

Kölle, Felix; Wenner, Lukas

Conference Paper

Time-Inconsistent Generosity: Present Bias across Individual and Social Contexts

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2019: 30 Jahre Mauerfall - Demokratie und Marktwirtschaft - Session: Experimental Economics III, No. B21-V1

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Kölle, Felix; Wenner, Lukas (2019) : Time-Inconsistent Generosity: Present Bias across Individual and Social Contexts, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2019: 30 Jahre Mauerfall - Demokratie und Marktwirtschaft - Session: Experimental Economics III, No. B21-V1, ZBW - Leibniz-Informationszentrum Wirtschaft, Kiel, Hamburg

This Version is available at:

<https://hdl.handle.net/10419/203505>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Time-Inconsistent Generosity: Present Bias across Individual and Social Contexts

Felix Kölle and Lukas Wenner*

University of Cologne

November 2018

Abstract

We investigate the presence and stability of dynamically inconsistent time preferences across contexts with and without interpersonal trade-offs. In a longitudinal experiment subjects make a series of intertemporal allocation decisions of real-effort tasks between themselves and another person. We find substantial time inconsistency in generosity: agents become disproportionately more selfish when decisions have immediate rather than delayed consequences. Based on our theoretical framework, structural estimations reveal that this is because agents exhibit present bias in own but not in others' consumption. We show that very similar differences in present bias are observed in the absence of any interpersonal trade-offs, when agents decide either for themselves or on behalf of another person. At the individual level, we find that present bias in own consumption is a stable behavioral trait which is correlated across individual and social contexts.

Keywords: Present bias; altruism; stability; real effort; dictator game; intertemporal choice.

JEL Classification Numbers: C91; D64; D90

*Kölle: Department of Economics, University of Cologne, Albertus Magnus Platz, 50923 Cologne, e-mail: *felix.koelle@uni-koeln.de*. Wenner: Department of Economics, University of Cologne, Albertus Magnus Platz, 50923 Cologne, e-mail: *lukas.wenner@uni-koeln.de*. We thank Antonio Cabrales, Paul Heidhues, Axel Ockenfels, Frederik Schwerter, Matthias Sutter, Nora Szech, Sebastian Tonke, Matthias Wibral, and various seminar and conference participants for helpful comments. Financial support from the Center for Social and Economic Behavior (C-SEB) at the University of Cologne is gratefully acknowledged.

1 Introduction

When faced with intertemporal trade-offs, many economic decision makers display a present bias, that is, their desire for immediate gratification leads them to become disproportionately more impatient when choices directly affect the present (Strotz, 1956; Loewenstein and Prelec, 1992; Laibson, 1997; O’Donoghue and Rabin, 1999; Frederick et al., 2002). Evidence for this comes from a variety of settings, such as financial decision-making (Ashraf et al., 2006), exercising (DellaVigna and Malmendier, 2006), and effort provision (Augenblick et al., 2015), supporting the notion that intertemporal decision-making is often time-inconsistent. The existing body of evidence, however, almost exclusively focuses on present bias in *individual decision contexts*, i.e., situations in which only own consumption is at stake. Yet, intertemporal trade-offs also play an important role in *social situations*, in which there is a conflict of interest between own and others’ well-being. For example, requests for supporting a friend or a colleague, donating to charity, or contributing to a group project, might generate very different degrees of generosity when carried out in advance rather than immediately. In these types of situations, evidence from static contexts documents that many people exhibit social preferences, i.e., they are willing to sacrifice their own payoffs for the benefit of others (see Sobel, 2005; Cooper and Kagel, 2009, for reviews of the literature). Much less is known, however, about how such social preferences unfold in dynamic contexts.

In this paper, we provide a first systematic analysis of time discounting in individual and social context within a unified framework. We report on a longitudinal experiment in which subjects make intertemporal allocation decisions of effort (i.e., negative leisure consumption) in contexts with and without a social dimension. We address the following two research questions.

First, do economic agents exhibit time-inconsistent generosity? That is, do people become more (or less) generous when consequences are delayed, and to what extent does this effect depend on whether the delay affects the present or not? To date, few papers have investigated decision making in situations which have both a social and an intertemporal dimension. Most of these papers focus on a single decision which is either implemented immediately or with delay. Kovarik (2009) and Dreber et al. (2016) study dictator game giving and find that giving decreases when delaying both own and the other person’s monetary payments. Breman (2011) studies charitable giving with delays of one and two months and finds an increase of charitable giving with delay. Unlike us, none of these studies analyze present bias directly.¹ The only paper we are aware that directly tests for present bias is Andreoni

¹A similar argument applies to the work by Rong et al. (2016) who study cases where dictator and recipient receive money at different points in time, either with a delay of one or five years. Also, their design

and Serra-Garcia (2017) who similar to Breman (2011) find an increase in charitable giving when agents are asked to donate in advance rather than on the spot. They, however, study present bias in the context of demand for commitment to donate, while we focus on time-inconsistency in a dictator game. Compared to these previous studies, the main contribution of this paper is to offer a more systematic and comprehensive analysis of the role of time on prosocial behavior. In particular, our design allows us to distinguish between the effects of short-run and long-run discounting, and to estimate time preference parameters structurally. Moreover, we study dictator game behavior in the effort domain, addressing the concern that the laboratory may not be well suited to capture present bias in money (Augenblick et al., 2015; Balakrishnan et al., 2017). We thus provide a direct link to the literature on time preferences in individual decision contexts (see Frederick et al., 2002; Cohen et al., 2016, for reviews of the literature).

Second, we ask whether time inconsistency is a behavioral phenomenon which is stable across individual and social contexts. That is, do people who display a present bias in individual contexts show a similar desire for immediate gratification in social contexts, where the costs of such behavior are borne by someone else rather than one's own future self? Studying to what extent economic behavior in different contexts is guided by stable underlying preferences, is a question that lies at the core of economic analysis (Stigler and Becker, 1977). Despite the importance of this question for the validity of economic research, surprisingly little is known about the stability of preferences across different domains. With respect to time preferences, Meier and Sprenger (2015) show that measures of time preferences are relatively stable across two consecutive years. Augenblick et al. (2015), in contrast, find no correlation of time preferences across the monetary and the effort domain. Other studies have looked at the correlation between experimental measures of time preference and real-world behavior outside the lab and report mixed results (see e.g., Ashraf et al., 2006; Chabris et al., 2008; Meier and Sprenger, 2010). Here, we provide an empirical test of the stability of present bias across contexts with and without interpersonal trade-offs, an exercise which, to the best of our knowledge, has not yet been made.²

Our design also allows for a direct comparison between intertemporal choices made for oneself and those made on behalf of someone else. Situations in which individuals make intertemporal decisions for others are frequent. Think, for instance, of asset managers investing on behalf of their clients, doctors choosing treatments for their patients, or parents

does not allow for a direct comparison with behavior when delay is symmetric.

²The stability of preferences has been investigated also in other domains. See, e.g., Andersen et al. (2008); Barseghyan et al. (2011); Dohmen et al. (2011) for studies on risk preferences and Blanco et al. (2011); Volk et al. (2012); Peysakhovich et al. (2013); Bruhin et al. (2017) for studies on other-regarding preferences.

deciding what is best for their children. Especially with regard to present bias, it is important to understand whether when deciding for another person, the desire for immediate gratification is equally strong compared to when deciding for oneself, or whether the greater personal distance mitigates time inconsistency. To date, the evidence on differences between intertemporal choices for oneself and those made on behalf of others is rather mixed and incomplete. Albrecht et al. (2011) study present bias in the choice of smaller-sooner versus larger-later monetary rewards, but find no aggregate effect of a difference in present bias for oneself and another person.³ Focusing on patience rather than present bias per-se, Shapiro (2010) and Howard (2013) find higher patience in choices made for oneself, while de Oliveira and Jacobson (2017) find the opposite.⁴ We thus contribute to this literature by providing a clean test for differences in discounting for self and others both in the long-run as well as in the short-run (in form of present bias).

To answer our research questions, we conduct a three-week longitudinal experiment in which participants are asked to make intertemporal allocation decisions of units of effort (i.e., negative leisure consumption) for varying prices using a convex budget set approach (Andreoni and Miller, 2002; Andreoni and Sprenger, 2012). The effort task is based on Erkal et al. (2011) and consists of encrypting a string of letters into numbers. Like in Augenblick et al. (2015), allocation decisions are made at two points in time—an initial allocation in week 1, and a subsequent allocation in week 2—while effort needs to be exerted in week 2 or in week 3. To incentivize all decisions, after subjects completed their week 2 decision, we randomly selected one decision—either from week 1 or from week 2—to be implemented and determine subjects’ workload. Differences between initial and subsequent allocation decisions allow for a precise measurement of dynamic inconsistency.

Each subject makes choices in two types of allocation decisions. In the first, subjects face intertemporal trade-offs in a social context in which they allocate tasks between themselves and another person. In contrast to choices in standard (static) dictator games, we systematically vary the timing of when the consequences for the decision maker and the consequences for recipient realize; either both immediately, both delayed, or one delayed and the other immediately. We refer to these decisions as *interpersonal choices*. In the second type of allocation decisions, subjects face intertemporal trade-offs that either only affect themselves or only affect another person, i.e., choices in which there is no conflict between own and others’ consumption. We call these decisions *intrapersonal choices*. The experimental variation of

³When analyzing the subset of subjects which are classified as having a strong present bias only, Albrecht et al. (2011) find that these participants are less present-biased when choosing on behalf of others.

⁴Relatedly, Rodriguez-Lara and Ponti (2017) study a situation in which own intertemporal trade-offs are imposed on another person. They find that when provided with information about the other’s discounting behavior, subjects adjust their own behavior to account for the other person’s time preferences.

consequences and decisions allows us to estimate time preference parameters for own and others' consumption in the individual and the social domain, and to compare their stability across domains.

Studying the interplay between time preferences and social preferences in a dynamic context requires a suitable theoretical framework. Paralleling the multi-attribute utility approach in Andersen et al. (2018) and Cheung (2015), we propose a utility function which allows for differences in discounting of own and others (atemporal) utility, while at the same time accounting for equality-efficiency trade-offs in own consumption versus another person's consumption. The key insight from our analysis is that if individuals discount own consumption to a larger extent than others' consumption, they should become less selfish when consequences are delayed. Moreover, if individuals exhibit stronger present bias for own compared to others' consumption, generosity is subject to time inconsistency. Intuitively, such a difference in present bias increases the relative weights of own vis-à-vis others' consumption when consequences are immediate rather than delayed. If, on the contrary, there are no differences in relative discounting between self and others, altruistic behavior should be unaffected by the timing of decisions and consequences as in this case, the relative weight of own compared to others' consumption is constant over time.

The results from our interpersonal decisions reveal a substantial time inconsistency in generosity. In allocations where both agents need to complete the task in week 2, subjects allocate 15.7% more tasks to themselves when choosing in advance (week 1) rather than in the present (week 2). When both agents need to work in week 3, in contrast, the number of tasks allocated to oneself only decreases by 5.6% between the two weeks. This implies a statistically significant decrease in generosity of 10.1% that is driven by the immediacy of consumption in the present. By including the data from those interpersonal choices in which the consequences for the decision-maker and the recipient occur at different points in time, we can structurally estimate time preference parameters. We find evidence for significant present bias in own but not in others' consumption. Depending on the exact specification, our estimates for present bias in own consumption, β_s , range from 0.883 to 0.910, which are all significantly lower than one. Our estimates for present bias in others' consumption, β_o , in contrast, lie between 1.043 and 1.060, which are significantly higher than our estimates for β_s and not significantly different from one. We find no evidence for significant long-run discounting, neither for own nor for others' consumption.

Very similar discounting patterns can be observed in our intrapersonal choices. We find that when deciding for themselves, subjects allocate 6.1% more tasks to the sooner date when deciding in advance rather than in the present. Our estimations reveal that this implies a

β_s of 0.842 to 0.863, which is statistically different from one, replicating the finding by Augenblick et al. (2015) for slightly different tasks and procedures. When subjects decide on behalf of someone else, instead, we find a decrease of only 2.2% across the two decision dates, which implies β_o estimates which are not significantly different from one. Again, we find no evidence for any long-run discounting in either of the two cases.

To test the stability of present bias across contexts, we structurally estimate time preference parameters at the individual level, separately for the interpersonal and intrapersonal choices. For present bias in own consumption, we find a significant positive correlation, suggesting that there is a stable underlying present bias trait across the two contexts. To our knowledge, this is the first paper which demonstrates that present bias extends from individual decision contexts to social contexts. For present bias in others' consumption, the correlation is much weaker and not significantly different from zero. Hence, while our aggregate result of no present bias in others' consumption is consistent across contexts, our individual-level analysis reveals that how an agent discounts another person's consumption seems to be conceptually different depending on whether there are trade-offs with own consumption, or not.

Our findings have important implications for the analysis and modeling of social preferences and time preferences. First, with regard to other-regarding behavior we show that the degree of generosity economic agents exhibit depends on the relative timing of decisions and consequences. More specifically, we show that altruism is time-inconsistent. Based on our theoretical framework, we identify the root of this inconsistency as differences in present bias in own compared to others' consumption. As such, our results provide important insights into the modeling of altruistic behavior in a dynamic context. Second, with regard to time preferences our results show that present bias in own consumption is a phenomenon that not only emerges in individual decision contexts, but one that extends to choices that also affect others. The fact that we find a positive correlation of present bias at the individual level suggests that the desire for immediate gratification is a stable underlying behavioral phenomenon, even though the type of trade-offs agents face are very different across the two contexts. To the best of our knowledge, this is the first paper showing that present bias is stable across contexts with and without a social dimension. Third, our results robustly demonstrate that there is no present bias when discounting others' consumption, which is consistent with neuro-economic evidence showing that different regions in the brain are active when deciding for oneself versus deciding for another person (McClure et al., 2004; Albrecht et al., 2011). Taken together with the evidence in Fedyk (2017) that agents correctly predict the present bias of others but are unaware of their own present bias, this result suggests that agents consider time-consistent rather than present-biased choices for

others as desirable and welfare-maximizing. Finally, insofar as present bias represents an impulsive, temptation-driven desire for immediate gratification, our results further corroborate the view that agents evaluate others’ consumption in a less biased, more controlled and analytical manner. As already argued by Schelling (1984), in many situations, casual observation suggests that agents might be willing to delegate choices to friends or family in the belief that when they choose on one’s behalf, they are to a lesser extent subject to temptations.⁵ Our findings are consistent with these observations.

The remainder of the paper is organized as follows. The next section presents the design of our experiment. In Section 3, we provide a theoretical framework for the analysis of the dictator games when consequences of decisions are delayed. Section 4 analyzes the data from our interpersonal choices. In Section 5, we first present the results from the intrapersonal choices at the aggregate level. We then structurally estimate time preference parameters at the individual level to investigate the stability of present bias across contexts. Section 6 investigates the robustness of our structural estimates to different specifications, while Section 7 concludes.

2 The Experiment

Our experiment investigates subjects’ allocation decisions about the completion of a real-effort encryption task. Similar to Augenblick et al. (2015), we implemented a longitudinal experiment that took place at three dates over three consecutive weeks. All meetings were conducted in the laboratory, and all subjects were required to participate at all dates of the experiment. In the first two weeks, subjects had to make a series of allocation decisions that could affect their own as well as another participant’s work load in week 2 and week 3. In the following, we present the experimental design in more detail. First, in Section 2.1, we describe the real-effort task participants had to work on. In Section 2.2, we present the decision environment in which effort allocations were made. Finally, in Section 2.3 we provide details about the general experimental procedures, payments, and recruitment.

⁵Schelling (1984) lists a number of examples, including handing over car keys to others when drinking, telling friends not to lend them money (when in a casino, for example), or relying on groups to commit to lose weight. We view these examples as plausibly supporting the notion that when evaluating others’ consumption, agents might be less (or not at all) present-biased, but do not delve deeper into the related, but separate, question of whether we should observe delegation of choices to others in addition or as an alternative to commitment devices provided by markets. We note, however, that implicit in the delegation argument is that one can trust the other person enough to “do the right thing”, an issue we will address in Section 5.

H	Z	C	M	A	V	S	U	Y	G	F	D	Q	P	L	W	T	E	K	R	X	J	B	O	I	N
835	890	242	321	585	963	881	288	652	125	477	448	760	982	183	833	212	858	561	686	246	692	105	366	583	222

Figure 1: Screenshot of the encryption task

2.1 Encryption Task

Our encryption task is based on Erkal et al. (2011). In this task, subjects have to encode a string of letters (a "word") to numbers. Each word consists of eight letters. The numbers are given by an encryption table, showing all 26 letters of the alphabet as well as corresponding three-digit numbers. The subjects' task is to type in the correct three-digit number corresponding to each letter into an empty textbox (see Figure 1 for a screenshot). After all eight letters are encrypted, subjects have to press a "submit" button. If the task is solved correctly, a new word appears, along with the information about the total number of correctly solved tasks so far and the remaining number of tasks to solve. In case of an incorrect entry, subjects are informed about their mistake.⁶ In this case, all entries are deleted and subjects have to encrypt the same word again. There is no time limit for correctly encrypting a word.

To mitigate learning effects over time and in order to make the exertion of effort as comparable as possible across the different dates of our experiment, we use a double randomization technique, introduced by Benndorf et al. (2014). After each correctly solved word, each letter is associated to a new, randomly allocated, three-digit number, and the position of all letters is randomly reshuffled.⁷

2.2 Effort Allocations

In both week 1 and week 2, subjects make a series of allocation decisions in which they have to allocate tasks between themselves and others as well as between week 2 and week 3. We distinguish between two types of decisions, *interpersonal* and *intrapersonal*.

⁶The overall level of mistakes was very low. 96.5% of all submitted answers were correct.

⁷It seems that we were largely successful in our attempt to mitigate learning effects. While in week 1 subjects took on average 39.1 seconds per task, in weeks 2 and 3 this number slightly drops to 36.5 and 35.9, respectively. These numbers are based on the minimum work of 10 tasks that every subject has to complete each week, as discussed below.

Decision Type	Block	X	Y
<i>Interpersonal</i>	SOONSOON	s_t	o_t
	LATELATE	s_{t+1}	o_{t+1}
	SOONLATE	s_t	o_{t+1}
	LATESOON	s_{t+1}	o_t
<i>Intrapersonal</i>	SELF	s_t	s_{t+1}
	OTHER	o_t	o_{t+1}

Table 1: Allocation decisions within each of the six blocks

In the *interpersonal allocation decisions* subjects make choices in four blocks. Here they have to decide, similar to standard dictator games, how many tasks they want to solve themselves and how many tasks have to be solved by another person. In two out of these four blocks, the time at which effort needs to be exerted is the same for the dictator and the receiver. In block SOONSOON agents decide about allocations of tasks which need to be completed in week 2, while in block LATELATE the decision environment is the same but the working date is week 3. In the following, we refer to these blocks as *symmetric dictator games*. In the other two blocks, the time at which the agents need to exert effort differs, we therefore call them *asymmetric dictator games*. In SOONLATE, the dictator has to work in week 2, while tasks allocated to the recipient have to be completed in week 3. In LATESOON, the roles are reversed such that the dictator has to work in week 3 and the recipient has to work in week 2.

In the *intrapersonal allocation decisions* subjects make choices in two blocks without any interpersonal trade-offs. In particular, in block SELF subjects choose how many tasks they want to solve in week 2 and how many tasks they want to solve in week 3. In block OTHER they face the exact same trade-off but now choose on behalf of another participant. The order in which subjects face these six blocks was randomized.⁸

Allocations are made in a convex time budget (CTB) environment (Andreoni and Sprenger, 2012). Subjects allocate tasks between two accounts, X and Y , whereby the exchange rate between X and Y differs from decision to decision. In particular, every task allocated to account Y reduces the number of tasks allocated to account X by R . Within each block, we use the following six rates: $R \in \{0.5, 0.75, 1, 1.25, 1.5, 2\}$. For example, a

⁸The randomization was as follows: Half of the subjects face the intrapersonal allocations first, followed by the symmetric dictator games and vice versa for the other half (always SELF before OTHER and SOONSOON before LATELATE). We then independently randomize whether these four blocks are followed by LATESOON or SOONLATE, leaving us with four different orderings. We do not find any evidence for systematic order effects.




Figure 2: Screenshot of the allocation environment

rate of 0.5 implies that each task allocated to account Y reduces the number of tasks allocated to account X by 0.5. Formally, a subject thus faces a budget constraint of the form $X + R \cdot Y = m$.

In each decision $m = 50$, hence, since negative number of tasks are not allowed, a subject can allocate at most 50 tasks to account X , while for account Y the maximum varies between 25 tasks ($R = 2$) and 100 tasks ($R = 0.5$). Depending on the block, account X and Y had different meanings. This is summarized in Table 1, where s stands for tasks allocated to oneself (*self*) and o stands for tasks allocated to someone else (*other*). The subscript indicates the time when the tasks have to be solved, t corresponds to week 2, and $t + 1$ corresponds to week 3. As an example, Figure 2 shows a screenshot of the allocation environment in block SOONSOON.

The real-effort task that we chose mandates that the number of allocated tasks is discrete. As Chakraborty et al. (2017) point out, in Augenblick et al. (2015) the authors chose a rounding method that leads to dominated choices being available to subjects, and subjects do indeed choose such dominated allocations. In our design this is not the case as we remove allocations in a way that no dominated allocations can be chosen.⁹ This approach seems most favorable as these violations may often be simply due to subjects being unaware that

⁹More precisely, we allow for $X \in \{0, 1, 2, \dots, 49, 50\}$ and, as a first step, round all Y to the closest integer. For $R > 1$, this leads to cases where two allocations (X, Y) and (X', Y) with $X > X'$ are both available. As a second step, we remove such “double appearances” in Y by keeping the allocation which does not contain a rounded value. For example, when $R = 2$ we have $(0, 25)$ and $(1, 25)$ and remove the latter. If both allocations contain rounded values, we remove the dominant alternative of the two, e.g., for $R = 1.25$ we remove $(2, 38)$ and keep $(3, 38)$.

	Minimum work	Allocation decisions	Allocation that counts chosen	Complete work
Week 1	✓	✓		
Week 2	✓	✓	✓	✓
Week 3	✓			✓

Table 2: Summary of the experiment

dominant options are available.

Each subject makes a total of 72 allocation decisions: 36 in week 1 and 36 in week 2 (six blocks with six different task rates each). Importantly, subjects in week 1 are informed that they will have to make allocation decisions in week 2 again, but they are not reminded of their initial week 1 allocations in week 2. After the week 2 decisions, subjects are randomly matched into pairs. Within each pair, one subject is randomly chosen as the decision maker.¹⁰ After that, one of the 72 allocations of the decision maker is chosen at random as the "allocation that counts". The allocated number of tasks from this decision then determines how many tasks each subject of the pair has to complete on the two work dates, in addition to a minimum requirement of 10 tasks that need to be completed at the beginning of every week (see below).¹¹ This procedure ensures that each decision is elicited in an incentive-compatible way.

In addition to their choices, in each week subjects are required to complete a "minimum work" of 10 encryption tasks prior to making their allocation decisions or completing their allocated tasks. As discussed in Augenblick et al. (2015, p.1077), this ensures that (i) at all dates subjects incur the cost of coming to the lab, (ii) in week 1 subjects get an idea how tedious the task is, and (iii) subjects have the same level of work experience at both allocation dates. Table 2 summarizes our experimental design, containing all tasks subjects face in each of the three weeks.

¹⁰To make the different roles more salient, we decided to use a physical randomization procedure. More precisely, subjects were asked to draw a colored card out of a bag containing the same number of blue and red cards. Red players were assigned the role of the decision maker.

¹¹In case a decision from block SELF or OTHER is selected, the respective other person only has to complete the minimum work. Similarly, in cases where the selected allocation decision does not specify any work by design, e.g., week 3 in block SOONSOON, only the minimum work has to be completed.

2.3 Recruitment, Payments, & Procedures

All sessions were computerized using the software Ztree (Fischbacher, 2007). We recruited subjects using ORSEE (Greiner, 2015). In the invitation email, subjects were informed about the longitudinal nature of the experiment. In particular, they were told that the experiment consists of three experimental sessions that each lie one week apart from each other. They were further told that they should only register if they can ensure that they participate at all three dates. The sessions took always place at the same day of the week, the same time of the day, and in the same laboratory. Before each session, subjects were send an email reminder about the remaining sessions. When invited for the experiment, participants were informed that the total average time of the experiment would be around 3 hours, but that the duration of each session could vary between 15 and 90 minutes.

If subjects showed up to all three experimental sessions and completed all tasks as specified by the randomly selected allocation, they received a completion payment of €40. If they failed to show up to one of the sessions in weeks 2 or 3, they were still eligible for a payment of €4, which corresponds to the usual show-up fee paid to subjects at the Cologne Laboratory of Experimental Research (CLER) where this study was run in August 2017. All payments were administered at the end of the third session in week 3 and subjects knew this in advance.

At the beginning of each experimental session, subjects received written instructions that were also read aloud by one of the experimenters. Instructions contained detailed information about the timeline of the experiment as well as the tasks to be solved in each of the three weeks.¹² After that, in each of the three weeks subjects had to complete the minimum work of 10 encryption tasks. Subsequently, in week 1 and week 2 subjects made their allocation decisions. In week 1, the session ended after the allocation decisions, followed by a short demographic questionnaire. In week 2 (after the allocation decisions) and week 3 (after the minimum work) subjects had to solve the number of tasks as specified by the allocation that counts. After completing all tasks, subjects could silently leave the lab without disturbing the other participants. In week 3, subjects received their payments immediately after completing their allocated tasks at their desk.

One concern with this procedure is that subjects may fear that others could draw conclusions about their allocation decisions. This is particularly relevant for the dictator games as previous literature has shown that social image concerns can increase pro-sociality (Benabou and Tirole, 2006; Charness and Gneezy, 2008; Andreoni and Bernheim, 2009). Note, however, that given our random implementation of one decision out of the six different blocks,

¹²A translated version of the instructions for all three weeks can be found in Appendix E.

by design, about half of the participants in each session are expected to only complete the minimum work in a given week. As a result, it is almost impossible for participants to infer others' degree of selfishness or impatience from the time they spend in the lab. We are hence confident that such concerns played no role in our setup.

Out of the $n = 110$ subjects who participated in our week 1 experiment, $n = 104$ showed up and completed all tasks in week 2.¹³ One crucial requirement for being able to identify an individual's time preference parameters is that we observe some variation in their allocation decisions. If in at least one week there is no variation in a subject's response to changes in the exchange rate R , behavior conveys limited information about time preferences. For example, in the interpersonal decisions, subjects who always allocate zero tasks to themselves can easily be identified as being completely selfish, but nothing can be said about their time preferences in this context. Hence, in our analyses we only focus on those subjects that do exhibit some positive amount of variation in their allocation decisions in both week 1 and week 2. For our dictator game decisions, we find a total of 33 subjects who do not exhibit any variation in at least one of the weeks, all of them because they do not allocate any tasks to themselves (20 out of these 33 subjects behave fully selfish in both weeks). Consequently, our remaining sample of $n = 71$ subjects is a selected sample that is more generous than the average. However, this selected sample is not more or less patient than the average, that is, we find no significant differences in choices in block SELF between those who are completely selfish (in at least one week) and those who are not. For our block SELF analysis, we have to drop four subjects who display no variation in at least one of the two weeks, leaving us with a sample of $n = 100$ subjects. For the same reason, when analyzing choices in block OTHER we have to drop six subjects, leaving us with $n = 98$ subjects.¹⁴

In Appendix D we check that our results are robust to these exclusion restrictions. In particular, for blocks SELF and OTHER we re-run all our main estimations for the full sample. For the dictator games, as an alternative restriction, we only exclude subjects who in both weeks always allocate zero tasks to themselves. In both cases, the estimates show no meaningful differences.

¹³An additional two subjects dropped out between week 2 and week 3. These subjects appear not to be different from others based on their allocation tasks, indicating that they did not know or plan to not show up in week 3 when making their week 1 or week 2 decisions. We hence do not drop these subjects from our analysis. All our results, however, are robust to dropping these two subjects.

¹⁴There is some overlap between our exclusion restrictions across the different blocks. One subject is excluded in both SELF and OTHER, leaving us with $n = 95$ subjects when analyzing both blocks jointly. One additional subject each is excluded in both the interpersonal choices and block SELF while three additional subjects are excluded in both the interpersonal choices and block OTHER. This implies that we use data from $n = 67$ subjects when analyzing decisions of all blocks combined.

3 Present Bias and Generosity: Some Theory

The goal of this section is to provide a coherent framework which captures social preferences when payoffs (or consumption, respectively) accrue at different points in time. In the standard discounted expected utility model, when taking the notion of present bias into account (Strotz, 1956; Laibson, 1997; O’Donoghue and Rabin, 1999; Frederick et al., 2002), an agent’s utility at time t can be written as:

$$U = u(c_t) + \beta \sum_{k=1}^T \delta^k u(c_{t+k}) \quad (1)$$

As is well known, if $\beta < 1$, the agent exhibits present bias, meaning that she discounts all future consumption by an additional factor which does not affect the relative discounting between any two future periods, but increases the importance of the present relative to all future periods.

In the settings we are interested in, agents not only decide (and care about) their own consumption, but also about the consumption of other agents. Hence, we also need to model how decision makers evaluate consumption of others. In light of the literature which analyzes decision making for others it seems natural to assume that preferences over own consumption differ from preferences over other people’s consumption.¹⁵ In the most general form of our model, we shall thus allow both the time preference parameters β and δ as well as the atemporal utility function $u(c_t)$ to differ depending on whether own or others’ consumption is evaluated.

Using the specification in (1), however, is only suitable for intrapersonal decisions, i.e., those decisions where there are no trade-offs between own consumption and another person’s consumption. In the altruistic choices we are concerned with here, however, these trade-offs are important and hence need to be properly taken account by the model. The few papers in the relevant literature provides little guidance on what the appropriate model should be. We therefore turn to the literature on multi-attribute utility in the domain of risk and time preferences. Andersen et al. (2018) and Cheung (2015) analyze intertemporal choices under risk and propose a model in which the (concave) intertemporal utility function takes the sum of atemporal utilities as its argument, and a standard expectation operator captures the

¹⁵For example, as discussed in the introduction there is some empirical evidence showing that agents discount very differently when deciding for themselves rather than on behalf of another person (Shapiro, 2010; Albrecht et al., 2011; de Oliveira and Jacobson, 2017). Similarly, in the domain of risky decision making, Andersson et al. (2014) show that agents exhibit lower degrees of loss aversion when deciding for others rather than for themselves.

weights of the different states. While we do not have any risk in our setting, consumption for oneself and consumption for another person can, for the purposes of the modeling approach, be treated analogously to different states of the world. In particular, we can capture the trade-offs between self and other by introducing α and $1 - \alpha$, with $0 \leq \alpha \leq 1$, as weights of own vis à vis others' consumption, and $\rho \geq 1$, which models the concavity of intertemporal utility.¹⁶ This yields to the following specification:

$$U = a \left(u_s(s_t) + \beta_s \sum_{k=1}^T \delta_s^k u_s(s_{t+k}) \right)^\rho + (1 - a) \left(u_o(o_t) + \beta_o \sum_{k=1}^T \delta_o^k u_o(o_{t+k}) \right)^\rho \quad (2)$$

In our setting agents decide about unpleasant consumption, which is why we assume that agents seek to minimize the expression in equation (2). To understand the intuition behind the role of ρ , note that for $\rho = 1$, the discounted utility from own and other's consumption are perfect substitutes (i.e., preferences are linear) but as ρ increases, the agent's desire to smooth consumption between herself and the other person becomes stronger, which increases equality between individuals at the expense of reduced efficiency.¹⁷ Moreover, the atemporal utility functions are best understood as cost functions capturing the disutility of exerting effort in our transcription task. We therefore assume that $u_s(\cdot)$ and $u_o(\cdot)$ are increasing and weakly convex. Note that if $u_i(\cdot)$ is linear and all consumption takes place in period t , this formulation is analogous to the constant elasticity of substitution (CES) functional form used by, for example, Andreoni and Miller (2002) and Fisman et al. (2007).¹⁸

In the dictator games of our experiment, agents allocate consumption between themselves and another person according to the budget constraint $s_{t,\tau} + Ro_{t,\tau} = m$. Since participants decide about a given allocation at time t at different points in time, the subscript τ indicates the period of decision. In a static framework where decisions have only immediate consequences, $\tau = t$. In the following, we allow for $\tau \leq t$ but maintain the assumption that

¹⁶Hence, compared to the specifications in Andersen et al. (2018) and Cheung (2015) for intertemporal choice behavior under risk, α and $1 - \alpha$ can be understood as analogous to states of the world which realize with probability p and $1 - p$, respectively. ρ is the analogous of a coefficient of relative intertemporal risk aversion, as it captures how consumption is smoothed between oneself and another person.

¹⁷To make the role of ρ precise, consider the case where $u_s(\cdot) = u_o(\cdot)$ and $\alpha = 0.5$, i.e., an agent who cares about her own workload exactly as much as about another person's workload. For $\rho = 1$, this agent is indifferent between the effort allocations $\{(s_t, s_{t+k}), (o_t, o_{t+k})\} = \{(10, 20), (40, 30)\}$ and $\{(s_t, s_{t+k}), (o_t, o_{t+k})\} = \{(40, 20), (10, 30)\}$. For $\rho > 1$, however, the agent prefers the second allocation because it allocates the work more equally across the two people.

¹⁸The formulation in (2) improves upon the specification proposed by Shapiro (2010) and Rodriguez-Lara and Ponti (2017) who simply use different weights for the discounted utility of own consumption and others' consumption, respectively. This restricts social preferences to be linear in the sums of discounted utility. Allowing for $\rho \geq 1$ can, thus, account for a broader class of social preferences.

consumption realizes at the same time for both agents (we discuss the asymmetric cases, i.e., cases in which own and other's consumption realize in different points in time, in more detail in Section 4.2). This is the most natural deviation from the static case, that also fits many of the real-world examples discussed in the introduction. For example, when agreeing to help a colleague with some future task, both the costs for oneself and the benefit for the other person accrue at the same time in the future. This leads to the following first-order condition:

$$\left(\frac{u_s(s_{t,\tau})}{u_o(o_{t,\tau})}\right)^{\rho-1} \frac{u'_s(s_{t,\tau})}{u'_o(o_{t,\tau})} = \frac{1}{R} \left(\frac{1}{\tilde{\beta}^{1\{t \neq \tau\}} \tilde{\delta}^{t-\tau}}\right)^\rho \frac{1-a}{a} \quad (3)$$

where $\tilde{\beta} = \frac{\beta_s}{\beta_o}$ and $\tilde{\delta} = \frac{\delta_s}{\delta_o}$ represent relative present bias and relative long-term discounting, respectively. We first note that if agents discount own and other's consumption to the same extent, i.e., if $\tilde{\beta} = 1$ and $\tilde{\delta} = 1$, any form of discounting only leads to a re-scaling of utility, making it irrelevant when deciding about optimal allocations. Intuitively, in this case discounting affects own and other's consumption in the same way, leaving relative preferences between the two unchanged.

To understand how changes in the timing affect the allocations when either $\tilde{\beta} \neq 1$ or $\tilde{\delta} \neq 1$ (or both), note first that the right hand side of equation (3) is decreasing in $\tilde{\beta}^{1\{t \neq \tau\}} \tilde{\delta}^{t-\tau}$, whereas the left hand side is—due to the (weak) convexity of $u_s(\cdot)$ and $u_o(\cdot)$ —increasing in $s_{t,\tau}$. Hence, compared to the static case ($\tau = t$), generosity increases or decreases, depending on whether $\tilde{\beta}^{1\{t \neq \tau\}} \tilde{\delta}^{t-\tau}$ is smaller or greater than one.¹⁹ To illustrate this, consider an agent who does not exhibit any relative present bias, i.e., $\tilde{\beta} = 1$, but may discount own and other's consumption differently in the long-run. In this case, delaying the consequences of the allocation decision to the future increases (if $\tilde{\delta} < 1$) or decreases (if $\tilde{\delta} > 1$) generosity at a constant rate, i.e., in a time-consistent manner.

On the contrary, if an agent is more or less present-biased when discounting own compared to other's consumption, but there are no differences in long-run discounting, i.e., $\tilde{\delta} = 1$, $\tilde{\beta} \neq 1$, then the change in generosity from delaying consumption by one period depends on whether this delay affects present or only future consumption. To illustrate this point, assume that there are two decision periods τ and $\tau + 1$, in which the agent decides about the allocation of consumption in periods t and $t + 1$. It follows that when deciding about relative consumption for oneself and another person to be realized in period t , if $\tilde{\beta} < 1$, generosity is larger when t is in the future (decision is made at time $\tau < t$), compared to when it is in the present (decision at time $\tau + 1 = t$). If, however, the same decisions are made for consumption to

¹⁹This result does not depend on whether we define an increase in generosity as an increase in $o_{t,\tau}$ or as a decrease in $s_{t,\tau}$.

be realized in period $t + 1$, generosity is unchanged because at both τ and $\tau + 1$ decisions only affect future consumption, and hence $\tilde{\beta}$ plays no role. As a consequence, generosity decreases for decisions that have immediate consequences, leading to time inconsistency in generosity as we move the periods of decision closer to the period of consumption. For $\tilde{\beta} > 1$ the effect is reversed.

Finally, when both $\tilde{\delta} \neq 1$ and $\tilde{\beta} \neq 1$, the effects described above are amplified or mitigated, depending on whether $\tilde{\beta}$ and $\tilde{\delta}$ point into the same or into opposite directions. Which of these effects is most relevant is ultimately an empirical question we will test with our data.

4 Effort Allocation in Interpersonal Choices

In this section, we present the results from the interpersonal choices in which decision makers have to allocate effort between themselves and another person, i.e., those decisions that can be considered generalized versions of dictator games. We start by analyzing the symmetric dictator games in blocks SOONSOON and LATELATE to investigate whether generosity is time-inconsistent. These blocks further allow for identification of a relative present bias as defined in Section 3. We then complement this analysis by incorporating the results from the asymmetric dictator games in blocks SOONLATE and LATESOON, because we can use them to identify concrete values for the discounting parameters β_s , β_o , δ_s and δ_o , rather than only their relative magnitudes.

Before analyzing the effects of timing on generosity, however, we briefly relate the overall level of generosity displayed by our subjects to the existing evidence on altruistic behavior. This is particularly interesting since we use effort rather than money allocations as in most previous dictator games, and so far there are only very few studies that have studied altruistic behavior in non-monetary domains (for exceptions see Noussair and Stoop, 2015; Danilov and Vogelsang, 2016). In a meta study of 131 standard monetary dictator games, Engel (2011, p. 607) reports that around 36% of the people give nothing, and that among those who give a positive amount to the receiver, the average amount given is 43% of the pie. The most comparable benchmark from our data is the case where consequences for both the dictator and the recipient are immediate, that is week 2 in SOONSOON, and $R = 1$. Using our whole sample, we find that 36% of our subjects allocate zero tasks to themselves. Among those who are not completely selfish, subjects allocate on average around 33% of the tasks to themselves. Hence, we find that while the fraction of completely selfish people is very similar across domains, conditional on giving, generosity in effort is somewhat weaker than




Figure 3: Effort allocations in symmetric dictator games ($n = 71$)

in the monetary domain.

4.1 Symmetric Dictator Games

Our main result is well summarized by Figure 3. It shows for each task rate R the number of tasks allocated to oneself. The left panel shows allocation decisions for block SOONSOON and the right panel shows the same data for block LATELATE. In both cases, we distinguish between initial allocation decisions made in week 1 (solid line with squares) and subsequent allocation decisions made in week 2 (dashed line with circles). Bars indicate standard errors of the mean.

As is apparent from Figure 3, all four lines are downward sloping, indicating that subjects' choices follow a basic law of demand: as R increases, it becomes “cheaper” to allocate more tasks to the other person. For example, in SOONSOON in week 1, at a task rate of of $R = 0.5$ participants allocate on average 25.93 tasks to themselves compared to 9.80 tasks when $R = 2$. Overall, we find that 92 (93) percent of choices in SOONSOON (LATELATE) are monotonically decreasing in R , suggesting that subjects understood our allocation environment.²⁰

²⁰At the individual level, in block SOONSOON (LATELATE), we find that 56 (63) percent of subjects do not exhibit any violations of monotonicity, and 27 (14) percent only violate monotonicity once. Furthermore,

Most importantly, as can be seen from the left panel of Figure 3, in block SOONSOON we find a large and significant difference between initial allocations made in week 1 and subsequent allocations made in week 2. The average number of tasks allocated to oneself decreases by 15.7% when work needs to be completed immediately (from 16.88 to 14.24; t-test, $p < 0.001$), indicating that generosity decreases when consequences are immediate. The left panel of Table 3 shows that this difference is statistically significant for each single task rate, except for $R = 2$. Recall from Section 3, that this result implies that for our subjects $\tilde{\beta}\tilde{\delta} < 1$.

We now consider the data from block LATELATE in order to investigate whether the decrease in generosity is due to differences in relative long-term discounting, i.e., $\tilde{\delta} < 1$, or driven by a relative present bias, i.e., $\tilde{\beta} < 1$. Our results support the latter. For LATELATE, we only find a (weakly significant) decrease in generosity by 5.6% (week 1: 15.28, week 2: 14.44; t-test, $p = 0.094$). As revealed by the right panel of Table 3, only for rates $R < 1$ this difference is significant at the 5%-level. Overall, this suggests that there is only weak evidence for relative differences in long-term discounting $\tilde{\delta}$. Consistent with this interpretation, the *difference-in-difference*, i.e., the difference between initial and subsequent allocation decisions between SOONSOON and LATELATE is large and significant, amounting to 10.1% or 1.80 tasks (t-test, $p = 0.015$). We thus find a much larger decrease in generosity when the decision in week 2 has immediate consequences (block SOONSOON) compared to when effort only needs to be exerted in the future (block LATELATE). These results provide strong indication that $\tilde{\beta}$ is significantly smaller than 1, both statistically and economically.

In order to quantify the values of $\tilde{\beta}$ and $\tilde{\delta}$, we estimate the preference parameters structurally. To do this, we revisit the first-order condition from equation (3) in Section 3. Close inspection of this expression reveals that from our dictator game data alone, we cannot separately identify the value of ρ from the atemporal utility functions $u_s(\cdot)$ and $u_o(\cdot)$. As we will show later in Section 6, this can, however, be done by combining the interpersonal decisions with the intrapersonal choices. But since the stability of time preferences between situations with and without interpersonal trade-offs is at the core of our paper, we first proceed by estimating time preferences separately for the two types of decisions. To do this, we make the simplifying assumption that the atemporal utility functions, or in this case the cost of effort functions are linear, i.e., $u_s(s_{t,\tau}) = s_{t,\tau}$ and $u_o(o_{t,\tau}) = o_{t,\tau}$. In Section 6, we evaluate whether this simplification leads to any systematic bias in our estimates. Foreshadowing our results from this robustness check, we find that this linearity assumption does not significantly affect our estimates, neither qualitatively nor quantitatively. The first order condition

deviations from monotonicity are typically very small with a median required allocation change of one task to restore monotonicity (see Table A1 in Appendix A for further details).

Rate R	SOONSOON ($n = 71$)			LATELATE ($n = 71$)			Diff-in-diff [t-test]
	$\tau = 1$ Task self	$\tau = 2$ Task self	t-test	$\tau = 1$ Task self	$\tau = 2$ Task self	t-test	
0.5	25.93 (10.29)	21.65 (10.46)	$p < 0.001$	23.87 (11.44)	22.14 (11.21)	$p = 0.035$	2.55 [$p = 0.049$]
0.75	21.83 (10.04)	17.99 (9.77)	$p = 0.001$	19.82 (10.77)	18.41 (10.56)	$p = 0.028$	2.34 [$p = 0.062$]
1	17.51 (9.32)	15.06 (8.67)	$p = 0.002$	16.04 (9.00)	14.87 (8.61)	$p = 0.084$	1.28 [$p = 0.089$]
1.25	13.96 (9.12)	11.92 (8.46)	$p = 0.002$	12.58 (8.61)	12.30 (8.23)	$p = 0.626$	1.76 [$p = 0.003$]
1.5	12.25 (9.05)	10.46 (8.13)	$p = 0.022$	10.85 (8.68)	10.52 (8.01)	$p = 0.580$	1.46 [$p = 0.061$]
2	9.80 (8.55)	8.37 (7.56)	$p = 0.111$	8.42 (780)	8.37 (6.98)	$p = 0.915$	1.38 [$p = 0.105$]
Overall	16.88 (10.90)	14.24 (9.94)	$p < 0.001$	15.28 (10.82)	14.43 (10.15)	$p = 0.094$	1.80 [$p = 0.015$]

Note: The table denotes the number of tasks allocated to oneself, separately for block SOONSOON (left panel) and block LATELATE (right panel). The p-values reported stem from t-tests with standard errors clustered at the individual level. The last column shows the difference-in-difference across week 1 and week 2 allocations between block SOONSOON and LATELATE.

Table 3: Symmetric dictator games: Aggregate behavior by task rate

can then be written as:

$$\frac{s_{t,\tau} + \omega}{o_{t,\tau} + \omega} = \left(\frac{1}{R} \left(\tilde{\beta} \mathbf{1}_{\{t \neq \tau\}} \tilde{\delta}^{t-\tau} \right)^{-\rho} \frac{1-a}{a} \right)^{\frac{1}{\rho-1}} \quad (4)$$

Also note that we add ω to the allocations for oneself and the other person, which can be interpreted as “background consumption”. This is relevant in our setting since subjects in each period have to complete the minimum work requirement of 10 tasks in addition to their allocated tasks, and subjects might take these into account when choosing their optimal allocation.

We present two different approaches that allow estimation of the parameters. In the first approach (“FOC”), we broadly follow Augenblick et al. (2015) and Andreoni and Sprenger (2012) and log-linearize the first-order condition to obtain:

$$\ln \left(\frac{s_{t,\tau} + \omega}{o_{t,\tau} + \omega} \right) = \ln(A) - \sigma \ln(R) - (\sigma + 1) \left[\ln \left(\tilde{\beta} \tilde{\delta} \right) \mathbf{1}_{\{t - \tau = 1\}} + \ln \left(\tilde{\beta} \tilde{\delta}^2 \right) \mathbf{1}_{\{t - \tau = 2\}} \right] \quad (5)$$

where we define $\sigma = \frac{1}{\rho-1}$ as the elasticity of substitution. $A = \left(\frac{1-a}{a} \right)^{\frac{1}{\rho-1}}$ describes a basic

measure of generosity in the sense that it corresponds to the ratio of tasks allocated to self and other when consequences are immediate and $R = 1$. From equation (5) it becomes apparent that we have obtained an expression that is linear in the parameters of interest. In particular, we can identify $\tilde{\beta}$ and $\tilde{\delta}$ from the coefficients of the two dummy variables indicating the difference between the period of decision and the period in which work has to be completed. We estimate this specification via two-limit Tobit by assuming that choices are made with some normally distributed error. We set $\omega = 10$ which corresponds to the minimum work requirement of 10 tasks in each week, which avoids the natural logarithm to be undefined for corner solutions. The exact details of the identification of the parameters and how we recover them from the regression coefficients can be found in Appendix B.

The second approach ("CFS") is based on a closed-form solution $s_{t,\tau}$, which is obtained as:

$$s_{t,\tau} = \frac{R^{-\sigma-1} \left[\tilde{\beta}^{\mathbf{1}\{t \neq \tau\}} \tilde{\delta}^{t-\tau} \right]^{-\sigma-1} + \omega \left(R^{-\sigma} \left[\tilde{\beta}^{\mathbf{1}\{t \neq \tau\}} \tilde{\delta}^{t-\tau} \right]^{-\sigma-1} - A^{-1} \right)}{A^{-1} + R^{-\sigma-1} \left[\tilde{\beta}^{\mathbf{1}\{t \neq \tau\}} \tilde{\delta}^{t-\tau} \right]^{-\sigma-1}} m \quad (6)$$

This specification can be estimated with two-limit Tobit maximum likelihood methods and has the advantage that we can estimate it for $\omega = 10$ and $\omega = 0$. Hence, this helps us to investigate the robustness of our estimates with respect to different estimation techniques as well as with regard to whether participants take the minimum work requirement into account when allocating tasks.

The estimation results can be found in Table 4 and confirm our reduced-form findings from above. Our estimates for relative present bias, $\tilde{\beta}$, range from 0.837 to 0.874, all significantly lower than one (all $p < 0.008$). The degree of relative weekly discounting, $\tilde{\delta}$, instead, is close to, and not significantly different from, one (all $p > 0.365$).²¹ We also find a relatively low elasticity of substitution, especially for the cases where we set $\omega = 10$. This indicates a substantial desire of subjects to smooth consumption between themselves and others, even if one option is relatively cheaper than the other. The value of A indicates that in a "standard" dictator game where consequences are immediate, our subjects allocate on

²¹We note that there is a slight inconsistency in the structural estimates for $\tilde{\delta}$ with our reduced-form results from above. While the former are (not significantly) larger than one, the latter indicate (weakly significant) evidence for $\tilde{\delta} < 1$. This is due to the fact that overall allocations are more generous in block SOONSOON than in LATELATE, which does not impact our "diff-in-diff" in the reduced-form analysis, but affects the structural estimates. Note, however, that identification of relative present bias does not rely on the social preference parameters to be identical for consumption in weeks 2 and 3. In particular, we can allow for the relative weight of own consumption a , to be different in weeks 2 and 3. In Table A2 in Appendix A we present the results from such an exercise which delivers estimates for $\tilde{\delta}$ which are below, but not significantly different from one, and leaves the estimates for $\tilde{\beta}$ virtually unchanged.

	(1) FOC $\omega = 10$	(2) CFS $\omega = 10$	(3) CFS $\omega = 0$
$\sigma = \frac{1}{\rho-1}$	0.081 (0.086)	0.014 (0.075)	0.201 (0.124)
$A = \left(\frac{1-a}{a}\right)^{\frac{1}{\rho-1}}$	0.491 (0.038)	0.513 (0.038)	0.369 (0.046)
$\tilde{\delta}$	1.040 (0.044)	1.034 (0.043)	1.040 (0.057)
$\tilde{\beta}$	0.873 (0.046)	0.874 (0.046)	0.837 (0.059)
Observations	1704	1704	1704
Participants	71	71	71
$H_0(\hat{\delta} = 1)$	$p = 0.366$	$p = 0.434$	$p = 0.481$
$H_0(\hat{\beta} = 1)$	$p = 0.006$	$p = 0.007$	$p = 0.006$

Note: The table reports the parameter estimates for the symmetric dictator games (blocks SOONSOON and LATELATE). Column (1) uses the log-linearized first order condition, while columns (2) and (3) use the closed form solution for the number of tasks allocated to oneself. Standard errors are clustered at the individual level and calculated via the delta method.

Table 4: Parameter estimates for symmetric dictator games

average about twice as many tasks to the other person than to themselves.²²

In summary, both the reduced-form as well as the structural estimates reveal strong evidence for differences in relative present bias, leading to time-inconsistent generosity. However, as pointed out previously, while the symmetric dictator games constitute a natural starting point for our analysis, we cannot make any statements about whether the decrease in generosity is due to a present bias for own consumption, or whether it is driven by a future bias for consumption of the other person (or a combination of both). In order to investigate this, in the following, we include the data from the asymmetric dictator games into our analysis, which allows for estimation of β_s , β_o , δ_s and δ_o .

4.2 Asymmetric Dictator Games

As in the previous section, before presenting the results from our structural estimation, we first describe the data and perform some reduced-form analysis. Analogous to Figure 3, Figure 4 shows for each task rate R the amount of tasks allocated to oneself in week 1 and week 2. The left panel shows allocation decisions for block SOONLATE and the right panel

²²We should point out here again that these estimates exclude subjects without any variation in their task allocations in at least one of the weeks. Since this restriction by and large only excludes subjects who behave perfectly selfish, our estimates for generosity are biased upwards.




Figure 4: Effort allocations in asymmetric dictator games ($n = 71$)

shows the same data for block LATESOON. The results reveal that for the case where the decision maker needs to exert effort at the sooner date and the recipient at the later date (SOONLATE), we see a small decrease for all six relative prices of giving. In week 1, agents allocate on average 15.41 tasks to themselves, compared to 14.69 tasks in week 2 (-5%). This decrease, however, does not reach statistical significance (t-test, $p = 0.272$). For the treatment LATESOON, where the timing of effort exertion is reversed, we obtain virtually no difference in allocation decision between weeks 1 and 2 (14.73 vs. 14.69; t-test, $p = 0.945$).

What do these effects tell us about our relative present bias, and, more specifically, about the magnitude of our coefficients of interest? In order to provide some intuition, we consider the first-order conditions for the two blocks. For SOONLATE we obtain:

$$\frac{s_{t,\tau} + \omega}{o_{t+1,\tau} + \omega} = \left(\frac{1}{R} (\beta_o \delta_o)^\rho \left(\frac{\delta_o}{\beta_s \delta_s} \right)^{\rho \cdot \mathbf{1}_{\{t \neq \tau\}}} \frac{1-a}{a} \right)^{\frac{1}{\rho-1}} \quad (7)$$

Equation (7) reveals that any differences in allocations between week 1 and week 2 can be accounted for by $\frac{\delta_o}{\beta_s \delta_s} \neq 1$. In particular, a decrease in tasks allocated to oneself from week 1 to week 2 is consistent with $\frac{\delta_o}{\beta_s \delta_s} > 1$. A similar exercise for LATESOON yields:

$$\frac{s_{t+1,\tau} + \omega}{o_{t,\tau} + \omega} = \left(\frac{1}{R} \left(\frac{1}{\beta_s \delta_s} \right)^\rho \left(\frac{\beta_o \delta_o}{\delta_s} \right)^{\rho \cdot \mathbf{1}\{t \neq \tau\}} \frac{1-a}{a} \right)^{\frac{1}{\rho-1}} \quad (8)$$

Hence, a decrease in tasks allocated to oneself when moving from week 1 to week 2 is consistent with $\frac{\beta_o \delta_o}{\delta_s} > 1$.

What becomes apparent from these considerations is that, without further assumptions, differences in allocations across weeks are not easily interpretable regarding their implications for subjects' time-preference parameters. If we are willing to assume that $\tilde{\delta} \approx 1$, which is in line with our previous results, we notice that the results from the asymmetric treatments are—at least directionally—consistent with an interpretation that the present bias found in the symmetric dictator games is due to $\beta_s < 1$, rather than $\beta_o > 1$. The decrease in SOONLATE indicates some present bias for own consumption, while the absence of any effect in LATESOON suggests that $\beta_o \approx 1$. Moreover, under the assumption that there is little difference in relative long-term discounting, the sum of the two decreases would correspond to a measure of relative present bias which is indeed in line with our previous findings.

A more compelling approach, however, is to combine the data from both the symmetric and the asymmetric dictator games, which allows us to directly estimate all parameters of interest, $\beta_s, \beta_o, \delta_s, \delta_o$. To this end, we again apply two different estimation approaches, one based on the first-order condition and the other based on the closed form solution. In both cases, the econometric specifications are very similar to the ones presented in the previous section. For the approach based on the closed form solution for effort allocated to oneself, we simply augment the log-likelihood function with the additional data. For the log-linearized first-order condition, we impose two linear constraints as to render the parameter just identified. The details of these procedures can be found in Appendix B.

The results from these estimations are presented in Table 5. The main finding is that we identify a present bias coefficient for own consumption, β_s , which is significantly lower than one. Depending on the specification, the actual estimate varies between 0.883 and 0.910 (all $p < 0.002$). We do not find any evidence for present bias in others' consumption. The estimated value for β_o is between 1.044 and 1.060, but not significantly different from one (all $p > 0.257$). Taken together, these results corroborate the findings from the symmetric dictator games as we can reject the hypothesis that $\tilde{\beta} = 1$, in favor of $\tilde{\beta} < 1$ (all $p < 0.016$). Furthermore, in line with the results from above, we cannot reject the hypothesis that $\tilde{\delta} = 1$ (all $p > 0.387$), indicating that there are no differences for long-run discounting.

In summary, the results from this section reveal that generosity is dynamically inconsistent. Subjects behave more altruistically towards others when deciding in advance rather

	(1) FOC $\omega = 10$	(2) CFS $\omega = 10$	(3) CFS $\omega = 0$
$\sigma = \frac{1}{\rho-1}$	0.067 (0.088)	-0.000 (0.076)	0.185 (0.125)
$A = \left(\frac{1-a}{a}\right)^{\frac{1}{\rho-1}}$	0.486 (0.038)	0.509 (0.038)	0.365 (0.045)
δ_s	1.048 (0.031)	1.046 (0.031)	1.056 (0.041)
β_s	0.910 (0.027)	0.910 (0.027)	0.883 (0.036)
δ_o	1.001 (0.027)	1.005 (0.027)	1.006 (0.035)
β_o	1.044 (0.040)	1.043 (0.040)	1.060 (0.053)
Observations	3408	3408	3408
Participants	71	71	71
$H_0(\hat{\beta}_s = 1)$	$p < 0.001$	$p = 0.001$	$p = 0.001$
$H_0(\hat{\beta}_o = 1)$	$p = 0.272$	$p = 0.277$	$p = 0.258$
$H_0(\hat{\beta}_s = \hat{\beta}_o)$	$p = 0.013$	$p = 0.014$	$p = 0.015$

Note: The table reports the parameter estimates from all dictator games (blocks SOONSOON, SOONLATE, LATELATE, and LATESOON). Column (1) uses the log-linearized first order condition, while columns (2) and (3) use the closed form solution for the number of tasks allocated to oneself. Standard errors are clustered at the individual level and calculated via the delta method.

Table 5: Parameter estimates from all dictator games

than in the present, while no such difference is observed when choices only affect the future. By disentangling discounting of own consumption from discounting of others' consumption, we show that only the former is subject to present bias while the latter is discounted in a time-consistent manner. As such, our results reveal that present bias in own consumption is not limited to individual decision contexts as studied in most of the previous literature, but also applies to social contexts in which there are trade-offs between own and others' consumption.

5 Present Bias across Individual and Social Contexts

In this section, we investigate the extent to which present bias (and the lack thereof) is correlated within individuals across individual and social contexts. A positive correlation would suggest that there is a stable underlying trait determining the degree to which individuals can resist the temptation of immediate gratification, irrespective of whether the

consequences of this have to be borne by the own future self or another person. The lack of any correlation, in contrast, would question the often made assumption that choices across different contexts are guided by some stable underlying primitives.

Before we present this analysis, however, we briefly describe choices made in the two intrapersonal blocks SELF and OTHER. This allows us to evaluate whether, at the aggregate level, the observed differences in present bias in the interpersonal choices translate into decision contexts without any interpersonal trade-offs. Further, this allows us to investigate differences in intertemporal allocation decisions made for oneself or on behalf of another person, a question previous literature has reported mixed evidence on (Shapiro, 2010; Albrecht et al., 2011; Howard, 2013; de Oliveira and Jacobson, 2017). After that, for each individual, we structurally estimate time preference parameters separately for the interpersonal and intrapersonal choices. We then compare the relationship between present bias across these two contexts at the individual level.

5.1 Aggregate Analysis

Formally, in block SELF, an individual chooses in periods $\tau \in \{1, 2\}$ (corresponding to week 1 and week 2) how many tasks to complete in periods $t = 2$ and $t = 3$. Following Augenblick et al. (2015), the optimal effort choices, denoted by $s_{t,\tau}$ and $s_{t+1,\tau}$, respectively, are found by minimizing

$$\beta_s^{1_{\{t \neq \tau\}}} \delta_s^{t-\tau} (s_{t,\tau} + \omega)^{\gamma_s} + \beta_s \delta_s^{t+1-\tau} (s_{t+1,\tau} + \omega)^{\gamma_s} \quad (9)$$

subject to the budget constraint $s_{t,\tau} + R s_{t+1,\tau} = m$. In the notation of Section 3, we have thus parameterized the atemporal cost of effort function as $u_s(s_{t,\tau}) = s_{t,\tau}^{\gamma_s}$. γ_s denotes the curvature of this function, i.e., the larger γ_s , the larger the agent's preference for smoothing consumption over the two periods. As before, δ_s represents long-term (exponential) discounting whereas present bias is captured by β_s . The first-order condition is given by:

$$\frac{s_{t,\tau} + \omega}{s_{t+1,\tau} + \omega} = \left(\frac{\beta_s^{1_{t=\tau}} \delta_s}{R} \right)^{\frac{1}{\gamma_s - 1}} \quad (10)$$

This implies that if $\beta_s < 1$, the agent allocates more tasks to the sooner date ($t = 2$) when she decides in advance ($\tau = 1$) rather than in the present ($\tau = 2$).

Our results are summarized by Figure 5. It depicts for each week and task rate the number of tasks allocated to the sooner work date in week 2. As can be seen from the left panel of Figure 5, we observe a systematic downward shift in the number of tasks




Figure 5: Effort allocations in intrapersonal decisions (SELF: $n = 100$, OTHER: $n = 98$)

allocated to the sooner date in week 2 compared to week 1. On average, subjects allocate 1.48 fewer tasks to the sooner work date when it is the present (-6.1%, 24.13 compared to 22.65, $p = 0.004$), indicating a significant and economically meaningful present bias for own consumption. These results are further corroborated by the left panel of Table 6, showing the number of tasks allocated to the sooner work date separately for each R . It also reveals that there is very little evidence for long-term discounting. This is most clearly seen for $R = 1$. In this case, subjects in week 1 allocate on average 25.86 tasks (or 51.7%) to the sooner date, thus splitting the workload almost evenly across weeks.

In order to estimate the time-preference parameters from these choices structurally, we can rely on the two different estimation approaches discussed in Section 4, as the first-order conditions have a very similar structure as the ones from the dictator games. The first approach is based on the log-linearization of the first-order condition ("FOC") in (10). The second approach uses the closed form solution for effort allocated to the sooner date ("CFS"), given by:

$$s_{t,\tau} = \frac{R^{-\frac{\gamma_s}{\gamma_s-1}} \left[\beta_s^{1\{t=\tau\}} \delta_s \right]^{\frac{1}{\gamma_s-1}} + \omega \left(R^{-\frac{1}{\gamma_s-1}} \left[\beta_s^{1\{t=\tau\}} \delta_s \right]^{\frac{1}{\gamma_s-1}} - 1 \right)}{1 + R^{-\frac{\gamma_s}{\gamma_s-1}} \left[\beta_s^{1\{t=\tau\}} \delta_s \right]^{\frac{1}{\gamma_s-1}}} m \quad (11)$$

Rate R	SELF ($n = 100$)			OTHER ($n = 98$)		
	$\tau = 1$ Tasks soon	$\tau = 2$ Tasks soon	t-test	$\tau = 1$ Tasks soon	$\tau = 2$ Tasks soon	t-test
0.5	37.84 (8.52)	35.71 (9.29)	$p = 0.008$	34.87 (12.36)	34.24 (10.26)	$p = 0.544$
0.75	33.31 (9.55)	31.26 (9.84)	$p = 0.018$	31.08 (11.73)	31.29 (10.10)	$p = 0.838$
1	25.86 (6.92)	24.07 (6.81)	$p = 0.031$	25.60 (7.20)	25.21 (6.19)	$p = 0.608$
1.25	18.58 (10.16)	17.16 (10.03)	$p = 0.037$	19.51 (11.87)	19.00 (10.93)	$p = 0.581$
1.5	15.58 (10.50)	15.06 (9.83)	$p = 0.446$	17.38 (12.43)	16.63 (11.07)	$p = 0.356$
2	13.62 (10.84)	12.66 (9.84)	$p = 0.173$	15.22 (12.79)	14.06 (11.06)	$p = 0.168$
Overall	24.13 (13.09)	22.65 (12.61)	$p = 0.004$	23.94 (13.58)	23.14 (12.52)	$p = 0.252$

Note: The table denotes the number of tasks allocated to the sooner date, separately for block SELF (left panel) and block OTHER (right panel). For each rate R , the p-value reported stems from a t-test with standard errors clustered at the individual level.

Table 6: Intrapersonal decisions: Aggregate behavior by task rate

which we estimate by two-limit Tobit maximum-likelihood. Further details on the estimation approach can be found in Appendix B.

The results of our estimations are shown in the left panel of Table 7. In line with our reduced-form results from above, the results reveal strong and significant evidence for present bias in own consumption. The estimates of β_s vary between 0.842 and 0.863 across specifications, and are always significantly lower than one (all $p < 0.006$). We find no evidence for long-term discounting; the weekly discount rate δ_s varies between 1.023 and 1.046, but it is never significantly different from 1 (all $p > 0.387$). Taken together, these results reveal that, at the aggregate level, present bias in own consumption is a robust phenomenon across individual and social contexts.

Given the similarity of our block SELF design to the one used in Augenblick et al. (2015), it is sensible to compare the findings of both studies, in particular as there are a few notable differences across the two studies. First of all, while in Augenblick et al. (2015) initial allocations were made in the lab and subsequent allocations were made online, all our allocations decisions took place in the the same lab at exactly the same time of the same day of the week. Furthermore, the encryption task we use is slightly different from theirs (they additionally use Tetris as a second, arguably more fun, real-effort task). Despite these differences, the

	SELF ($j = s$)			OTHER ($j = o$)		
	(1)	(2)	(3)	(4)	(5)	(6)
	FOC	CFS	CFS	FOC	CFS	CFS
	$\omega = 10$	$\omega = 10$	$\omega = 0$	$\omega = 10$	$\omega = 10$	$\omega = 0$
γ_j	2.284 (0.256)	2.667 (0.402)	2.083 (0.277)	2.748 (0.551)	3.534 (1.050)	2.688 (0.726)
δ_j	1.045 (0.052)	1.046 (0.063)	1.023 (0.059)	0.989 (0.055)	0.991 (0.074)	0.967 (0.069)
β_j	0.863 (0.045)	0.842 (0.056)	0.844 (0.055)	0.931 (0.059)	0.912 (0.078)	0.919 (0.076)
Observations	1200	1200	1200	1176	1176	1176
Participants	100	100	100	98	98	98
$H_0(\hat{\delta}_j = 1)$	$p = 0.388$	$p = 0.464$	$p = 0.692$	$p = 0.850$	$p = 0.901$	$p = 0.623$
$H_0(\hat{\beta}_j = 1)$	$p = 0.003$	$p = 0.005$	$p = 0.005$	$p = 0.245$	$p = 0.259$	$p = 0.282$

Note: The table reports the parameter estimates for the choices made in blocks SELF (left panel) and OTHER (right panel), respectively. Columns (1) and (4) use the log-linearized first order condition, while the other columns use the closed form solution for the number of tasks allocated to the sooner date. Standard errors are clustered at the individual level and calculated via the delta method.

Table 7: Parameter estimates for blocks SELF and OTHER

results from both studies are remarkably similar. Augenblick et al. (2015) estimate a β of 0.888, compared to our β_s estimate of 0.863 (see model (1) in Table 7, which is the approach that Augenblick et al. (2015) use for their structural estimation). The strong similarity of the results suggests that present bias in own non-monetary consumption is a robust finding across different subject pools, experimental procedures, and tasks.²³

We now turn to the analysis of choices made on behalf of someone else in block OTHER. As pointed out in the introduction, there are many situations in which agents have to make intertemporal decisions for others (e.g., asset managers investing on behalf of their clients, doctors choosing treatments for their patients, parents deciding what is best for their children, etc.). These situations are further interesting as they can help to understand some of the underlying principles of present bias. In particular, they can reveal whether when deciding on behalf of others, the desire for immediate gratification is equally strong compared to when deciding for oneself, or whether the greater personal distance mitigates this effect. The latter effect would be consistent with neuro-economic evidence (McClure et al., 2004; Albrecht et al., 2011) that links present bias to the more affective and more impulsive system, compared to a more deliberative and reasoned system which may play a more central

²³Another paper that uses a similar environment is Augenblick and Rabin (2017) where agents choose how many tasks to complete for varying wages and (future) dates. The authors estimate individual present bias to be between 0.81 and 0.84, which is also close to our numbers.

role when discounting others' consumption.

The results from block OTHER are summarized in the right panel of Figure 5 and Table 6. Compared to the choices in block SELF, a somewhat different picture emerges. In particular, the differences between initial allocations in week 1 and subsequent allocations in week 2 are now much less pronounced. On average, subjects allocate 0.54 fewer tasks to the sooner work date when consequences are immediate. This corresponds to a decrease of only 2.2%, which is not statistically significant (week 1: 23.94, week 2: 23.41, $p = 0.252$).

Using the same approach as for block SELF, we corroborate the reduced-form findings by structurally estimating the time preference parameters for others' consumption. As shown in the right panel of Table 7, we find little evidence for intertemporal discounting, neither in the form of present bias, nor in form of long-run discounting. We estimate a β_o between 0.912 and 0.931 and a δ_o ranging from 0.967 to 0.991, none of these estimates are significantly different from one (all $p > 0.244$ and $p > 0.622$, respectively). Hence, in line with our results from the interpersonal choices reported in the previous section, also in the intrapersonal choices we find little evidence for present bias in others' consumption.²⁴

An important general question that arises when analyzing decision-making on behalf of others is to what extent subjects take this seriously. After all, these decisions have no bearing on the number of tasks they have to solve themselves, and thus purely self-interested subjects may have no incentive to make reasonable choices. To investigate this, we analyze the decision quality of choices in block OTHER compared to decisions made in block SELF. In the latter, 92 percent of choices are monotonically decreasing in R and 60 percent of subjects have no monotonicity violation in their effort choices.²⁵ In OTHER, 90 percent of choices are monotonically decreasing in R and 64 percent of subjects are fully consistent. While these numbers suggest a similarly high level of decision quality for decisions in blocks SELF and OTHER, a closer inspection of the data reveals that this is not the case. In particular, as we demonstrate in more detail in Table A1 in Appendix A, conditional on violating monotonicity, the minimum number of tasks that need to be reallocated within a block to bring the data in line with monotonicity is significantly higher in block OTHER than in block SELF (4.26 vs. 0.75; paired t-test, $p = 0.042$). That is, while we find no difference in the likelihood of violating monotonicity, the magnitude of these violations is much larger in

²⁴Note, however, that while our estimates of β_o from the dictator games were slightly above one, the ones obtained from the intrapersonal choices are slightly below one.

²⁵The numbers are comparable to the ones reported in Augenblick et al. (2015) who find 95 percent of effort choices to be monotonically decreasing in R . In addition, we find about 20% of the choices being corner solutions (19% in SELF and 21% in OTHER), which is somewhat lower than the 31% observed in Augenblick et al. (2015) and much lower than the numbers typically observed in monetary discounting (e.g., 70% in Andreoni and Sprenger, 2012 and 86% in Augenblick et al., 2015).

	(1) FOC $\omega = 10$	(2) CFS $\omega = 10$	(3) CFS $\omega = 0$
γ	2.410 (0.309)	2.884 (0.507)	2.236 (0.351)
δ_s	1.025 (0.057)	1.009 (0.066)	0.994 (0.062)
β_s	0.847 (0.053)	0.821 (0.066)	0.825 (0.065)
δ_o	1.008 (0.048)	1.025 (0.064)	0.996 (0.059)
β_o	0.947 (0.048)	0.940 (0.059)	0.943 (0.057)
Observations	2280	2280	2280
Participants	95	95	95
$H_0(\hat{\beta}_s = 1)$	$p = 0.004$	$p = 0.007$	$p = 0.007$
$H_0(\hat{\beta}_o = 1)$	$p = 0.270$	$p = 0.308$	$p = 0.319$
$H_0(\hat{\beta}_s = \hat{\beta}_o)$	$p = 0.087$	$p = 0.106$	$p = 0.098$

Note: The table reports the parameter estimates for the choices made in blocks SELF and OTHER under the restriction that $\gamma_s = \gamma_o = \gamma$. Column (1) uses the log-linearized first order condition, while the columns (2) and (3) use the closed form solution for the number of tasks allocated to the sooner date. The estimation uses the data from those 95 subjects who have sufficient variation in block SELF and block OTHER (see footnote 14). Standard errors are clustered at the individual level and calculated via the delta method.

Table 8: Parameter estimates for blocks SELF and OTHER combined

OTHER compared to SELF. Importantly, this difference is entirely driven by subjects who in the dictator games behave (in at least one week) completely selfishly. For these subjects, we need to reallocate on average 14.95 tasks to restore monotonicity, compared to 1.09 tasks for the non-selfish subjects (two sample t-test, $p = 0.032$). These results suggest that there may be important differences in the decisions made on behalf of others, depending on whether a subject exhibits some degree of other-regarding concerns or not. In particular, our results indicate that one should be cautious with reading too much into decisions made on behalf of others by fully selfish subjects, as decision quality may be low.

Taken together, in line with our findings from the interpersonal choices, we find evidence for stronger present bias in own compared to others' consumption when there are no trade-offs between own and others' consumption. This result is further corroborated when, similar to the analysis of the dictator games, estimating all four discounting parameters jointly. To do so, we constrain the curvature of the cost of effort function to be the same for own and other's consumption ($\gamma = \gamma_s = \gamma_o$). The results from this estimation, shown in Table 8, provide a very similar picture regarding the differences in present bias from above. Specifically, we estimate β_s to be between 0.821 and 0.847 and β_o to be between 0.940 and 0.947. Moreover,

we can use this joint estimation to directly test whether β_s and β_o are the same, and reject this hypothesis at the 10% level.²⁶

5.2 Individual-level Analysis

Our results so far have revealed that, at the aggregate level, there are systematic differences in present bias in own consumption compared to others' consumption, both in interpersonal as well as in intrapersonal choices. However, aggregate analyses may disguise important heterogeneity at the individual level. In particular, the previous findings do not reveal anything about the extent to which present bias in individual and social contexts is correlated within the individual, i.e., whether present bias is a behavioral phenomenon that is stable across contexts. To investigate this, we estimate individual-level discounting parameters separately for each of the two types of decisions.

To estimate individual-level present bias, we use the approach based on the closed-form solution for $s_{t,\tau}$ (see equation (11)), and concentrate on the case with $\omega = 10$. Compared to the log-linearized first-order condition approach, it has the advantage that it allows us to place a restriction on the curvature parameters γ and ρ , which, for the analytic solution to be an interior optimum, need to be larger than one. Following the aggregate analysis, we obtain separate estimates from the dictator games and the intrapersonal choices (combining blocks SELF and OTHER).²⁷ We obtain reasonable individual-level estimates for about 93% of the subjects (intrapersonal choices: 88 out of 95 subjects, dictator games: 66 out of 71 subjects).²⁸ See Appendix C for a more detailed description of our procedures and the full

²⁶Given the results of low decision quality of selfish subjects when deciding on behalf of others from above, as a robustness check, we re-estimate time preference parameters from the intrapersonal decision blocks by excluding selfish subjects that we also remove from the analysis of the interpersonal decisions. The results from the structural estimations can be found in Table A3 in Appendix A. The main result that emerges from this analysis is that the coefficient for β_o gets closer to one (now between 0.971 and 0.977). At the same time, β_s slightly decreases. As a result, we can reject the equality of β_s and β_o with higher confidence than before (all $p < 0.072$). These results suggest that including the choices of subjects without a relevant concern for the well-being of others may have underestimated the observed differences in present bias in our intrapersonal choices.

²⁷For the latter, we jointly estimate the discounting parameters, restricting $\gamma = \gamma_s = \gamma_o$, corresponding to the aggregate estimation presented in Table 8. This reduces the number of parameters to be estimated from a given number of observations, thereby increasing the precision of the estimation.

²⁸The behavior of five subjects in the intrapersonal choices and one subject in the interpersonal choices is fully consistent with utility maximization, but we can only identify bounds on β_s and β_o , i.e., whether they are (weakly) above or below one, because they have insufficient variation across weeks. One subject in the intrapersonal choices displays behavior which is too noisy to yield convergence. For the remaining subjects, following Augenblick and Rabin (2017), we use Grubb's outlier test with a confidence level of 99.99%. For the intrapersonal choices, the test is rejected for three subjects with very large β_o estimates (and very small β_s estimates). For the interpersonal choices we have to remove two subjects, one because of a very high β_s and the other because of a very high β_o . Tables C1 to C4 in Appendix C list the estimates for each subject




Figure 6: Individual estimates for present bias in own and others' consumption from intrapersonal and interpersonal choices

list of individual estimates.

Figure 6 plots the distributions of the individual estimates for β_s and β_o , separately for the intrapersonal and interpersonal choices. It reveals that in all cases there is a big spike around 1 indicating (close to) dynamically consistent discounting behavior, but that there is also pronounced heterogeneity across individuals. Table 9 highlights different moments of these distributions. In line with our aggregate results from above, we find that for intrapersonal choices individuals exhibit a stronger present bias for own compared to others' consumption; the mean β_s is significantly lower than the mean β_o (0.930 vs. 0.990; paired t-test, $p = 0.041$). For the estimates from interpersonal choices, we find a mean β_s of 0.956, which, again, is significantly lower than the 1.012 for β_o (paired t-test, $p = 0.036$).

We are now in a position to test whether present bias is correlated across interpersonal and intrapersonal decision situations. Figure 7 shows this relationship, separately for present bias estimates for own and others' consumption. We find a strong and significant positive correlation for β_s ($\rho = 0.41$, $p < 0.001$), while for β_o the correlation is much lower and not statistically significant ($\rho = 0.11$, $p = 0.371$). The positive correlation for present bias in own consumption indicates that the desire for immediate gratification can be seen as a trait

separately and highlight the excluded cases.

	N	Mean (s.d.)	Proportion present biased ($\beta < 0.99$)	Proportion dynamically consistent ($0.99 \leq \beta \leq 1.01$)	Proportion future biased ($\beta > 1.01$)
β_s^{Intra}	88	0.930 (0.188)	0.523	0.239	0.239
β_o^{Intra}	88	0.990 (0.217)	0.443	0.273	0.284
β_s^{Inter}	66	0.956 (0.143)	0.515	0.167	0.318
β_o^{Inter}	66	1.012 (1.187)	0.364	0.227	0.409

Table 9: Summary statistics of individual-level estimates for β_s and β_o .

that is relatively stable across contexts in which there are interpersonal trade-offs or not.

A very similar pattern can be observed when using the individual-level estimates to classify subjects into different "discounting types", as done in previous empirical studies (see e.g., Ashraf et al., 2006; Meier and Sprenger, 2010). We follow Augenblick et al. (2015) and classify a participant as "present-biased" if her estimated $\beta < 0.99$, as "future-biased" if $\beta > 1.01$, and as "dynamically consistent" otherwise. The distributions of these types are shown in Table 9. In line with the results above, it reveals that, for both intrapersonal and interpersonal choices, more subjects are classified as present-biased when own rather than others' consumption is at stake.²⁹ When using these classifications to analyze stability across contexts, we find that 69% of the subjects who display a present bias in block SELF also display a present bias in own consumption in the dictator games. This implies a correlation of discounting types which is positive and significant across the two contexts ($\rho = 0.28$, $p = 0.030$). For present bias in others' consumption, in contrast, only 38% of subjects classified as present-biased in block OTHER display the same pattern in the interpersonal choices. Compared to present bias in own consumption, the correlation of discounting types across contexts is much weaker and does not reach statistical significance ($\rho = 0.09$, $p = 0.482$).

Finally, to check the validity of our structural estimates, we compare them to a simple reduced-form measure of present bias. For the intrapersonal choices, a direct measure for present bias is the difference between allocations made in week 1 and week 2. For block SELF, the average difference is -1.82 tasks, which is highly correlated with the structural estimates for β_s^{Intra} ($\rho = 0.983$, $p < 0.001$). Similarly, for block OTHER, our direct measure yields -0.75, which is also strongly correlated with our estimates for β_o^{Intra} ($\rho = 0.973$, $p < 0.001$). For the dictator games, the construction of a similar measure for present bias in own and others' consumption is a little less straightforward, since the identification relies on differences-in-

²⁹The results for β_s^{Intra} are again very much in line with those of Augenblick et al. (2015) who find 56% of people being present-biased, compared to 29% who are future-biased.




Figure 7: Correlation of present bias in own and others' consumption across intrapersonal and interpersonal choices. The line indicates a linear fit from a OLS regression

differences. By appropriately combining the differences in allocations between weeks 1 and 2, we obtain two separate measures of present bias in own and others' consumption for each case. We then use the average of the two to obtain our reduced-form measure of present bias.³⁰ Again, we find a high degree of consistency with our structural estimates. For present bias in own consumption, we find a diff-in-diff of -0.64 tasks, whereas for present bias in others' consumption, the corresponding difference is only -0.18 tasks. In both cases, our reduced-form measure is highly correlated with our structural estimates (β_s^{Inter} : $\rho = 0.931, p < 0.001$, β_o^{Inter} : $\rho = 0.961, p < 0.001$). Overall, these results show a very high level of consistency of our structurally estimated parameters. Unsurprisingly, we thus reach a very similar conclusion regarding the stability of present bias when calculating correlations based on our reduced-form measure. For present bias in own consumption, we find a correlation of $\rho = 0.351$ ($p = 0.006$), compared to $\rho = 0.068$ ($p = 0.605$) for present bias in others' consumption.

The stability of present bias in own consumption is remarkable as previous studies have

³⁰More precisely, define Δ_k as the difference between allocations in weeks 1 and 2 for block k , where $k \in \{\text{SOONSOON}, \text{LATELATE}, \text{SOONLATE}, \text{LATESOON}\}$. Based on the first-order conditions in Section 4, for present bias in own consumption, we calculate our measure as the average of $\Delta_{\text{SOONSOON}} - \Delta_{\text{LATESOON}}$ and $\Delta_{\text{SOONLATE}} - \Delta_{\text{LATELATE}}$, and for present bias in others' consumption, we calculate our measure as the average of $\Delta_{\text{LATELATE}} - \Delta_{\text{LATESOON}}$ and $\Delta_{\text{SOONLATE}} - \Delta_{\text{SOONSOON}}$.

shown that experimentally elicited time preferences often lack a strong correlation across different domains (see e.g., Chabris et al., 2008; Augenblick et al., 2015). The fact that similar conclusions do not hold for present bias in others’ consumption reveals that discounting of others’ consumption is more malleable and context-specific. In particular, it shows that the evaluation of others’ consumption streams is very different between social settings in which also own consumption is at stake, and situations without trade-offs between own and others’ consumption. In the former, agents may engage in relative comparisons which may trigger feelings of envy, spite, or guilt, while in the latter, they may base their behavior on what they think is best for the other person. Whether this is what the agent thinks the other person wants, or should want, is an interesting question we will return to in our discussion. One possibility that we can rule out based on our data, however, is that a majority of subjects simply implement their own discounting pattern when choosing for others. Only 8% of subjects reveal $\beta_s = \beta_o$, and for an additional 9% β_s and β_o differ by less than 0.01.

6 Testing the robustness of the structural estimates

In Section 3, we proposed a functional form which allows us to capture intertemporal social preferences. In Section 4, however, we made a simplifying assumption and constrained the curvature of the atemporal utility/cost-of-effort function to be linear. The upside of this assumption was that it allowed us to estimate time preferences for own and others’ consumption separately for the interpersonal and the intrapersonal decisions. Since a-priori there is no guarantee that preferences will be the same, we considered it a useful empirical exercise to test the stability of time preferences across these two decision contexts. The downside, however, is that by essentially neglecting the fact that there is substantial curvature in the cost-of-effort function—as is evident from the estimates for γ obtained above (compare Table 7 and 8)—we may have produced biased estimates for the time preferences estimated from the dictator games. To test for this possibility, as a robustness check, in this section we provide results from an estimation approach in which we estimate time preferences using data from all decision blocks, without imposing any linearity restriction on the utility function. We show that all our previous main results hold, indicating that such a bias is small.

In order to estimate the parameters of the utility specification in equation (2) (see Section 3), we rely on the estimation approach based on the closed-form solution for $s_{t,\tau}$ (or $o_{t,\tau}$, for decisions in block OTHER, respectively).³¹ In line with the theory, this estimation imposes

³¹An estimation using the log-linearized first-order condition is not feasible because this would require non-linear constraints on the parameters (see Appendix B for details). For the interpersonal decisions, however, such a closed-form solution only exists if we constrain the parameter γ , which measures the curvature of the

	(1) CFS $\omega = 10$	(2) CFS $\omega = 0$		(3) FOC $\omega = 10$	(4) CFS $\omega = 10$	(5) CFS $\omega = 0$
$\sigma = \frac{1}{\rho-1}$	-0.088 (0.195)	0.340 (0.318)		0.096 (0.219)	-0.088 (0.197)	0.349 (0.324)
$\tilde{A} = \left(\frac{1-a}{a}\right)^{\frac{1}{\gamma\rho-1}}$	0.549 (0.038)	0.415 (0.045)		0.517 (0.038)	0.540 (0.038)	0.404 (0.044)
γ	3.048 (0.625)	2.355 (0.436)		2.577 (0.406)	3.087 (0.645)	2.375 (0.444)
δ_s	1.081 (0.092)	1.057 (0.089)	δ_s^{Inter}	1.136 (0.094)	1.158 (0.118)	1.152 (0.119)
			δ_s^{Intra}	1.053 (0.081)	1.063 (0.100)	1.036 (0.094)
			β_s^{Inter}	0.823 (0.062)	0.791 (0.078)	0.796 (0.079)
β_s	0.809 (0.071)	0.815 (0.071)	β_s^{Intra}	0.849 (0.065)	0.813 (0.083)	0.819 (0.081)
			δ_o^{Inter}	1.006 (0.070)	1.019 (0.084)	1.019 (0.086)
			δ_o^{Intra}	0.957 (0.051)	0.946 (0.062)	0.925 (0.058)
δ_o	0.975 (0.057)	0.958 (0.056)	β_o^{Inter}	1.076 (0.106)	1.095 (0.130)	1.102 (0.133)
			β_o^{Intra}	0.982 (0.055)	0.977 (0.066)	0.983 (0.064)
Observations	4824	4824		4824	4824	4824
Participants	67	67		67	67	67
$H_0(\hat{\beta}_s = 1)$	$p = 0.007$	$p = 0.009$	$H_0(\hat{\beta}_s^{Inter} = 1)$	$p = 0.005$	$p = 0.007$	$p = 0.010$
			$H_0(\hat{\beta}_s^{Intra} = 1)$	$p = 0.020$	$p = 0.024$	$p = 0.026$
$H_0(\hat{\beta}_o = 1)$	$p = 0.902$	$p = 0.860$	$H_0(\hat{\beta}_o^{Inter} = 1)$	$p = 0.470$	$p = 0.464$	$p = 0.441$
			$H_0(\hat{\beta}_o^{Intra} = 1)$	$p = 0.742$	$p = 0.725$	$p = 0.788$
$H_0(\hat{\beta}_s = \hat{\beta}_o)$	$p = 0.033$	$p = 0.034$	$H_0(\hat{\beta}_s^{Intra} = \hat{\beta}_o^{Intra})$	$p = 0.070$	$p = 0.071$	$p = 0.088$
			$H_0(\hat{\beta}_s^{Inter} = \hat{\beta}_o^{Inter})$	$p = 0.062$	$p = 0.082$	$p = 0.066$
			$H_0(\hat{\beta}_s^{Inter} = \hat{\beta}_s^{Intra})$	$p = 0.719$	$p = 0.798$	$p = 0.789$
			$H_0(\hat{\beta}_o^{Inter} = \hat{\beta}_o^{Intra})$	$p = 0.323$	$p = 0.725$	$p = 0.316$

Note: The table reports the parameter estimates from all the blocks, using the utility specification introduced in equation (2). We use data from those 67 subjects who have sufficient variation in the interpersonal and intrapersonal decisions. Columns (1) and (2) restrict the β 's and δ 's to be the same across interpersonal and intrapersonal decisions. Columns (3) to (5) allow them to differ. Column (3) uses the approach via the log-linearized first order condition, all others use the closed form solution. The estimation uses the data from those 67 subjects who have sufficient variation in block SELF and block OTHER, as well as in the dictator game choices (see footnote 14). Standard errors are clustered at the individual level and calculated via the delta method.

Table 10: Parameter estimates from all blocks

cost-of-effort function, to be the same for *self* and *other*.

the restriction that the β 's and δ 's are the same across contexts. Thus, it is the most direct estimation of the model, even though we know from the results above that especially for β_o this restriction may not be warranted.

The results from this estimation can be found in columns (1) and (2) of Table 10. We estimate a β_s between 0.809 and 0.815. Hence, while this estimate is slightly lower than the β_s estimated from the dictator game choices alone, it is well in line with our main results. In particular, we find β_s to be significantly smaller than one (both $p < 0.010$), and to be significantly lower than β_o (both $p < 0.035$). For the latter we estimate values between 1.009 and 1.012, which are not significantly different from one (both $p > 0.859$).

As we show in columns (3) to (5) of Table 10, similar conclusions hold when we allow time preferences to differ between the two contexts, i.e., if we allow for $\beta_s^{Intra} \neq \beta_s^{Inter}$, $\beta_o^{Intra} \neq \beta_o^{Inter}$, $\delta_s^{Intra} \neq \delta_s^{Inter}$, and $\delta_o^{Intra} \neq \delta_o^{Inter}$. Our results reveal a significant present bias in own consumption for both the interpersonal as well as the intrapersonal choices. The estimates for β_s^{Inter} and β_s^{Intra} range between 0.791 and 0.823 and 0.813 and 0.849, respectively, all significantly lower than one (all $p < 0.027$). Moreover, for none of the specifications we find the estimates for β_s^{Inter} and β_s^{Intra} to be significantly different from each other (all $p > 0.718$). Importantly, when comparing these estimates with the ones reported in Section 4 (Table 5) and 5 (Table 8) where we estimate time preferences separately for each context, we find that the estimates for β_s^{Intra} remain virtually the same, while the estimates for β_s^{Inter} become even somewhat smaller when accounting for the curvature in the utility function. In any case, however, we find the estimates from the separate estimations to fall into the 95%-confidence interval of the joint estimation. This suggests that our linearity assumption in Section 4 introduced (if at all) only a small bias and that, if anything, we underestimated the degree of present bias for own consumption in the previous estimation. Similar conclusions can be drawn for present bias in others' consumption, β_o^{Inter} and β_o^{Intra} , as we find the estimates of the joint estimation to be very similar to the ones reported in the previous sections. In particular, in no case we find β_o to be significantly different from one (all $p > 0.440$). As a consequence, in all cases we find β_s to be smaller than β_o , both for the interpersonal as well as the intrapersonal choices (all $p < 0.089$). Overall, these results show that the functional form we proposed to model intertemporal social preferences (see equation (2)) provides reasonable parameter estimates and organizes behavior in dictator games with delayed consequences well.

7 Discussion and Conclusion

In this paper we provide a systematic analysis of time discounting across contexts with and without interpersonal trade-offs. We show that time-inconsistent behavior in form of present bias is not limited to individual decision contexts, but also extends to social contexts in which there is a straightforward trade-off between own and others' consumption. In particular, we find that agents' generosity is subject to dynamic inconsistency, i.e., they behave significantly more altruistically towards others when deciding in advance rather than immediately, while no such difference is observed for choices which only involve decisions about the future. As such, our paper provides important insights into the understanding of social behavior in situations in which consequences play out over time, which, in turn, may have important implications for the modeling of social preferences in dynamic environments.

Our results provide new evidence on the context-dependency of other-regarding concerns. Some previous studies have demonstrated the malleability of prosocial behavior for a variety of manipulations of the decision environment, such as when giving people the possibility to avoid information about the consequences of their actions for others (Dana et al., 2006; 2007) or avoid situations that involve giving decisions (Andreoni and Rao, 2011; DellaVigna et al., 2012), framing (List, 2007; Bardsley, 2008), the inclusion of risk (Exley, 2015), or by diffusing pivotality (Falk and Szech, 2013). Our results add another layer to this central aspect of human behavior by showing that people behave more selfish when consequences are immediate.

We further contribute to the literature on time preferences that, so far, has mainly focused on individual decision situations. Here, we show that people not only exhibit present bias in individual decision contexts (e.g., as in Augenblick et al., 2015 and Augenblick and Rabin, 2017), but that this translates into social contexts in which choices have consequences for someone else. In contrast, no such time inconsistency is observed when consumption of others is concerned. Importantly, we show that present bias in own consumption across these two contexts is correlated within individuals, suggesting that the desire for immediate gratification is a robust phenomenon which is stable across contexts. As such, our paper further relates to the literature investigating the extent to which individual preferences are stable across time and contexts, a topic that is becoming increasingly popular within economics. While some papers have investigated the intertemporal stability of time preferences (Meier and Sprenger, 2015), the stability of time preferences across the monetary and the effort domain (Augenblick et al., 2015), and the predictive power of experimental measures of time preferences for real-world behavior outside the lab (Ashraf et al., 2006; Chabris et al., 2008;

Meier and Sprenger, 2010), we are not aware of any study that compares time preferences across individual and social contexts.

Our results also reveal that agents resolve intertemporal trade-offs very differently, depending on whether they decide about own consumption or on behalf of others. The observation that only the former choices reveal a present bias allows for two different interpretations. Either, agents behave as if they choose what they believe the other person would have chosen for themselves, but mistakenly believe that the other person is time-consistent in their choices. Alternatively, decision makers hold correct beliefs about the present bias of others, but decide to implement time-consistent allocations because they believe that this is the intertemporal allocation of consumption which, from a normative perspective, *should* be implemented for the other agent. While an in-depth investigation of this question is not the focus of this paper, we note that recent work by Fedyk (2017) shows that, in a setting similar to Augenblick and Rabin (2017), agents are unable to foresee their own present bias, but are relatively accurate in predicting the present bias of others. If similar results would hold also in setting, this would suggest that choices made on behalf of others reflect paternalism and that when not affected directly, agents treat present-biased choices as temptation-driven and in need of correction. This is in line with neuro-economic evidence (Albrecht et al., 2011; McClure et al., 2004) which links present bias to the more affective and more impulsive system compared to a more deliberative and reasoned system which may play a more central role when discounting others' consumption.³² Yet, more research is needed to gain a deeper understanding of the underlying psychological mechanisms when discounting own and others' consumption. Evidence from this research, however, would provide important insights about how to incorporate present bias when trying to form a welfare function.

Finally, the current paper only provides a first step in analyzing the link between social and time preferences. Here, we investigate a setting where interactions among players are limited to only one of the two parties making choices as this has the advantage that we can isolate preferences for generosity from strategic motivations. Many situations in which social preferences play a crucial role, such as (ultimatum) bargaining, public good provision, or fostering and maintaining trust, however, have an important strategic component. We hence believe that our study can provide a good starting point to encourage more research that looks at the interaction of social preferences and time preferences more generally.

³²Andersson et al. (2016) make a similar case when they study the role of loss aversion when deciding for others. They find that agents are more loss averse in own than others' choices and therefore argue that loss aversion should be treated as a bias in decision making.

References

- ALBRECHT, K., K. G. VOLZ, M. SUTTER, D. I. LAIBSON, AND D. Y. VON CRAMON (2011): “What is for me is not for you: brain correlates of intertemporal choice for self and other,” *Social Cognitive and Affective Neuroscience*, 6, 218–225.
- ANDERSEN, S., G. W. HARRISON, M. I. LAU, AND E. ELISABET RUTSTRÖM (2008): “Lost in state space: are preferences stable?” *International Economic Review*, 49, 1091–1112.
- ANDERSEN, S., G. W. HARRISON, M. I. LAU, AND E. E. RUTSTRÖM (2018): “Multiattribute Utility Theory, Intertemporal Utility, and Correlation Aversion,” *International Economic Review*, 59, 537–555.
- ANDERSSON, O., H. J. HOLM, J.-R. TYRAN, AND E. WENGSTRÖM (2016): “Deciding for Others Reduces Loss Aversion,” *Management Science*, 62, 29–36.
- ANDERSSON, O., H. J. HOLM, J.-R. TYRAN, AND E. WENGSTRÖM (2014): “Deciding for others reduces loss aversion,” *Management Science*, 62, 29–36.
- ANDREONI, J. AND B. D. BERNHEIM (2009): “Social image and the 50–50 norm: A theoretical and experimental analysis of audience effects,” *Econometrica*, 77, 1607–1636.
- ANDREONI, J. AND J. MILLER (2002): “Giving according to GARP: An experimental test of the consistency of preferences for altruism,” *Econometrica*, 70, 737–753.
- ANDREONI, J. AND J. M. RAO (2011): “The power of asking: How communication affects selfishness, empathy, and altruism,” *Journal of Public Economics*, 95, 513–520.
- ANDREONI, J. AND M. SERRA-GARCIA (2017): “Time-Inconsistent Charitable Giving,” *NBER Working Paper No. 22824*.
- ANDREONI, J. AND C. SPRENGER (2012): “Estimating time preferences from convex budgets,” *American Economic Review*, 102, 3333–3356.
- ASHRAF, N., D. KARLAN, AND W. YIN (2006): “Tying Odysseus to the Mast: Evidence From a Commitment Savings Product in the Philippines,” *Quarterly Journal of Economics*, 121, 635–672.
- AUGENBLICK, N., M. NIEDERLE, AND C. SPRENGER (2015): “Working over time: Dynamic inconsistency in real effort tasks,” *Quarterly Journal of Economics*, 130, 1067–1115.
- AUGENBLICK, N. AND M. RABIN (2017): “An Experiment on Time Preference and Misprediction in Unpleasant Tasks,” *Review of Economic Studies* (forthcoming).
- BALAKRISHNAN, U., J. HAUSHOFER, AND P. JAKIELA (2017): “How Soon Is Now? Evidence of Present Bias from Convex Time Budget Experiments,” *NBER Working Paper No. 23558*.
- BARDSLEY, N. (2008): “Dictator game giving: altruism or artefact?” *Experimental Economics*, 11, 122–133.
- BARSEGHYAN, L., J. PRINCE, AND J. C. TEITELBAUM (2011): “Are Risk Preferences Stable across Contexts? Evidence from Insurance Data,” *American Economic Review*, 101, 591–631.
- BENABOU, R. AND J. TIROLE (2006): “Incentives and prosocial behavior,” *American Economic Review*, 96, 1652–1678.
- BENNDORF, V., H. A. RAU, AND C. SÖLCH (2014): “Minimizing learning behavior in experiments with repeated real-effort tasks,” *Working Paper*.
- BLANCO, M., D. ENGELMANN, AND H. T. NORMANN (2011): “A within-subject analysis of other-regarding preferences,” *Games and Economic Behavior*, 72, 321–338.
- BREMAN, A. (2011): “Give more tomorrow: Two field experiments on altruism and intertemporal choice,” *Journal of Public Economics*, 95, 1349 – 1357.
- BRUHIN, A., E. FEHR, AND D. SCHUNK (2017): “The Many Faces of Human Sociality: Uncovering the Distribution and Stability of Social Preferences,” *Journal of the European Economic Association* (forthcoming).

- CHABRIS, C. F., D. LAIBSON, C. L. MORRIS, J. P. SCHULDT, AND D. TAUBINSKY (2008): “Individual laboratory-measured discount rates predict field behavior,” *Journal of risk and uncertainty*, 37, 237.
- CHAKRABORTY, A., E. M. CALFORD, G. FENIG, AND Y. HALEVY (2017): “External and internal consistency of choices made in convex time budgets,” *Experimental Economics*, 20, 687–706.
- CHARNESS, G. AND U. GNEEZY (2008): “What’s in a name? Anonymity and social distance in dictator and ultimatum games,” *Journal of Economic Behavior & Organization*, 68, 29–35.
- CHEUNG, S. L. (2015): “Risk preferences are not time preferences: on the elicitation of time preference under conditions of risk: comment,” *American Economic Review*, 105, 2242–60.
- COHEN, J. D., K. M. ERICSON, D. LAIBSON, AND J. M. WHITE (2016): “Measuring time preferences,” *NBER Working Paper No. 22455*.
- COOPER, D. AND J. H. KAGEL (2009): “Other regarding preferences: a selective survey of experimental results,” *Handbook of Experimental Economics*, 2.
- DANA, J., D. M. CAIN, AND R. M. DAWES (2006): “What you don’t know won’t hurt me: Costly (but quiet) exit in dictator games,” *Organizational Behavior and Human Decision Processes*, 100, 193–201.
- DANA, J., R. A. WEBER, AND J. X. KUANG (2007): “Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness,” *Economic Theory*, 33, 67–80.
- DANILOV, A. AND T. VOGELSANG (2016): “Time for helping,” *Journal of the Economic Science Association*, 2, 36–47.
- DE OLIVEIRA, A. AND S. JACOBSON (2017): “(Im)patience by Proxy: Making Intertemporal Decisions for Others,” *Working Paper*.
- DELLAVIGNA, S., J. A. LIST, AND U. MALMENDIER (2012): “Testing for altruism and social pressure in charitable giving,” *Quarterly Journal of Economics*, 127, 1–56.
- DELLAVIGNA, S. AND U. MALMENDIER (2006): “Paying Not to Go to the Gym,” *American Economic Review*, 96, 694–719.
- DOHMEN, T., A. FALK, D. HUFFMAN, U. SUNDE, J. SCHUPP, AND G. G. WAGNER (2011): “Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences,” *Journal of the European Economic Association*, 9, 522–550.
- DREBER, A., D. FUDENBERG, D. K. LEVINE, AND D. G. RAND (2016): “Self-control, social preferences and the effect of delayed payments,” *Working Paper*.
- ENGEL, C. (2011): “Dictator games: a meta study,” *Experimental Economics*, 14, 583–610.
- ERKAL, N., L. GANGADHARAN, AND N. NIKIFORAKIS (2011): “Relative earnings and giving in a real-effort experiment,” *American Economic Review*, 101, 3330–3348.
- EXLEY, C. L. (2015): “Excusing selfishness in charitable giving: The role of risk,” *Review of Economic Studies*, 83, 587–628.
- FALK, A. AND N. SZECH (2013): “Organizations, diffused pivotality and immoral outcomes,” *Working Paper*.
- FEDYK, A. (2017): “Asymmetric Naiveté: Beliefs about Self-Control,” *Working Paper*.
- FISCHBACHER, U. (2007): “Z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental Economics*, 10, 171–178.
- FISMAN, R., S. KARIV, AND D. MARKOVITS (2007): “Individual Preferences for Giving,” *American Economic Review*, 97, 1858–1876.
- FREDERICK, S., G. F. LOEWENSTEIN, AND T. O’DONOGHUE (2002): “Time discounting and time preference: A critical review,” *Journal of Economic Literature*, 40, 351–401.
- GREINER, B. (2015): “Subject pool recruitment procedures: organizing experiments with ORSEE,” *Journal of the Economic Science Association*, 1, 114–125.

- HOWARD, G. (2013): “Discounting for personal and social payments: Patience for others, impatience for ourselves,” *Journal of Environmental Economics and Management*, 66, 583–597.
- KOVARIK, J. (2009): “Giving it now or later: Altruism and discounting,” *Economics Letters*, 102, 152–154.
- LAIBSON, D. (1997): “Golden eggs and hyperbolic discounting,” *Quarterly Journal of Economics*, 112, 443–477.
- LIST, J. A. (2007): “On the interpretation of giving in dictator games,” *Journal of Political Economy*, 115, 482–493.
- LOEWENSTEIN, G. AND D. PRELEC (1992): “Anomalies in intertemporal choice: Evidence and an interpretation,” *Quarterly Journal of Economics*, 107, 573–597.
- MCCLURE, S. M., D. I. LAIBSON, G. LOEWENSTEIN, AND J. D. COHEN (2004): “Separate neural systems value immediate and delayed monetary rewards,” *Science*, 306, 503–507.
- MEIER, S. AND C. SPRENGER (2010): “Present-biased preferences and credit card borrowing,” *American Economic Journal: Applied Economics*, 2, 193–210.
- MEIER, S. AND C. D. SPRENGER (2015): “Temporal stability of time preferences,” *Review of Economics and Statistics*, 97, 273–286.
- NOUSSAIR, C. N. AND J. STOOP (2015): “Time as a medium of reward in three social preference experiments,” *Experimental Economics*, 18, 442–456.
- O’DONOGHUE, T. AND M. RABIN (1999): “Doing it now or later,” *American Economic Review*, 89, 103–124.
- PEYSAKHOVICH, A., M. NOWAK, AND D. RAND (2013): “Humans display a cooperative phenotype that is domain general and temporally stable,” *Nature Communications*, 5, 4939–4939.
- 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.
- RONG, R., T. C. GRIJALVA, J. L. LUSK, AND W. D. SHAW (2016): “Interpersonal Discount Rates: An Experiment,” *Working Paper*.
- SCHELLING, T. C. (1984): “Self-Command in Practice, in Policy, and in a Theory of Rational Choice,” *American Economic Review*, 74, 1–11.
- SHAPIRO, J. (2010): “Discounting for you, me, and we: Time preference in groups and pairs,” *Working Paper*.
- SOBEL, J. (2005): “Interdependent Preferences and Reciprocity,” *Journal of Economic Literature*, 43, 392–436.
- STIGLER, G. J. AND G. S. BECKER (1977): “De gustibus non est disputandum,” *American Economic Review*, 67, 76–90.
- STROTZ, R. H. (1956): “Myopia and inconsistency in dynamic utility maximization,” *Review of Economic Studies*, 23, 165–180.
- VOLK, S., C. THÖNI, AND W. RUIGROK (2012): “Temporal stability and psychological foundations of cooperation preferences,” *Journal of Economic Behavior & Organization*, 81, 664–676.