intervals) from OLS regressions using contribution decisions as the dependent variable and the number of late choices as independent variable (see Table A2 in Appendix A for the regression results).[16] The figure shows that, in line with Hypothesis 2, the treatment differences in contributions between *NowLater* (*LaterNow*) and *NowNow* are particularly large among the most impatient subjects. To test Hypothesis 2 more rigorously, within each treat-

[16]As a robustness check, we also run Tobit regressions to account for the censoring of the data. Since the results are very similar, for the ease of interpretation, we only report the OLS results.

ment we divide our sample into three equally large subgroups depending on their degree of (im)patience. In particular we classify a subject as a low, medium, or high patience type if the number of her larger-later choices lies in the bottom, intermediate, or upper tercile of the overall distribution. Redoing our analysis from above reveals that cooperation differences between our two asymmetric treatment and *NowNow* are large and highly significant for subjects classified as having low patience (*NowLater*: -6.0 tokens, *LaterNow*: +11.0 tokens, t-tests: both $p < 0.003$, Mann-Whitney U tests, both $p < 0.005$), but that these differences become much smaller and insignificant for subjects classified as having high patience (*NowLater*: -3.3 tokens, *LaterNow*: +1.7 tokens; t-tests: both $p > 0.164$, Mann-Whitney U tests: both $p > 0.204$). We summarize these findings in our second result:

> **Result 2a:** *When the costs and benefits of cooperation are realized at the same time (irrespective of whether delayed or not), an individual's degree of impatience is uncorrelated with her contribution behavior. When the costs and benefits of cooperation occur at different points in time, in contrast, cooperation systematically varies with elicited time preferences.*

> **Result 2b:** *The cooperation gap between our two asymmetric treatment (NowLater and LaterNow) and our baseline treatment without any delay (NowNow) is largest for subjects who are very impatient, and smallest for subjects who are very patient.*

4.2 Explaining cooperation differences

Our results so far have revealed substantial differences in cooperation rates depending on when the costs and benefits of cooperation are realized (Result 1), and that the magnitude of these differences systematically varies according to subjects' degree of impatience (Result 2b). Our next step is to further understand these results by determining some of the underlying behavioral mechanisms. To this end, we investigate to what extent the observed differences are driven by a shift in beliefs about others' cooperativeness, a shift in attitudes towards cooperation, or a combination of both.

We start by investigating how the timing of costs and benefits affects beliefs about others' contributions. The results of our belief elicitation are summarized in Figure 3. It shows for each treatment the average belief interval - determined by the upper and lower bound - as well as the average midpoint of the interval as indicated by the horizontal line. The figure reveals several interesting observations. First, we observe that the location of the beliefs follows the ordering of contributions (compare Figure 1). While beliefs in *NowNow* and *LaterLater* are very similar, subjects in *NowLater* are more pessimistic about others' contributions and in *LaterNow* they are more optimistic about others' contributions. This visual impression is corroborated by non-parametric tests. The null hypothesis of equality of the distribution of beliefs is clearly rejected using the lower bound, the upper bound, or the midpoint of the belief interval (Kruskal-Wallis tests, all $p < 0.001$). Pairwise treatment comparisons reveal highly significant differences across all comparisons (t-tests, all $p < 0.005$; Mann-Whitney U test, all $p < 0.001$), except for comparisons between *NowNow* and *LaterLater*, which are always insignificant (all $p > 0.302$).

Second, as in the case of contributions, we find these effects to be attenuated or amplified depending on subjects' time preferences. That is, while in our two symmetric treatments (*NowNow* and *LaterLater*) we find no systematic relationship between a subject's time preferences and her beliefs about others' contributions - in both cases the Spearman's rank correlation coefficient is close to zero and not statistically significant (both $p > 0.189$) - in our two asymmetric treatments (when either the costs or the benefits of cooperation are delayed) we observe a significant relationship between an individual's degree of impatience and her beliefs about others' cooperativeness. More specifically, in *NowLater* we find that more patient subjects are more optimistic about others' contributions (Spearman's rank correlation: $\rho = 0.31, p = 0.013$). In *LaterNow*, we observe the opposite; the more patient a subjects, the lower her beliefs about others' contributions ($\rho = -0.32, p = 0.017$). One potential explanation for these findings might be the so-called *false consensus effect* (Ross et al., 1977) according to which people tend to overestimate the extent to which others think the same way as they do. In particular, if subjects expect others to be similarly (im)patient as themselves, given the significant relationship between time preferences and contribution behavior (see Result 2a), this can explain why beliefs about others' contributions and impatience are correlated, too.

Third, the width of the belief interval, which we take as a proxy for subjects' uncertainty about others' contribution behavior, is very similar across treatments (Kruskal-Wallis test, $p = 0.277$). Hence, we find little evidence that adding an additional layer of complexity to the decision problem by delaying the costs and/or the benefits of cooperation significantly affects strategic uncertainty.

Finally, we also find the link between beliefs and contributions (displayed by the diamonds) to be systematically different across treatments. While in both *NowNow* and *LaterLater* average contributions closely match the average midpoint of the belief interval (the ratio between contributions and beliefs amounts to 0.95 and 0.94, respectively; t-test, $p = 0.554$; Mann-Whitney U test, $p = 0.494$), in *NowLater* and *LaterNow* contributions sys-

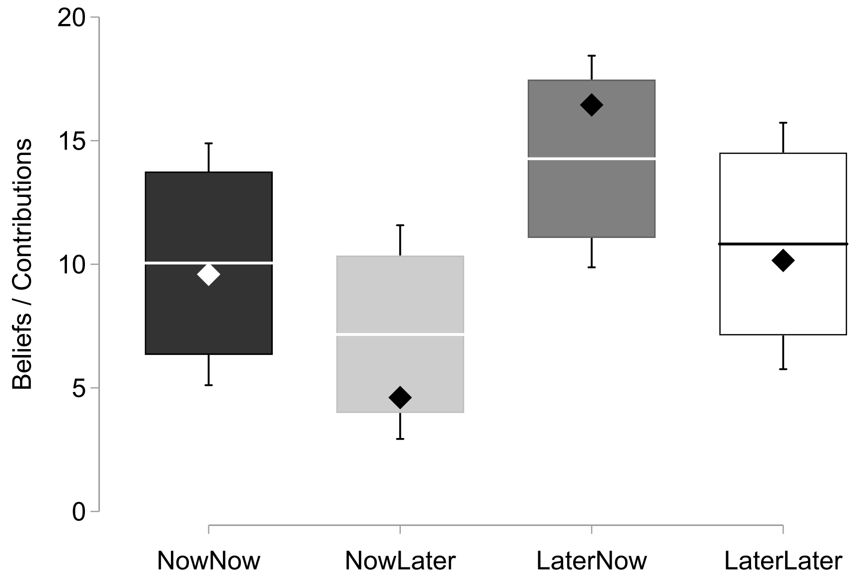

Figure 3: Belief intervals (± 1 s.e.m.) and contributions (diamonds) in Part 2.

tematically deviate from that midpoint. Instead, in *NowLater* we observe contributions close to the lower bound of the belief interval, yielding a contribution-belief ratio of 0.64, which is significantly lower than in *NowNow* (t-test, $p < 0.001$; Mann-Whitney U test, $p < 0.001$). In *LaterNow*, in contrast, we observe contributions to be closer to the upper bound of the belief interval. Here, subjects contribute on average 115% of what they belief others will contribute, significantly more than in *NowNow* (t-test, $p < 0.004$; Mann-Whitney U test, $p = 0.004$).[17]

One potential explanation for these differences could be that introducing an asymmetry to the timing when payments realize makes it harder for subjects to predict others' behavior. That is, while in the two symmetric treatments subjects have well-calibrated beliefs about others' behavior, in the two asymmetric treatments, despite anticipating the direction of the shift in contributions, subjects underestimate the magnitude of this effect. An alternative explanation for these effects is that delaying the costs or benefits of cooperation has a direct effect on subjects' cooperation attitudes, i.e., their willingness to reciprocate others' contributions. While both of these explanations may operate at the same time and based on this

[17]At the individual level, we find that the fraction of subjects contributing less than what they believe others will contribute significantly differs across treatments. In both *NowNow* and *LaterLater*, we find that 45% of the subjects contribute less than what they believe others will contribute, while 55% contribute more. When only benefits are delayed (*NowLater*), the fraction of people undercutting their beliefs significantly increases to 77% ($\chi^2 - test, p < 0.001$). The opposite picture is observed when only costs are delayed (*LaterNow*). In this case, 80% of the subjects contribute more than what they believe others will contribute, and only 20% undercut. This shift is highly significant ($\chi^2 - test, p < 0.001$).

data alone we cannot rule out one or the, we believe that the second explanation has at least some bite. The are two reasons for why we believe this to be the case. First, we find no pronounced treatment differences with regard to the accuracy of beliefs as measured by the number of cases in which the actual average contributions of the other group members lay within the stated belief interval. Overall, beliefs are correct in 52% of the cases, with no significant differences across treatments (*NowNow*: 50%, *NowLater*: 53%, *LaterNow*: 61%, and *LaterLater*: 47%; $\chi^2(3) = 2.50$, $p = 0.476$). Second, if the mismatch between contributions and beliefs were solely driven by inaccurate beliefs, this effect should be uncorrelated with individuals' time preferences. Instead, in *NowLater* we find a significant positive correlation between a subject's degree of impatience and the ratio between her contributions and her beliefs about others' cooperativeness (Spearman's rank correlation: $\rho = 0.41, p = 0.001$), while in *LaterNow* this relationship is significantly negative ($\rho = -0.33, p = 0.014$).[18] Together with the existing pervasive evidence on (imperfect) conditional cooperation (Keser and Van Winden, 2000; Fischbacher et al., 2001; Frey and Meier, 2004; Fischbacher and Gächter, 2010; Chaudhuri, 2011; Gächter et al., 2017), these results indicate that delaying the costs or benefits of cooperation might have a direct effect on subject's degree of self-serving bias when deciding about their willingness to match others' contributions. Before testing this conjecture more rigorously (see below), we summarize these findings in our third result:

> **Result 3:** *Delaying the costs (benefits) of cooperation significantly increases (decreases) subjects' beliefs about others' cooperativeness. This shift in beliefs can only partly explain the shifts in contributions.*

To test for the effect of timing on people's willingness to conditionally cooperate more directly, we analyze the data from Part 1 of our experiment, in which we elicited subjects' attitudes towards cooperation using the strategy method. The advantage of this method is that by allowing subjects to condition their own contributions on the other group members' contributions, the decision becomes (from a game-theoretic perspective) sequential and, hence, does not exhibit any strategic uncertainty. This allows us to investigate the role of "cooperation preferences" without the confounding factor of endogenous beliefs.

The results from the strategy method are summarized in Figure 4. Panel A depicts the average conditional contribution schedule across the four treatments. The results reveal that the slopes of the contribution schedule are very similar across treatments, but that there is a vertical shift in the contribution schedules across treatments. To quantify these effects,

[18] Again, no such effects are observed in the other two treatments (*NowNow*: $\rho = 0.01, p = 0.928$; *LaterLater*: $\rho = -0.02, p = 0.852$).

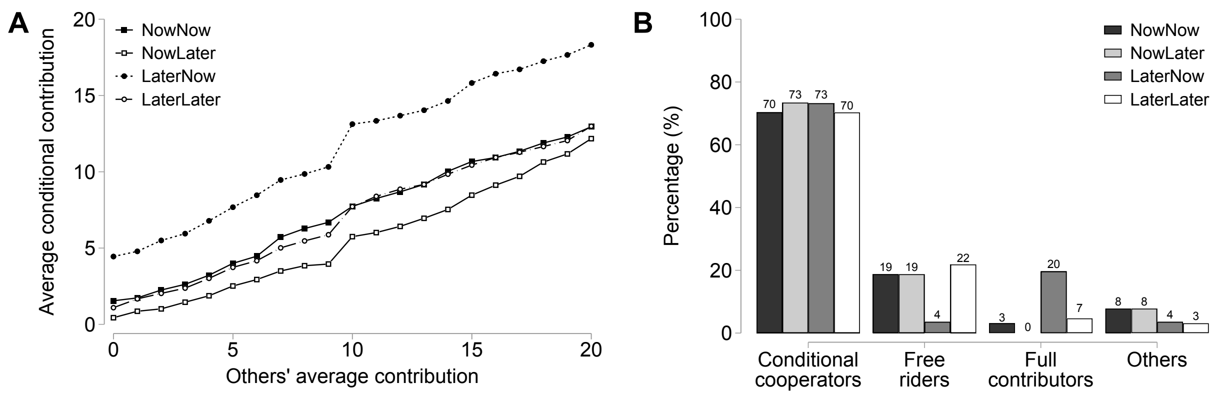

Figure 4: Panel A: Average conditional contribution schedule in Part 1 of the experiment. Panel B: Distribution of cooperation types.

we use regression analysis in which we regress subjects' conditional contribution decisions on the average contribution of the other group members, treatment dummies, as well as interaction terms between the treatment dummies and the average contributions of others. The results from this regression are reported in Table A3 in Appendix A. They confirm that there are no significant differences across treatments with regard to the slopes of the average schedule (as indicated by the insignificant interaction terms), but that there are significant shifts in contribution levels. In particular, subjects in *NowLater* contribute on average 1.6 tokens less than subjects in *NowNow*, while subjects in *LaterNow* contribute on average 3 tokens more. Both of these effects are statistically significant ($p = 0.015$ and $p = 0.020$, respectively). In *LaterLater*, in contrast, the treatment dummy is close to zero and statistically insignificant ($p = 0.666$), confirming that when both costs and benefits of cooperation are delayed, behavior is very similar compared to the case without any delay.

Taken together, these results indicate that while on the one hand the principle of reciprocity, i.e., the general willingness to contribute more the more others contribute, is similarly strong across treatments, on the other hand the timing of costs and benefits significantly affects subjects' degree of self-serving bias; they are less (more) willing to match others' contributions when benefits (costs) of cooperation are delayed.

Furthermore, in line with our findings from above, we find that our time preference measure significantly affects average conditional contributions in the asymmetric but not in the symmetric treatments. In particular, the more patient an individual is the higher her average conditional contributions in *NowLater* (Spearman's rank correlation, $\rho = 0.30$, $p = 0.017$) and the lower her average conditional contributions in *LaterNow* ($\rho = -0.53$, $p < 0.001$). In *NowNow* and *LaterLater*, the correlation coefficient is close to zero and not statistically significant ($\rho = 0.03$, $p = 0.841$; and $\rho = 0.05$, $p = 0.692$, respectively).

The strategy method further allows us to investigate the heterogeneity in individuals' cooperation attitudes that underlies Figure 4A, and to test to what extent the observed treatment differences are driven by a shift in the distribution of cooperation attitudes. We follow previous literature (Fischbacher et al., 2001; Fischbacher and Gächter, 2010; Fischbacher et al., 2012; Gächter et al., 2017) and classify subjects into cooperation types depending on their contribution schedule. We classify a subject as *'conditional cooperator'* if the entries in her contribution vector are increasing in others' average contributions, as a *'free rider'* if she never contributes anything, or as a *'full contributor'* if she contributes maximally irrespective of others' contributions.[19] We refer to the remaining subjects as *'others'*.

The results of this classification are shown in Panel B of Figure 4. The findings reveal a remarkably similar distribution of types across treatments. In particular, the fraction of subjects classified as *'conditional cooperators'* is almost identical in all four treatments, varying between 70% and 73% ($\chi^2(3) = 0.28$, $p = 0.964$). Also the fraction of *'others'* is very similar across treatments, varying between 3% and 8% ($\chi^2(3) = 2.34$, $p = 0.504$). The only notable difference we observe is a significant shift in the fraction of *'free riders'* ($\chi^2(3) = 8.74$, $p = 0.033$) and *'full contributors'* ($\chi^2(3) = 22.06$, $p < 0.001$). A closer look reveals that this difference is entirely driven by the treatment *LaterNow*, in which compared to all other treatments we observe a significantly lower fraction of *'free riders'* (χ^2-tests, all pairwise comparisons $p < 0.011$) and a significantly higher share of *'full contributors'* (χ^2-tests, all pairwise comparisons $p < 0.012$). When comparing the remaining three treatments, *NowNow*, *NowLater*, and *LaterLater*, in contrast, we observe no difference in the fraction of *'free riders'* ($\chi^2(2) = 0.26$, $p = 0.877$) and *'full contributors'*($\chi^2(2) = 2.87$, $p = 0.238$).

These results highlight that the downward shift in the average contribution schedule in *NowLater* (compare Figure 4A) is not driven by a shift in the distribution of types, but rather by an adjustment of contribution behavior among conditional cooperators; on average they contribute 1.6 tokens less than in *NowNow*, a difference that is highly significant (7.2 vs. 8.8; t-test, $p = 0.017$; Mann-Whitney U test, $p = 0.009$).[20] The upward shift in the average contribution schedule in *LaterNow*, in contrast, is to a large extent driven by a shift from 'free riders' towards 'full contributors', although 'conditional cooperators' adjust their behavior, too. Relative to *NowNow*, they increase their average conditional contributions by 1.2 tokes to 10.0, but this increase is only marginally significant (t-test, $p = 0.094$; Mann-Whitney U

[19]Following Fischbacher et al. (2001), we classify a subject as a 'conditional cooperator' either if her contribution schedule is weakly monotonically increasing in the average contributions of others, or if the Spearman correlation between own and others' contributions is positive and significant at the 1% level.

[20]The effect of 'others' is negligible as there are only very few ($n = 14$) subjects within this category across all treatments.

test, $p = 0.095$). We summarize our findings as follows:

> **Result 4:** *Delaying the benefits (costs) of cooperation causes a downward (upward) shift in conditional contributions, while leaving the slope of the contribution schedule unchanged. In NowLater this can be explained by a lower willingness of conditional cooperators to match others' contributions, while in LaterNow this can be explained by a shift in the fraction of 'free riders' and 'full contributors'.*

What can explain this pronounced shift in cooperation attitudes in *LaterNow*? Recall that when the costs of cooperation are delayed, if subjects discount future payments very strongly, contributing to the full extent might become individually optimal even when subjects only care about their own monetary payoff (compare Section 2). Our results are in line with this prediction. That is, we find that subjects classified as *'full contributors'* in *LaterNow* are significantly more impatient than any other type in the same treatment; in our time preference elicitation task they choose the larger-later reward on average in only 2.1 out of the 20 cases, compared to 8.3 we observe for all the remaining subjects (t-test, $p = 0.005$; Mann-Whitney U test, $p = 0.006$). No such difference is observed in the other treatments.[21] Hence, it seems to be indeed the case that subjects in *LaterNow* who contribute maximally irrespective of the contributions of others do so because of their strong degree of impatience rather than, e.g., their concern for social efficiency.

Note that while the model presented in Section 2 can well capture the shift in cooperation types from free-riders to full contributors in *LaterNow*, the effects observed in *NowLater* (no change in the distribution of types but a decreased willingness to match others' contributions among conditional cooperators) seem inconsistent with that model. The reason is that due to its linearity, the model is inconsistent with any type of imperfect conditional cooperation; it can only predict shifts between full free-riding, perfect conditional cooperation, and unconditional full contribution. The effects observed in *NowLater*, however, are reconcilable with a more general version of that model, which allows payoff inequalities to enter utility in a non-linear way. In this case, instead of having sharp cut-off values for the degree of advantageous inequity aversion, different levels of β would determine different "degrees of conditional cooperation", which in the case of delayed benefits (as in *NowLater*) might then be attenuated or accelerated depending on subjects' degree of impatience.

In summary, our results so far have shown that varying the timing of when costs and benefits of cooperation are realized has a systematic effect on the level of public goods provision

[21] Note, however, that the number of *'full contributors'* is very small in these treatments ($n = 5$ in all other treatments combined). Yet, if anything *'full contributors'* in these treatments seem to be more patient than all other types (average number of late choices 13.0 vs. 7.9, respectively).

as well as on its underlying behavioral mechanisms (beliefs and cooperation attitudes). While the direction of the observed effects are well in line with a standard model of discounting in combination with some other-regarding concerns that can explain positive contributions to the public good, an open question that remains is whether also the magnitude of these effects are in line with what one would expect given a common degree of time discounting. In the following, we report the results from a second experiment that was specifically designed to address this question.

5 What economic incentives are needed to close the cooperation gap?

Our second study aims at quantifying the degree of discounting observed in our strategic public goods setting. Specifically, in the case of delayed benefits we ask by how much economic incentives have to be increased in order to raise contributions up to the level without any delay. Similarly, in the case of delayed costs we study by how much economic incentives can be decreased such that contributions are at the level of the no delay case. This allows us to calculate discount factors in a strategic context in which choices also affect others, and to compare it to the degree of discounting typically observed in standard individual time preferences elicitation tasks.

The basic design of our second study is similar to the one in Study 1 (compare Section 3). A new set of $n = 188$ subjects that had not participated in our first study were randomly assigned to one out of two treatments, *NowLater* ($n = 92$) or *LaterNow* ($n = 96$), using a between-subjects design. Both treatments consist of three parts. In the first two parts, subjects played two versions of a one-shot linear public goods game. As before, the game was played in groups of four, groups were randomly re-shuffled between parts using a perfect stranger-matching protocol, and no information about behavior of others was revealed until the very end of the experiment. In Part 1, subjects either played the game with delayed benefits (*NowLater*) or with delayed costs (*LaterNow*). In both cases, subjects had to state their contributions for varying levels of the marginal per capita returns, θ, using the strategy method. We had a total of 13 different MPCR levels, ranging from 0.3 to 0.9 (increasing in increments of 0.05), presented to subjects in a simple table on one decision screen (see online Appendix D for a screenshot and a copy of the experimental instructions). To incentivize all choices, at the end of the experiment, for each group one MPCR was randomly selected to determine subjects' contributions and payoffs. After all subjects made their contribution decisions, we also elicited subjects' beliefs about the other group members' contributions for

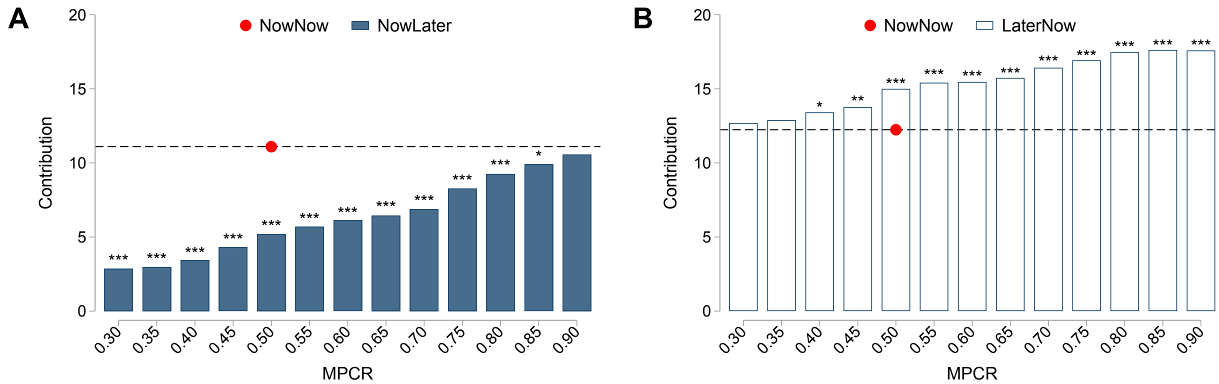

Figure 5: Contributions as a function of the MPCR in *NowLater* (Panel A) and *LaterNow* (Panel B). ***, **, and * indicate significance at the 1%, 5%, and 10% level using non-parametric Wilcoxon Signrank tests.

each level of the MPCR. As in Study 1, we elicited beliefs using the interval method. In Part 2, subjects in both treatments then played a one-shot public goods game without any delay and a fixed MPCR of 0.5 (as in our *NowNow* treatment in Study 1). Finally, in Part 3 we elicited a proxy for subjects' time preferences in an individual decision context using the same multiple price list design as in Study 1. All other procedures including the recruitment and payment of subjects were identical to the ones in our first study. As before, subjects had to correctly answer a set of control questions before the start of the experiment.

Our main results are summarized in Figure 5. For both treatments, it shows the average contributions in Part 1 of the experiment for each level of the MPCR. Panel A shows the results from treatment *NowLater*, and Panel B shows the same data for treatment *LaterNow*. As a comparison, in both cases we also display the level of cooperation from the second part of the experiment in which there was no delay and a fixed MPCR of 0.5 (*NowNow*), as indicated by the dot and the dashed line.

Figure 5 reveals several interesting observations. First, it replicates (in a within-subjects design) the main finding from our first study of significantly lower (higher) contributions at a MPCR-level of 0.5 in *NowLater* (*LaterNow*) relative to *NowNow* (Paired t-test, both $p < 0.001$; Signrank tests, both $p < 0.011$). Second, it shows that in both treatments contributions are monotonically increasing in the MPCR, indicating that subjects positively react to economic incentives in the expected way; the higher the MPCR (i.e., the lower the costs of cooperation) the higher the contributions.[22] Third and most importantly, it shows for

[22]At the individual level, in *NowLater* we find that 66% of all subjects display a (weakly) monotonically increasing contribution pattern, while 22% are free riders who never contribute anything, and 5% are full

both treatments by how much economic incentives (in terms of the MPCR) need to change in order to close the cooperation gap compared to the case without any delay. For *NowLater*, our results reveal that the MPCR has to be increased up to a level of 0.9 in order for the cooperation difference to become insignificant (Signrank test, $p = 0.247$). For all lower levels of the MPCR, contributions are significantly lower in the former than in the latter (Signrank tests, all $p < 0.065$). For *LaterNow*, we find that the MPCR can be decreased to a level of 0.35 in order for the cooperation difference relative to *NowNow* to become insignificant (Signrank test, $p = 0.218$). For all higher levels of the MPCR, contributions are significantly higher when the costs of cooperation are delayed (Signrank tests, all $p < 0.060$).

What do these results reveal about the degree of discounting in our context of cooperation? Similar to the commonly used approach to infer an individual's discount rate from the money amount needed to make her indifferent between a smaller-sooner and a larger-later reward, in our public goods setting we can use the variation in the MPCR to calculate a similar point of indifference. In particular, using linear regressions, we can estimate a threshold value the MPCR has to reach in order to equalize contributions between *NowLater* (*LaterNow*) and *NowNow*. At the aggregate level, in *NowLater* this break-even point is given by $\hat{\theta} = 0.962$. Relative to the case without any delay, this corresponds to a necessary increase in economic incentives of 92.4%, yielding an implicit discount factor of $\delta = \frac{0.5}{0.962} = 0.52$. In *LaterNow*, instead, we find that the MPCR can be lowered to a level of $\hat{\theta} = 0.247$ in order to equalize contributions with the ones observed in *NowNow*. This corresponds to a decrease of 50.7%, or an implicit discount factor of $\delta = 0.49$.[23]

To evaluate the magnitude of these effects, we can compare them to the revealed discount rates in the individual time preference elicitation task in Part 3 of our experiment. Using the midpoint of the interval in which an individual switches from the smaller-sooner to the larger-later reward, we find an average discount rate of 30.2%, which (under the assumption of locally linear utility) corresponds to an implicit discount factor of $\delta = 0.78$, which is

contributors. The remaining 7% display some non-monotonic contribution pattern. In *LaterNow*, we find 43% of all subjects exhibiting an increasing contribution pattern, 2% never contribute anything, and 46% contribute their whole endowment for each level of the MPCR. The remaining 11% fall under neither of these categories.

[23] Very similar results are obtained when instead of using the point in which contributions to the public are the same, we use the point in which expected payoffs are the same. To do this, we incorporate the elicited beliefs from Part 1 and 2 into our analysis. In particular, using linear regressions, we first make a point prediction of subjects' beliefs about others' level of contributions at the point in which estimated contributions are the same, i.e., at $\hat{\theta} = 0.962$ and $\hat{\theta} = 0.247$ in *NowLater* and *LaterNow*, respectively. We then use these predicted beliefs and the threshold values for θ to calculate expected payoffs. We then compare these to the expected earnings in *NowNow*, which we calculate using an MPCR of 0.5 and the elicited beliefs. The implied discount factors from these comparison yield $\delta = 0.58$ and $\delta = 0.56$ in *NowLater* and *LaterNow*, respectively, which are very similar to the ones above.

well in line with evidence from previous studies (see e.g., Harrison et al., 2002; Dohmen et al., 2010).[24] Although the procedures with which we compute discount rates in the public goods game and the individual decision-contexts slightly differ, the results reveal strong and pronounced differences in the degree of discounting between the two. In particular, our results indicate that in strategic environments—like the one considered here— delaying the consequences of actions can lead to exacerbated discounting, much higher than the levels typically observed in individual, non-strategic, environments. To the best of our knowledge, this is the first paper that demonstrates such an effect.

As a final step, we can check to what extent discounting is correlated across individual and strategic contexts at the individual level. To do that, as a simple reduced-form measure of discounting in our public goods setting, we calculate for each subject the cooperation gap at a MPCR level of 0.5 between the case with or without delay (compare Figure 5). We then correlate this measure with a subject's degree of impatience as measured by the number of late choices in the individual time preference task in Part 3 of the experiment. For both *NowLater* and *LaterNow*, we find a significant relationship between these two measures; the more patient a subject in Part 3 of the experiment, the lower her adjustment in contributions when the costs or the benefits of cooperation are delayed. The Spearman's rank correlation coefficient between the number of late choices and the absolute adjustment in contributions is -0.22 in *NowLater* ($p = 0.034$) and -0.23 in *LaterNow* ($p = 0.027$). This indicates that discounting behavior is correlated across individual and strategic contexts, at least to some degree. We note, however, that the correlation coefficients are far from one, indicating that other factors apart from one's own degree of impatience, such as the own attitude towards cooperation or the beliefs about others' degree of cooperativeness and impatience, matter for cooperation in intertemporal collective action problems, too. Of course, given the different techniques we used to measure discounting across the two contexts, measurement error might affect these results, too. We summarize these findings in our fifth result:

> **Result 5:** *The degree of discounting implied by the amount of monetary incentives needed to close the cooperation gap between NowLater (LaterNow) and NowNow is much larger than the degree of discounting observed in our individual time preferences elicitation task, indicating that discounting in strategic contexts is exacerbated. At the individual level, we find some moderate correlation of discounting between the two contexts.*

[24]The average lower bound is 28.9 and the average upper bound is 31.4, corresponding to $0.76 \leq \delta \leq 0.78$. As in our first study, we find no significant differences in individual time preferences between *NowLater* and *LaterNow* (Mann-Whitney U test, $p = 0.499$, see also see Figure A2 in Appendix A) indicating that randomization into treatments worked and that having experienced different treatments in Part 1 of the experiment had no systematic effect on Part 3 behavior.

6 Conclusion

Many economic decision situations involve a social dimension and a time dimension, leading agents to make choices that affect both themselves and others at different points in time. While previous literature has provided many important insights about how agents solve intertemporal and interpersonal trade-offs in isolation, so far relatively little is known about agents' behavior when both of these trade-offs exist at the same time, especially in strategic decision situations in which agents' payoffs are interdependent.

In this paper, we provide a systematic and comprehensive analysis of how the timing of payments affects behavior in a strategic context by studying how delaying the costs or the benefits of cooperation (or both) affects collective action in a linear public goods game. Overall, our paper makes two novel contributions. First, our paper contributes to the literature on time discounting, showing that in strategic environments delaying the consequences of actions can lead to exacerbated discounting. Compared to the 30% discount rate we observe in a standard individual decision environment (without any externalities on others), in our strategic environment discount rates increase considerably to about 50%. To the best of our knowledge, this is the first paper that demonstrates such an effect. Second, with regard to the literature on cooperation and voluntary provision of public goods, our study highlights that the factor time is an important but so far largely ignored determinant of cooperation. While delaying both the costs and the benefits of cooperation to the future had no effect on contributions indicating that cooperation is time-consistent, delaying only the benefits (cots) lead to a sharp decrease (increase) in voluntary contributions. With regard to applications outside the laboratory such as the problem of climate change in which the benefits of today's costly efforts can only be reaped in the far future, our results reveal that the expected level of voluntary contributions is rather low, although some significant positive amount of cooperation remains even in this setting. As such, evidence from previous literature should be seen as an upper bound for the level of achievable cooperation in these type of situations.[25]

Our results may have also some policy implications. First of all, the results from our second study provide some estimate of how much money a government or principal would need to invest in order to offset the effect of delayed payments. Second, our result that the general principle of reciprocity, i.e., the willingness to contribute more the more others' contribute, is maintained even when the benefits of cooperation are delayed is good news for policy makers as it demonstrates a point potential solution mechanisms can build on. For example, finding ways of engineering trust and optimism, or decreasing uncertainty about

[25]Of course, in many real-world social dilemmas additional factors such as technological complexities, as well as risk and uncertainty about the impact of one's contributions are important, too.

others' cooperativeness, e.g., via communication or other coordination devices, can already offset a large part of the observed negative effect on cooperation. Last, with regard to situations in which the timing of costs and benefits can be determined endogenously, our results show that costs should be delayed while benefits should be immediate. While this is of course something that is often hard to achieve (e.g., as in the case of climate change where the timing of payments cannot be chosen endogenously), it might be relevant, for example, in firms or organizations that rely on performance-contingent team incentives and which have some discretion about when to pay these bonuses. Our results reveal that paying bonuses only at the end of year might yield to lower collaborative team efforts than paying them, e.g., quarterly or immediately after the end of the project.

Obviously, our study only constitutes a first step towards an understanding of strategic decision making in intertemporal contexts. More research is needed to test the robustness of these results across different domains. For example, one interesting next step is to understand to what extent our results are driven by subjects' beliefs about the impatience of others (Rodriguez-Lara and Ponti, 2017; Fedyk, 2017). It is further important to analyze the sensitivity of our results with regard to the length of the delay. In particular, by systematically varying by how much the costs and benefits of cooperation are delayed, one could investigate whether, similar to the evidence from choices in individual decision situations (Frederick et al., 2002), people exhibit a present-bias also in strategic settings. Furthermore, while in our experiment we focused on a series of one-shot interactions to cleanly identify the different channels through which the timing of costs and benefits affects cooperation, many human interactions outside the lab take place in repeated contexts. Future research should therefore investigate the extent to which our results translate into such repeated interactions, in which the effects of learning, reputation, and other strategic incentives matter, too. Finally, another interesting avenue for future research is to investigate whether institutions that have been shown to foster cooperation in static contexts without any delay, such as communication (Isaac and Walker, 1988; Brosig et al., 2003; Balliet, 2010), punishment (Fehr and Gächter, 2000; 2002; Boyd et al., 2003; Herrmann et al., 2008), or leading by example (Güth et al., 2007; Potters et al., 2007; Gächter et al., 2012), are equally effective in promoting social efficiency in intertemporal social dilemmas. This is non-trivial, as in a situation in which the benefits of cooperation are delayed, for instance, punishment might be less effective in increasing contributions because agents might be less willing to sanction free riders as the costs of punishment might outweigh the discounted benefits from increased cooperation. In this case, different institutions might be needed in order to sustain cooperation.

References

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2008). Eliciting risk and time preferences. *Econometrica 76*(3), 583–618.

Anderson, S., J. Goeree, and C. Holt (1998). A theoretical analysis of altruism and decision error in public goods games. *Journal of Public Economics 70*(2), 297–323.

Andreoni, J. (1989). Giving with impure altruism: Applications to charity and ricardian equivalence. *Journal of Political Economy 97*(6), 1447–1458.

Andreoni, J. (1990). Impure altruism and donations to public goods: A theory of warm-glow giving. *Economic Journal 100*, 464–477.

Andreoni, J., M. A. Kuhn, and C. Sprenger (2015). Measuring time preferences: A comparison of experimental methods. *Journal of Economic Behavior & Organization 116*, 451–464.

Andreoni, J. and M. Serra-Garcia (2017). Time-inconsistent charitable giving. *NBER Working Paper No. 22824*.

Andreoni, J. and C. Sprenger (2012a). Estimating time preferences from convex budgets. *American Economic Review 102*(7), 3333–3356.

Andreoni, J. and C. Sprenger (2012b). Risk preferences are not time preferences. *American Economic Review 102*(7), 3357–3376.

Balakrishnan, U., J. Haushofer, and P. Jakiela (2017). How soon is now? evidence of present bias from convex time budget experiments. Technical report, National Bureau of Economic Research.

Balliet, D. (2010). Communication and cooperation in social dilemmas: A meta-analytic review. *Journal of Conflict Resolution 54*(1), 39–57.

Bolton, G. E. and A. Ockenfels (2000). ERC - A theory of equity, reciprocity and competition. *American Economic Review 100*(1), 166–93.

Boyd, R., H. Gintis, S. Bowles, and P. J. Richerson (2003). The evolution of altruistic punishment. *Proceedings of the National Academy of Sciences 100*(6), 3531–3535.

Breman, A. (2011). Give more tomorrow: Two field experiments on altruism and intertemporal choice. *Journal of Public Economics 95*(11), 1349 – 1357.

Brosig, J., J. Weimann, and A. Ockenfels (2003). The effect of communication media on cooperation. *German Economic Review 4*(2), 217–241.

Charness, G. and M. Rabin (2002). Understanding Social Preferences with Simple Tests. *Quarterly Journal of Economics 117*(3), 817–869.

Chaudhuri, A. (2011). Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature. *Experimental Economics 14*(1), 47–83.

Cheung, S. L. (2016). Recent developments in the experimental elicitation of time preference. *Journal of Behavioral and Experimental Finance 11*, 1–8.

Coller, M. and M. B. Williams (1999). Eliciting individual discount rates. *Experimental Economics 2*(2), 107–127.

Cornes, R. and T. Sandler (1996). *The theory of externalities, public goods, and club goods.* Cambridge University Press.

Davis, D., A. Ivanov, and O. Korenok (2016). Individual characteristics and behavior in repeated games: an experimental study. *Experimental Economics 19*(1), 67–99.

Deck, C. and S. Jahedi (2015a). An experimental investigation of time discounting in strategic settings. *Journal of Behavioral and Experimental Economics 54*, 95–104.

Deck, C. and S. Jahedi (2015b). Time discounting in strategic contests. *Journal of Economics & Management Strategy 24*(1), 151–164.

Dohmen, T., A. Falk, D. Huffman, and U. Sunde (2010). Are risk aversion and impatience related to cognitive ability? *American Economic Review 100*(3), 1238–60.

Dreber, A., D. Fudenberg, D. K. Levine, and D. G. Rand (2016). Self-control, social preferences and the effect of delayed payments. *Working Paper*.

Dufwenberg, M. and G. Kirchsteiger (2004). A theory of sequential reciprocity. *Games and Economic Behavior 47*(2), 268–298.

Falk, A. and U. Fischbacher (2006). A theory of reciprocity. *Games and Economic Behavior 54*(2), 293–315.

Fedyk, A. (2017). Asymmetric naiveté: Beliefs about self-control. *Working Paper*.

Fehr, E. and S. Gächter (2000). Cooperation and punishment in public goods experiments. *American Economic Review 90*(4), 980–994.

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

Fehr, E. and A. Leibbrandt (2011). A field study on cooperativeness and impatience in the tragedy of the commons. *Journal of Public Economics 95*(9), 1144–1155.

Fehr, E. and K. Schmidt (1999). A Theory of Fairness, Competition, and Cooperation. *Quarterly Journal of Economics 114*(3), 817–868.

Fehr, E. and K. M. Schmidt (2006). The economics of fairness, reciprocity and altruism - experimental evidence and new theories. In S. Kolm and J. M. Ythier (Eds.), *Handbook of the Economics of Giving, Altruism, and Reciprocity*, Volume 1, pp. 615–691. Elsevier.

Fischbacher, U. (2007). Z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics 10*(2), 171–178.

Fischbacher, U. and S. Gächter (2010). Social preferences, beliefs, and the dynamics of free riding in public goods experiments. *American Economic Review 100*(1), 541–556.

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

Fischbacher, U., S. Gächter, and S. Quercia (2012). The behavioral validity of the strategy method in public good experiments. *Journal of Economic Psychology 33*(4), 897–913.

Frederick, S., G. F. Loewenstein, and T. O'Donoghue (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature 40*(2), 351–401.

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

Gächter, S. and B. Herrmann (2009). Reciprocity, culture and human cooperation: previous insights and a new cross-cultural experiment. *Philosophical Transactions of the Royal Society B: Biological Sciences 364*(1518), 791–806.

Gächter, S., F. Kölle, and S. Quercia (2017). Reciprocity and the tragedies of maintaining and providing the commons. *Nature Human Behaviour 1*(9), 650.

Gächter, S., D. Nosenzo, E. Renner, and M. Sefton (2012). Who makes a good leader? cooperativeness, optimism, and leading-by-example. *Economic Inquiry 50*(4), 953–967.

Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association 1*(1), 114–125.

Güth, W., M. V. Levati, M. Sutter, and E. Van Der Heijden (2007). Leading by example with and without exclusion power in voluntary contribution experiments. *Journal of Public Economics 91*(5-6), 1023–1042.

Hardin, G. (1968). Tragedy of the commons. *Science 162*, 1243–1248.

Harrison, G. W., M. I. Lau, and M. B. Williams (2002). Estimating individual discount rates in denmark: A field experiment. *American Economic Review 92*(5), 1606–1617.

Hartig, B., B. Irlenbusch, and F. Kölle (2015). Conditioning on what? heterogeneous contributions and conditional cooperation. *Journal of Behavioral and Experimental Economics 55*, 48–64.

Hauser, O. P., D. G. Rand, A. Peysakhovich, and M. A. Nowak (2014). Cooperating with the future. *Nature 511*(7508), 220.

Herrmann, B., C. Thöni, and S. Gächter (2008). Antisocial punishment across societies. *Science 319*(5868), 1362.

Isaac, R. M. and J. M. Walker (1988). Communication and free-riding behavior: The voluntary contribution mechanism. *Economic Inquiry 26*(4), 585–608.

Jacquet, J., K. Hagel, C. Hauert, J. Marotzke, T. Röhl, and M. Milinski (2013). Intra-and intergenerational discounting in the climate game. *Nature climate change 3*(12), 1025.

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

Kölle, F. and L. Wenner (2018). Present-biased generosity: Time inconsistency across individual and social contexts. Technical report.

Kovarik, J. (2009). Giving it now or later: Altruism and discounting. *Economics Letters 102*(3), 152–154.

Laibson, D. (1997, May). Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics 112*(2), 443–477.

Ledyard, J. O. (1995). Public goods: A survey of experimental research. In J. H. Kagel and A. E. Roth (Eds.), *The Handbook of Experimental Economics*, pp. 111–194. Princeton University Press.

O'Donoghue, T. and M. Rabin (1999). Doing it now or later. *American Economic Review 89*(1), 103–124.

Olson, M. (1965). *The Logic of Collective Action: Public Goods and the Theory of Groups.* Harvard University Press.

Ostrom, E. (2000, February). Collective action and the evolution of social norms. *Journal of Economic Perspectives 14*(3), 137–158.

Potters, J., M. Sefton, and L. Vesterlund (2007). Leading-by-example and signaling in voluntary contribution games: an experimental study. *Economic Theory 33*(1), 169–182.

Rodriguez-Lara, I. and G. Ponti (2017). Social motives vs social influence: An experiment on interdependent time preferences. *Games and Economic Behavior 105*, 177–194.

Ross, L., D. Greene, and P. House (1977). The false consensus effect: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology 13*(3), 279–301.

Schelling, T. (1995). Integenerational discounting. *Energy Policy 23*(4), 395–401.

Schlag, K. H. and J. van der Weele (2015). A method to elicit beliefs as most likely intervals. *Judgment and Decision Making 10*(5), 456.

Selten, R. (1967). Die strategiemethode zur erforschung des eingeschrnkt rationalen verhaltens im rahmen eines oligopolexperimentes. In *Beitrge zur experimentellen Wirtschaftsforschung*, pp. 136–168. Sauermann, H.

Sobel, J. (2005). Interdependent preferences and reciprocity. *Journal of Economic Literature 43*(2), 392–436.

Zelmer, J. (2003). Linear public goods experiments: A meta-analysis. *Experimental Economics 6*(3), 299–310.

Appendix

A Additional Tables and Figures

Decision #	Option A	Option B	Annual interest rate (percent)
1	€50	€51.25	2.5
2	€50	€52.50	5.0
3	€50	€53.75	7.5
4	€50	€55.00	10.0
5	€50	€56.25	12.5
6	€50	€57.50	15.0
7	€50	€58.75	17.5
8	€50	€60.00	20.0
9	€50	€61.25	22.5
10	€50	€62.50	25.0
11	€50	€63.75	27.5
12	€50	€65.00	30.0
13	€50	€66.25	32.5
14	€50	€67.50	35.0
15	€50	€68.75	37.5
16	€50	€70.00	40.0
17	€50	€71.25	42.5
18	€50	€72.50	45.0
19	€50	€73.75	47.5
20	€50	€75.00	50.0

Table A1: Multiple price list for eliciting individual time preferences in Part 3 of the experiment.

	NowNow	NowLater	LaterNow	LaterLater
Number of late choices	0.077	0.281**	-0.463***	-0.100
	(0.135)	(0.116)	(0.113)	(0.144)
Constant	9.004***	2.605***	19.724***	11.080***
	(1.418)	(0.967)	(0.573)	(1.646)
N	64	64	56	64
R^2	0.005	0.114	0.224	0.007

Note: The dependent variable is the number of tokens contributed to the public good. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A2: OLS regression. Effects of time preferences on contributions by treatment.

	(1)
Avg. contrib. of others	0.605***
	(0.055)
Avg. contrib. of others x *NowLater*	-0.015
	(0.076)
Avg. contrib. of others x *LaterNow*	0.137
	(0.086)
Avg. contrib. of others x *LaterLater*	0.015
	(0.083)
NowLater	-1.569**
	(0.642)
LaterNow	2.996**
	(1.275)
LaterLater	-0.371
	(0.859)
Constant	1.212**
	(0.563)
N	5208
R^2	0.355

Robust standard errors (clustered at the individual level) are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A3: OLS regression. Conditional contributions in strategy-method experiment.

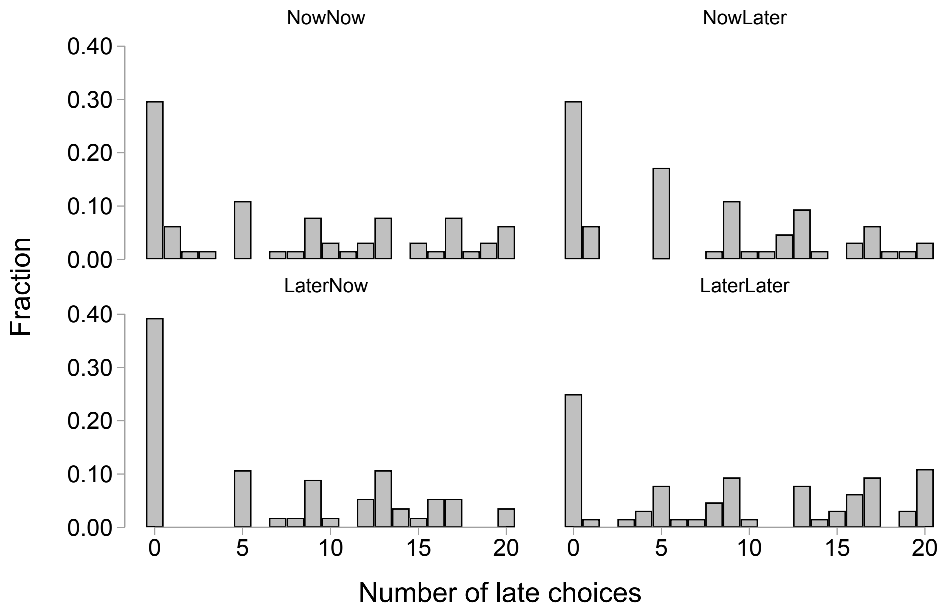

Figure A1: Distribution of number of late choices in the time preference elicitation task in Part 3 of Study 1, separated by treatment.

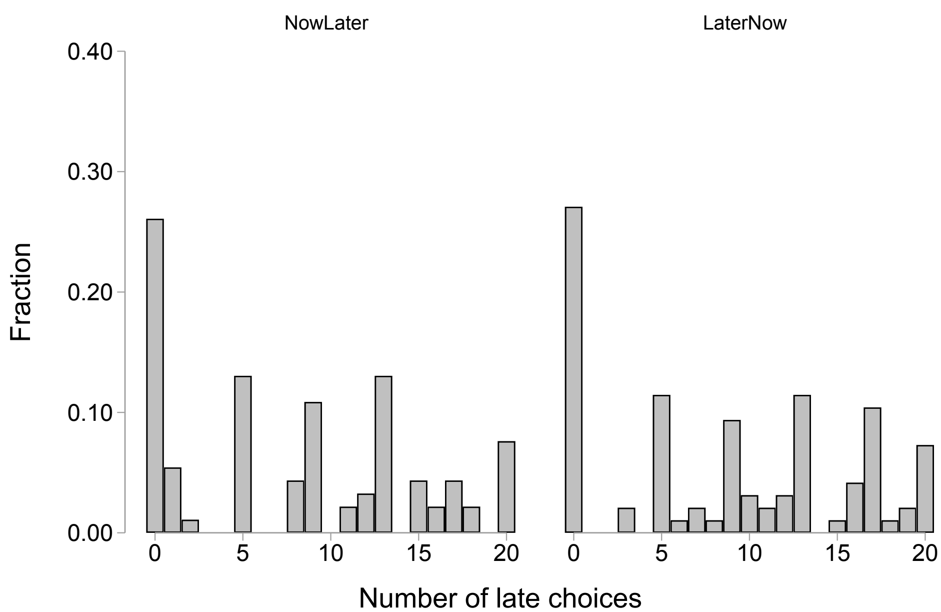

Figure A2: Distribution of number of late choices in the time preference elicitation task in Part 3 of Study 2, separated by treatment.