Patience Is Power: Bargaining and Payoff Delay*

Jeongbin Kim† Wooyoung Lim‡

Sebastian Schweighofer-Kodritsch§

August 14, 2024

Abstract

We provide causal evidence that patience is a significant source of bargaining power. Generalizing the Rubinstein (1982) bargaining model to arbitrarily non-stationary discounting, we first show that dynamic consistency across bargaining rounds is sufficient for a unique equilibrium, which we characterize. We then experimentally implement a version of this game where bargaining delay is negligible (frequent offers, so dynamic consistency holds by design), while payoff delay is significant (a week or month per round of disagreement, with or without front-end delay). Our treatments induce different time preferences between subjects by randomly assigning individuals different public payoff delay profiles. The leading treatment allows to test for a general patience advantage, predicted independent of the shape of discounting, and it receives strong behavioral support. Additional treatments show that this advantage hinges on the availability of immediate payoffs and reject exponential discounting in favor of present-biased discounting.

Keywords: Alternating-Offers Bargaining, Time Preferences, Present Bias, Laboratory Experiments

***JEL* Classification**: C78, C91, D03

*We are grateful to Andreas Blume, Anujit Chakraborty, Dongkyu Chang, Pedro Dal Bó, Michael Böhm, John Duffy, Yoram Halevy, Christina Gravert, Emin Karagözoğlu, Shih En Lu, Johannes Maier, Philip Neary, Thomas Palfrey, Ariel Rubinstein, Klaus Schmidt, Peter Norman Sørensen, Charles Sprenger, Egor Starkov, and Emanuel Vespa for valuable comments and suggestions, and we thank our various seminar audiences for helpful discussion. This study is supported by a grant from the Research Grants Council of Hong Kong (Grant No. GRF-16500318), and Sebastian Schweighofer-Kodritsch also gratefully acknowledges financial support from the *Deutsche Forschungsgemeinschaft* through CRC TRR 190 (project number 280092119). The paper was previously circulated and presented under the title "Bargaining and Time Preferences: An Experimental Study."

†Florida State University, Department of Economics, Tallahassee, Florida 32306-2160. Email: `jkim33@fsu.edu`. Phone: (850) 644-3835.

‡The Hong Kong University of Science and Technology, Department of Economics, Kowloon, Hong Kong. Email: `wooyoung@ust.hk`. Phone: (852) 2358-7628. Fax: (852) 2358-2084.

§Berlin School of Economics and Humboldt-Universität zu Berlin, School of Business and Economics, Spandauer Str. 1, 10178 Berlin. Email: `sebastian.kodritsch@hu-berlin.de`. Phone: (49) 30-2093-99497. Fax: (49) 30-2093-99491.

1 Introduction

How will two parties share an economic surplus? This classic distributional question known as the bargaining problem arises in numerous settings.[1] To theoretically resolve this problem with a clear prediction boils down to developing a theory of bargaining power.

The seminal work of Rubinstein (1982) that initiated modern non-cooperative bargaining theory achieved this by explicitly modeling the dynamic process of bargaining as a game, in which disagreement leads to costly payoff delay, and it identified patience as a general source of an individual's bargaining power. Greater patience means greater willingness to delay agreement for a better deal, and in recognition of this, the opponent is led to offering a better deal right away. The advantage due to greater patience extends to incomplete information about time preferences in the sense that it is advantageous to be *perceived as* more patient.[2] In looser terms, the basic claim that being more patient or being perceived as more patient confers an advantage in bargaining also appears in consultants' guides to negotiation.[3] If true, it would add a strategic perspective on the observed positive correlation between individuals' patience and their long-run economic success (e.g., Mischel, Shoda, and Rodriguez, 1989; Epper, Fehr, Fehr-Duda, Kreiner, Lassen, Leth-Petersen, and Rasmussen, 2020; Sunde, Dohmen, Enke, Falk, Huffman, and Meyerheim, 2022), implying that policy makers concerned with economic inequality may for instance consider regulating opportunities for individual wage bargaining (see the recent work of Biasi and Sarsons, 2022).[4]

Yet, to the best of our knowledge, there exists no direct evidence to substantiate this basic prediction or claim. Besides the scarce indirect field evidence, which is suggestive and at best only weakly favorable to it (e.g., Ambrus, Chaney, and Salitskiy, 2018; Backus, Blake, Larsen, and Tadelis, 2020), the only controlled bargaining study in which disagreement results in actual time delay of payoffs finds that participants do not strategically respond to information on the opponent's measured discount factor and concludes that time preferences do not matter in bargaining, in stark contradiction to theory (Manzini, 2001, for a discussion of related literature see below).

In this paper, we offer *causal* evidence on the effect of time preferences on bargaining. We achieve this by experimentally inducing differences in time preferences between otherwise identical groups of participants whom we match to bargain. Following this causal approach, we obtain two main

[1]It arises within households (e.g., Browning and Chiappori, 1998), between workers and firms (e.g., Hall and Milgrom, 2008), as well as between firms (e.g., Ho and Lee, 2017) or between nations (see Powell, 2002, for a survey of bargaining theory in political science analyses of international conflict).

[2]For instance, see Rubinstein (1985), Chatterjee and Samuelson (1987), Bikhchandani (1992) and Watson (1998); a patience advantage similarly prevails under reputational incentives (e.g., Abreu and Gul, 2000; Compte and Jehiel, 2002).

[3]For instance, as in "Be patient—and show it" (Korda, 2011, p. 107) or "Patience is a key characteristic of the good negotiator" (Forsyth, 2009, p. 160).

[4]The importance of bargaining for individuals' long-run economic outcomes has received particular attention in the literature relating gender inequality and wage bargaining (e.g., Bowles, Babcock, and Lai, 2007; Sin, Stillman, and Fabling, 2022). Babcock and Laschever (2003, p. 5) provides a drastic numerical example to illustrate how important even a single wage bargain can potentially be in generating inequality.

results. First, we find that patience is indeed a significant source of bargaining power. Thus, we empirically substantiate the aforementioned general prediction and claim. Second, we find that also in this strategic context the notion of patience has to be qualified to distinguish between the immediate short run and the longer run: Based on what observed bargaining behavior in our experiment reveals, exponential discounting is rejected in favor of present-biased time discounting.

To structure our experimental manipulation, we adopt Rubinstein's classic indefinite alternating offers protocol.[5] Our key innovation is to disentangle *bargaining delay* (i.e., the time delay in bargaining due to disagreement in a round) from *payoff delay* (i.e., the time delay of payoffs due to disagreement in a round), which allows us to *induce* different time preferences among bargainers. Specifically, we let all bargaining take place in a single session, so that bargaining delay is negligible (frequent offers), while at the same time imposing significant payoff delay, of either a week or a month per round of disagreement. Importantly, we exogenously and transparently vary this payoff delay at the individual level (including also whether someone additionally faces a front-end delay): These payoff delay types are randomly assigned and made common knowledge within every bargaining match. Thus, we create groups of bargainers that are essentially identical in every respect other than their *effective time preferences*, and we can compare bargaining behavior and outcomes between different matches to identify causal effects due to people's *underlying time preferences*.

Our experimental method, which we call *effective discounting procedure*, thus permits clean tests of effects due to actual time preferences.[6] Our choice of specific treatments (corresponding to pairings of payoff delay types) is guided by two objectives: First, we aim to obtain and test general predictions that essentially rely only on positive time discounting, to establish whether greater patience indeed confers a strategic advantage; second, we aim to additionally obtain and test discounting-specific predictions to discriminate in particular between exponential discounting (EXD) and the most commonly considered alternative of quasi-hyperbolic (β, δ)-discounting (QHD, see Phelps and Pollak, 1968; Laibson, 1997), as well as more generally present-biased discounting (cf. Chakraborty, 2021).[7]

Our leading treatment *WM* achieves our first objective of testing for a patience advantage independent of details of underlying time preferences. It matches bargainers whose payoff is delayed by one week per round of disagreement ("weekly bargainers") with bargainers for whom this is one month ("monthly bargainers"), and we observe both versions of the game, differing in the type of the initial proposer. As we show theoretically, weekly bargainers are *generally* predicted to be at an advantage

[5] Thus we are able to directly relate to this important theoretical benchmark, including its arguably natural feature that there is always a chance of a counteroffer. For practical as well as theoretical reasons (see Section 2.1), we additionally impose a commonly known 25% chance of random termination after any disagreement throughout all matches, so overall time discounting includes this risk (cf. Halevy, 2008; Chakraborty, Halevy, and Saito, 2020).

[6] We would like to thank John Duffy for helping us coin this term. A version of the method varying discounting between but not within matches was introduced by one of us in Kim (2023) to study the effect of time preferences on cooperation in an indefinitely repeated prisoners' dilemma, which theoretically exhibits equilibrium multiplicity, however. As there, we use the convenient mobile app *Venmo* for all payments, including immediate payments.

[7] Negligible bargaining delay implies that only a single dated self of any individual gets to make all decisions. While somewhat artificial, this has important practical as well as theoretical advantages (see Section 2.1).

over monthly bargainers, holding constant their initial role. While the modal proposal is an equal split (around 50% of all initial proposals, and this is roughly similar also in the other treatments), in line with existing evidence highlighting the importance of fairness concerns (starting with Güth, Schmittberger, and Schwarze, 1982), we nonetheless strongly as well as robustly confirm the generally predicted patience advantage.[8] Hence, time preferences matter, and—in the broad sense of our manipulation—we confirm that patience is a significant source of bargaining power. At the same time, due to the attractiveness of equal splits, the effects we observe in our laboratory study are small in size; in plain averages, they remain below 3.5% of the cake.

The remaining two treatments *WM2D* and *WW1D* further allow us to determine the robustness of this result and discriminate between discounting models based on a revealed preference argument. Treatment *WM2D* is similar to *WM*, except that every bargainer's payoff comes with an additional front-end delay of one week (hence, this delay applies to immediate agreements, and we call these bargainers "delayed"). In contrast to *WM*, but also to the predictions from both EXD and QHD, the common front-end delay removes the significant asymmetry in bargaining power favoring weekly over monthly bargainers, though various comparisons do point in the predicted direction. Altogether this treatment's findings indicate that the significant patience advantage observed in *WM* hinges on the availability of immediate payoffs.

Treatment *WW1D* matches a weekly and a delayed weekly bargainer. Under EXD, the front-end delay is irrelevant, and outcomes should be the same, irrespective of which type gets to make the initial proposal. However, if discounting exhibits a present bias (as is the key feature of QHD), we should observe that delayed weekly bargainers enjoy a significant advantage over non-delayed weekly bargainers. This is indeed what we find—again, with a variety of tests, hence robustly—and from the perspective of (selfish) time preferences as a driving force in bargaining, it is evidence that participants not only expect but also strategically exploit a present bias in others.[9]

In addition to these main results, our treatments permit comparisons between treatments, fixing a given payoff type against two different opponent types (weekly bargainers in Treatments *WM* vs. *WW1D*, and delayed weekly bargainers in Treatments *WM2D* vs. *WW1D*). Analyzing these, on the one hand, we are able to establish robustness of our leading result, confirming a generally predicted patience advantage; on the other hand, however, we also find some suggestive evidence (against both EXD and QHD) that, beyond present bias, bargainers perceive and respond to diminishing impatience, as in general hyperbolic discounting (Loewenstein and Prelec, 1992), though its predictions in our experiment are also more permissive than those of QHD.

Since our design does not itself induce incomplete information, we derive and test comparative

[8] Our main test for this advantage compares the distributions of initial proposals for (strict) first-order stochastic dominance. We complement this with additional analyses such as analogously comparing the distributions of *accepted* initial proposals or those of *selfish* proposals demanding more than half the cake, but also simply comparing the rates at which the types propose as well as accept equal splits.

[9] We discuss potential confounds due to the interaction of our manipulation with social preferences in Section 3.3.

predictions from a complete information theory.[10] In doing so, we essentially assume that the behavioral effects due to any "natural" incomplete information—in particular, about individuals' likely heterogeneous underlying time preferences—are not systematically different between the groups of matches/types we compare, analogous to "noise" in behavior. To address this issue, we consider the rate(s) of immediate agreement: This rate is overall high, close to 75%, and, importantly, it is similar across all kinds of matches/games we observe. This is evidence that incomplete information is non-negligible, but that our design is successful in keeping its effects both relatively mild and roughly constant.

Nonetheless, from this more general perspective, our experimental manipulation may be mainly one of *beliefs about patience*. To investigate this question, we also measured time preferences using standard methods for a subsample of our participants (after bargaining). These measures' correlations with bargaining behavior (in particular, initial demands) have the expected signs, but almost none of them are statistically significant. This highlights the critical importance of beliefs in strategic interaction, and of controlling them experimentally, supporting our approach. Regarding the substantial interpretation of our main results this adds only a minor twist, however, because if beliefs about patience matter strategically, then so does (knowledge of) patience.

Overall, we conclude that time preferences are certainly not all that matters in bargaining, but they do matter significantly. Moreover, they do so in a manner that is theoretically predicted by and consistent with what we know from the large body of work that has researched them: Present bias and diminishing impatience. Though the notion of patience is therefore more complex than under EXD, it is generally a significant source of bargaining power.

Related Literature. Our main contribution is to the experimental economics literature on sequential bargaining, which we focus on here.[11] We propose a novel design that leverages/manipulates people's true time preferences on variously delayed payoffs, to deliver the first causal insights into whether and how they affect bargaining power. In this spirit, the most closely related work is Manzini (2001), which is the only other bargaining experiment with actual time delay of payoffs. Other than that, her design and conclusion are radically different from ours. She first measures relevant discount factors for a two-rounds alternating-offers game—so in the second round it becomes an ultimatum game—where immediate agreements result in payoffs the next day and delayed agreements result in payoffs a month later. Initial proposers then get to learn their opponents' discount factors, and she finds that their opening offers are basically uncorrelated with those, concluding that the task of bargaining distracts attention completely away from time considerations. Our results qualify this negative conclusion: Manzini's very careful design quite compellingly shows that time preferences are

[10]Explicitly modeling incomplete information about time preferences to capture their observed heterogeneity seems elusive.

[11]In particular, we omit a discussion of ultimatum game experiments, see Camerer (2003, Chapter 2) or Güth and Kocher (2014) for surveys.

not *all* that matters in bargaining; and yet, they do once differences in patience are induced in a transparent *ceteris paribus* manner between otherwise identical groups of participants. Indeed, we see a major strength in our approach that it does not rely on accurately measuring individual time preferences, which is well-known to be problematic.[12] Our own correlational results support this view and demonstrate the importance of controlling participants beliefs' about patience.

While the theoretical bargaining literature—following the seminal work by Rubinstein (1982, 1985)—emphasizes costs of disagreement due to time delay of payoffs, all other of the numerous existing sequential bargaining experiments would only "simulate" these costs either via a "shrinking pie" (for the most closely related work, implementing indefinite-horizon games, see Weg, Rapoport, and Felsenthal, 1990; Rapoport, Weg, and Felsenthal, 1990; Binmore, Swierzbinski, and Tomlinson, 2007), or via exogenously imposed breakdown risk (here, see Zwick, Rapoport, and Howard, 1992);[13] all payoffs occur immediately at the end of the experimental session, however, in all these studies (for a classic survey of the experimental bargaining literature during its most active period, see Roth, 1995). First, it is unclear how behavioral responses to these simulations compare with responses to actual time delay of payoffs: individual shrink rates require computational discounting, which appears psychologically very different from how we naturally/intuitively deal with time delay, and breakdown risk is not only about a different domain (relatedly, see Andreoni and Sprenger, 2012b) but also constrains the design to symmetric costs, limiting its scope for testing predictions. Second, all of these simulations assume exponential discounting, at odds with the evidence on time preferences, in particular robustly documenting present bias (for a recent survey, see Ericson and Laibson, 2019). By contrast, our design directly operates on people's natural time preferences as the cost of disagreement, and our findings show that simulating exponential discounting may indeed be a serious limitation.

At the same time, while significant statistically, the causal effects on relative bargaining power due to time preferences we find with our novel method are small, due to the prevalence of equal splits that theoretically advantaged as well as disadvantaged proposers hardly move away from by much. In this particular quantitative respect, our evidence is broadly in line with the classic experimental literature studying different (including asymmetric) shrink rates in definite-horizon bargaining over small numbers of rounds (most prominently Ochs and Roth, 1989).[14] As with any novel method, however, ours also raises new issues regarding this literature's central theme of "gamesmen or fairmen" (Binmore, Shaked, and Sutton, 1985, p. 1179)—reviewed authoritatively by Roth (1995)—which we discuss in Section 3.3.

[12]See the different approaches, arguments, and conclusions in, e.g., Andersen, Harrison, Lau, and Rutström (2008); Andreoni and Sprenger (2012a); Halevy (2015); Augenblick, Niederle, and Sprenger (2015); Chakraborty, Calford, Fenig, and Halevy (2017). Moreover, time preferences may well correlate with other potentially relevant traits and hence potential confounds, such as social or risk preferences, or cognitive ability.

[13]Rapoport et al. (1990) actually implement fixed costs per round of disagreement rather than constant shrink rates; see Schweighofer-Kodritsch (2022) for a recent clarification of this specification as costs of disagreement.

[14]This is further confirmed by Heggedal and McKay (2024), who very recently compared our method with a shrinking-pie design. Relative to the aforementioned experiments implementing indefinite-horizon games, our findings are more supportive of theoretical predictions.

Outline. The rest of this paper is organized as follows. We first present our experimental design and behavioral predictions for the empirically most important classes of time preferences in Section 2. We then report and discuss the findings from our experiment in Section 3, and in the final Section 4 offer concluding remarks, in particular with regards to future research. All formal proofs are relegated to this paper's Appendix, which also contains a general result of equilibrium uniqueness under dynamic consistency, as foundation for our approach. An Online Appendix provides the following supplemental material: additional figures and tables that complement the main body of the paper (part B.1); experimental instructions and selected screenshots for all treatments (part B.2); all details of our additional time preference elicitation and results on how measured time preferences relate to bargaining behavior (part B.3). Throughout, we simply use OA in references to materials in part B.1 of the Online Appendix and specify further only in references to its other parts.

2 Experimental Design and Behavioral Predictions

As the basic structure for our experimental manipulation of bargainers' relative patience we employ the canonical alternating-offers protocol without a deadline, as in the seminal theory of Rubinstein (1982). Two individuals $i \in \{1, 2\}$ decide on how to share a fixed monetary amount. For simplicity, normalize the amount to one, so divisions correspond to shares, and assume it is perfectly divisible. In any round $n \in \mathbb{N}$, one individual i proposes a division $x \in \{(x_1, x_2) : x_1 \in [0, 1] \text{ and } x_2 = 1 - x_1\}$ to the other individual $j = 3 - i$, who can then either accept or reject. If the proposal is accepted, there is agreement, and the game ends; if the proposal is rejected, then the game continues to round $n + 1$, where this protocol is repeated with reversed roles such that j proposes and i responds. Player 1 makes the initial proposal in round 1, and the game continues until a proposal is accepted.

Our main innovation in experimentally implementing this protocol is to disentangle bargaining delay—i.e., the time delay in bargaining upon disagreement, which will be negligible—from payoff delay—i.e., the time delay in payoffs upon disagreement, which will be substantial. This allows us to derive implications of non-exponential discounting, but without the theoretical as well as practical complications due to the implied dynamic inconsistency.[15] The reason is that negligible bargaining delay essentially means that a single dated self of any player makes all decisions. Accordingly, let i's preferences at any point in the game be represented by a single utility function, hence dynamically consistent, and of the form

$$U_i(q, n) = d_i(n - 1) \cdot u_i(q),$$

[15] Significant bargaining delay and the resulting dynamic inconsistency would introduce the issue of potentially imperfect self-knowledge, which gets compounded through the entire hierarchy of beliefs in strategic interaction. See O'Donoghue and Rabin (1999) for pioneering theoretical work on "sophistication vs. naïvetè" in individual decision making, and Augenblick and Rabin (2019) for related evidence of naïveté; Akin (2007) and Haan and Hauck (2023) offer related theoretical bargaining analyses (Ok and Masatlioglu, 2007; Noor, 2011; Pan, Webb, and Zank, 2015; Lu, 2016; Schweighofer-Kodritsch, 2018, analyze bargaining under "full sophistication" about dynamic inconsistency). A practical complication we avoid is attrition, which is likely selective (see Sprenger, 2015; Kim, 2020).

where (q,n) is the outcome where agreement is reached in round n and such that i obtains share q of the surplus; $u_i(\cdot) \geq 0$ is the (continuously increasing) instantaneous utility associated with any payoff q, and $d_i(\cdot) \in [0,1]$ is the (decreasing) discount factor due to the payoff delay corresponding to any bargaining delay. In Appendix A.1 (Proposition 1), we show that, given dynamic consistency, this extension of the Rubinstein (1982) model has a unique subgame-perfect Nash equilibrium under otherwise very general conditions (e.g., u_i concave and $\sup_n d_i(n) < 1$ for both i is sufficient; see also Binmore, 1987, and Coles and Muthoo, 2003, for general analyses of non-stationary bargaining games). The unique equilibrium maintains the familiar property that, in any round, the proposer's offer makes the respondent indifferent between accepting to end the bargaining with this agreement and rejecting to make a counter-offer in the next round. (This offer is also the minimally accepted offer by the respondent, hence accepted indeed.) This uniqueness and characterization result provides the theoretical foundation for our experimental design and behavioral predictions below, but also various alternative treatment implementations one could consider in future research.

2.1 Experimental Design

In line with the theory just described, our experiment implements indefinitely alternating-offers bargaining games with frequent offers and significant payoff delay. The monetary surplus to be divided is fixed and always amounts to US$50. All bargaining takes place within a standard experimental session. However, any round of disagreement in this bargaining causes the two individuals' respective payoffs from any later agreed division to be delayed (further) either by a week or by a month. This delay profile, which additionally includes whether there is a front-end delay of one week to all payoffs (whereby also immediate agreements result in payoffs delayed by a week), is randomly assigned at the individual level, following our effective discounting procedure, according to the specific treatment.

Table 1: Experimental Treatments

Bargainer 1	Bargainer 2		
	Monthly (M)	Delayed Monthly (MD)	Delayed Weekly (WD)
Weekly (W)	***WM***	N/A	***WW1D***
Delayed Weekly (WD)	N/A	***WM2D***	N/A

Notes: Delay (D) = 1 week (front-end delay to payment).

Table 1 presents our experimental design, which consists of three such treatments. Each of these corresponds to a particular pairing of "bargainer types," which are the payoff delay profiles that any two matched participants face. In Treatment WM, one bargainer faces a payoff delay of one week (W) per round of disagreement, while the other faces a payoff delay of one month (M). In this treatment, immediate agreements correspond to payoffs received immediately after the experimental session ends.

Treatment *WM2D* is similar, but both bargainers additionally face a front-end delay (D) of one week, such that immediate agreements produce payoffs received one week after the experiment (hence type labels *WD* and *MD*). Finally, in Treatment *WW1D*, both bargainers face identical delays of one week (W) per round of disagreement, but one of them additionally faces a front-end delay (D) of one week (hence type labels W and *WD*). In the rest of the paper, we will call a bargainer whose payment window is weekly/monthly/delayed a *weekly/monthly/delayed* bargainer.[16]

Importantly, the treatment is always public, whereby the types of any two matched participants are common knowledge, in line with the theory's stylized assumption of perfect information. Moreover, who is assigned to be the initial proposer is randomized at the match level, so we observe both kinds of games of any treatment; thus, e.g., Treatment *WM* allows us to compare initial proposals by weekly bargainers to their monthly opponents with the initial proposals by monthly bargainers to their weekly opponents. Indeed, as we show in the next section, this will allow us to test for a patience advantage that should theoretically obtain independent of the details of time preferences. The purpose of Treatment *WM2D* is then to permit a robustness check concerning the availability of immediate payoffs, and Treatment *WW1D* will allow us to discriminate between exponential and present-biased discounting, which make different predictions here.[17]

While disentangling the timing of bargaining—in particular, that of agreements—from the timing of payoffs renders the experiment somewhat artificial, this has several important advantages for the purpose of our study. Besides those mentioned already earlier (footnote 15 just above), it allows us to transparently induce *different* effective discounting between otherwise identical (groups of) participants matched to bargain with each other, which is key (cf. Kim, 2023). Moreover, it means a tight connection to the related evidence on time preferences, which comes from static choices over differently delayed payoffs (for a rare longitudinal study see Halevy, 2015); the discounting measures thus obtained have been shown to predict a variety of field behaviors in line with theory (Chabris, Laibson, Morris, Schuldt, and Taubinsky, 2008; Meier and Sprenger, 2010; Sutter, Kocher, Glätzle-Rützler, and Trautmann, 2013; Castillo, Jordan, and Petrie, 2019; Backes-Gellner, Herz, Kosfeld, and Oswald, 2021).

We couple this treatment design with a fixed, commonly known termination probability of 25% that was transparently applied to all rounds of all games in all treatments (so it could not cause any systematic differences). This serves several quite related purposes. First, it ensures that our bargainers are actually impatient regarding when to reach agreement, despite also the basically zero interest rates. (For a theoretical foundation, see Property 3 of our preference assumptions in Appendix A.1, and—

[16] In total we thus have four types *(W, M, WD, MD)*. For a concrete illustration of their different payoff delays, suppose an individual's bargaining game resulted in agreement in Round 3, so after two rounds of disagreement. They then receive their corresponding payoff two weeks from the experiment if their type is W, two months if M, three weeks if *WD* (one week of front-end delay on top of the two of W), and one week plus two months if *MD*.

[17] Testing these "within-treatment predictions" was the main motivation for our design. Accordingly, all treatments match different payoff types. However, we will additionally consider "between-treatments predictions" that compare a given type in the same initial role against different opponent types; e.g., weekly bargainers against a monthly opponent (in *WM*) or against a delayed weekly opponent (in *WW1D*), see below for details.

from the perspective of incomplete information—Fanning and Kloosterman, 2022.) Second, every bargaining game, while indefinite, is thus still expected to end after a reasonable amount of time, which is important for the credibility and smooth running of our experiment. Finally, it limits the potential complications due to incomplete information by making screening and signaling additionally costly. Of course, in terms of the model, discounting should therefore be interpreted as also including this constant risk, which we accordingly make explicit in the following derivation of behavioral predictions (assuming expected utility, for simplicity; see Appendix A.3 for extensions).[18]

2.2 Behavioral Predictions

We now operationalize our aforementioned general uniqueness result and equilibrium characterization to derive the behavioral predictions related to time preferences that our experiment is designed to test.[19] We focus on the benchmark of exponential discounting (EXD; Rubinstein, 1982) and its leading alternative of quasi-hyperbolic discounting (QHD), which has a single additional parameter to explain present bias. (Throughout, we use QHD assuming a strict present bias, so it does not generalize EXD but pose an actual alternative, see below.) In each case, to capture the implied typical behavior, we impose preference symmetry in terms of *underlying/natural* preferences over delayed payoffs: i.e., both individuals have the same atemporal utility function, $u_1 = u_2 = u$, and for the same future delay $\Delta_{t,t'}$ from some given date in time t to some later date $t' > t$, discount utility with the same discount factor $\delta_{t,t'}$. Our effective discounting procedure induces different *effective* time preferences by implementing idiosyncratic payoff delay profiles (types).

To see this, take the following utility specification, which nests all such types and covers both EXD and (present-biased) QHD, and where we make the constant exogenous breakdown risk explicit. Let then $\delta \in (0,1)$ denote the common basic/exponential discount factor for a weekly delay such that δ^k is that for a monthly delay ($k \approx 4$), $\beta \in (0,1]$ the common additional present-bias discount factor so that EXD has $\beta = 1$ and QHD has $\beta < 1$, and $\gamma \in (0,1)$ the continuation probability upon disagreement ($\gamma = 3/4$). To capture a bargainer i's payoff type, let $\mathbb{I}_i(W)$, $\mathbb{I}_i(M) \equiv 1 - \mathbb{I}_i(W)$, and $\mathbb{I}_i(D)$ be indicator functions returning one if i is, respectively, a weekly (W), a monthly (M) and a delayed (D) bargainer, and zero otherwise in each case. Then, defining i's effective (basic/exponential) discount factor as $\delta_i \equiv \delta^{\mathbb{I}_i(W)+k\mathbb{I}_i(M)}$ and using another indicator $\mathbb{I}(n > 1)$ to distinguish between delayed agreements

[18] Heggedal and McKay (2024) recently adopted our procedure to implement our baseline Treatment *WM* without any exogenous risk and instead fixed/known deadlines after two or three rounds, and their findings related to time preferences replicate ours.

[19] Since we leverage our bargainers' unobserved underlying/natural time preferences, point predictions are unavailable.

($n > 1$, then returning one) and immediate agreements ($n = 1$, then returning zero), take:[20]

$$U_i(q,n) = \begin{cases} \beta^{\mathbb{I}(n>1)} \cdot (\gamma\delta_i)^{n-1} \cdot u(q), & \text{if } \mathbb{I}_i(D) = 0, \\ \beta\delta \cdot (\gamma\delta_i)^{n-1} \cdot u(q), & \text{if } \mathbb{I}_i(D) = 1. \end{cases}$$

In terms of implications for i's willingness to agree in round n rather than $m > 0$ rounds later, this means

$$U_i(q,n) \geq U_i(q', n+m) \Leftrightarrow u(q) \geq \begin{cases} \beta \cdot (\gamma\delta_i)^m \cdot u(q'), & \text{if } \mathbb{I}_i(D) = 0 \text{ and } n = 1, \\ (\gamma\delta_i)^m \cdot u(q'), & \text{if } \mathbb{I}_i(D) = 1 \text{ or } n > 1. \end{cases}$$

Under EXD ($\beta = 1$), whether a bargainer is delayed has no effect, and the only effective asymmetry our treatments would induce concerns whether $\delta_i = \delta$ (weekly bargainers) or $\delta_i = \delta^k < \delta$ (monthly bargainers). Assuming QHD, present bias due to $\beta < 1$ would never matter once bargaining is past the initial round ($n > 1$)—from then on, all trade-offs involve only payoffs in the future—but an additional effective asymmetry concerns whether it would matter in the very initial round ($n = 1$), where it does not for delayed bargainers.

The only universal prediction from the general model in Appendix A.1 (Proposition 1) is immediate agreement. We now formulate the behavioral predictions in terms of relative bargaining power, as reflected in this immediate agreement. For this purpose, we denote by $x_{type}^{treatment}$ and $y_{type}^{treatment}$ the equilibrium shares of a given *type (W, M, WD, MD)* in a given *treatment (WM, WM2D, WW1D)*, as the initial proposer (x) and as the initial respondent (y), respectively. Our main predictions compare bargaining power within a given treatment matching types A and B, where we have that A obtains a greater proposer share than B if and only if A obtains a greater respondent share than B; e.g., take Treatment *WM* and simply use that $x_W^{WM} = 1 - y_M^{WM}$ as well as $x_M^{WM} = 1 - y_W^{WM}$. Additional predictions compare the same type's bargaining power between two treatments matching this very type with two different opponent types, in the same initial role; e.g., type W as initial proposer against type M in Treatment *WM* vs. against type *WD* in Treatment *WW1D*, and analogously for type W as initial respondent. Here, the ranking of proposer shares does not generally—i.e., for any time preferences—imply a similar ranking of respondent shares (e.g., see QHD's prediction part (B2) in Table 2 below). Since our main empirical analysis will be based on initial proposals in terms of demands, we formulate all predictions in terms of comparing these; e.g., concerning W as initial respondent in Treatment *WM* vs. Treatment *WW1D*, predictions compare $x_M^{WM} = 1 - y_W^{WM}$ vs. $x_{WD}^{WW1D} = 1 - y_W^{WW1D}$ as opposed to the equivalent y_W^{WM} vs. y_W^{WW1D}.

Table 2 summarizes and contrasts all behavioral predictions under EXD and QHD, which correspond to Predictions 1 and 2, respectively, as formally stated and proven in Appendix A.2. Prediction

[20]We assume that payoffs received at the end of the experimental session are perceived as immediate payoffs, which is standard also in time preference elicitations using money.

Table 2: Predictions under EXD and QHD

	Treatments	EXD (Prediction 1)	QHD (Prediction 2)
(A1)	*WM*	$x_W^{WM} > x_M^{WM}$	$x_W^{WM} > x_M^{WM}$
(A2)	*WM2D*	$x_{WD}^{WM2D} > x_{MD}^{WM2D}$	$x_{WD}^{WM2D} > x_{MD}^{WM2D}$
(A3)	*WW1D*	$x_{WD}^{WW1D} = x_W^{WW1D}$	$x_{WD}^{WW1D} > x_W^{WW1D}$
(B1)	*WM* and *WW1D*	$x_W^{WM} > x_W^{WW1D}$	$x_W^{WM} > x_W^{WW1D}$
		$x_{WD}^{WW1D} > x_M^{WM}$	$x_{WD}^{WW1D} > x_M^{WM}$
(B2)	*WM2D* and *WW1D*	$x_{WD}^{WM2D} > x_{WD}^{WW1D}$	No general pred.
		$x_W^{WW1D} > x_{MD}^{WM2D}$	$x_W^{WW1D} > x_{MD}^{WM2D}$

Notes: Table shows behavioral predictions in our experiment, depending on whether one assumes EXD or QHD, where QHD assumes $\beta < 1$ and therefore does not nest EXD. Predictions 1 for EXD and 2 for QHD refer to the formal statements in Appendix A.2, and $x_{type}^{treatment}$ denotes the (immediate-agreement) equilibrium share of a given *type (W, M, WD, MD)* in a given *treatment (WM, WM2D, WW1D)* as the initial proposer.

1 for EXD is straightforward. Simply note that under EXD front-end delay is irrelevant, and weekly bargainers have a higher effective discount factor than monthly bargainers, i.e., they are effectively more patient.

Prediction 2 for QHD—recall this is taken to mean an actual present bias, $\beta < 1$—differs only in parts (A3) and (B2). To understand this prediction, note that (i) Prediction 1 applies to the Round-2 subgame, where bargaining is only about delayed payoffs and QHD coincides with EXD, and (ii) the immediate Round-1 agreement has the initial respondent indifferent to the Round-2 agreement, so only the respondent's discounting for the first round's disagreement matters. In particular, a present bias in the sense of $\beta < 1$ enters the actual equilibrium agreement if and only if the initial respondent is not delayed, so *ceteris paribus* front-end delay makes an initial respondent stronger.

Within Treatment *WM*, the weekly bargainer is therefore stronger than the monthly bargainer in the Round-2 subgame, and since present bias applies equally to both types, this carries over to the immediate Round-1 agreement. (Present bias here simply reinforces the proposer advantage.) The QHD prediction within Treatment *WM2D* is immediate from that under EXD because present bias is irrelevant. This is in stark contrast to Treatment *WW1D*, which is symmetric under EXD, but not under QHD: Whereas the Round-2 subgame is symmetric, the delayed weekly bargainer is the stronger initial respondent due to the effective absence of a present bias. (Equivalently, this type is stronger as the initial proposer because it faces a weaker respondent.)

The observation that under QHD front-end delay is advantageous as initial respondent implies that the EXD prediction between Treatments *WM* and *WW1D* is only reinforced under QHD regarding the weekly bargainer as initial proposer, since the delayed weekly respondent then faces no present bias whereas the monthly one does. For the weekly bargainer as the initial respondent, present bias equally weakens this type irrespective of the type of proposer and does not affect the comparison relative to EXD.

Between Treatments *WM2D* and *WW1D*, when the delayed weekly bargainer is the initial respondent, the game under QHD is the same as that under EXD, so the prediction immediately carries over. However, when the delayed weekly bargainer is the initial proposer, present bias is effective in *WW1D* but not in *WM2D*; while the Round-2 agreement is less favorable with a weekly opponent than a delayed monthly one, the former is therefore weakened by present bias, whereas the latter is not. This means that the comparison depends on how strong present bias is relative to the difference in long-run discounting, so there is no general prediction under QHD.

Due to the tractability they afford, EXD and QHD are, by far, the most important models of time preferences for theoretical analyses. However, empirical studies, especially from psychology, suggest hyperbolic discounting (HYD)—a form of diminishing impatience, which implies present bias—as the "universal" form of discounting (for discussion see Frederick et al., 2002). We discuss the implications of diminishing impatience and generalizations (especially by Chakraborty, 2021) as well as alternatives more extensively in Appendix A.3. Here, we only briefly sketch them for HYD, in the form of the model by Loewenstein and Prelec (1992), which imposes the structure of $d(t) = (1 + \alpha \cdot t)^{-\beta/\alpha}$ (with $\alpha, \beta > 0$ and $t \geq 0$ time delay). While impatience is diminishing, a weekly bargainer then still always remains more patient than a monthly bargainer, and this extends also to when *both* face the same front-end delay. Hence, within-treatment predictions (A1) and (A2) for *WM* and *WM2D* remain the same as with both EXD and QHD. Concerning (A3) for *WW1D*, the diminishing impatience only adds to the advantage of the delayed bargainer from removing present bias, thus reinforcing the prediction under QHD. (B1) similarly carries over, while with the HYD model above it still depends on parameters whether the weekly or delayed monthly bargainer is always more patient, such that it not only makes no general (B2) prediction about x_{WD}^{WM2D} vs. x_{WD}^{WW1D}, but also not about x_{MD}^{WM2D} vs. x_{W}^{WW1D}. This additional permissiveness in (B2) is the only difference in predictions from QHD.[21]

Overall, regarding (A1), (A2), and (B1), all discounting models considered here make the same predictions regarding relative bargaining power. As such, they allow us to test for a *basic patience advantage*, and jointly let us determine its robustness. Provided this is established, (A3) will then allow us to test whether bargaining behavior is better explained by EXD or by *present bias*, as under QHD or HYD. Finally, in case we indeed find evidence of present bias here, (B2) will allow us to further determine whether the simple version of present bias in QHD suffices for explaining behavior, or this requires the more permissive version of *diminishing impatience* under HYD.

[21] Similar to how we use QHD to mean $\beta < 1$, thereby not generalizing EXD, the HYD model neither generalizes EXD nor QHD. However, in terms of predictions in our experiment, HYD is still more permissive than QHD, meaning our design could possibly make only a weak case for HYD over QHD.

2.3 Administrative Details

Our experiment was conducted using z-Tree (Fischbacher, 2007) at the University of California, Irvine. A total of 348 subjects who had no prior experience with our experiment were recruited from the graduate and undergraduate student population of the university. Upon arrival at the laboratory, the participants were instructed to sit at separate computer terminals. Each received a copy of the experiment's instructions (see Online Appendix B.2). To ensure that the information contained in the instructions was induced as public knowledge, these instructions were read aloud, and the reading was accompanied by slide illustrations followed by a comprehension quiz.

Each session employed a single treatment, and we conducted 6 sessions for each treatment, for a total of 18 sessions (6 sessions × 3 treatments).[22] In all sessions, the participants anonymously played 10 games under the corresponding treatment condition, say matching bargainer types A and B, where bargaining was over how to divide 500 tokens worth \$50. At the beginning of the experiment, one half of the participants were randomly assigned to be Type A and the other half to be Type B. Individual participants' types remained fixed throughout the session. We used random rematching across subsequent games, subject to the treatment condition of always matching a Type A and a Type B. Any participant therefore always had the same type and always faced the same opponent type to avoid any confusion regarding payoff delay profiles. However, the identity of the initial proposer was always determined by chance, so we observe both kinds of games of any treatment, and every participant would sometimes be the initial proposer and sometimes the initial respondent. Each session had 16–20 participants and hence involved 8–10 simultaneous games.

At the end of the experiment, one of the 10 matches a participant had played was randomly selected for payment (see Azrieli, Chambers, and Healy, 2018, who offer a theoretical justification of such incentives). For the selected match, if agreement was reached, the agreed number of tokens for this participant was converted into US dollars at a fixed and commonly known exchange rate of \$0.1 per token, and the delay of the participant's dollar payment was determined according to (1) his/her bargainer type and (2) the round of the agreement.

After all ten bargaining matches were over, we additionally measured the participants' time preferences using a version of the BDM (Becker, DeGroot, and Marschak, 1964) method. We elicited switching points (indifferences) between sooner and later money amounts. One decision was randomly selected for actual payment.[23]

In addition, participants received a show-up payment of \$10. Any amount a participant was due to receive, including immediate payments, was paid electronically via the popular mobile payment system *Venmo*, and when recruiting our participants we clearly announced that those without a Venmo

[22] We conducted 6 sessions in May and June 2018, and 12 sessions in October and December 2018.

[23] We implemented the elicitation task in 4 sessions per treatment. While initially focusing on our causal approach of inducing variation in time preferences, we afterwards decided to also measure time preferences. This allows us to check (the subsample) whether the random assignment was successfully implemented in terms of participants' underlying time preferences and to compare our causal behavioral findings with correlational ones. See Online Appendix B.3 for details.

account were not eligible to participate in the experiment.[24] Earnings were \$37.90 on average, and the average duration of a session was approximately 1.5 hours.

3 Experimental Results

In line with the literature, we conduct our main tests of the behavioral predictions based on observed initial proposals, as they reveal the proposers' perceptions of relative bargaining power. We have this data for every match. Acknowledging potential preference heterogeneity and incomplete information, we take the predictions to concern shifts in the distribution of bargaining power between the different kinds of games/matches that are compared. We therefore conduct our comparisons based on the entire observed distributions (CDFs) of initial proposals, always in terms of the proposer's claimed share (demand).[25] In line with the predictions in Table 2, $F(x_{type}^{treatment})$ denotes the distribution of initial demands by a given type in a given treatment. We test pairs of such distributions for *strict* first-order stochastic dominance (in what follows, simply "dominance"), using the procedure proposed by Barrett and Donald (2003), which involves testing both "weak" Null hypotheses and hence yields two p-values (the two-fold Null corresponds to the distributions being the same); in particular, for two distributions/CDFs F and G, F is inferred to dominate G if and only if Null $F \leq G$ is accepted *and* Null $G \leq F$ is rejected.[26] Since demands and offers add up to a constant amount (500 tokens), if type A's demands dominate type B's demands as initial proposer, type A faces an unambiguously more favorable distribution of initial proposals than type B as initial respondent.

As a preliminary, Table 3 shows average initial proposals—in terms of the share demanded—by each of the two types in each of the three treatments, as well as aggregating all of these. It does so for all ten matches, as well as splitting into the first five and the last five matches, to compare inexperienced and experienced proposers. Overall, the average initial demand is 53% of the cake, hence very close to an equal split. Furthermore, there is limited variation by treatment and type, and also by experience: all averages are between 51.5% and 54.7%. While, as explained, our main tests compare distributions, this means that any effects of our effective discounting manipulation on bargaining power—even if significant in shifting proposal distributions—are quantitatively small. The largest difference between types' average demands amounts to little more than 3% of the cake, see last five matches in Treatment *WM*, whose two types' averages coincide also with the aforementioned bounds.

At the same time, the fact that all initial demand averages are above 50% suggests a *proposer*

[24] At the end of the experiment, the participants were asked to report their account information for payment, including username and email address details. None of the participants reported any error or difficulty in providing this information, suggesting that all our participants were sufficiently familiar with Venmo in their daily lives. The longest delay among the matches selected for payment was 7 months, and the corresponding amount was paid on May 17, 2019.

[25] The CDF figures plot the shares distributions always over the range $[0.50, 0.62]$, for ease of graphical representation. This range contains, on average, more than 95% of the data. It corresponds to $[250, 310]$ in tokens bargained over.

[26] For the p-values in each test, we employ a bootstrap of size 1,000.

Table 3: Average Initial Proposals

Matches	Treatment/Type					
	All	*WM*		*WM2D*		*WW1D*	
		W	*M*	*WD*	*MD*	*W*	*WD*
All ten	0.530 (0.061)	0.541 (0.076)	0.521 (0.062)	0.532 (0.051)	0.532 (0.063)	0.519 (0.048)	0.533 (0.060)
First five	0.531 (0.064)	0.535 (0.061)	0.527 (0.081)	0.528 (0.048)	0.535 (0.064)	0.519 (0.049)	0.539 (0.074)
Last five	0.529 (0.058)	0.547 (0.089)	0.515 (0.033)	0.535 (0.054)	0.529 (0.061)	0.519 (0.047)	0.527 (0.040)
Observations (all ten)	1,740	290	300	284	296	299	271

Notes: Table shows average initial proposals in terms of demanded shares for all ten matches, the first five matches, and the last five matches, respectively, with standard deviations in parentheses below. Observations refer to all ten matches. While there are equally many such observations from the first and last five matches (870 = 1740/2), the numbers differ for the two types in a given treatment, due to match-level randomization of the initial proposer.

(first-mover) advantage. This prediction concerns a within-treatment comparison that is not specific to time preferences as the source of disagreement costs, and it neither implies nor is implied by a patience advantage. While it is possible to construct examples of time preferences that even violate it, the prediction obtains under EXD, and a present bias as in QHD only reinforces it. Since we employ a novel method, we also test for it here, using the procedure described above. Indeed, in every treatment, for each of the two types, we find that the distribution of initial *demands by* this type dominates that of initial *offers to* this type (see also OA Figures 16–18).

With this background, we now turn to the main tests of the predictions in Table 2, comparing distributions of various types' initial demands. We throughout complement them with reference to additional analysis for which the details can be found in Online Appendix B.1; in particular, this concerns the distributions of individual mean proposals and of accepted proposals, as well as learning over the ten matches and rejection behavior.

3.1 Within-Treatment Comparisons

The most general prediction of a patience advantage as in (A1) concerns our leading treatment *WM*. Comparing average proposals by the two types in Table 3 is suggestive of confirmation. Figure 1(a) now presents the CDF of all Round-1 proposals/demands in this treatment, aggregating all 10 matches/games of any session, by bargainer type. The solid line indicates the CDF of weekly demands, and the dotted line indicates the CDF of monthly demands. Consistent with prior findings, fairness

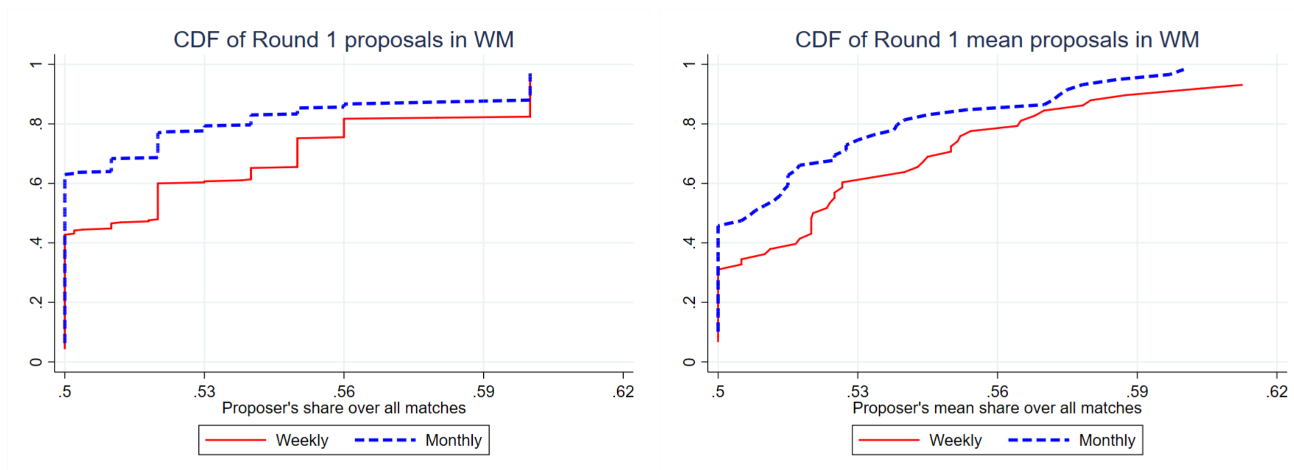

(a) All Proposals

(b) Individual Mean Proposals

Figure 1: Round-1 Proposals in Treatment *WM*

concerns appear important, as approximately 50% of proposals are equal splits (corresponding to .5 in the figure; see OA Table 4 for details). However, the rate of equal-split proposals varies strongly by type, with weekly bargainers much less likely to propose an equal split than monthly bargainers. Indeed, regarding the key comparative prediction, the weekly CDF clearly lies below the monthly CDF, and we have statistically highly significant dominance (p-values of 0.961 and 0.001). Since participants make several initial proposals, however, Figure 1(b) also presents the CDFs of individuals' average Round-1 proposals (demands). It confirms our finding from considering all proposals, and it additionally shows that this finding is not driven only by those that, by chance, get to make relatively many initial proposals. Moreover, the fraction of individuals always proposing equal splits is much lower than the fraction of equal-split proposals overall.

Result 1 (Basic Patience Advantage (A1))**.** *$F(x_W^{WM})$ dominates $F(x_M^{WM})$, confirming the generally predicted basic patience advantage of weekly over monthly bargainers in Treatment WM.*

OA Figures 7 and 10 show further that this result holds up also for accepted proposals and that it obtains rather quickly after the initial match, despite a fair amount of noise. Comparing the first and second half of the experiment, Table 3 further shows that experience reinforces the patience advantage of weekly over monthly bargainers. We can also exploit the prevalence of equal splits for comparing *respondent* behavior: While such proposals are generally hardly ever rejected, weekly bargainers reject them somewhat more often, though this difference is not statistically significant (7% vs. 4%, $p = 0.378$).

Overall, we therefore conclude that weekly and monthly bargainers share a common perception about their relative bargaining power, favoring the former. Accordingly, we strongly confirm the basic prediction (A1) that patience is a source of bargaining power.

Treatment *WM2D* adds a front-end delay of one week to both types and therefore allows us to investigate to what extent this patience advantage is driven by short-run patience vs. long-run patience.

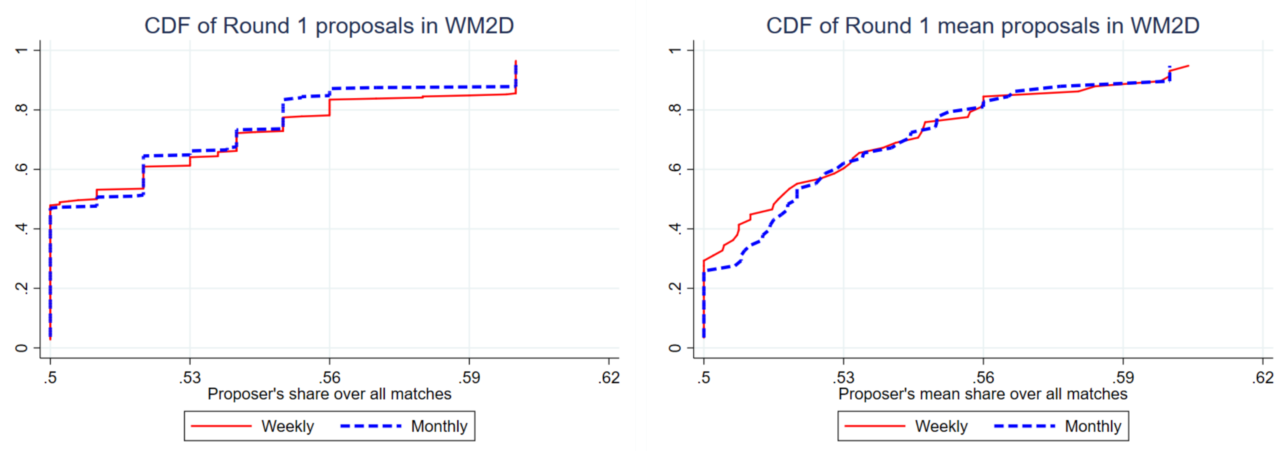

(a) All Proposals (b) Individual Mean Proposals

Figure 2: Round-1 Proposals in Treatment *WM2D*

Figure 2(a) presents the CDFs of all Round-1 proposals in this treatment by bargainer type. The solid line indicates the CDF of delayed weekly demands, and the dotted line indicates the CDF of delayed monthly demands. Again, close to 50% of proposals are equal splits. Here, however, these are equally likely for both types, and also the entire distributions of proposals in Figure 2(a) are quite obviously not significantly different (both p-values > 0.131). Figure 2(b) confirms this in terms of individual average demands.

Result 2 (Basic Patience Advantage (A2)). *$F(x_{WD}^{WM2D})$ and $F(x_{MD}^{WM2D})$ are not statistically different, rejecting the generally predicted patience advantage of delayed weekly over delayed monthly bargainers in Treatment WM2D.*

OA Figures 8 and 11 qualify this finding somewhat: In terms of accepted proposals the delayed weekly proposers tend to do better than the delayed monthly ones—in particular, delayed weekly *respondents* reject equal splits significantly more often than delayed monthly ones (8% vs. 2%, $p = 0.018$)—and also in the later phase of the experiment learn to demand slightly more than their monthly counterparts (as can be seen also in Table 3).

Notwithstanding these qualifications, the data show that—in contradiction to the general prediction (A2)—there is no clear common perception of the predicted basic patience advantage of weekly over monthly bargainers as in *WM* once a front-end delay is added to all payoffs. Recalling that the effects in *WM* are quantitatively small, this Null finding is arguably not too surprising. We may speculate that the additional and *symmetric* front-end delays in *WM2D* makes the asymmetry in payoff delay profiles somewhat less salient, and any confusion (or expected confusion) pushes demands towards the safe equal split. However, notably, under QHD, the front-end delays in *WM2D* also theoretically reduce all initial demands relative to *WM*, in which respondents' present bias matters, and given the small scope for effects to begin with, this reduction could easily render differences insignificant.

The third treatment *WW1D* allows us to directly test for present-biased QHD (or HYD) against EXD, based on what initial proposals reveal. There is again a clear asymmetry, but now it concerns only the bargainers' front-end delay, which only one of them faces.

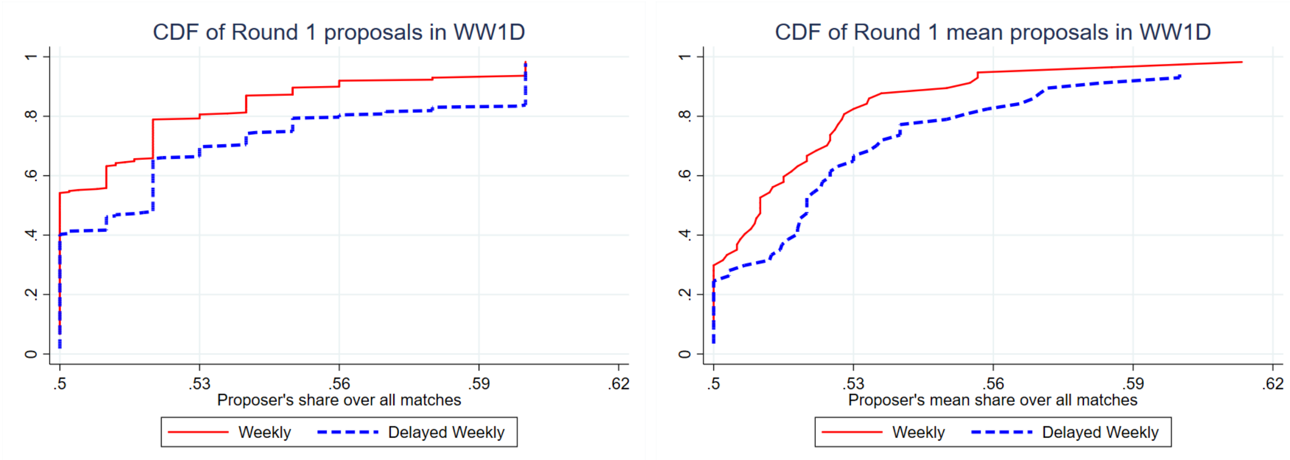

(a) All Proposals

(b) Individual Mean Proposals

Figure 3: Round-1 Proposals in Treatment *WW1D*

Figure 3(a) shows the CDFs of all Round-1 proposals in Treatment *WW1D* by bargainer type. The solid line indicates the CDF of weekly demands, and the dotted line indicates the CDF of *delayed* weekly demands. Once again, approximately 50% of proposals are equal splits, though the proportion is higher for weekly than delayed weekly bargainers. Indeed, whereas EXD predicts no difference, it is visually clear that the delayed weekly demands dominate the weekly ones (without front-end delay), as alternatively predicted under QHD (with p-values 0.971 and 0.001, so high statistical significance). We again obtain the same result for individuals' average proposals shown in Figure 3(b).

Result 3 (Present Bias (A3)). *$F(x_{WD}^{WW1D})$ dominates $F(x_{W}^{WW1D})$, rejecting EXD in favor of present bias as under QHD or HYD in Treatment WW1D.*

As for Treatment *WM*, here, OA Figures 9 and 12 again provide further confirmation: The result holds up also for accepted proposals, and it obtains immediately after the initial match though here the difference becomes smaller towards the end (see also Table 3). Regarding *respondent* behavior, delayed weekly bargainers also reject equal splits at a significantly greater rate (8% vs. 2%, $p = 0.044$).

Hence, we conclude that the two types of weekly bargainers of this treatment—one delayed, the other not—share a common perception that the front-end delay increases bargaining power. Regarding (A3), the EXD prediction is thus rejected in favor of the alternative prediction under present-biased QHD (or also HYD).[27]

Overall, while average demands showed that potential effects due to time preferences are quantitatively small, our within-treatment comparisons show that: (i) patience is nonetheless a significant

[27]Note that this result most strongly rejects near-future bias (see Appendix A.3).

source of bargaining power, (ii) this patience advantage's significance may hinge on the availability of immediate payoffs, and (iii) it is in line with exploiting present bias.[28] The fact that all of these results are reflected in types' relative propensities to propose equal splits raises the question whether they are entirely driven by these or there are also wedges in the distributions of more selfish proposals that demand more than half the cake. OA Figures 13–15 analogously compare CDFs of only such proposals, and tests on these confirm the earlier results with one notable exception: Treatment *WM2D* then confirms the predicted basic patience advantage (A2) with the corresponding statistically significant dominance. This lends support to the discussion following (Null-) Result 2 concerning its reasons, to conclude that *WM2D* provides weak support rather than rejection of the predicted advantage.[29]

3.2 Between-Treatment Comparisons

Our comparisons between treatments always concern two treatments that have one bargainer type in common. They therefore involve testing two predictions, one for this type as the initial proposer and another for this type as the initial respondent. For simplicity, we now only present the results based on all initial proposals, since—as seen for all within-treatment comparisons—they are similar for individual mean proposals, and we also focus on the main tests. The above observations about Treatment *WM2D* indicate that participants' initial responses and subsequent learning from their interaction may practically depend on treatment in ways that the theory does not capture (relatedly, see Azrieli et al., 2018, pp. 1489–90). This is a potential caveat of between-treatments tests and also the reason behind our focus on within-treatment tests when designing the experiment.

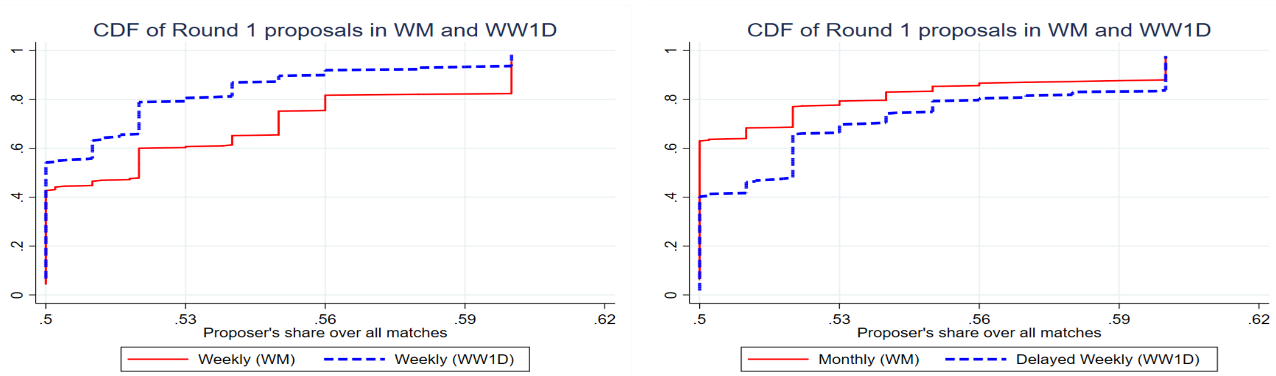

(a) Weekly as Proposer

(b) Weekly as Respondent

Figure 4: Round-1 Proposals in Treatments *WM* and *WW1D*

We first revisit the generally predicted patience advantage as in Prediction (B1). Figure 4(a) com-

[28]OA Tables 5 and 6 corroborate these qualitative conclusions with alternative tests controlling for heterogeneity and session effects.

[29]The p-values are: 0.570 and 0.019 in *WM*; 0.639 and 0.030 in *WM2D*; 0.936 and 0.017 in *WW1D*. Further supporting evidence that both types in *WM2D* propose symmetric/equal splits "too often," as opposed to demanding more (weekly) or less (monthly) instead, can be discerned from OA Table 4.

pares weekly demands from monthly bargainers, as in Treatment *WM* (solid), and from delayed weekly bargainers, as in Treatment *WW1D* (dashed). Consistent with the general theoretical prediction, the former distribution clearly dominates the latter. This dominance is also highly statistically significant (p-values 0.845 and 0.001).

Figure 4(b) compares initial proposals to weekly bargainers, as the initial respondents. It shows monthly demands from weekly bargainers, as in Treatment *WM* (solid), and delayed weekly demands from weekly bargainers, as in Treatment *WW1D* (dashed). Again, consistent with the general theoretical prediction, the former very clearly dominate the latter, and this difference is highly statistically significant (p-values 0.885 and 0.001). Hence, we fully and strongly confirm the general prediction (B1) of both EXD and QHD (and also HYD).

Result 4 (Basic Patience Advantage (B1))**.** *Comparing weekly bargainers against different opponent types, (i) as initial proposers,* $F(x_W^{WM})$ *dominates* $F(x_W^{WW1D})$*, and (ii) as initial respondents,* $F(x_{WD}^{WW1D})$ *dominates* $F(x_M^{WM})$*, confirming both parts of the generally predicted basic patience advantage between Treatments WM and WW1D.*

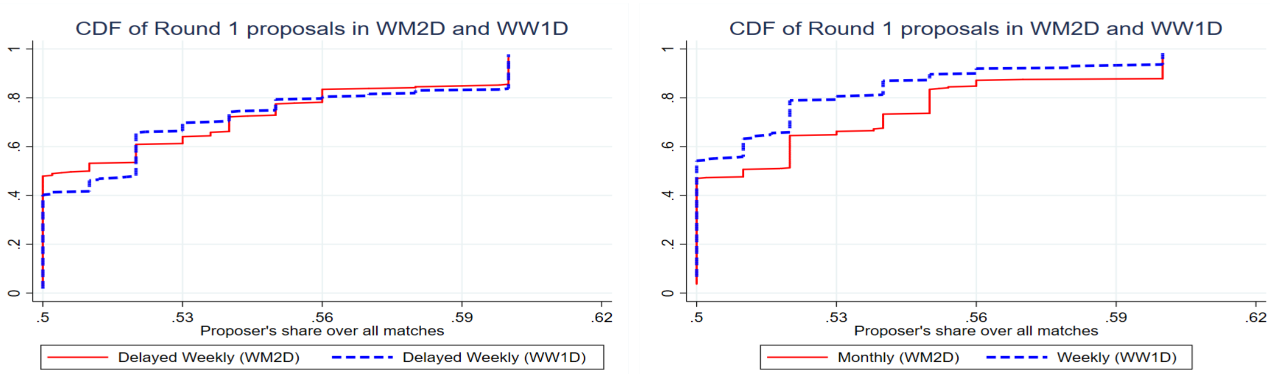

(a) Delayed Weekly as Proposer

(b) Delayed Weekly as Respondent

Figure 5: Round-1 Proposals in Treatments *WM2D* and *WW1D*

Finally, we turn to Prediction (B2) under EXD, vs. its (in one part) more permissive qualification under QHD, vs. the yet more permissive HYD. Figure 5(a) compares delayed weekly demands from delayed monthly bargainers, as in Treatment *WM2D* (solid), and those from weekly bargainers (with no delay), as in Treatment *WW1D* (dashed). In this case, EXD predicts an unambiguously greater advantage against delayed monthly bargainers, whereas QHD makes no general prediction, as weekly bargainers with no delay may also be weaker respondents than delayed monthly bargainers if present bias is sufficiently strong. We find the EXD prediction rejected: while there is no obvious qualitative dominance relationship between these distributions, the opposite dominance—greater delayed weekly demands from weekly bargainers than from delayed monthly bargainers—is marginally statistically significant (p-values 0.239 and 0.090).

Figure 5(b) then compares demands *from* delayed weekly bargainers, those by delayed monthly bargainers, as in Treatment *WM2D* (solid), and those by weekly bargainers (with no delay), as in Treatment *WW1D* (dashed). Since an initial proposer's present bias is irrelevant to equilibrium, both EXD and QHD imply that weekly bargainers should claim more than delayed monthly bargainers from the same opponent type, here a delayed weekly bargainer. Our data reject this, however. While the distributions are related via statistically significant dominance, it is the opposite of the predicted direction (p-values 0.748 and 0.001), meaning delayed monthly bargainers demand more from delayed weekly ones than the weekly bargainers (with no delay) do.

Result 5 (Diminishing Impatience (B2))**.** *Comparing delayed weekly bargainers against different opponent types, (i) as initial proposers,* $F(x_{WD}^{WM2D})$ *does not dominate* $F(x_{WD}^{WW1D})$*, and (ii) as initial respondents,* $F(x_{MD}^{WM2D})$ *dominates* $F(x_{W}^{WW1D})$*, both rejecting EXD and (ii) rejecting also the more permissive QHD, while the yet more permissive HYD could not be rejected.*

We therefore find that delayed weekly bargainers perceive their bargaining power to be roughly similar against delayed monthly bargainers and against weekly bargainers, whereas delayed monthly bargainers perceive their bargaining power against delayed weekly bargainers to be greater than weekly bargainers. The former finding rejects prediction (B2) under EXD but is consistent with QHD, suggesting once again a pronounced present bias; the latter finding, however, rejects that part of prediction (B2) under QHD where it agrees with EXD. Now recall that HYD makes the very same predictions as QHD that the previous results confirm, and it can additionally rationalize this last finding (cf. Section 2.2). When impatience diminishes at the appropriate speed, the front-end delay of the monthly bargainer may put this type into an even stronger position than the weekly one as the initial proposer.

Of course, this latter conclusion from between-treatments tests involving *WM2D* is especially subject to the earlier caveat, hence rather suggestive. Indeed, when analogously comparing proposals demanding more than half the cake, we confirm this latter dominance but also find statistically significant dominance as predicted by EXD (B2) for delayed weekly bargainers as initial proposers.[30]

3.3 Additional Results and Discussion

Our experimental design and interpretation of empirical findings are informed by bargaining theory that assumes perfect information and also selfish preferences. We now discuss the potential limitations of our conclusions, as they depend on theoretical assumptions, especially concerning what bargaining behavior reveals about people's time preferences. In doing so, we also offer some additional results.

[30]The p-values are: 0.836 and 0.002 for delayed weekly bargainers as initial proposers in *WM2D* vs. *WW1D*; 0.963 and 0.001 for delayed weekly bargainers as initial respondents in *WM2D* vs. *WW1D*. At the same time, we then only confirm Result 4 for weekly bargainers (with no delay) as initial proposers in *WM* vs. *WW1D* (p-values 0.885 and 0.001), while there is no significant difference in the distributions of proposals made to them as initial respondents in the two treatments (both p-values >0.339).

3.3.1 Immediate vs. Delayed Agreement, and Incomplete Information

A general implication of the perfect information theory that informs our bargaining predictions is immediate agreement (see Proposition 1 in Appendix A.1). This counterfactual prediction in Rubinstein (1982) has led to the development of the theory of bargaining under incomplete information to explain delay (see, e.g., the surveys Ausubel, Cramton, and Deneckere, 2002; Fanning and Wolitzky, 2022). While the focus of this paper lies on predictions about *relative* bargaining power, it is nonetheless instructive to consider the incidence of delay to better understand our results with regard to the potential role of incomplete information in generating them.

Recall for this purpose that our design makes disagreement costly via both significant time delays of payoffs and a sizeable 25% chance of exogenous termination resulting in zero payoffs. Thus, we pushed participants towards trying to reach immediate agreement rather than exploit incomplete information. Basically, when testing the theoretical predictions, we assume that any effects due to incomplete information are constant in our comparisons, analogous to independent noise in behavior.

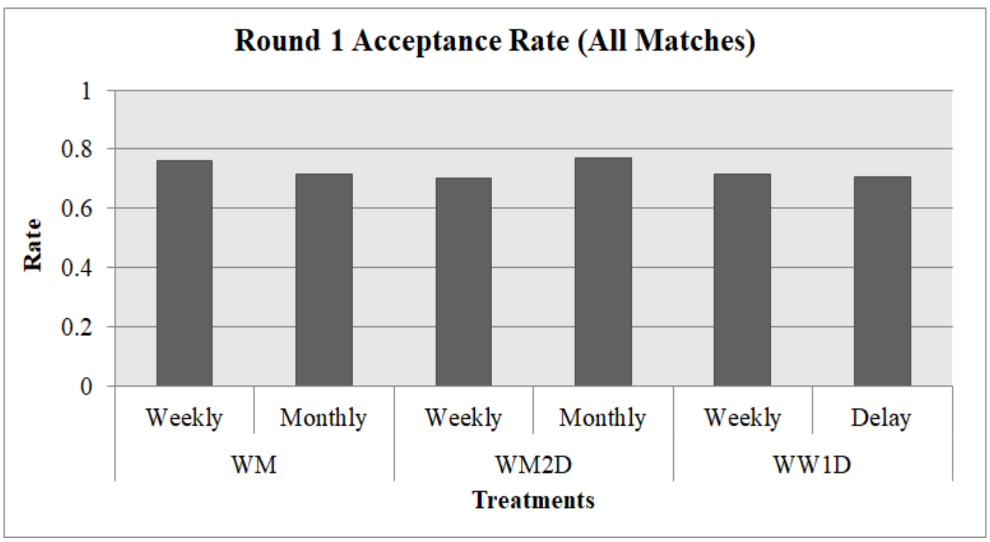

Figure 6: The Proportions of Immediate Agreements – All Matches

Figure 6 shows that the rate of immediate agreement is (i) relatively high overall, approximately 75%, and (ii) similar both across treatments and across the two versions of the game within treatments, always remaining strictly between 70% and 80%. Notably—and unsurprisingly, given the above—Treatment *WM2D* has both the highest acceptance rate, by delayed monthly respondents, and the lowest acceptance rate, by delayed weekly respondents. Hence, there is a fair amount of initial disagreement, and the role of incomplete information appears non-negligible in this particular respect. However, our design appears to have been successful in keeping its effects relatively mild overall and roughly constant for the purpose of all our comparisons. Recall also that we obtain similar within-treatment results for accepted proposals as for all proposals, so they are not driven by rejected proposals.

OA Figures 19 and 20 provide further detail on the proportions of agreements including later rounds (aggregating over the first 5 and last 5 matches, respectively). These are again similar between treatments. Overall, the proportions of agreement before random termination are 91.8%, 89.3% and 90.5% for the three treatments *WM*, *WM2D* and *WW1D*, respectively. The average number of rounds for agreement is only slightly above 1.3 overall and does not differ between treatments (Mann-Whitney test, p-values > 0.5). These observations indicate that initial proposals are indeed informative about perceptions of relative bargaining power, incorporating similar trade-offs between obtaining a greater share and a greater risk of rejection due to incomplete information.

3.3.2 Equal Splits and Fairness Concerns

To what extent fairness concerns versus strategic considerations drive bargaining behavior has been a central as well as topical issue in the experimental bargaining literature from its very beginnings (see Roth, 1995). While some progress has recently been made with incomplete information theory (Fanning and Kloosterman, 2022; Fanning, 2022; Keniston, Larsen, Li, Prescott, Silveira, and Yu, 2024), it remains so to this date.[31] Our procedure induces differences in *effective* time preferences between ex ante identical groups of participants, which holds constant in particular their (beliefs about others') social preferences. Thus, for the purpose of testing predictions concerning relative bargaining power, our method permits causal inference based on participants' true time preferences, while sidestepping this issue.[32] The large proportion of equal-split proposals is notable, of course, and in line with the evidence from the large related literature employing different methods. However, it is relevant to our results only insofar as pro-social behavior limits the size (as opposed to sign) of effects due to time preferences. In view of the debate about the generalizability of laboratory evidence on pro-social behavior to the field (Levitt and List, 2007; Camerer, 2015), our results' contribution is qualitative in nature, a proof of concept.

Yet, the above is true only as long as our manipulation of time preferences does not interact with participants' fairness concerns. Observe now that an equal split is unambiguously the most fair division in Treatments *WM* and *WM2D*, where such an immediate agreement would result not only in equal payoffs but also equal delay of these payoffs. This is not the case in Treatment *WW1D*, however,

[31]This perspective raises the question of whether an equal split is actually the optimal proposal, see Fanning and Kloosterman (2022); e.g., based on overall averages (OA Table 4), weekly bargainers in our Treatment *WM* would thus obtain an expected payoff bounded below by the expected immediate payoff of \$25 times the Round-1 acceptance rate for equal splits of 0.955, which equals 23.875 and is thus even slightly above their realized average payoff (23.758) *without* accounting for delay discounting (OA Figure 17, while the overall Round-1 acceptance rate is only around 72%). What speaks against suboptimal experimentation with selfish proposals, however, is that weekly bargainers' proposals become more selfish over time: They propose fewer equal splits and also fewer altruistic proposals in the second half of the experiment, and even the conditional average of their selfish demands increases. Together with the fact that monthly respondents' acceptance of equal splits also falls somewhat over time, this suggests signaling as well as inferring (heterogeneous) time preferences.

[32]Our stylized theory therefore assumes "selfish" preferences. Since shares add up to one, this here only means utility increases more in own than in other's share, which covers a fairly large subset of plausible social preferences (see Charness and Rabin, 2002; Engelmann and Strobel, 2004).

because the payoff delay under immediate agreement differs by type; specifically, an equal split would have equal payoffs but one bargainer would receive it immediately after the session whereas the other would receive it with a delay of one week. What is evidence for present bias in Result 3 under the assumption of selfish preferences or payoff fairness concerns may therefore reflect (discounted) utility fairness concerns that only coincidentally operate like payoff fairness concerns in the other treatments.

Our design does not allow us to directly identify the extent to which this is true. In fact, the time dimension of social preferences has not received any attention until most recently. However, we see good reasons to doubt that such fairness concerns would fully explain the observed difference. First, recall that this difference is one of first-order stochastic dominance and highly significant, that we know from Treatment *WM* that time preferences as such do matter significantly and that present bias is an empirically well established property of time preferences. For fairness concerns of the kind just described to produce the entire difference in Treatment *WW1D*, they would therefore have to be widespread. The large fraction of equal-split proposals in the other two treatments suggests this may be true, but it would predict a similarly sizeable fraction of proposals demanding *strictly less than half the cake* by weekly bargainers in *WW1D*. This is not what we observe: Overall, the fraction initially offering more than half of the surplus tends to indeed be highest for the weekly bargainers in this treatment, but it is still small and decreases sharply with experience, from almost 9% in the first to just 4% in the last five matches; remarkably, the very same type of (weekly) bargainers in the different Treatment *WM*, where it could not be driven by fairness concerns of this kind, follows the opposite dynamic pattern with similar rates, and comparing this type's propensity to propose equal splits in the two treatments shows it is higher in *WW1D* than in *WM* (see OA Table 4 for further detail). These observations are in line with the findings on self-serving bias in the bargaining context (see Babcock and Loewenstein, 1997, for a survey), which imply that fair-minded weekly bargainers faced with the two conflicting norms of payoff fairness vs. utility fairness tend towards the norm that yields themselves a greater share, hence payoff fairness.

Regarding delayed weekly bargainers, on the other hand, the recent findings by Kölle and Wenner (2023) from studying time preferences over both individual and social payoffs imply that front-end delay of payoffs should make these bargainers more pro-social (see also Chopra, Falk, and Graeber, 2024). Social preferences would then tend to work in the opposite direction of the observed difference in Treatment *WW1D*. However, this pro-social effect of front-end delay could well play a role in explaining why the predicted difference in bargaining power favoring weekly over monthly bargainers is observed significantly only in Treatment *WM* but not in Treatment *WM2D*.

3.3.3 Elicited Time Preferences and Behavior

We elicited discounting measures for a subsample of our participants in all three treatments, always after all bargaining games were completed (see Online Appendix B.3 for details of the task, measurement and analysis). This allows us, on the one hand, to (at least partially) check whether treatment assignment was random in terms of underlying time preferences, which our data confirm, and, on the other hand, to also relate measured discounting to bargaining behavior.

Basic regression analysis indeed shows the expected correlations between discounting and initial proposals, such that individuals that discount less demand more (for β as well as δ). These correlations are very weak and hardly significant, however. While this may also be due to correlations of time preferences with other relevant aspects of preferences, such as attitudes towards risk or fairness, that work against each other, one likely reason is incomplete information. Especially in our design, where disagreement is rather costly, beliefs about the opponent are likely to be a major determinant of behavior (especially one's initial proposal), and these beliefs are controlled by the public payoff types, independent of one's own actual preferences.

4 Concluding Remarks

We see two approaches to the contribution of our paper, depending on prior beliefs regarding our results. To the extent that the latter would "have to be true," our main contribution consists in offering a method that successfully delivers them, against the background of the related literature's often quite negative findings. At a general level, we propose an approach to deriving empirical content from intentionally stylized models and identify causal effects of preferences related to a particular sub-domain, acknowledging the presence of various confounds and incomplete information. This approach is demonstrated here for time preferences, but may be fruitfully developed for other domains, such as risk. Moreover, for the particular setting of sequential bargaining and the role of time preferences studied here, it is straightforward to see how it may also be used to structurally induce and investigate incomplete information about time preferences.

Our theoretically cleanest and empirically most robust result establishes that being (perceived as) transparently more patient significantly increases one's bargaining power. This result contributes a fundamentally positive message to the large body of theoretical analyses of dynamic strategic interaction, where time preferences—with few exceptions, this means simply "the" (exponential) discount factor—are a key driver of behavior. It lends empirical support to the basic idea behind theoretical comparative statics exercises in this discount factor as reflecting comparative statics in patience.

Moreover, this result practically implies that more patient individuals (or those perceived as more patient, though we think these would ultimately coincide) will benefit more from bargaining opportunities. How important a role this plays in generating or exacerbating inequality depends on how important those bargaining opportunities are for individuals' long-run economic success. As far as we are aware of, this question has not received much attention in empirical economics research, except in relation to gender inequality.[33] We hope our work will help raise awareness of this question's

[33] See Footnote 4 in the introduction. The general issue of inequality of bargaining power is discussed prominently in classic works such as Adam Smith's *Wealth of Nations* and Alfred Marshall's *Principles of Economics* (see Dunlop and

importance and promote future empirical research that quantitatively addresses it.

Our other main findings are that the size/significance of this patience advantage depends on the availability of immediate payoffs and that front-end delay is beneficial in payoff terms. These results from bargaining behavior are well in line with the large amount of empirical evidence that patience importantly depends on whether one is subject to present bias, most parsimoniously captured by the quasi-hyperbolic (β, δ)-discounting model. This indirect/behavioral evidence that people not only recognize but also seem to exploit the present bias of others supports theoretical analyses of dynamic strategic interaction incorporating present-biased individuals. Indeed, it raises the interesting question of whether present-biased individuals could improve their bargaining outcomes by publicly committing to a front-end delay of any payoffs prior to bargaining.[34] This could be directly investigated via revealed preference in future research, with an extension of our design around Treatment *WW1D*, by adding a symmetric weekly-weekly *WW* treatment and a variant where one of the two players could publicly choose the front-end delay on payoffs.

However, as discussed earlier, the latter two findings from the first implementation of our novel method to bargaining are subject to some remaining caveats. In particular, they highlight the importance of carefully controlling beliefs in the face of incomplete information—in particular, of ensuring common knowledge of delayed payoff consequences with more intricate delay profiles (see Treatment *WM2D* and between-treatments comparisons)—and of better understanding the time dimension of social preferences, which is itself a novel and only emerging area of research but highly complementary to our method for experimentally identifying causal effects of time preferences in dynamic games (see Treatment *WW1D*). For instance, a natural avenue in this direction could be to revisit also the classic ultimatum game with variously delayed payoffs and link decisions to corresponding individual and dictator game choices (Forsythe, Horowitz, Savin, and Sefton, 1994).

Furthermore, while our experimental design with payoff delay but no bargaining delay has clear advantages for the purpose of testing theoretical predictions, it removes potentially important psychological aspects of real bargaining, like dynamic inconsistency (for theoretical analysis, see Schweighofer-Kodritsch, 2018). A natural extension to identify its effects, exemplifying the potential of our method in future research, would employ our design alongside an analogous longitudinal version with actual bargaining delay (e.g., our Treatment *WW1D* with literally weekly bargaining, where one player's payoff would come immediately upon agreement and the other's one week later).

Generally, our laboratory evidence indicates rather small such effects of basic patience as well as present bias in bargaining. Reliably quantifying these effects—relative to other drivers of behavior, such as fairness concerns—in natural bargaining settings seems to us an important and natural next step in future research, which will require careful design combining laboratory and field measures, ideally in a structural approach.

Higgins, 1942) and subject to debate among legal scholars (e.g., Barnhizer, 2005).

[34] We thank an anonymous reviewer for pointing this out to us.

References

Abreu, Dilip and Faruk Gul (2000), "Bargaining and reputation." *Econometrica*, 68, 85–117.

Akin, Zafer (2007), "Time inconsistency and learning in bargaining games." *International Journal of Game Theory*, 36, 275–299.

Ambrus, Attila, Eric Chaney, and Igor Salitskiy (2018), "Pirates of the Mediterranean: An empirical investigation of bargaining with asymmetric information." *Quantitative Economics*, 9, 217–246.

Andersen, Steffen, Glenn W. Harrison, Morten I. Lau, and E. Elisabet Rutström (2008), "Eliciting risk and time preferences." *Econometrica*, 76, 583–618.

Andreoni, James and Charles Sprenger (2012a), "Estimating time preferences from convex budgets." *American Economic Review*, 102, 3333–56.

Andreoni, James and Charles Sprenger (2012b), "Risk preferences are not time preferences." *American Economic Review*, 102, 3357–3376.

Augenblick, Ned, Muriel Niederle, and Charles Sprenger (2015), "Working over time: Dynamic inconsistency in real effort tasks." *The Quarterly Journal of Economics*, 130, 1067–1115.

Augenblick, Ned and Matthew Rabin (2019), "An experiment on time preference and misprediction in unpleasant tasks." *The Review of Economic Studies*, 86, 941–975.

Ausubel, Lawrence M., Peter Cramton, and Raymond J. Deneckere (2002), "Bargaining with incomplete information." In *Handbook of Game Theory with Economic Applications*, volume 3, chapter 50, 1897–1945, Elsevier.

Azrieli, Yaron, Christopher P. Chambers, and Paul J. Healy (2018), "Incentives in experiments: A theoretical analysis." *Journal of Political Economy*, 126, 1472–1503.

Babcock, Linda and Sara Laschever (2003), *Women Don't Ask: Negotiation and the Gender Divide*. Princeton University Press.

Babcock, Linda and George Loewenstein (1997), "Explaining bargaining impasse: The role of self-serving biases." *Journal of Economic Perspectives*, 11, 109–126.

Backes-Gellner, Uschi, Holger Herz, Michael Kosfeld, and Yvonne Oswald (2021), "Do preferences and biases predict life outcomes? evidence from education and labor market entry decisions." *European Economic Review*, 134, 103709.

Backus, Matthew, Thomas Blake, Brad Larsen, and Steven Tadelis (2020), "Sequential bargaining in the field: Evidence from millions of online bargaining interactions." *The Quarterly Journal of Economics*, 135, 1319–1361.

Barnhizer, Daniel D. (2005), "Inequality of bargaining power." *University of Colorado Law Review*, 76, 139–241.

Barrett, Garry F. and Stephen G. Donald (2003), "Consistent tests for stochastic dominance." *Econometrica*, 71, 71–104.

Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak (1964), "Measuring utility by a single-response sequential method." *Behavioral Science*, 9, 226–232.

Biasi, Barbara and Heather Sarsons (2022), "Flexible wages, bargaining, and the gender gap." *The Quarterly Journal of Economics*, 137, 215–266.

Bikhchandani, Sushil (1992), "A bargaining model with incomplete information." *The Review of Economic Studies*, 59, 187–203.

Binmore, Ken (1987), "Perfect equilibria in bargaining models." In *The Economics of Bargaining* (Ken Binmore and Partha Dasgupta, eds.), chapter 5, 77–105, Basil Blackwell.

Binmore, Ken, Avner Shaked, and John Sutton (1985), "Testing noncooperative bargaining theory: A preliminary study." *American Economic Review*, 75, 1178–1180.

Binmore, Ken, Joseph Swierzbinski, and Chris Tomlinson (2007), "An experimental test of Rubinstein's bargaining model.", URL https://discovery.ucl.ac.uk/id/eprint/14439/. Mimeo.

Bleichrodt, Han, Kirsten I. M. Rohde, and Peter P. Wakker (2009), "Non-hyperbolic time inconsistency." *Games and Economic Behavior*, 66, 27–38.

Bowles, Hannah Riley, Linda Babcock, and Lei Lai (2007), "Social incentives for gender differences in the propensity to initiate negotiations: Sometimes it does hurt to ask." *Organizational Behavior and Human Decision Processes*, 103, 84–103.

Browning, M. and P. A. Chiappori (1998), "Efficient intra-household allocations: A general characterization and empirical tests." *Econometrica*, 66, 1241–1278.

Camerer, Colin F. (2003), *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.

Camerer, Colin F. (2015), "The promise and success of lab–field generalizability in experimental economics: A critical reply to levitt and list." In *Handbook of Experimental Economic Methodology* (Guillaume R. Fréchette and Andrew Schotter, eds.), chapter 14, Oxford University Press.

Castillo, Marco, Jeffrey L. Jordan, and Ragan Petrie (2019), "Discount rates of children and high school graduation." *The Economic Journal*, 129, 1153–1181.

Chabris, Christopher F., David Laibson, Carrie L. Morris, Jonathon P. Schuldt, and Dmitry Taubinsky (2008), "Individual laboratory-measured discount rates predict field behavior." *Journal of Risk and Uncertainty*, 37, 237–269.

Chakraborty, Anujit (2021), "Present bias." *Econometrica*, 89, 1921–1961.

Chakraborty, Anujit, Evan M. Calford, Guidon Fenig, and Yoram Halevy (2017), "External and internal consistency of choices made in convex time budgets." *Experimental Economics*, 20, 687–706.

Chakraborty, Anujit, Yoram Halevy, and Kota Saito (2020), "The relation between behavior under risk and over time." *American Economic Review: Insights*, 2, 1–16.

Charness, Gary and Matthew Rabin (2002), "Understanding social preferences with simple tests." *The Quarterly Journal of Economics*, 117, 817–869.

Chatterjee, Kalyan and Larry Samuelson (1987), "Bargaining with two-sided incomplete information: An infinite horizon model with alternating offers." *The Review of Economic Studies*, 54, 175–192.

Chopra, Felix, Armin Falk, and Thomas Graeber (2024), "Intertemporal altruism." *American Economic Journal: Microeconomics*, 16, 329–357.

Coles, Melvyn G. and Abhinay Muthoo (2003), "Bargaining in a non-stationary environment." *Journal of Economic Theory*, 109, 70–89.

Compte, Olivier and Philippe Jehiel (2002), "On the role of outside options in bargaining with obstinate parties." *Econometrica*, 70, 1477–1517.

Dunlop, John T. and Benjamin Higgins (1942), ""Bargaining power" and market structures." *Journal of Political Economy*, 50, 1–26.

Ebert, Jane E. J. and Drazen Prelec (2007), "The fragility of time: Time-insensitivity and valuation of the near and far future." *Management Science*, 53, 1423–1438.

Engelmann, Dirk and Martin Strobel (2004), "Inequality aversion, efficiency, and maximin preferences in simple distribution experiments." *American Economic Review*, 94, 857–869.

Epper, Thomas, Ernst Fehr, Helga Fehr-Duda, Claus Thustrup Kreiner, David Dreyer Lassen, Søren Leth-Petersen, and Gregers Nytoft Rasmussen (2020), "Time discounting and wealth inequality." *American Economic Review*, 110, 1177–1205.

Ericson, Keith Marzilli and David Laibson (2019), "Intertemporal choice." In *Handbook of Behavioral Economics – Foundations and Applications 2* (B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds.), chapter 1, 1–67, Elsevier.

Fanning, Jack (2022), "Fairness and the Coase conjecture." *Journal of Economic Psychology*, 93, 102571.

Fanning, Jack and Andrew Kloosterman (2022), "An experimental test of the Coase conjecture: Fairness in dynamic bargaining." *The RAND Journal of Economics*, 53, 138–165.

Fanning, Jack and Alexander Wolitzky (2022), "Reputational bargaining." In *Bargaining: Current Research and Future Directions* (Emin Karagözoğlu and Kyle B. Hyndman, eds.), chapter 3, 35–60, Springer.

Fischbacher, Urs (2007), "z-tree: Zurich toolbox for ready-made economic experiments." *Experimental Economics*, 10, 171–178.

Forsyth, Patrick (2009), *Negotiation Skills for Rookies*. Marshall Cavendish Limited.

Forsythe, Robert, Joel L. Horowitz, Nathan E. Savin, and Martin Sefton (1994), "Fairness in simple bargaining experiments." *Games and Economic Behavior*, 6, 347–369.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue (2002), "Time discounting and time preference: A critical review." *Journal of Economic Literature*, 40, 351–401.

Güth, Werner and Martin G. Kocher (2014), "More than thirty years of ultimatum bargaining experiments: Motives, variations, and a survey of the recent literature." *Journal of Economic Behavior & Organization*, 108, 396–409.

Güth, Werner, Rolf Schmittberger, and Bernd Schwarze (1982), "An experimental analysis of ultimatum bargaining." *Journal of Economic Behavior and Organization*, 3, 367–388.

Haan, Marco A. and Dominic Hauck (2023), "Games with possibly naive present-biased players." *Theory and Decision*, 95, 173–203.

Halevy, Yoram (2008), "Strotz meets Allais: Diminishing impatience and the certainty effect." *American Economic Review*, 98, 1145–62.

Halevy, Yoram (2015), "Time consistency: Stationarity and time invariance." *Econometrica*, 83, 335–352.

Hall, Robert E. and Paul R. Milgrom (2008), "The limited influence of unemployment on the wage bargain." *American Economic Review*, 98, 1653–1674.

Heggedal, Tom-Reiel and Thomas McKay (2024), "Discounting in finite-time bargaining experiments.", URL https://drive.google.com/file/d/1R35v_QwvfC8fCO4D33qnpN7OTehxh7nF/view. Mimeo (June).

Ho, Kate and Robin S. Lee (2017), "Insurer competition in health care markets." *Econometrica*, 85, 379–417.

Keniston, Daniel, Bradley J. Larsen, Shengwu Li, J.J. Prescott, Bernardo S. Silveira, and Chuan Yu (2024), "Splitting the difference in incomplete-information bargaining: Theory and widespread evidence from the field." *International Economic Review (forthcoming)*.

Kim, Jeongbin (2020), "Discounting, dynamic consistency, and cooperation in an infinitely repeated game experiment.", URL https://www.dropbox.com/scl/fi/wwnej1rym46mwkrhxluhr/Dratf_5.0.pdf?rlkey=c8ng2spht52fmgal5tx8qxtwm&e=1&dl=0. Mimeo (June).

Kim, Jeongbin (2023), "The effects of time preferences on cooperation: Experimental evidence from infinitely repeated games." *American Economic Journal: Microeconomics*, 15, 618–637.

Kölle, Felix and Lukas Wenner (2023), "Is generosity time-inconsistent? Present bias across individual and social contexts." *The Review of Economics and Statistics*, 105, 683–699.

Korda, Philippe (2011), *The five golden rules of negotiation*. Business Expert Press, LLC.

Laibson, David (1997), "Golden eggs and hyperbolic discounting." *The Quarterly Journal of Economics*, 112, 443–478.

Levitt, Steven D. and John A. List (2007), "What do laboratory experiments measuring social preferences reveal about the real world?" *Journal of Economic Perspectives*, 21, 153–174.

Loewenstein, George and Drazen Prelec (1992), "Anomalies in intertemporal choice: Evidence and an interpretation." *The Quarterly Journal of Economics*, 107, 573–597.

Lu, Shih En (2016), "Self-control and bargaining." *Journal of Economic Theory*, 165, 390–413.

Manzini, Paola (2001), "Time preferences: Do they matter in bargaining?" Queen Mary University of London Working Paper No. 445.

Meier, Stephan and Charles Sprenger (2010), "Present-biased preferences and credit card borrowing." *American Economic Journal: Applied Economics*, 2, 193–210.

Mischel, W., Y. Shoda, and M. I. Rodriguez (1989), "Delay of gratification in children." *Science*, 244, 933–938.

Noor, Jawwad (2011), "Intertemporal choice and the magnitude effect." *Games and Economic Behavior*, 72, 255–270.

Ochs, Jack and Alvin E. Roth (1989), "An experimental study of sequential bargaining." *American Economic Review*, 355–384.

O'Donoghue, Ted and Matthew Rabin (1999), "Doing it now or later." *American Economic Review*, 89, 103–124.

Ok, Efe A. and Yusufcan Masatlioglu (2007), "A theory of (relative) discounting." *Journal of Economic Theory*, 137, 214–245.

Osborne, Martin J. and Ariel Rubinstein (1990), *Bargaining and Markets*. Academic Press.

Pan, Jinrui, Craig S. Webb, and Horst Zank (2015), "An extension of quasi-hyperbolic discounting to continuous time." *Games and Economic Behavior*, 89, 43–55.

Phelps, Edmund S. and Robert A. Pollak (1968), "On second-best national saving and game-equilibrium growth." *The Review of Economic Studies*, 35, 185–199.

Powell, Robert (2002), "Bargaining theory and international conflict." *Annual Review of Political Science*, 5, 1–30.

Rapoport, Amnon, Eythan Weg, and Dan S. Felsenthal (1990), "Effects of fixed costs in two-person sequential bargaining." *Theory and Decision*, 28, 47–71.

Roth, Alvin E. (1995), "Bargaining experiments." In *The Handbook of Experimental Economics* (John H. Kagel and Alvin E. Roth, eds.), chapter 4, Princeton University Press.

Rubinstein, Ariel (1982), "Perfect equilibrium in a bargaining model." *Econometrica*, 50, 97–109.

Rubinstein, Ariel (1985), "A bargaining model with incomplete information about time preferences." *Econometrica*, 53, 1151–1172.

Saito, Kota (2011), "Strotz meets Allais: Diminishing impatience and the certainty effect: Comment." *American Economic Review*, 101, 2271–75.

Schweighofer-Kodritsch, Sebastian (2018), "Time preferences and bargaining." *Econometrica*, 86, 173–217.

Schweighofer-Kodritsch, Sebastian (2022), "The bargaining trap." *Games and Economic Behavior*, 136, 249–254.

Sin, Isabelle, Steven Stillman, and Richard Fabling (2022), "What drives the gender wage gap? Examining the roles of sorting, productivity differences, bargaining and discrimination." *The Review of Economics and Statistics*, 104, 636–651.

Sprenger, Charles (2015), "Judging experimental evidence on dynamic inconsistency." *American Economic Review: Papers & Proceedings*, 105, 280–85.

Sunde, Uwe, Thomas Dohmen, Benjamin Enke, Armin Falk, David Huffman, and Gerrit Meyerheim (2022), "Patience and comparative development." *The Review of Economic Studies*, 89, 2806–2840.

Sutter, Matthias, Martin G. Kocher, Daniela Glätzle-Rützler, and Stefan T. Trautmann (2013), "Impatience and uncertainty: Experimental decisions predict adolescents' field behavior." *American Economic Review*, 103, 510—-31.

Takeuchi, Kan (2011), "Non-parametric test of time consistency: Present bias and future bias." *Games and Economic Behavior*, 71, 456–478.

van Damme, Eric, Reinhard Selten, and Eyal Winter (1990), "Alternating bid bargaining with a smallest money unit." *Games and Economic Behavior*, 2, 188–201.

Watson, Joel (1998), "Alternating-offer bargaining with two-sided incomplete information." *The Review of Economic Studies*, 65, 573–594.

Weg, Eythan, Amnon Rapoport, and Dan S. Felsenthal (1990), "Two-person bargaining behavior in fixed discounting factors games with infinite horizon." *Games and Economic Behavior*, 2, 76–95.

Zwick, Rami, Amnon Rapoport, and John C. Howard (1992), "Two-person sequential bargaining behavior with exogenous breakdown." *Theory and Decision*, 32, 241–268.

A Appendix: Theory and Proofs

This Appendix formally establishes the equilibrium uniqueness and characterization result referred to at the beginning of Section 2, underlying our experimental design and behavioral predictions. Furthermore, it has the full formulations and proofs for Predictions 1 and 2, as well as a discussion of other forms of discounting and generalizations.

A.1 Equilibrium Uniqueness and Characterization

Based on the main text's model description, we define an individual i's preferences over the domain of her agreement outcomes $(q,n) \in ([0,1] \times \mathbb{N}) \cup \{(0,\infty)\}$, where q denotes i's share under the agreement and n denotes the round in which it is reached, and where $(0,\infty)$ subsumes any infinite history (perpetual disagreement). We assume that i's preferences at any point in the game are represented by a single utility function

$$U_i(q,n) = d_i(n-1) \cdot u_i(q),$$

consisting of a delay discounting function d_i and an atemporal utility function u_i such that

1. (Delay Discounting) $d_i(0) = 1 > d_i(n) > d_i(n+1) > 0 = d_i(\infty)$ for all $n \in \mathbb{N}$;

2. (Atemporal Utility) $u_i : [0,1] \to [0,1]$ is continuous and strictly increasing from $u(0) = 0$ to $u(1) = 1$;[35]

3. (Intertemporal Utility) There exists $\alpha_i < 1$ such that for all $n \in \mathbb{N}$, and for all $q \in [0,1)$ and $q' \in (q,1]$,

$$u_i^{-1}\left(\delta_i(n) \cdot u_i\left(q'\right)\right) - u_i^{-1}\left(\delta_i(n) \cdot u_i(q)\right) \leq \alpha_i \cdot \left(q' - q\right),$$

   where $\delta_i(n) \equiv d_i(n)/d_i(n-1)$.

The discounting function $d_i(n-1)$ gives the discount factor for the total payoff delay associated with agreement being reached in round n, i.e., after $(n-1)$ rounds of disagreement. The expression $\delta_i(n)$ is the discount factor for the specific period of payoff delay caused by disagreement in round n; by property 1, it lies between zero and one. Note that $d_i(n) = \prod_{m=1}^{n} \delta_i(m)$ holds true.

Properties 1 and 2 define the bargaining problem: On the one hand, any round of disagreement causes (further) payoff delay, which is costly to both individuals because they are impatient, and on the other hand, each of them always wants more of the cake for herself.

Property 3 will guarantee uniqueness of equilibrium by ensuring that backwards-induction dynamics are well-behaved. It says that i's willingness to pay to avoid another round's payoff delay is always increasing in the amount that she would obtain in case of this delay. This property extends what has been termed "increasing loss to delay" (see the axiomatic formulation of Rubinstein, 1982 and its treatment in Osborne and Rubinstein, 1990, Chapter 3) or "immediacy" (see the utility formulation of Schweighofer-Kodritsch,

[35] The assumption that $u(1) = 1$ is a mere normalization and without loss of generality.

2018) to the non-stationary setting studied here, and it is implied by standard assumptions; e.g., u_i concave and $\sup_n \delta_i(n) < 1$.[36]

Our equilibrium notion for this extensive-form game of perfect information is that of subgame perfect Nash equilibrium (SPNE). SPNE outcomes of a more general version of this game, where bargaining is over a general time-varying surplus, are geometrically analyzed by Binmore (1987), who shows that the extreme utilities are obtained in history-independent SPNE. Coles and Muthoo (2003) establish existence for a version that also contains our model. We contribute here a uniqueness result and a characterization for general discounted utility where non-stationary discounting is the source of time-varying surplus, and we provide algebraic proofs.[37] The key result is the following lemma, from which equilibrium uniqueness and its general properties are immediate; r_n will denote the responding player in any round $n \in \mathbb{N}$ (so $r_n = 2$ for n odd, and $r_n = 1$ for n even).

Lemma 1. *There exists a unique sequence x_n such that, for all $n \in \mathbb{N}$,*

$$x_n = 1 - u_{r_n}^{-1}\left(\delta_{r_n}(n) \cdot u_{r_n}(x_{n+1})\right). \tag{A.1}$$

Proof. Define, for each player i, the function $f_i : [0,1] \to [0,1]$ as $f_i(U) = 1 - u_j^{-1}(U)$. If player j is the respondent and could obtain a fixed utility U by rejecting, then $1 - u_j^{-1}(U)$ is the maximal share of proposer i that j is willing to accept. Equation (A.1) then says that $x_n = f_{r_{n+1}}(\delta_{r_n}(n) \cdot u_{r_n}(x_{n+1}))$, whereby any sequence x_n corresponds to a history-independent equilibrium: in any round n, the proposing player offers share $1 - x_n$, thus keeping x_n for herself, and this is the smallest offer accepted by the responding player, who upon rejection would similarly capture x_{n+1}. (Note the indifference of the responding player, $u_{r_n}(1 - x_n) = \delta_{r_n}(n) \cdot u_{r_n}(x_{n+1})$.)

Take now any odd-numbered round N in which player 1 is the proposer, and consider the two extreme cases for responding player 2's continuation utility upon rejection: first, when it is minimal and equals zero, and second, when it is maximal and equals one. For each of these two cases, compute the implied backwards induction solution for the thus truncated game. Clearly, it has immediate agreement in every round, and starting from the respective extreme terminal values, it is characterized by the recursive equation (A.1) for all rounds up through round N. (The extreme shares $x_{N+1} = 0$ and $x_{N+1} = 1$ correspond to the extreme continuation utilities $U_2 = 0$ and $U_2 = 1$.) Define these two finite sequences as a_n^N and b_n^N, and—using

[36] Let u be concave, $q_0 < q_1$ and $\varepsilon > 0$. Then

$$\frac{u(q_0+\varepsilon) - u(q_0)}{\varepsilon} \geq \frac{u(q_1+\varepsilon) - u(q_1)}{\varepsilon} > \frac{\delta u(q_1+\varepsilon) - \delta u(q_1)}{\varepsilon}$$

for any $\delta < 1$. Moreover, if $u(q_0) = \delta u(q_1)$, then $u(q_0+\varepsilon) > \delta u(q_1+\varepsilon)$ follows immediately from the above. This is equivalent to $\varepsilon > u^{-1}(\delta u(q_1+\varepsilon)) - q_0$ and upon substituting $q_0 = u^{-1}(\delta u(q_1))$ to $\varepsilon > u^{-1}(\delta u(q_1+\varepsilon)) - u^{-1}(\delta u(q_1))$. Denoting $q \equiv q_1$ and $q' \equiv q_1 + \varepsilon$, and applying this to individual i's preferences, the third assumed property follows for any given n; $\sup_n \delta_i(n) < 1$ ensures boundedness away from equality across all n by ruling out that $\lim_{n\to\infty} \delta(n) \to 1$.

[37] In their axiomatic extension of quasi-hyperbolic discounting to continuous time, Pan, Webb, and Zank (2015) consider a bargaining application, which is a special case of our model.

assumption 3 with $\alpha \equiv \max\{\alpha_1, \alpha_2\}$—observe that

$$\begin{aligned}
\left|a_N^N - b_N^N\right| &= a_N^N - b_N^N \\
&= f_1(0) - f_1(\delta_2(N)) \\
&= u_2^{-1}(\delta_2(N)) - u_2^{-1}(0) \\
&\leq \alpha \cdot \delta_2(N) \\
\left|a_{N-1}^N - b_{N-1}^N\right| &= b_{N-1}^N - a_{N-1}^N \\
&= f_2(\delta_1(N-1) \cdot u_1(f_1(\delta_2(N)))) - f_2(\delta_1(N-1) \cdot u_1(f_1(0))) \\
&= u_1^{-1}(\delta_1(N-1) \cdot u_1(f_1(0))) - u_1^{-1}(\delta_1(N-1) \cdot u_1(f_1(\delta_2(N)))) \\
&\leq \alpha \cdot (f_1(0) - f_1(\delta_2(N))) \\
&\leq \alpha^2 \cdot \delta_2(N) \\
&\vdots \\
\left|a_1^N - b_1^N\right| &\leq \alpha^N \cdot \delta_2(N).
\end{aligned}$$

Clearly, $\left|a_1^{2n-1} - b_1^{2n-1}\right| \to_{n\to\infty} 0$ (recall that we use only odd-numbered rounds), and hence $\lim_{n\to\infty} a_1^{2n-1} = \lim_{n\to\infty} b_1^{2n-1}$, which proves the claim, since $a_1^{2n-1} \geq x_1 \geq b_1^{2n-1}$ for all n. □

Proposition 1. *There exists a unique equilibrium. This unique equilibrium is in history-independent strategies that imply immediate agreement in every round. It is characterized by the unique sequence x_n of Lemma 1 as follows: in round n, the respective proposer demands share x_n, and the respective respondent accepts a demand q if and only if $q \leq x_n$.*

Proof. Consider any odd-numbered round N in which player 1 is the proposer, and suppose the supremal equilibrium continuation utility of player 2 takes the highest possible value of 1. Then, there exists an equilibrium with the outcome that players agree in round 1, and proposing player 1 obtains share a_1^N, defined in the proof of Lemma 1. Similarly, supposing the infimal equilibrium continuation utility of player 2 takes the lowest possible value of 0, there exists an equilibrium with the outcome that players agree in round 1 and proposing player 1 obtains share b_1^N, defined in the proof of Lemma 1. Now, any equilibrium utility value U_1 of player 1 (as of round 1) satisfies $u_1(a_1)^N \geq U_1 \geq u_1\left(b_1^N\right)$, whereby Lemma 1 proves its uniqueness. A similar argument proves the uniqueness of player 2's equilibrium utility. Both are uniquely obtained in the immediate-agreement equilibrium characterized by the sequence of Lemma 1. □

Proposition 1 delivers a general characterization of SPNE. It has the familiar property that in each round, the proposer makes the smallest acceptable offer to the respondent, given the unique continuation agreement that results upon rejection. Hence, in terms of time preferences as of a given round n, only the respondent's discount factor for that round's delay $\delta_{r_n}(n)$ enters the equilibrium outcome. It is easily verified that in the special case where the payoff delay per round of agreement is constant and both players discount exponentially, the infinite sequence in (A.1) indeed reduces to two equations as in Rubinstein (1982):

$$x_1 = 1 - u_2^{-1}\left(\delta_2 \cdot u_2\left(x_2\right)\right),$$
$$x_2 = 1 - u_1^{-1}\left(\delta_1 \cdot u_1\left(x_1\right)\right).$$

The only substantial assumption we impose on preferences is dynamic consistency *across rounds* of bargaining, i.e., that there is a single utility function representing an individual's preferences at any point in the game. This implies that only the *payoff delay* due to disagreement matters, the *bargaining delay* (i.e., the time delay until the next round) is irrelevant.[38] Given this, our abstract formulation of preferences in terms of rounds of agreement actually captures a huge variety of protocols and preferences. For instance, even assuming symmetric exponential discounting, if the payoff delay due to the first round of bargaining is longer than that due to any later round where it is constant, then preferences take the quasi-hyperbolic form of $d_i(n-1) = \beta\delta^{n-1}$ $(n > 1)$, even though time preferences are dynamically consistent. Under this assumption, we therefore essentially cover any combination of time preferences and payoff (as well as bargaining) timings, and we establish a very general result regarding equilibrium uniqueness and structure.[39]

In view of the vast body of evidence on time preferences, dynamic consistency of our experimental participants' time preferences would be hardly tenable as an assumption.[40] What we impose, however, is only dynamic consistency of preferences across bargaining rounds, which is satisfied in the limiting case of frequent offers where bargaining delay is negligible.[41] Then, a single dated self of any individual makes all the strategic decisions and only this one temporal snapshot of preferences matters (sometimes called "commitment preferences").

A.2 Predictions 1 and 2

Here we first formulate and then prove Predictions 1 and 2, summarized in the main text in Table 2. In the formulation, we say "any symmetric EXD (resp., QHD)" for any symmetric time preferences over delayed payoffs/shares representable as $U(q,t) = \delta^t u(q)$ (resp., $U(q,t) = \beta^{\mathbb{I}_{t>0}}\delta^t u(q)$) with function $u : [0,1] \to \mathbb{R}_+$ continuously increasing from $u(0) = 0$ and concave, and parameter $\delta \in (0,1)$ (resp., parameters $(\beta,\delta) \in (0,1)^2$). Thus preferences satisfy all three properties in A.1 above, and the general results there apply. (Recall that with regards to the exogenous breakdown risk, we assume expected utility, but see A.3 below for robustness.)

In the proofs, we employ the following notation: First, define $f(U) \equiv 1 - u^{-1}(U)$ for any $U \in [0,1]$; second,

[38]Formalizing bargaining and payoff delay requires explicitly accounting for time. Suppose round n takes place at date τ_n and agreement in round n results in payoffs at date t_n, where both τ_n and t_n are increasing sequences, such that $\tau_n \leq t_n$ holds (bargaining is never about past payoffs). The bargaining and the payoff delay due to disagreement in round n are, respectively, the delay from date τ_n to date τ_{n+1} and the delay from date t_n to date t_{n+1}. The statement that bargaining delay is irrelevant formally says that, for given t_n, any τ_n such that $\tau_n \leq t_n$ yields the same game.

[39]We focus on the separable case of *discounted* utility merely to simplify the notation. It is relatively straightforward to formulate the three assumed properties for non-separable preferences and to then generalize our uniqueness and characterization result using the same line of proof.

[40]See, however, Halevy (2015) for evidence that some violations of exponential discounting may be due to *time variance* rather than dynamic inconsistency of discounting.

[41]As long as payoff delay remains significant, the model is not susceptible to the "smallest-units" critique of van Damme, Selten, and Winter (1990).

as in Section 2.2, we let δ denote the basic weekly discount factor, but we introduce $\phi_i \equiv \gamma\cdot(\mathbb{I}_i(W)+\mathbb{I}_i(M)\delta^{k-1})$ so that individual i's effective basic discount factor accounting also for the breakdown risk equals $\phi_i\delta$; finally, when considering QHD and $\beta<1$ for Prediction 2, we let $\beta_i \equiv \mathbb{I}_i(D)+(1-\mathbb{I}_i(D))\beta$.

A.2.1 Prediction 1

Prediction 1. *Any symmetric EXD, for any $\gamma\in(0,1)$ (in particular $\gamma=3/4$), implies:*

(A1) Treatment WM: $x_W^{WM} > x_M^{WM}$ (equivalently, $y_W^{WM} > y_M^{WM}$).

(A2) Treatment WM2D: $x_{WD}^{WM2D} > x_{MD}^{WM2D}$ (equivalently, $y_{WD}^{WM2D} > y_{MD}^{WM2D}$).

(A3) Treatment WW1D: $x_{WD}^{WW1D} = x_W^{WW1D}$ (equivalently, $y_{WD}^{WW1D} = y_W^{WW1D}$).

(B1) Treatments WM and WW1D, with common type W: $x_W^{WM} > x_W^{WW1D}$ and $x_{WD}^{WW1D} > x_M^{WM}$ (equivalently, $y_{WD}^{WW1D} > y_M^{WM}$ and $y_W^{WM} > y_W^{WW1D}$).

(B2) Treatments WM2D and WW1D, with common type WD: $x_{WD}^{WM2D} > x_{WD}^{WW1D}$ and $x_W^{WW1D} > x_{MD}^{WM2D}$ (equivalently, $y_W^{WW1D} > y_{MD}^{WM2D}$ and $y_{WD}^{WM2D} > y_{WD}^{WW1D}$).

Proof. Proposition 1 implies that the unique equilibrium is characterized by

$$x_1^E = f\left(\phi_2\delta u\left(f\left(\phi_1\delta u\left(x_1^E\right)\right)\right)\right) \text{ and } x_2^E = f\left(\phi_1\delta u\left(x_1^E\right)\right), \tag{A.2}$$

where x_i^E is the share that individual i obtains in immediate agreement whenever she gets to propose. This share x_i^E obtains as the unique (and interior) fixed point of the function $g_i(q) \equiv f(\phi_j\delta u(f(\phi_i\delta u(q))))$, defined for any $q\in[0,1]$.[42] The characterization covers all matches of all treatments.

Observe now that $\phi_1 > \phi_2$ implies $g_1(q) > g_2(q)$ for all $q\in[0,1]$, and therefore $x_1^E > x_2^E$ (comparison of proposer shares), which is equivalent to $1-x_2^E > 1-x_1^E$ (comparison of respondent shares). Given A.2, this covers all parts except for (A3). The latter follows directly from the irrelevance of front-end delay under EXD. □

A.2.2 Prediction 2

Prediction 2. *Any symmetric QHD, for any $\gamma\in(0,1)$ (in particular, $\gamma=3/4$), implies the same as any symmetric EXD, except:*

(A3) Treatment WW1D: $x_{WD}^{WW1D} > x_W^{WW1D}$ (equivalently, $y_{WD}^{WW1D} > y_W^{WW1D}$).

(B2) Treatments WM2D and WW1D, with common type WD: $x_W^{WW1D} > x_{MD}^{WM2D}$, but no general prediction concerning x_{WD}^{WM2D} vs. x_{WD}^{WW1D} (equivalently, $y_{WD}^{WM2D} > y_{WD}^{WW1D}$, but no general prediction concerning y_W^{WW1D} vs. y_{MD}^{WM2D}).

[42] Our preference assumptions imply that each g_i is continuous and increasing from $g_i(0)>0$ through $g_i(1)<1$, whereby a fixed point exists and any fixed point is interior. Moreover, by our third preference assumption, each g_i has a slope less than one, so there is a unique fixed point.

Proof. The second-round continuation equilibrium is characterized by the shares x_i^E solving the two equations (A.2). Backward induction then yields immediate agreement in the first round, with the initial proposer's share given by

$$x_1^Q = f\left(\beta_2\phi_2\delta u\left(x_2^E\right)\right).$$

Regarding (A1), observe that *WM* has $\beta_1 = \beta_2 = \beta$ and that the respondent's continuation share is smaller for the monthly than the weekly bargainer from EXD. Hence, the initial proposer's share x_1^Q is greater (equivalently, the initial respondent's share $1 - x_1^Q$ is smaller) when the weekly bargainer initially proposes against the monthly bargainer than when the monthly bargainer initially proposes against the weekly bargainer.

Regarding (A2), observe that *WM2D* has $\beta_1 = \beta_2 = 1$, whereby predictions are as under EXD.

Regarding (A3), observe that when the weekly bargainer is the initial proposer, then $x_1^Q = x_1^E$, while when the weekly bargainer is the initial respondent, then $x_1^Q > x_1^E$.

Regarding (B1), observe that the weekly bargainer's continuation share is greater against the monthly bargainer (*WM*) than against the delayed weekly bargainer (*WW1D*), both as the initial proposer and as the initial respondent, from EXD. Hence, when the weekly bargainer is the initial respondent, $(\phi_2, \beta_2) = (p, \beta)$, $1 - x_1^Q$ is greater against the monthly bargainer, $(\phi_1, \beta_1) = \left(p\delta^{k-1}, \beta\right)$, than against the delayed weekly bargainer, $(\phi_1, \beta_1) = (p, 1)$. When the weekly bargainer is the initial proposer, a responding delayed weekly bargainer is unaffected by present bias, whereas a responding monthly bargainer is additionally weakened by it; this implication also follows for the between-treatment comparison of the weekly bargainer's shares as the initial proposer.

Regarding (B2), first observe that with the initial respondent's type equal to $(\phi_2, \beta_2) = (p, 1)$, her continuation share—hence also $1 - x_1^Q$—is smaller against the weekly than the monthly bargainer, as under EXD. Second, fixing $(\phi_1, \beta_1) = (p, 1)$, it should be clear from continuity that a violation of the prediction under EXD—meaning x_1^Q becomes smaller when $(\phi_2, \beta_2) = \left(p\delta^{k-1}, 1\right)$ than when $(\phi_2, \beta_2) = (p, \beta)$—is obtained as δ approaches one (so $\phi_2\delta = p\delta^k$ approaches $p = 3/4$) while β approaches zero. □

A.3 Other Forms of Discounting and Generalizations

A.3.1 Other Forms of Discounting

First, consider diminishing impatience, meaning $\delta_i(n)$ increases in n (cf. Section A.1 above). This implies a present bias, so a front-end delay increases such a discounter's bargaining power as the respondent. However, disagreement in round n adds a shorter basic payoff delay to a shorter existing delay for a weekly bargainer than for a monthly bargainer, meaning that for n large enough, even a monthly bargainer may in general become more patient than a weekly bargainer. This could resonate through the entire recursion of equation (A.1), thereby affecting the equilibrium outcome. Based on the intuition that discounting for the same additional delay would not change too quickly with the preceding delay (except for the immediate present) and in view of the sizable termination probability, we would assume that the effect of pushing a basic delay further into the future does not outweigh that of the basic delay being longer in determining the immediate equilibrium agreement. Indeed, the most general model of HYD proposed by Loewenstein and Prelec (1992), which we will identify with HYD here, imposes the structure of $d(t) = (1+\alpha\cdot t)^{-\beta/\alpha}$ for discounting time delays $t \geq 0$ with parameters $\alpha, \beta > 0$, and thus nesting neither EXD nor QHD.[43] However, this functional form still implies that a weekly bargainer always remains more patient than a monthly bargainer, and this extends also to when *both* are delayed.[44] We therefore immediately obtain the same predictions within Treatments *WM* and *WM2D*. Regarding Treatment *WW1D*, the delayed weekly bargainer is only further strengthened by diminishing impatience, so the prediction under QHD extends to HYD. For a similar reason—since the weekly bargainer is always more patient than the monthly bargainer, the delayed weekly bargainer is so—the prediction between Treatments *WM* and *WW1D*, which is common to both EXD and QHD, carries over with this model of HYD. However, it still depends on parameters whether the weekly bargainer or the delayed monthly bargainer is always more patient (though one can show that one is). This renders HYD yet more permissive with respect to the comparison between Treatments *WM2D* and *WW1D*.

Experimental studies from economics also document the opposite of present bias, namely (near-) future bias (see, e.g., Ebert and Prelec, 2007; Bleichrodt, Rohde, and Wakker, 2009; Takeuchi, 2011). Somewhat loosely, this means that the discounting function is initially concave (hump-shaped), in contrast to the convex discounting functions under EXD, QHD or HYD. While empirically documented, it is neither known how prevalent this bias is (hence, whether it could be reasonably expected to guide typical behavior) nor how far the *near* future extends from the immediate present (hence, whether a week's front-end delay would mute it). In view of these open issues, we omit a detailed analysis but note that if a near-future bias operates like an "inverted" present bias in the QHD model—i.e., $1 < \beta < 1/\delta$—then a front-end delay would make the initial respondent weaker rather than stronger. Hence, in Treatment *WW1D*, the weekly bargainer rather than the delayed weekly bargainer would be stronger, $x_W^{WW1D} > x_{WD}^{WW1D}$. Using that $\beta\delta^k < \delta$, which follows from $\beta < 1/\delta$ and $k > 1$, it is straightforward to show that the between-treatment predictions under such near-future bias coincide with those under EXD.

A.3.2 Generalizations

Chakraborty (2021) characterizes *all* (weakly) present-biased time preferences over payoffs/shares q with delays t in terms of their representability as

$$F(q,t) = \min_{u \in \mathcal{U}} u^{-1}(\delta^t u(q)),$$

where $\mathcal{U}$ is a set of (continuous and increasing) utility functions, and $F(q,t)$ is the immediate payoff/share that makes the decision maker indifferent to (q,t). When $\mathcal{U}$ is a singleton, we have EXD, while having multiple

[43] To be precise, EXD is only the limiting case where $\alpha \to 0$; as $\alpha \to \infty$, discounting approaches a step function, which would approximate only QHD subject to $\delta = 1$.

[44] The reason is that the different delays per round have a constant ratio, which also equals the ratio of total delays the two bargainers face in any agreement. Measuring time t in the unit that is the shorter delay per round and letting the corresponding type be type A, A's discount factor for round n is $\delta_A(n) \equiv d_A(n)/d_A(n-1) = [(1+\alpha\cdot n)/(1+\alpha\cdot(n-1))]^{-\beta/\alpha}$; letting the longer delay be $k > 1$ times the shorter delay with corresponding type B, B's discount factor for round n is $\delta_B(n) = [(1+\alpha\cdot kn)/(1+\alpha\cdot k(n-1))]^{-\beta/\alpha}$. Basic algebra yields $\delta_A(n) > \delta_B(n)$, and it is straightforward to check that the same holds true if both A and B face the same front-end delay. Breakdown risk is the same for both and anyways cancels out.

functions u with different slopes in the set $\mathcal{U}$ allows the representation to cover QHD and HYD as above. Given it essentially characterizes any form of present bias, it covers also more general discounting with diminishing impatience, and also time preferences where discounting is not well-defined due to non-separability, such as the magnitude-dependent discounting according to $U(q,t) = \delta(q)^t u(q)$, with $\delta(\cdot)$ an increasing function of the reward, proposed by Noor (2011). In addition to the issue explained above for general diminishing impatience, non-separability introduces another issue, namely multiplicity of (stationary) equilibrium, even given present bias and standard concavity of (all) u. This follows from the analysis by Schweighofer-Kodritsch (2018, Section 5.1), which covers all time preferences considered by Chakraborty (2021), and which more or less straightforwardly extends to the game with negligible bargaining delay considered here;[45] see especially his analysis of Noor's model, violating the key "immediacy" property (the analogue in the present non-stationary setting but separable utility is "intertemporal utility" in A.1). With such generalizations, therefore, either relative bargaining powers depend on various details of time preferences, or there is multiplicity of equilibrium. Arguably, however, the essence of present bias concerning behavior in our experiment is captured by QHD (and/or HYD), whereas those details of time preferences or how equilibrium would be selected are of relatively minor importance.[46]

While we have argued that our design rules out dynamic inconsistency through negligible bargaining delay, having an exogenous probability of breakdown introduces another potential source of dynamic inconsistency, namely risk preferences that violate expected utility. More specifically, Halevy (2008) shows how well-documented deviations from expected utility in decision making under risk translate into present bias and diminishing impatience in decision making over time, once the future is inherently uncertain (see also Saito, 2011; Chakraborty et al., 2020). Specifically, Halevy adds uncertainty to EXD with the model

$$U(q,t) = g(\gamma^t)\delta^t u(q),$$

where $g : [0,1] \to [0,1]$ is a probability weighting function. To allow for the certainty effect requires in particular that $g(\gamma)g(\gamma^{m-1}) \leq g(\gamma^m)$ for all $m \in \mathbb{N}$. Then, however, each individual's time preferences accounting for breakdown risk satisfy weak present bias, and Schweighofer-Kodritsch (2018, Proposition 1) directly applies to show that the unique equilibrium with fully sophisticated players then is as if there was EXD and expected utility, but where the breakdown probability equals $1 - g(\gamma)$ instead of $1 - \gamma$. This immediately extends to QHD, of course, so Predictions 1 and 2 remain unchanged (they are proven to hold for any $\gamma \in (0,1)$, and since $g(\gamma) \in (0,1)$ we can indeed simply replace γ by $g(\gamma)$). It is relatively easy to extend this result also to HYD.[47] For the purpose of predicting initial proposals, full naïveté is equivalent to dynamic consistency, thus covered by our analysis as $d(n) = g(\gamma^n)\delta^n$, whereby all our results can be applied, again depending on the shape of $g(\cdot)$. For rare (because conceptually challenging) approaches to bargaining implications with intermediate cases, see Akin (2007) and Haan and Hauck (2023).

[45]This applies also when accounting for breakdown risk, as with the corresponding extension in Chakraborty (2021, Section 6.).

[46]Our data confirm this, in the sense that—due to the prevalence of equal splits—there is anyways little scope for time preferences to matter via our manipulation.

[47]Given the above assumption on g, we may take any discounting such that the implied per-round discount factor of each player i, given by $\delta_i(n) \equiv d_i(n)/d_i(n-1)$ with d_i taken to exclude any breakdown risk, is weakly increasing.

B Online Appendix: Supplemental Material

This Online Appendix consists provides the following supplemental material: Part B.1 provides additional figures and tables that complement those provided in the main body of the paper; part B.2 provides full experimental instructions for all of our treatments and selected screenshots (Treatment WM); final part B.3 presents all details of our additional time preference elicitation and results on how measured time preferences relate to bargaining behavior.

B.1 Additional Figures and Tables

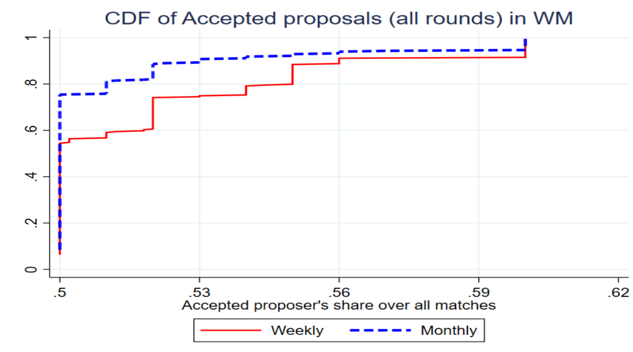

(a) Accepted Proposals

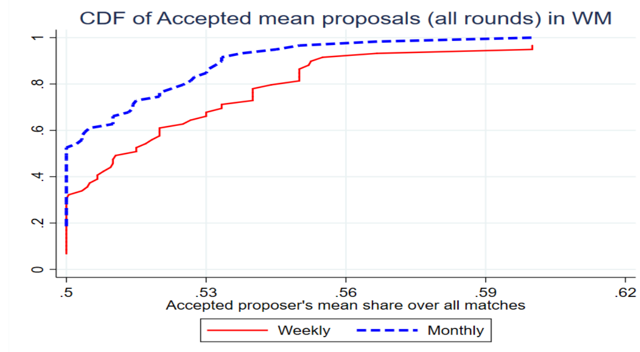

(b) Accepted Mean Proposals

Figure 7: Accepted Proposals over All Matches in Treatment *WM*

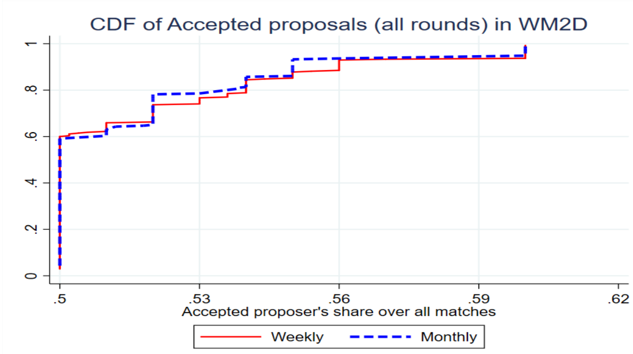

(a) Accepted Proposals

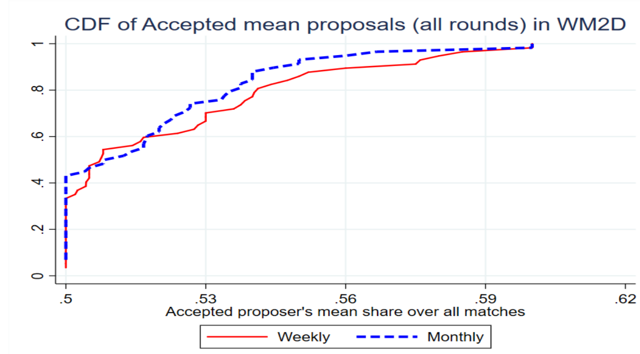

(b) Accepted Mean Proposals

Figure 8: Accepted Proposals over All Matches in Treatment *WM2D*

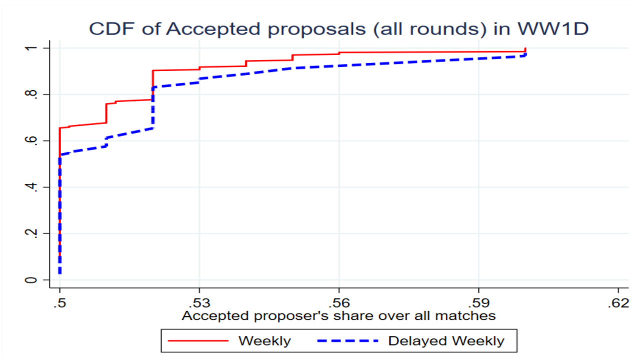

(a) Accepted Proposals

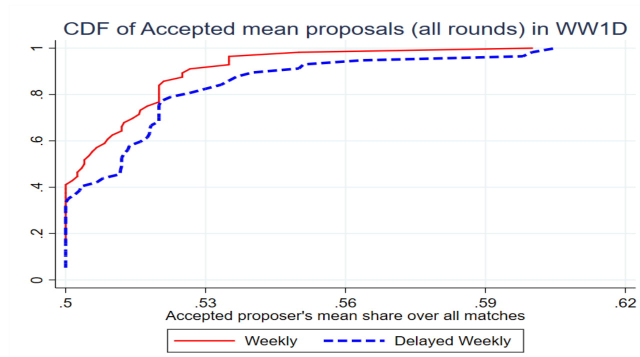

(b) Accepted Mean Proposals

Figure 9: Accepted Proposals over All Matches in Treatment *WW1D*

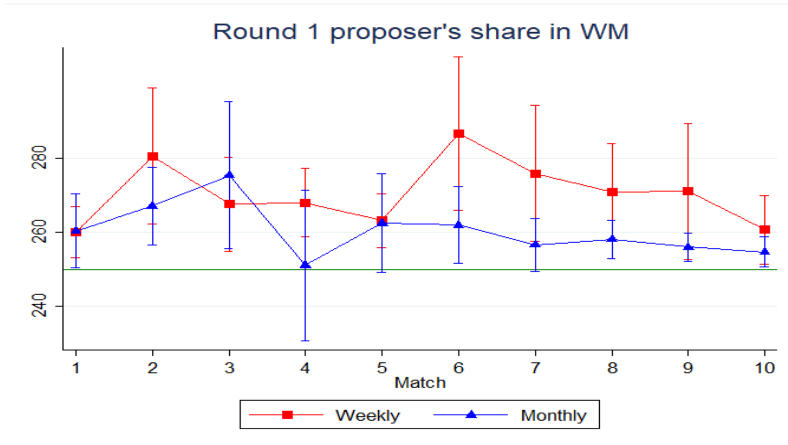

Figure 10: Round-1 Proposals over Matches in Treatment *WM*

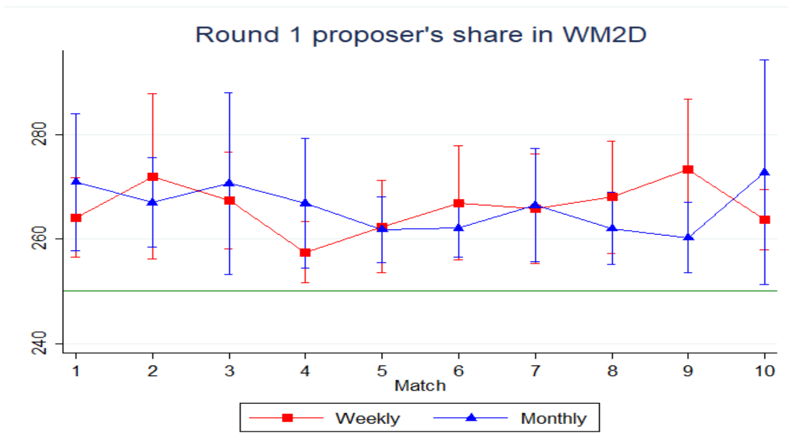

Figure 11: Round-1 Proposals over Matches in Treatment *WM2D*

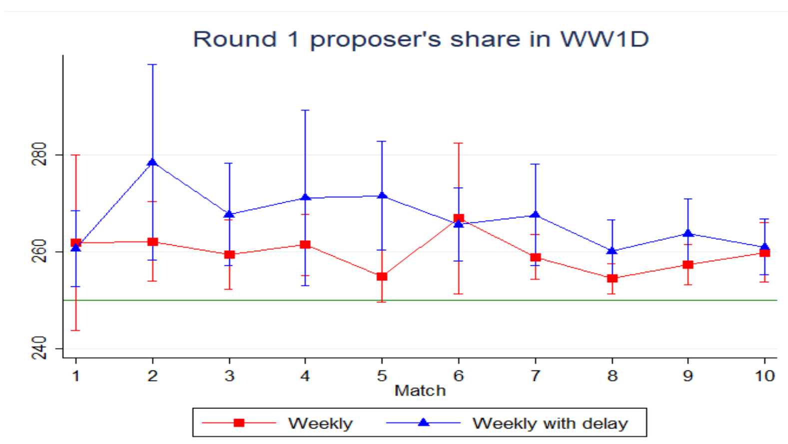

Figure 12: Round-1 Proposals over Matches in Treatment *WW1D*

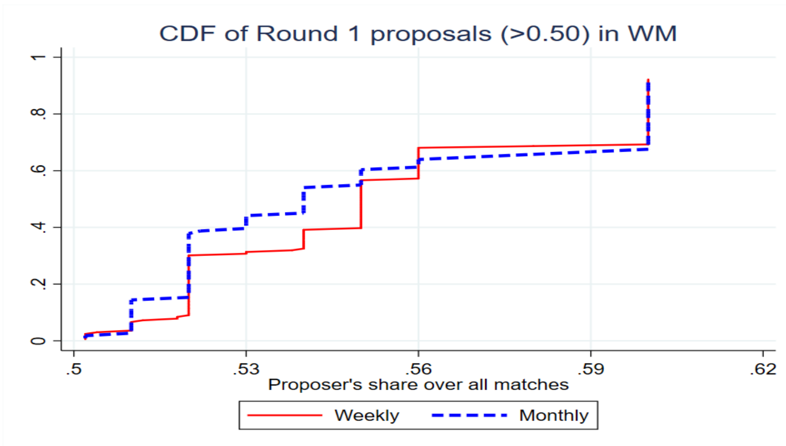

Figure 13: (All) Round-1 Proposals demanding > 0.50 in Treatment *WM*

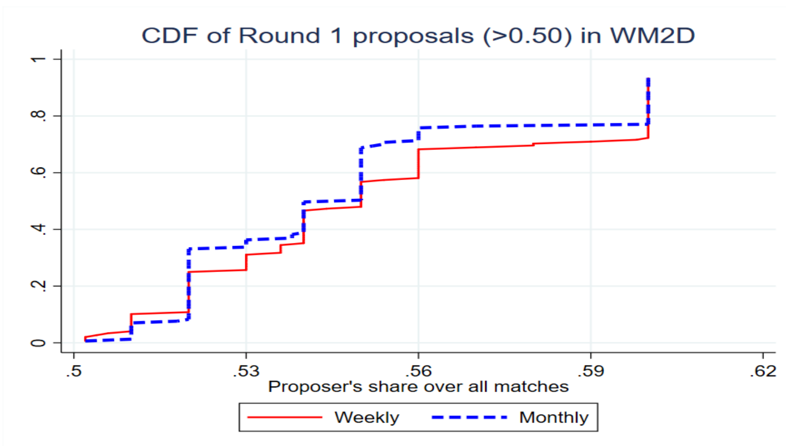

Figure 14: (All) Round-1 Proposals demanding > 0.50 in Treatment *WM2D*

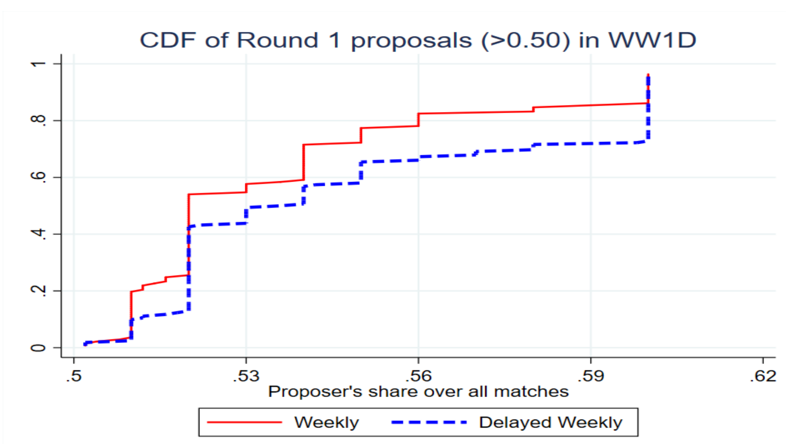

Figure 15: (All) Round-1 Proposals demanding > 0.50 in Treatment *WW1D*

Proposer Advantage

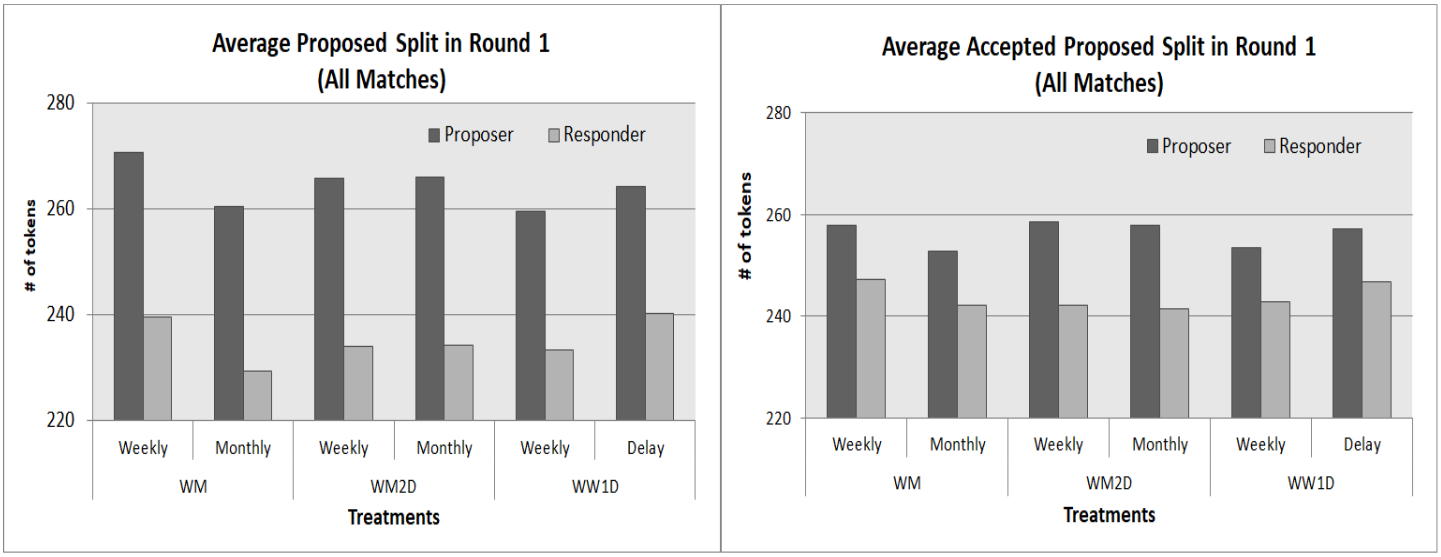

Figure 16: Initial Proposals and Accepted Proposals – All Matches

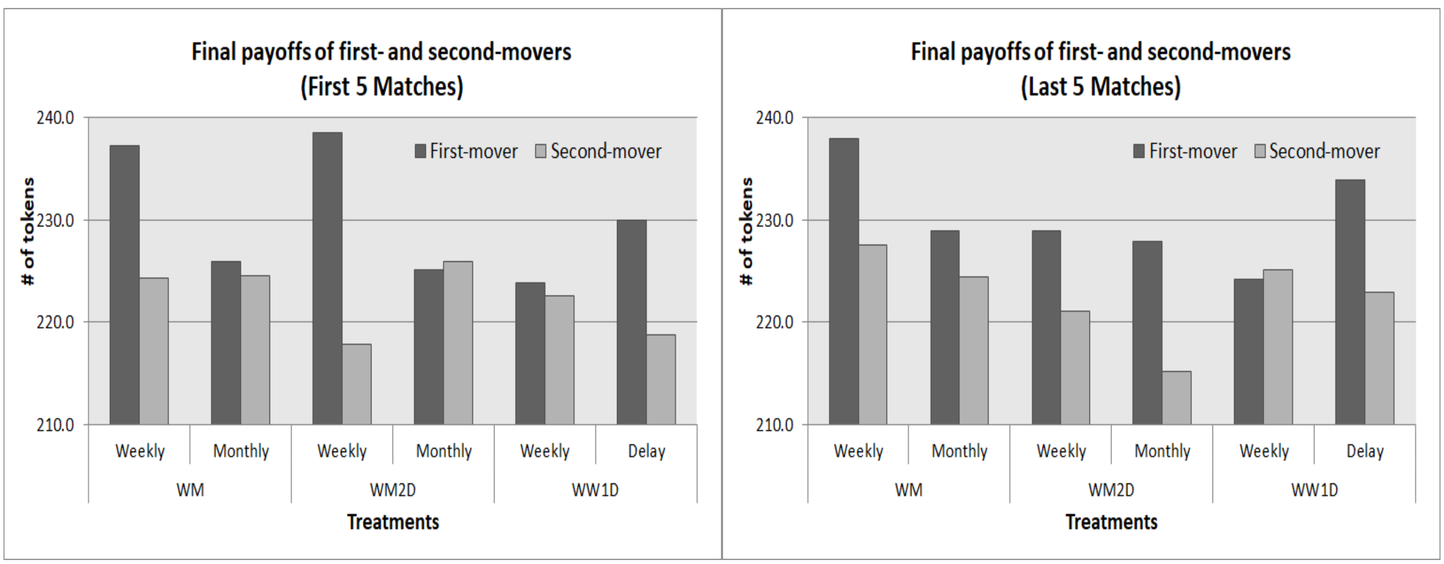

Figure 17: Final Payoffs (All, incl. Random Terminations) – First and Last 5 Matches

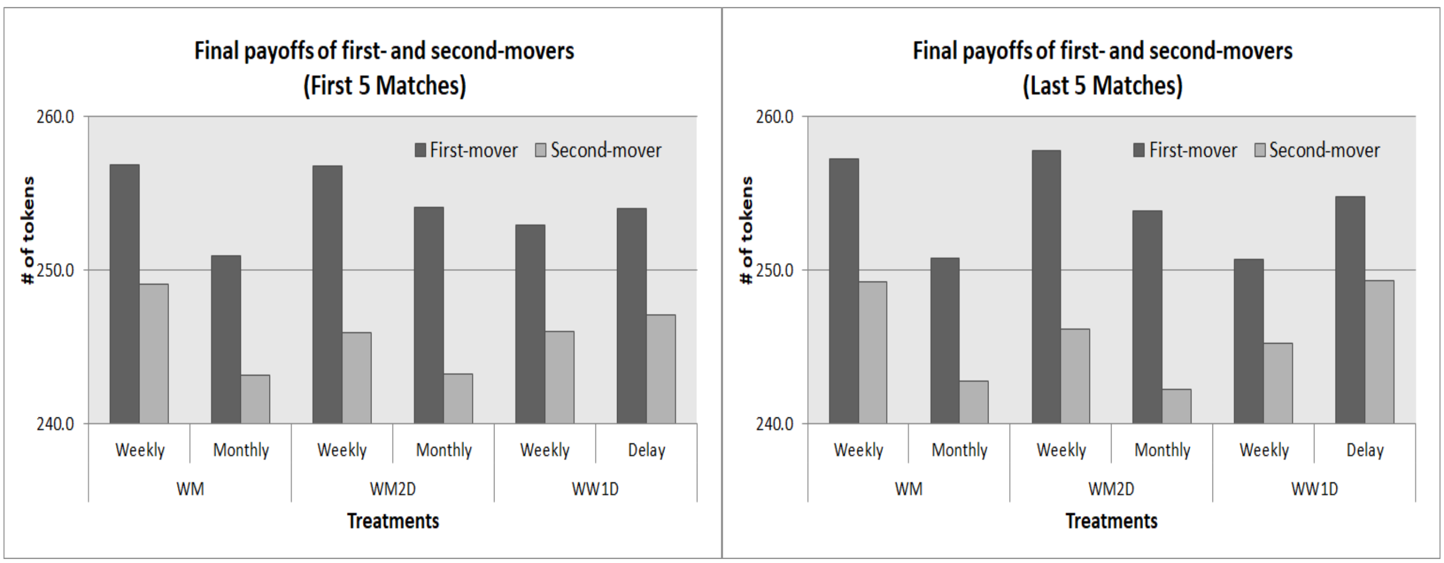

Figure 18: Final Payoffs (excl. Random Terminations) – First and Last 5 Matches

Immediate vs. Delayed Agreements

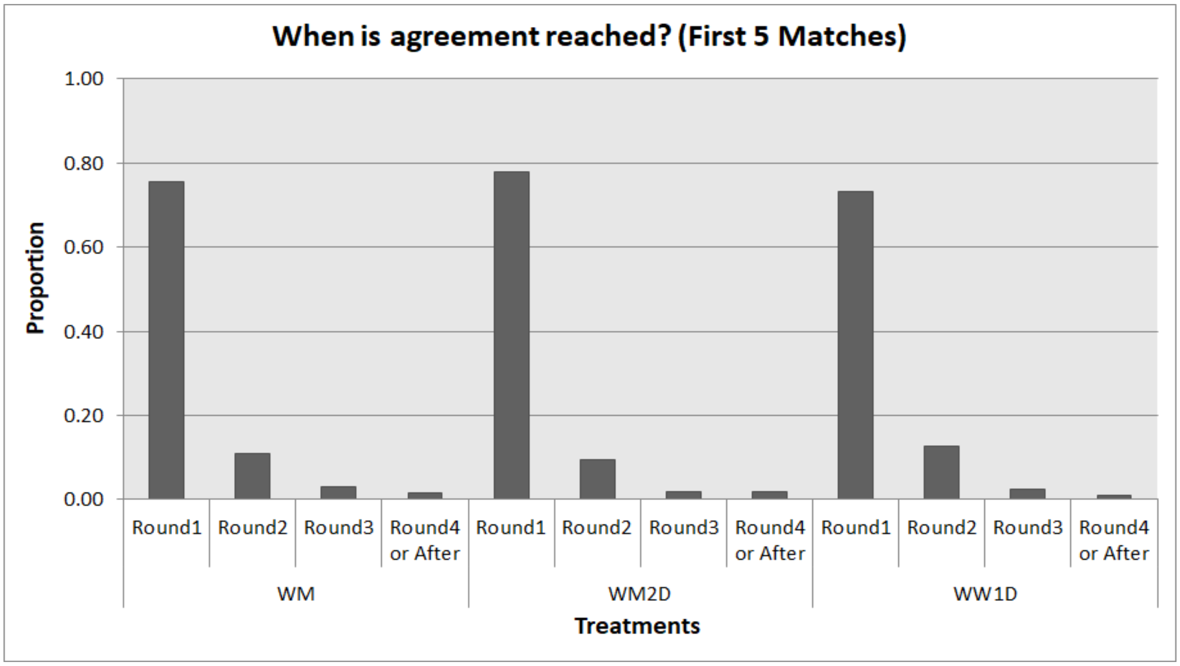

Figure 19: The Proportions of Agreements over Rounds – First 5 Matches

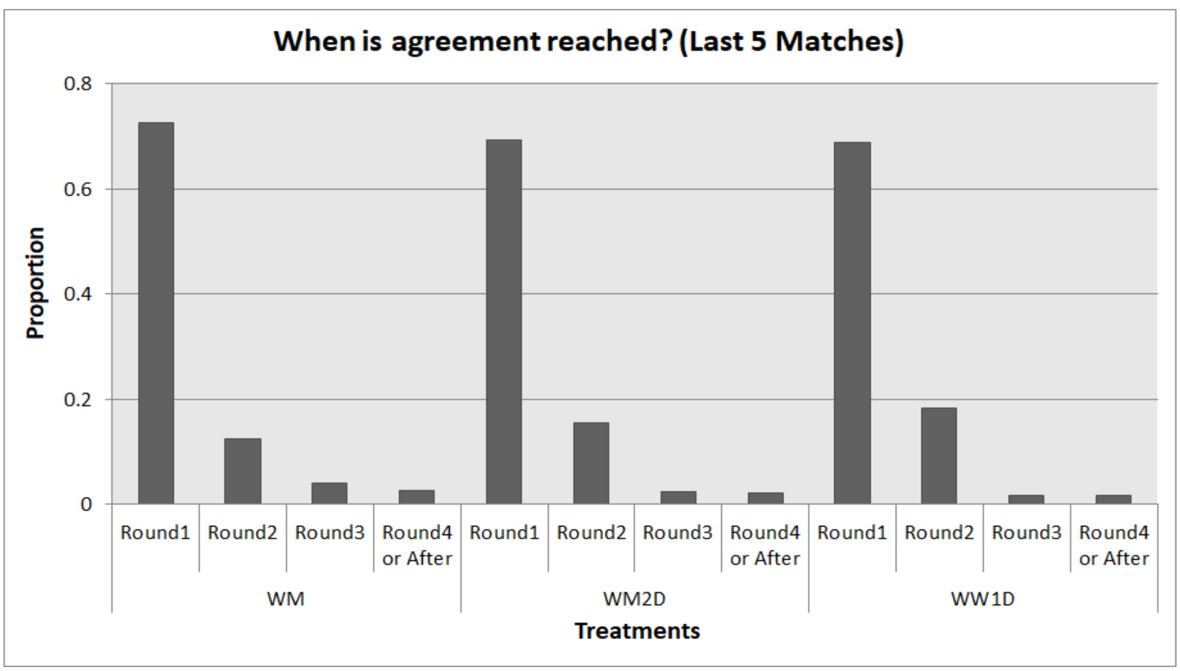

Figure 20: The Proportions of Agreements over Rounds – Last 5 Matches

Table 4: Further Descriptives: Initial Proposals and Acceptance Rates

Statistic/Matches	Treatment/Type						
	All	*WM*		*WM2D*		*WW1D*	
		W	*M*	*WD*	*MD*	*W*	*WD*
% equal-split proposals							
All ten	0.453	0.386	0.570	0.454	0.435	0.478	0.387
	(0.498)	(0.488)	(0.496)	(0.499)	(0.497)	(0.500)	(0.488)
First five	0.475	0.438	0.596	0.475	0.443	0.486	0.401
	(0.500)	(0.498)	(0.492)	(0.501)	(0.498)	(0.502)	(0.492)
Last five	0.432	0.336	0.544	0.434	0.429	0.470	0.373
	(0.496)	(0.474)	(0.500)	(0.497)	(0.497)	(0.501)	(0.485)
% proposals demanding < 0.50							
All ten	0.040	0.041	0.060	0.025	0.034	0.064	0.015
	(0.197)	(0.200)	(0.238)	(0.155)	(0.181)	(0.244)	(0.121)
First five	0.044	0.042	0.040	0.035	0.034	0.088	0.022
	(0.204)	(0.201)	(0.196)	(0.186)	(0.181)	(0.284)	(0.147)
Last five	0.037	0.041	0.081	0.014	0.034	0.040	0.007
	(0.188)	(0.199)	(0.279)	(0.182)	(0.182)	(0.196)	(0.086)
% accepted of all proposals							
All ten	0.729	0.717	0.763	0.771	0.703	0.716	0.705
	(0.444)	(0.451)	(0.426)	(0.421)	(0.458)	(0.452)	(0.457)
First five	0.756	0.771	0.742	0.801	0.758	0.764	0.701
	(0.430)	(0.422)	(0.439)	(0.400)	(0.430)	(0.426)	(0.460)
Last five	0.702	0.664	0.785	0.741	0.646	0.669	0.709
	(0.458)	(0.474)	(0.412)	(0.439)	(0.480)	(0.472)	(0.456)
% accepted of equal-split proposals							
All ten	0.945	0.955	0.930	0.984	0.922	0.923	0.981
	(0.225)	(0.207)	(0.256)	(0.124)	(0.268)	(0.267)	(0.137)
First five	0.954	0.968	0.922	1.000	0.924	0.931	1.000
	(0.210)	(0.177)	(0.269)	(0.000)	(0.267)	(0.256)	(0.000)
Last five	0.939	0.938	0.938	0.968	0.921	0.915	0.960
	(0.240)	(0.242)	(0.242)	(0.178)	(0.272)	(0.280)	(0.198)
Observations (all ten)	1,740	290	300	284	296	299	271
Observations (first five)	870	144	151	141	149	148	137
Observations (last five)	870	146	149	143	147	151	134

Notes: Table shows further descriptives statistics for all (ten) matches, the first five matches, and the last five matches, respectively, and with standard deviations in parentheses below. Observations refer to initial proposals by each type in each treatment. While there are equally many such observations from the first and last five matches (870 = 1740/2), the numbers differ for the two types in a given treatment, due to match-level randomization of the initial proposer. "% accepted of equal-split proposals" shows conditional rates, using as basis the corresponding number of equal-split proposals.

Table 5: Alternative Statistical Tests

	K-S test		OLS		OLS with controls		OLS with session FE	
	Panel (a)	Panel (b)	Panel (a)	Panel (b)	Panel (a)	Panel (b)	Panel (a)	Panel (b)
Figure 1	0.218**	0.248**	10.210**	8.891[b]	9.444*	8.933	9.701**	8.911[c]
			(4.958)	(5.614)	(5.534)	(6.002)	(4.851)	(5.594)
Figure 2	0.066	0.138	-0.251	-1.747	-0.646	-1.883	-0.748	-1.747
			(3.791)	(3.826)	(3.852)	(4.030)	(3.454)	(3.749)
Figure 3	0.183***	0.228[a]	-7.052**	-6.631*	-6.480*	−6.219[c]	-7.018**	-6.631*
			(3.252)	(3.558)	(3.335)	(3.756)	(3.266)	(3.594)
Figure 4	0.218***	0.228***	10.920**	-6.340*	9.499**	-4.344	16.542***	11.517
			(4.544)	(3.809)	(4.228)	(4.488)	(5.969)	(10.377)
Figure 5	0.083	0.149***	-0.880	6.422	-2.476	5.544[d]	-4.623	10.962
			(3.781)	(3.265)	(3.662)	(3.398)	(5.687)	(11.783)

Notes: Kolmogorov–Smirnov test, OLS with standard errors clustered at the individual level, OLS with standard errors clustered at the individual level and control variables (gender, economics major, race, years in college), and OLS with standard errors clustered at the individual level and session fixed effects. For the Kolmogorov–Smirnov test, the largest difference between the two distributions is presented. For the OLS, the independent variable is the treatment or type dummy variable and standard errors are reported in parentheses. For panel (b) of Figures 1–3 (individual mean proposals), OLS standard errors are not clustered.
***Significant at 1%; **5%; *10%.
a: p-value=0.103, *b*: p-value=0.115, *c*: p-value=0.101, *d*: p-value=0.105, *e*: p-value=0.114.

Table 6: Rank-Sum Tests (p-values) for Within-Treatment Comparisons on Session-Averaged Initial Proposals

Averaging Level (Observations)	Treatment		
	WM	*WM2D*	*WW1D*
Session (12)	0.055	0.873	0.055
Session×Match (120)	0.000	0.543	0.001

Notes: Recall that there were 6 sessions run per treatment, in which participants played 10 matches/games. For each of the two types in a treatment, Round 1 (initial) proposals are here averaged at (i) the session-level, averaging also over all matches, or at (ii) the session×match-level, yielding one session-average per match.

B.2 Experimental Instructions and Screenshots

This part provides full experimental instructions, followed by example screenshots of proposers as well as respondents, for all of our three treatments – *WM*, *WM2D*, and *WW1D*.

B.2.1 Treatment *WM*

Welcome to the experiment. Please read these instructions carefully; the payment you will receive from this experiment depends on the decisions you make. The amount you earn will be paid through **VENMO**.

Your Payment Type and Match

At the beginning of the experiment, one-half of the participants will be randomly assigned to be Payment **Type A** and the other half to be Payment **Type B**. Your payment type will remain fixed throughout the experiment. Your payment type will affect when you will be paid, which will be explained below.

The experiment consists of 10 **matches**. At the beginning of each match, one Type A participant and one Type B participant are randomly paired. The pair is fixed **within the match**. After each match, participants will be randomly repaired, and new pairs will be formed. You will not learn the identity of the participant you are paired with, nor will that participant learn your identity—even after the end of the experiment.

Your Decisions in Each Match

Round 1: At the beginning of Round 1, one participant will be randomly assigned to the role of a **proposer** and the other participant to the role of a **responder**. Each participant in a match has 50-50 chance to be the proposer and to be the responder regardless of his/her payment type.

The proposer is then asked to propose how to split 500 tokens (= $50) between the two participants as:

"_________ tokens for yourself and _________ tokens for the other person."

After observing the split proposed by the proposer, the responder decides whether to accept or reject the proposed split.

Outcome, Termination, and Transition to Next Round: The outcome of Round 1 depends on whether the split proposed by the proposer is accepted or rejected.

1. If the responder **accepts** the proposed split, both participants will receive the amounts of tokens as proposed, and the match will be terminated.
2. If the responder **rejects** the proposed split, then the match will proceed to the next round with 75% (3/4) chance or be terminated with 25% (1/4) chance. This is as if we were to roll a 100-sided die and continue if the selected number is less than or equal to 75 and end if the number chosen is larger than 75.
   (a) If a match is **terminated** after a rejection of a proposed split, both participants will receive 0 tokens for the match.
   (b) If the match **proceeds** to the next round, then the proposer-responder roles are alternated. That is, the participant who is the proposer in the current round will become the responder in the next round, and vice versa. The number of tokens the participants receive will be determined by the outcome of the subsequent rounds.

Round $K > 1$: In Round $K > 1$, the participant who was the proposer in Round $(K - 1)$ becomes the responder, and the participant who was the responder in Round $(K-1)$ becomes the proposer. The proposer is then asked to propose how to split 500 tokens (= $50) between the two participants. After observing the split proposed by the proposer, the responder decides whether to accept or reject the proposed split.

The rest of the procedures determining the outcome, termination of the round, and transition to next round, is the same as those in Round 1.

Information Feedback

- At the end of each **round**, you will be informed about the proposal made by the proposer and the accept/reject decision made by the responder.
- At the end of each **match**, you will be informed **when and how much** you are going to be paid.

Your Monetary Payments

At the end of the experiment, one match out of 10 will be randomly selected for your payment. Every match has an equal chance to be selected for your payment so that it is in your best interest to take each match seriously. Participants will receive the amounts of tokens according to the outcome from the selected match with the exchange rate of 1 token = $0.1.

When you are going to be paid depends on (1) your payment type and (2) the round in which the proposed split is accepted.

If you are **Type A**, you may be paid today or in a few **weeks**. If a proposed split is accepted in Round 1, you will be paid today right after the experiment. If a proposed split is accepted in Round 2, you will be paid in one week. If a proposed split is accepted in Round $K > 1$, you will be paid in $(K - 1)$ weeks.

If you are **Type B**, you may be paid today or in a few **months**. If a proposed split is accepted in Round 1, you will be today right after the experiment. If a proposed split is accepted in Round 2, you will be paid in one month. If a proposed split is accepted in Round $K > 1$, you will be paid in $(K - 1)$ months.

The following table summarizes the schedule of payment for each type:

If a proposed split is accepted in	Type A will be paid	Type B will be paid
Round 1	Today	Today
Round 2	In 1 week	In 1 month
Round 3	In 2 weeks	In 2 months
Round 4	In 3 weeks	In 3 months
Round 5	In 4 weeks	In 4 months
.....	.....	.....
Round K	In $(K-1)$ weeks	In $(K-1)$ months

Any amount you are supposed to receive will be paid electronically via VENMO.

In addition to your earnings from the selected match, you will receive a **show-up fee of $10** through VENMO, right after the experiment.

A Practice Match

To ensure your comprehension of the instructions, you will participate in a practice match. The practice match is part of the instructions and is not relevant to your cash payment; its objective is to get you familiar with the computer interface and the flow of the decisions in each round of a match. Once the practice match is over, the computer will tell you "The official matches begin now!"

Rundown of the Study

1. At the beginning of the experiment, your payment type will be randomly determined. Your payment type will remain fixed throughout the experiment.
2. At the beginning of each match, one Type A participant and one Type B participant are randomly paired.
3. At the beginning of Round 1, one participant will be randomly assigned to the role of a proposer and the other to the role of a responder.
4. The proposer then proposes how to split 500 tokens (= $50).
5. If the responder accepts the proposed split, both participants will receive the amounts of tokens as proposed, and the match will be terminated.
6. If the responder rejects the proposed split, then the match will proceed to the next round with 75% (3/4) chance or be terminated with 25% (1/4) chance. If a match is terminated after the rejection of a proposed split, both participants will receive 0 tokens for the match.
7. If the match proceeds to the next round, then the proposer-responder roles are alternated.
8. At the end of the experiment, one of 10 matches will be randomly selected for payment. For the selected match, the timing of your payment depends on (1) your payment type and (2) the round in which the proposed split was accepted. All your earnings will be paid to you through VENMO.
9. For Type A, you may be paid today or in a few weeks. For Type B, you may be paid today or in a few months.
10. In addition to your earnings from the selected match, you will receive a show-up fee of $10 right after the experiment.

Administration

Your decisions, as well as your monetary payment, will be kept confidential. Remember that you have to make your decisions entirely on your own; please do not discuss your decisions with any other participants. Upon finishing the experiment, you will be asked to sign your name to acknowledge your receipt of the payment. You are then free to leave. If you have any question, please raise your hand now. We will answer your question individually.

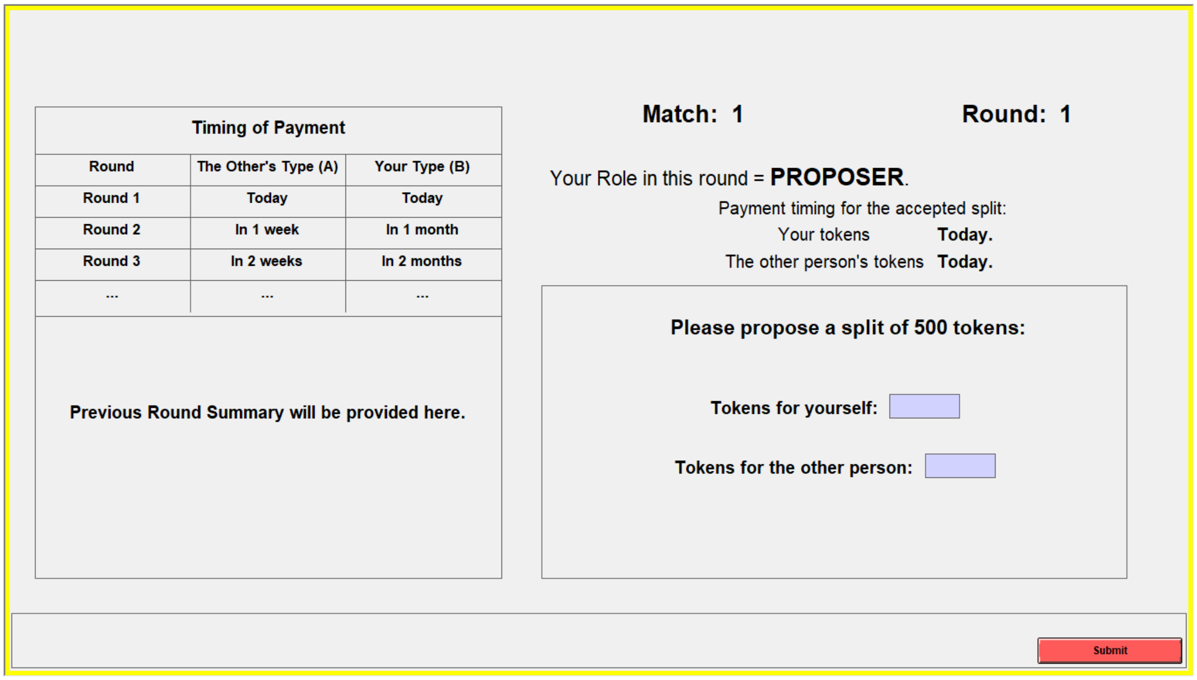

Figure 21: *Proposer*'s Screen

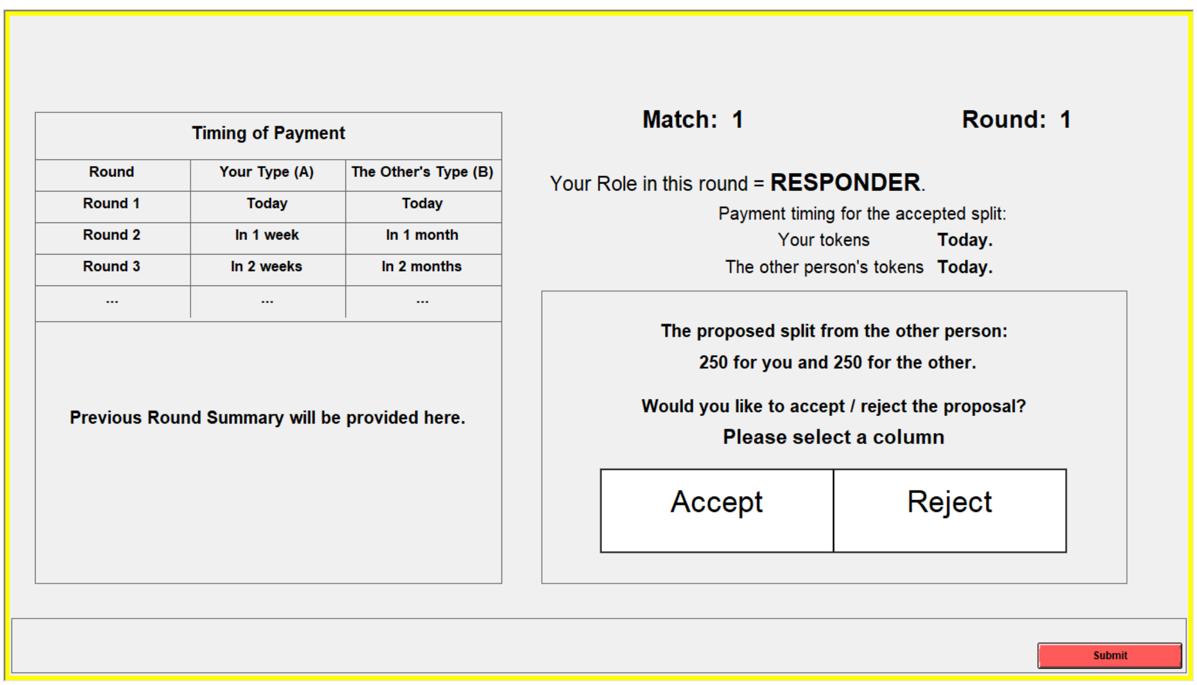

Figure 22: *Responder*'s Screen

B.2.2 Treatment *WM2D*

Welcome to the experiment. Please read these instructions carefully; the payment you will receive from this experiment depends on the decisions you make. The amount you earn will be paid through **VENMO**.

Your Payment Type and Match

At the beginning of the experiment, one-half of the participants will be randomly assigned to be Payment **Type A** and the other half to be Payment **Type B**. Your payment type will remain fixed throughout the experiment. Your payment type will affect when you will be paid, which will be explained below.

The experiment consists of 10 **matches**. At the beginning of each match, one Type A participant and one Type B participant are randomly paired. The pair is fixed **within the match**. After each match, participants will be randomly repaired, and new pairs will be formed. You will not learn the identity of the participant you are paired with, nor will that participant learn your identity—even after the end of the experiment.

Your Decisions in Each Match

Round 1: At the beginning of Round 1, one participant will be randomly assigned to the role of a **proposer** and the other participant to the role of a **responder**. Each participant in a match has 50-50 chance to be the proposer and to be the responder regardless of his/her payment type.

The proposer is then asked to propose how to split 500 tokens (= $50) between the two participants as:

"_________ tokens for yourself and _________ tokens for the other person."

After observing the split proposed by the proposer, the responder decides whether to accept or reject the proposed split.

Outcome, Termination, and Transition to Next Round: The outcome of Round 1 depends on whether the split proposed by the proposer is accepted or rejected.

1. If the responder **accepts** the proposed split, both participants will receive the amounts of tokens as proposed, and the match will be terminated.

2. If the responder **rejects** the proposed split, then the match will proceed to the next round with 75% (3/4) chance or be terminated with 25% (1/4) chance. This is as if we were to roll a 100-sided die and continue if the selected number is less than or equal to 75 and end if the number chosen is larger than 75.

   (a) If a match is **terminated** after a rejection of a proposed split, both participants will receive 0 tokens for the match.

   (b) If the match **proceeds** to the next round, then the proposer-responder roles are alternated. That is, the participant who is the proposer in the current round will become the responder in the next round, and vice versa. The number of tokens the participants receive will be determined by the outcome of the subsequent rounds.

Round $K > 1$: In Round $K > 1$, the participant who was the proposer in Round $(K - 1)$ becomes the responder, and the participant who was the responder in Round $(K-1)$ becomes the proposer. The proposer is then asked to propose how to split 500 tokens (= $50) between the two participants. After observing the split proposed by the proposer, the responder decides whether to accept or reject the proposed split.

The rest of the procedures determining the outcome, termination of the round, and transition to next round, is the same as those in Round 1.

Information Feedback

- At the end of each **round**, you will be informed about the proposal made by the proposer and the accept/reject decision made by the responder.
- At the end of each **match**, you will be informed **when and how much** you are going to be paid.

Your Monetary Payments

At the end of the experiment, one match out of 10 will be randomly selected for your payment. Every match has an equal chance to be selected for your payment so that it is in your best interest to take each match seriously. Participants will receive the amounts of tokens according to the outcome from the selected match with the exchange rate of 1 token = $0.1.

When you are going to be paid depends on (1) your payment type and (2) the round in which the proposed split is accepted.

If you are **Type A**, you may be paid in one week from today or in a few **weeks**. If a proposed split is accepted in Round 1, you will be paid in one week. If a proposed split is accepted in Round 2, you will be paid in two weeks. If a proposed split is accepted in Round $K > 1$, you will be paid in K weeks.

If you are **Type B**, you may be paid in one week from today, or in one week and a few **months**. If a proposed split is accepted in Round 1, you will be paid in one week. If a proposed split is accepted in Round 2, you will be paid in one week and one month. If a proposed split is accepted in Round $K > 1$, you will be paid in one week and $K - 1$ months.

The following table summarizes the schedule of payment for each type:

If a proposed split is accepted in	Type A will be paid	Type B will be paid
Round 1	In 1 week	In 1 week
Round 2	In 2 weeks	In 1 week and 1 month
Round 3	In 3 weeks	In 1 week and 2 months
Round 4	In 4 weeks	In 1 week and 3 months
Round 5	In 5 weeks	In 1 week and 4 months
.....	.....	.....
Round K	In K weeks	In one week and $K - 1$ months

Any amount you are supposed to receive will be paid electronically via VENMO.

In addition to your earnings from the selected match, you will receive a **show-up fee of $10** through VENMO, right after the experiment.

A Practice Match

To ensure your comprehension of the instructions, you will participate in a practice match. The practice match is part of the instructions and is not relevant to your cash payment; its objective is to get you familiar with the computer interface and the flow of the decisions in each round of a match. Once the practice match is over, the computer will tell you "The official matches begin now!"

Rundown of the Study

1. At the beginning of the experiment, your payment type will be randomly determined. Your payment type will remain fixed throughout the experiment.
2. At the beginning of each match, one Type A participant and one Type B participant are randomly paired.
3. At the beginning of Round 1, one participant will be randomly assigned to the role of a proposer and the other to the role of a responder.
4. The proposer then proposes how to split 500 tokens (= $50).
5. If the responder accepts the proposed split, both participants will receive the amounts of tokens as proposed, and the match will be terminated.
6. If the responder rejects the proposed split, then the match will proceed to the next round with 75% (3/4) chance or be terminated with 25% (1/4) chance. If a match is terminated after the rejection of a proposed split, both participants will receive 0 tokens for the match.
7. If the match proceeds to the next round, then the proposer-responder roles are alternated.
8. At the end of the experiment, one of 10 matches will be randomly selected for payment. For the selected match, the timing of your payment depends on (1) your payment type and (2) the round in which the proposed split was accepted. All your earnings will be paid to you through VENMO.
9. For Type A, you may be paid in one week, or a few weeks. For Type B, you may be paid in one week, or in one week and a few months.
10. In addition to your earnings from the selected match, you will receive a show-up fee of $10 right after the experiment.

Administration

Your decisions, as well as your monetary payment, will be kept confidential. Remember that you have to make your decisions entirely on your own; please do not discuss your decisions with any other participants. Upon finishing the experiment, you will be asked to sign your name to acknowledge your receipt of the payment. You are then free to leave. If you have any question, please raise your hand now. We will answer your question individually.

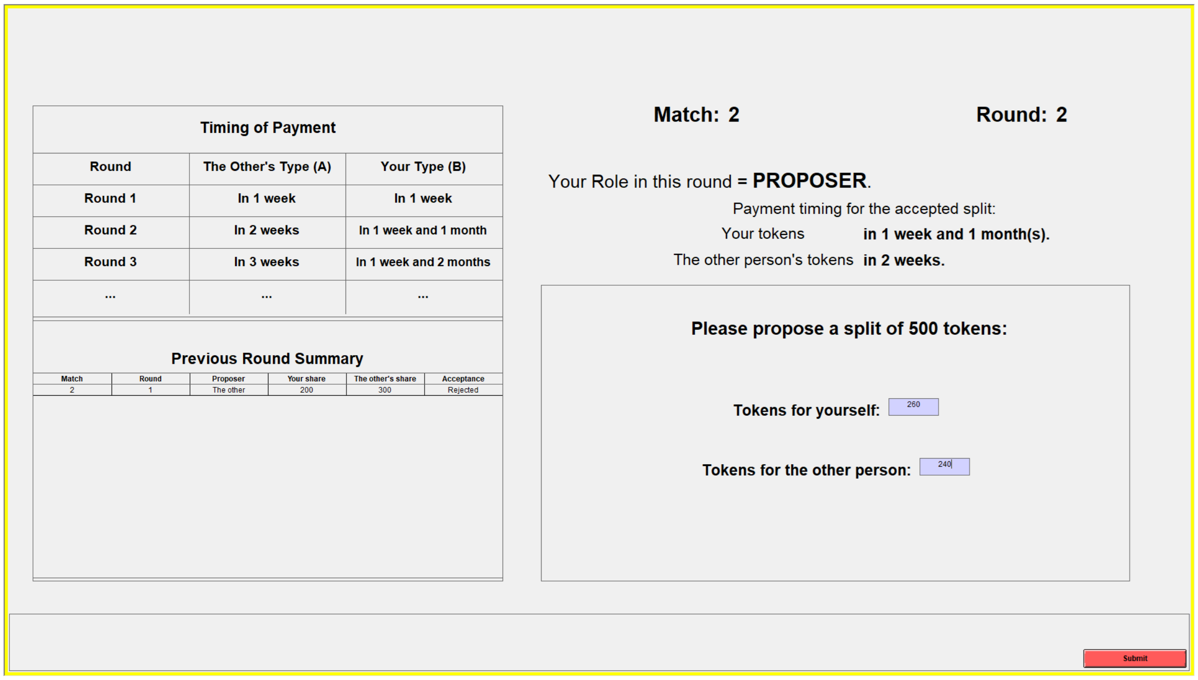

Figure 23: *Proposer*'s Screen

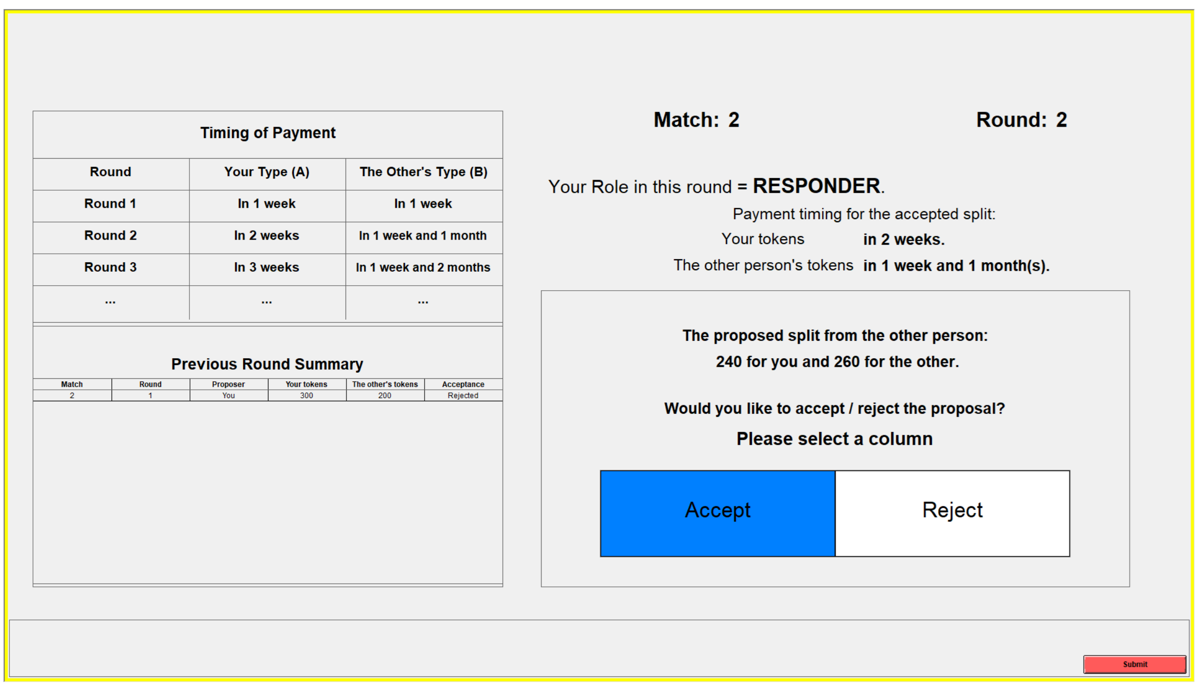

Figure 24: *Responder*'s Screen

B.2.3 Treatment *WW1D*

Welcome to the experiment. Please read these instructions carefully; the payment you will receive from this experiment depends on the decisions you make. The amount you earn will be paid through **VENMO**.

Your Payment Type and Match

At the beginning of the experiment, one-half of the participants will be randomly assigned to be Payment **Type A** and the other half to be Payment **Type B**. Your payment type will remain fixed throughout the experiment. Your payment type will affect when you will be paid, which will be explained below.

The experiment consists of 10 **matches**. At the beginning of each match, one Type A participant and one Type B participant are randomly paired. The pair is fixed **within the match**. After each match, participants will be randomly repaired, and new pairs will be formed. You will not learn the identity of the participant you are paired with, nor will that participant learn your identity—even after the end of the experiment.

Your Decisions in Each Match

Round 1: At the beginning of Round 1, one participant will be randomly assigned to the role of a **proposer** and the other participant to the role of a **responder**. Each participant in a match has 50-50 chance to be the proposer and to be the responder regardless of his/her payment type.

The proposer is then asked to propose how to split 500 tokens (= $50) between the two participants as:

"_________ tokens for yourself and _________ tokens for the other person."

After observing the split proposed by the proposer, the responder decides whether to accept or reject the proposed split.

Outcome, Termination, and Transition to Next Round: The outcome of Round 1 depends on whether the split proposed by the proposer is accepted or rejected.

1. If the responder **accepts** the proposed split, both participants will receive the amounts of tokens as proposed, and the match will be terminated.

2. If the responder **rejects** the proposed split, then the match will proceed to the next round with 75% (3/4) chance or be terminated with 25% (1/4) chance. This is as if we were to roll a 100-sided die and continue if the selected number is less than or equal to 75 and end if the number chosen is larger than 75.

    (a) If a match is **terminated** after a rejection of a proposed split, both participants will receive 0 tokens for the match.

    (b) If the match **proceeds** to the next round, then the proposer-responder roles are alternated. That is, the participant who is the proposer in the current round will become the responder in the next round, and vice versa. The number of tokens the participants receive will be determined by the outcome of the subsequent rounds.

Round $K > 1$: In Round $K > 1$, the participant who was the proposer in Round $(K-1)$ becomes the responder, and the participant who was the responder in Round $(K-1)$ becomes the proposer. The proposer is then asked to propose how to split 500 tokens (= \$50) between the two participants. After observing the split proposed by the proposer, the responder decides whether to accept or reject the proposed split.

The rest of the procedures determining the outcome, termination of the round, and transition to next round, is the same as those in Round 1.

Information Feedback

- At the end of each **round**, you will be informed about the proposal made by the proposer and the accept/reject decision made by the responder.
- At the end of each **match**, you will be informed **when and how much** you are going to be paid.

Your Monetary Payments

At the end of the experiment, one match out of 10 will be randomly selected for your payment. Every match has an equal chance to be selected for your payment so that it is in your best interest to take each match seriously. Participants will receive the amounts of tokens according to the outcome from the selected match with the exchange rate of 1 token = \$0.1.

When you are going to be paid depends on (1) your payment type and (2) the round in which the proposed split is accepted.

If you are **Type A**, you may be paid **today** or in a few **weeks**. If a proposed split is accepted in Round 1, you will be paid today right after the experiment. If a proposed split is accepted in Round 2, you will be paid in one week. If a proposed split is accepted in Round $K > 1$, you will be paid in $K-1$ weeks.

If you are **Type B**, you may be paid in **one week from today**, or in a few **weeks**. If a proposed split is accepted in Round 1, you will be paid in one week. If a proposed split is accepted in Round 2, you will be paid in two weeks. If a proposed split is accepted in Round $K > 1$, you will be paid in K weeks.

The following table summarizes the schedule of payment for each type:

If a proposed split is accepted in	Type A will be paid	Type B will be paid
Round 1	Today	In 1 week
Round 2	In 1 week	In 2 weeks
Round 3	In 2 weeks	In 3 weeks
Round 4	In 3 weeks	In 4 weeks
Round 5	In 4 weeks	In 5 weeks
.....	.....	.....
Round K	In $K-1$ weeks	In K weeks

Any amount you are supposed to receive will be paid electronically via VENMO.

In addition to your earnings from the selected match, you will receive a **show-up fee of \$10** through VENMO, right after the experiment.

A Practice Match

To ensure your comprehension of the instructions, you will participate in a practice match. The practice match is part of the instructions and is not relevant to your cash payment; its objective is to get you familiar with the computer interface and the flow of the decisions in each round of a match. Once the practice match is over, the computer will tell you "The official matches begin now!"

Rundown of the Study

1. At the beginning of the experiment, your payment type will be randomly determined. Your payment type will remain fixed throughout the experiment.
2. At the beginning of each match, one Type A participant and one Type B participant are randomly paired.
3. At the beginning of Round 1, one participant will be randomly assigned to the role of a proposer and the other to the role of a responder.
4. The proposer then proposes how to split 500 tokens (= $50).
5. If the responder accepts the proposed split, both participants will receive the amounts of tokens as proposed, and the match will be terminated.
6. If the responder rejects the proposed split, then the match will proceed to the next round with 75% (3/4) chance or be terminated with 25% (1/4) chance. If a match is terminated after the rejection of a proposed split, both participants will receive 0 tokens for the match.
7. If the match proceeds to the next round, then the proposer-responder roles are alternated.
8. At the end of the experiment, one of 10 matches will be randomly selected for payment. For the selected match, the timing of your payment depends on (1) your payment type and (2) the round in which the proposed split was accepted. All your earnings will be paid to you through VENMO.
9. For Type A, you may be paid today or in a few weeks. For Type B, you may be paid in one week, or in a few weeks.
10. In addition to your earnings from the selected match, you will receive a show-up fee of $10 right after the experiment.

Administration

Your decisions, as well as your monetary payment, will be kept confidential. Remember that you have to make your decisions entirely on your own; please do not discuss your decisions with any other participants. Upon finishing the experiment, you will be asked to sign your name to acknowledge your receipt of the payment. You are then free to leave. If you have any question, please raise your hand now. We will answer your question individually.

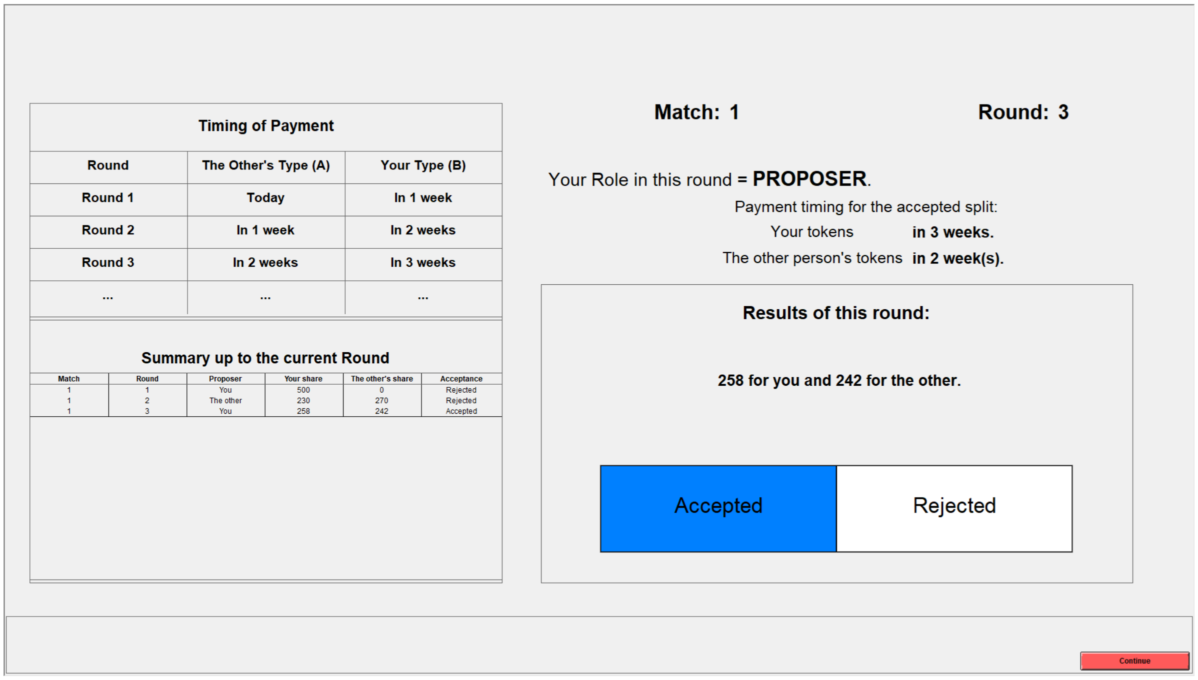

Figure 25: *Proposer*'s Screen

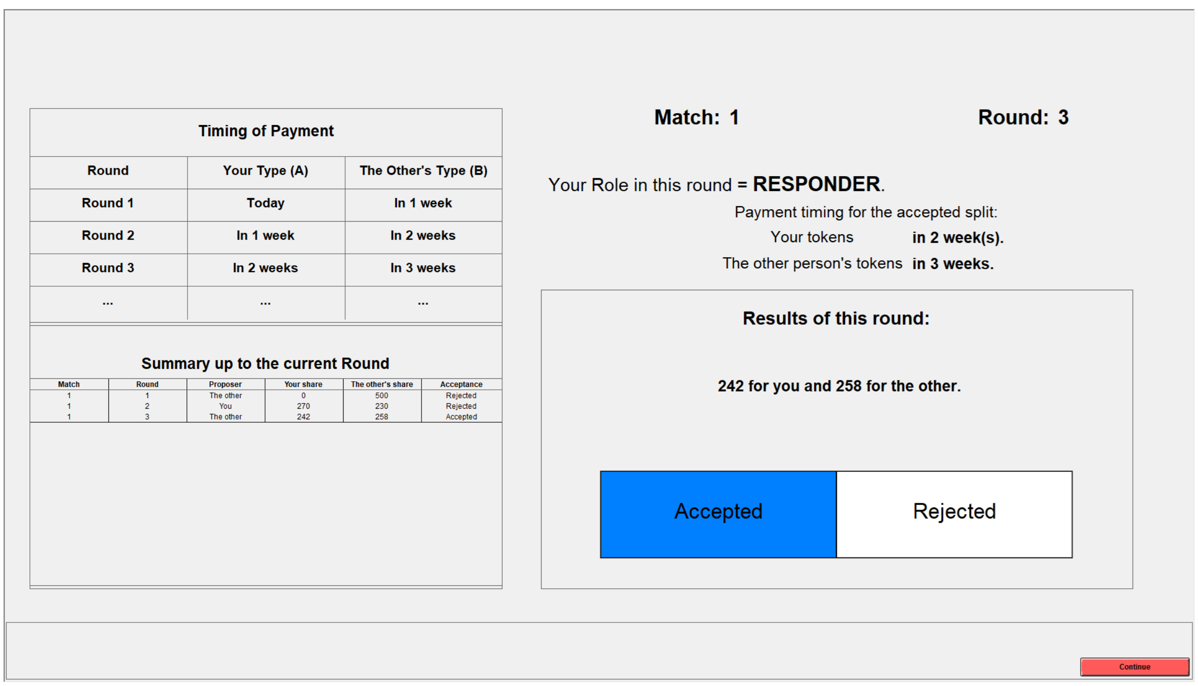

Figure 26: *Responder*'s Screen

B.3 Elicited Time Preferences and Behavior

We also elicited conventional measures of time preferences from our participants. This served two purposes: First, we can thereby test whether the random assignment to treatment and also bargainer type was indeed successful with regards to the underlying time preferences (for the subsample with elicitation), and second, we can also relate those conventional measures to behavior, as a complement to our main analysis.

Elicitation Procedure. We administered our elicitation task in only 4 out of the 6 sessions in each treatment (228 out of 348 participants), where it followed the bargaining games. Participants were not informed about this elicitation task beforehand, and they received all payoff-relevant information from their choices only at the very end of the experiment. The elicitation task asked participants to make 8 blocks of binary decisions between a sooner payment (option A) and a later payment (option B). In each block, one of the two was a fixed amount (either \$4 or \$10), and the other amount increased from \$0.01 in minimal steps of \$0.01 to \$10.00, resulting in effectively 1,000 binary decisions (rows) per block. Participants were asked for their switching point in terms of the varying option's amount, which they had to enter. The computer would automatically select the fixed option in all rows with a smaller varying amount and the varying option in all rows with a larger such amount. One row would be selected at random and the decision implemented, for one randomly drawn block. In essence, this is a version of the BDM (Becker, DeGroot, and Marschak, 1964) method, hence incentive compatible, but explained via a price list. The full instructions and a screenshot are available at the end of this section.

Table 7: Description of the Elicitation Task

<table>
<tr><td>Switching</td><td colspan="4">Sooner ⇒ Later</td></tr>
<tr><td>Block</td><td>(1)</td><td>(2)</td><td>(3)</td><td>(4)</td></tr>
<tr><td>Sooner</td><td>$4
Today</td><td>$4
Today</td><td>$4
1 month</td><td>$4
1 month</td></tr>
<tr><td>Later</td><td>$$X$
1 week</td><td>$$X$
1 month</td><td>$$X$
1 month and 1 week</td><td>$$X$
2 months</td></tr>
<tr><td>Switching</td><td colspan="4">Sooner ⇐ Later</td></tr>
<tr><td>Block</td><td>(5)</td><td>(6)</td><td>(7)</td><td>(8)</td></tr>
<tr><td>Sooner</td><td>$$X$
Today</td><td>$$X$
Today</td><td>$$X$
1 month</td><td>$$X$
1 month</td></tr>
<tr><td>Later</td><td>$10
1 week</td><td>$10
1 month</td><td>$10
1 month and 1 week</td><td>$10
2 months</td></tr>
</table>

*Note: X denotes the amounts that vary from 0.01 to 10.

Table 7 provides an overview of the details of the task. The block numbers correspond to their order in the task. There were four different sooner and later payment combinations: (1) sooner payment today and later payment in 1 week, (2) sooner payment today and later payment in 1 month, (3) sooner payment in 1 month and later payment in 1 month plus 1 week, and (4) sooner payment in 1 month and later payment in 2 months. For the first 4 blocks, the sooner payment was fixed at \$4.00 while the later payment ranged from \$0.01 to \$10.00. For the last 4 blocks, the later payment was fixed at \$10.00, and the sooner payment ranged from \$0.01 to \$10.00.

Distributions of Switching Points. We first compare the distributions of switching points X_k, where $k \in \{1, 2, \ldots, 8\}$ refers to the block number, by treatment and bargainer type, to check whether our randomization in terms of underlying time preferences was successful. Figure 27 provides the corresponding box plots. We employ Kolmogorov-Smirnov tests for equality of the switching point distributions on all 8 blocks. Since we test bargaining predictions both concerning comparisons between the two bargainer types within any treatment and between treatments for a given bargainer type, we also carry out distributional tests on the time preference task responses. Comparing, first, the switching points between the two bargainer types within any treatment—e.g., weekly vs. montly in Treatment *WM*—we find no significant differences (8 binary comparisons per treatment times 3 treatments, hence 24 binary comparisons, all p-values greater than 0.239). Second, and given this finding, we compare responses between various pairs of treatments—e.g., *WM* vs. *WM2D*—with a similar result (8 binary comparisons per treatment pairing times 3 treatment pairings, hence 24 binary comparisons, all p-values greater than 0.226).[48] Overall, we therefore conclude that our randomization into treatments and types in terms of underlying time preferences was successful indeed, certainly for the subsample considered.

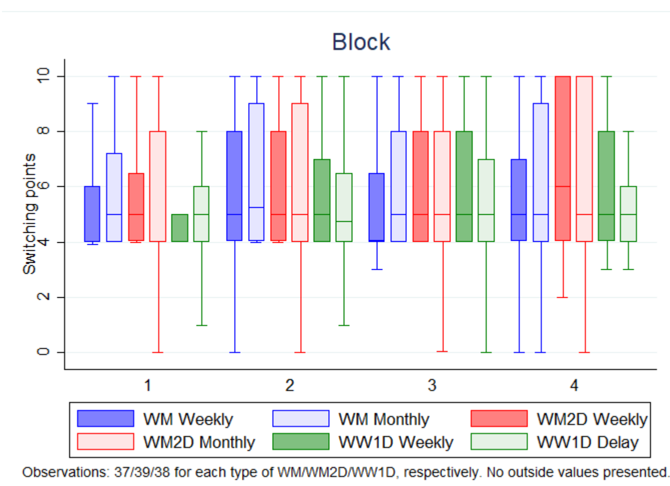

(a) Blocks 1-4

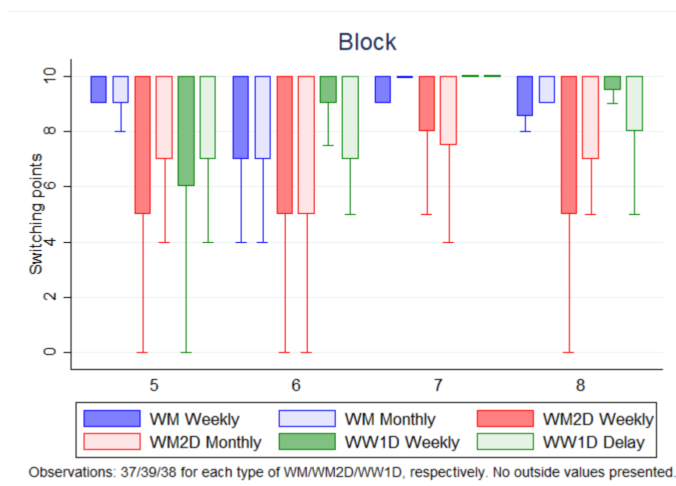

(b) Blocks 5-8

Figure 27: Distribution of Switching Points by Type/Treatment/Block

[48] We run the same test for weekly types only, where there are three treatment comparisons (there are weekly types in all treatments) and for monthly types only, where there is one treatment comparion (*WM* vs. *WM2D*). This results in $(3+1) \cdot 8 = 32$ binary comparisons, and all except three of them have p-values greater than 0.375. The smallest three equal 0.117, 0.123 and 0.167, so may be considered borderline. However, all of them concern comparisons of weekly types for trade-offs with a month's delay, namely X_4 and X_8, which are not the relevant ones for their bargaining.

Relation to Bargaining Behavior. We next relate our elicitation to bargaining behavior. The elicitation task is designed to infer parameters of (β, δ)-discounting, under the assumption that the participants are approximately risk neutral together with the standard narrow bracketing assumption (recall here the small stakes of at most \$10). We first estimate these for every participant, using the switching points for indifference equations—e.g., $4 = \beta\delta X_1$ and $4 = \delta X_3$, or $X_5 = \beta\delta 10$ and $X_7 = \delta 10$; details below—and then relate proposer as well as respondent behavior to the parameter estimates using regressions.

To estimate the two parameters we use for each participant the responses to all blocks; i.e., for weekly parameters we consider X_1, X_3, X_5, and X_7, and for monthly parameters we consider the other four. We then exclude participants whose responses are inconsistent or do not allow us to infer indifference.[49] For the remaining participants, we compute $(\beta_w, \beta_m, \delta_w, \delta_m)$ once from the relevant sooner-to-later switching points among the first four blocks and again from the relevant later-to-sooner switching point among the last four blocks, and we then take the average of the two for each parameter to reduce measurement error. For instance, we compute δ_w as the average of $\delta_{w(1)} = 4/X_3$ and $\delta_{w(2)} = X_7/10$, and then β_w as the average of $\beta_{w(1)} = 4/\delta_{w(1)}X_1 = X_3/X_1$ and $\beta_{w(2)} = X_5/\delta_{w(2)}10 = X_5/X_7$; similarly, for monthly parameters, where we denote estimates by (β_m, δ_m). Given the computed four parameters, β and δ take the average of the relevant parameters, i.e., $\beta = (\beta_w + \beta_m)/2$ and $\delta = (\delta_w + \delta_m)/2$, again to reduce measurement error. The results are summarized in Table 8 in terms of averages with standard deviations, and in Figure 28 in terms of box-plots, by types and treatments.

Table 8: Average Elicited Time Preferences by Type

	Treatments					
	WM		WM2D		WW1D	
	Weekly	Monthly	Weekly	Monthly	Delay	Weekly
β	1.03 (0.20)	1.03 (0.11)	1.00 (0.08)	1.03 (0.12)	0.99 (0.13)	1.03 (0.14)
δ	0.86 (0.15)	0.85 (0.15)	0.86 (0.13)	0.90 (0.13)	0.89 (0.11)	0.88 (0.12)
Obs.	21	25	21	16	24	25
# excluded	16	12	18	23	14	13

*Note: Standard deviations in parentheses.

Table 8 shows that around 40% of participants per type and treatment had to be excluded. The average β is very similar in all six cases, ranging from 0.99 to 1.03. Moreover, the standard deviations are of similar sizes, except for weekly types in Treatment *WM*. Also δ is very similar in all six cases, ranging from 0.85 to 0.90, and all standard deviations are of similar size.

Figure 28 presents the underlying distributional information as box-plots, also including outside values. The median values of β are all equal to one, and most of the mass lies around one in all cases, so the

[49] Inconsistency refers to assumed impatience and transitivity. It means here (i) $X_k < 4$, or (ii) $X_k = 10$ and $X_{k+4} > 4$, for at least one of the two relevant $k \in \{1, 2, 3, 4\}$; moreover, while $X_k = 10$ together with $X_{k+4} < 4$ is not inconsistent, it does not allow to establish indifference because a highly impatient person may *strictly* prefer the fixed sooner amount of \$4 in block k over the maximal possible switching point of \$10.

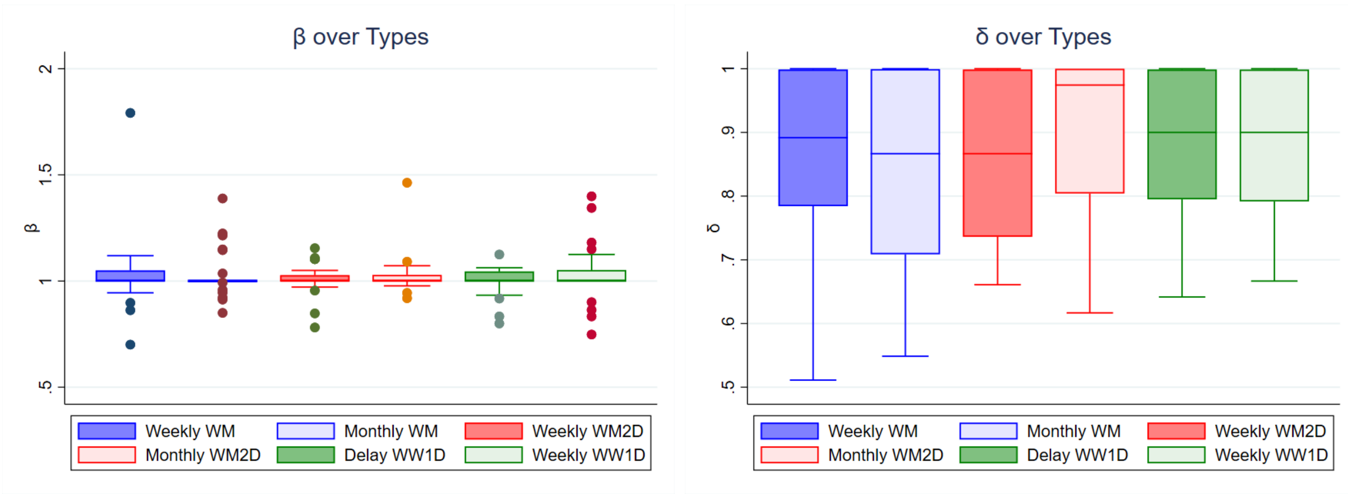

(a) β (b) δ

Figure 28: Distribution of β and δ by Type/Treatment

distributions are rather similar. The median values of δ are around 0.90 in all cases except for the monthly types of Treatment *WM2D* for whom the median equals 0.97, and the distributions are quite similar too.

We now use our estimates of the two discounting parameters as regressors in two basic regression specifications regarding bargaining behavior, one regarding proposer behavior (Round-1 proposals/demands) and another regarding respondent behavior (Round-1 acceptance vs. rejection of equal split proposals). Table 9 presents the results of OLS regressions of Round-1 average proposals of all types in all treatments on the proposer's discounting parameters (and a constant). Column 1 presents the result from considering all matches, and columns 2 and 3 present the results from considering the first and the last five matches, respectively.[50]

Overall, β and δ appear positively correlated with average Round-1 demands, but only one out of the corresponding six estimates is significantly different from zero, statistically. In other words, conventional time preference measures partially but only very weakly explain overall proposer behavior.

Additionally, we relate the discounting estimates also to respondent behavior, for which we take average Round-1 acceptance of equal-split proposals. Table 10 presents the results of analogous OLS regressions. Overall, we find similar results regarding statistical significance. While we don't find any significant relationships when considering all matches, δ is negatively and significantly correlated with average Round-1 acceptance of equal splits in the first five matches, in line with more patient respondents going after a better deal in the next round. With one exception, all other estimates are rather close to zero (here also the signs are not as consistently in line with patience being an advantage). The exception concerns the last five matches, where it appears that less present biased respondents are more likely to accept equal splits. The corresponding estimate comes with a large standard error, however, and is not statistically significant.

Potential reasons for the weak relationships observed include behaviorally relevant confounds (social preferences and risk attitudes, belief formation about the opponent) or also a relatively low signal-to-noise ratio

[50]The number of observations overall is 131, because one of the final 132 participants with estimated discount factors (see Table 8) happened to never be selected as initial proposer.

Table 9: β, δ, and Round-1 Average Proposer behavior (OLS)

	All Matches	First Five	Last Five
	(1)	(2)	(3)
β	14.78	2.00	20.46
	(16.85)	(11.27)	(20.36)
δ	26.33	36.11**	7.11
	(14.98)	(15.09)	(19.54)
Constant	224.5***	230.5***	232.8***
	(23.66)	(18.81)	(31.57)
Obs.	131	128	126
R-squared	0.019	0.024	0.017

Notes: Dependent variable: Proposer's Round-1 average share. Clustered standard errors at the session level in parentheses.
***Significant at the 1%-level.
**Significant at the 5%-level.
*Significant at the 10%-level.

Table 10: β, δ, and Round-1 Average Acceptance of the Equal Splits (OLS)

	All Matches	First Five	Last Five
	(1)	(2)	(3)
β	0.14	-0.06	0.50
	(0.16)	(0.12)	(0.29)
δ	-0.12	-0.46**	0.16
	(0.16)	(0.17)	(0.15)
Constant	0.90***	1.39**	0.32
	(0.26)	(0.21)	(0.38)
Obs.	116	94	86
R-squared	0.02	0.05	0.10

Notes: Dependent variable: Responder's Round-1 average acceptance of the equal splits. Clustered standard errors at the session level in parentheses.
***Significant at the 1%-level.
**Significant at the 5%-level.
*Significant at the 10%-level.

of such measures. As such, the findings lend further support to our study's design and analysis.

Instructions for Elicitation Task and Selected z-Tree Screenshot.

<u>**Instructions**</u>

In this task, we will ask you to make decisions for 8 blocks of questions. In each block, there are 1,000 questions. For each question, you can choose one of two options - Option A, which pays you sooner, and Option B, which pays you later.

After you answer all questions, one question will be randomly selected and the option you chose on that question will determine your earnings. Each question is equally likely to be chosen for payment. Obviously, you have no reason to misreport your preferred choice for any question, because if that question gets chosen for payment, then you would end up with the option you like less.

For example, the questions in one block are as follows. Note that each row corresponds to a question so that you have to choose one option for each row.

Questions	Option A Today	Option B in 1 month
1	\$4.00	\$0.01
2	\$4.00	\$0.02
3	\$4.00	\$0.03
⋮	⋮	⋮
999	\$4.00	\$9.99
1,000	\$4.00	\$10.00

It is natural to expect that you will choose Option A for at least the first few questions, but at some point switch to choosing Option B. In order to save time, you can report at which dollar value of Option B you'd switch. The computer program can then 'fill out' your answers to all 1,000 questions based on your reported switching point (choosing Option A for all questions before your switching point, and Option B for all questions at and after your switching point).

Timing of payment: The 8 blocks will differ in the following two ways: (1) the timings of sooner and later payments:

- Between payment **today** and payment in **1 week**.
- Between payment **today** and payment in **1 month**.
- Between payment in **1 month** and payment in **1 month and 1 week**.
- Between payment in **1 month** and payment in **2 months**.

and (2) whether you are asked to switch from <u>Option A to Option B</u>, or from <u>Option B to Option A</u>.

Payment: At the end of the experiment, one question in one of the blocks will be randomly selected for payment. The selected question and the block as well as your choice for the question will be displayed on your screen. Then the payment will be made on the designated date through VENMO. For example, 1. If your choice in the randomly selected question was to receive a payment today, then you will be paid through

VENMO right after the experiment. 2. If your choice in the randomly selected question was to receive a payment in the future, you will be paid on the designated date through VENMO.

Rundown of the Study

1. There are 8 blocks of questions, each of which you will be asked to report your switching point.
2. Only one question in one of the eight blocks will be randomly selected for payment.
3. You will be paid on the designated date through VENMO.

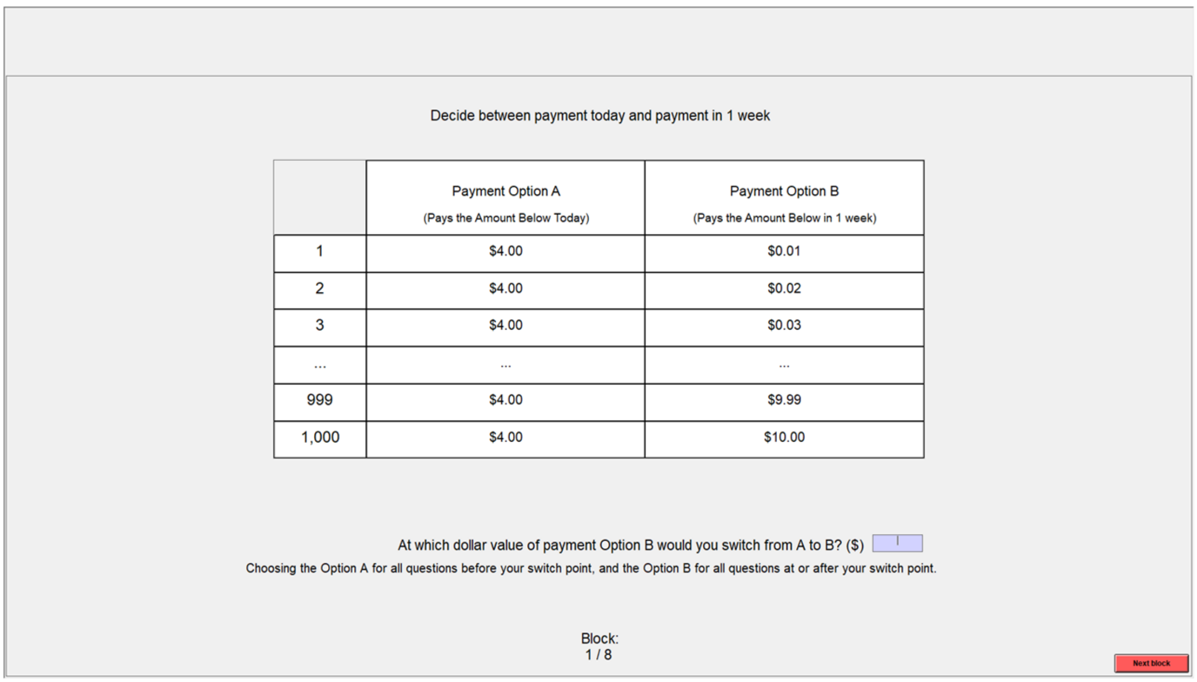

Decide between payment today and payment in 1 week

	Payment Option A (Pays the Amount Below Today)	Payment Option B (Pays the Amount Below in 1 week)
1	$4.00	$0.01
2	$4.00	$0.02
3	$4.00	$0.03
...	...	...
999	$4.00	$9.99
1,000	$4.00	$10.00

At which dollar value of payment Option B would you switch from A to B? ($)

Choosing the Option A for all questions before your switch point, and the Option B for all questions at or after your switch point.

Block:
1 / 8

Next block

Figure 29: Elicitation Task Screen-shot Block 1