



BERLIN SCHOOL OF ECONOMICS
DISCUSSION PAPERS

Discussion Paper #15

May 2023

**Patience Is Power: Bargaining and
Payoff Delay**

Jeongbin Kim

Wooyoung Lim

Sebastian Schweighofer-Kodritsch

Patience Is Power: Bargaining and Payoff Delay*

Jeongbin Kim[†]

Wooyoung Lim[‡]

Sebastian Schweighofer-Kodritsch[§]

December 7, 2022

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 and Egor Starkov 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.

[§]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, 2021](#)), 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](#)), there exists a sizeable experimental bargaining literature, of course. However, none of these experiments involve time delay of payoffs (all of which occur at the end of the session); rather, they could only be interpreted as “simulating” discounting via shrinking cakes or breakdown risk (see our discussion in [Section 5](#)). This, however, still restricts potential conclusions to exponential discounting, at odds with the bulk of evidence on actual human time preferences (see [Frederick, Loewenstein, and O’Donoghue, 2002](#); [Ericson and Laibson, 2019](#), for a classic and a recent review, respectively). Moreover, while the theoretical prediction of [Rubinstein \(1982\)](#) relies on a fair amount of sophisticated strategic reasoning

¹It 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).

²For 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](#)).

³For 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).

⁴The 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, 2020](#)). [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.

(e.g., under mere Nash equilibrium any immediate division is an equilibrium outcome, independent of time preferences; relatedly, see also [Vannetelbosch, 1999](#); [Friedenberg, 2019](#)), bargaining is also inherently a distributional problem, potentially giving rise to fairness concerns (the related behavioral findings from ultimatum bargaining of [Güth, Schmittberger, and Schwarze, 1982](#), actually sparked the study of social preferences in economics). Even people recognizing that patience is a source of bargaining power may find its strategic exploitation unacceptable as a cause of inequality and be willing to sacrifice some payoff to prevent this. Indeed, 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](#)).

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 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. This constitutes the first evidence we are aware of that people strategically exploit others’ present bias.

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*.

We provide a theoretical foundation for this design, by extending the [Rubinstein \(1982\)](#) model to arbitrarily non-stationary discounting. Under the assumption that the parties’ preferences are

⁵Thus 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. To ensure participants’ impatience (in view of essentially zero interest at the time) and hence the credibility of our experiment, but also to limit potential effects due to incomplete information, we additionally impose a commonly known 25% chance of random termination after any disagreement, throughout all matches. Discounting should therefore be taken to include this risk (arguably, time preferences inherently include attitudes to uncertainty; see [Chakraborty, Halevy, and Saito, 2020](#)).

dynamically consistent across *bargaining rounds*—as is true by design in our experiment—the strong results for exponential discounting (EXD) carry over: We prove general existence and uniqueness of (perfect) equilibrium, which always implies immediate agreement, and we characterize it in terms of the two parties’ (potentially very complex) time preferences. Thus, bargaining power is generally well-defined, and we can use the characterization to obtain comparative statics predictions for our experiment: Assuming identical *underlying time preferences*, our experimental manipulation turns these into potentially different *effective time preferences*. How this affects bargaining power generally depends on what underlying time preferences are assumed, which will allow us to discriminate between leading classes of time preferences based on what our participants’ bargaining behavior reveals.

Our experimental method, which we call *effective discounting procedure*, thus permits clean comparative statics tests in 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 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.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. Weekly bargainers are *generally* predicted to be at an advantage 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, we nonetheless strongly confirm the *comparative statics prediction* of a patience advantage.⁸ In fact, we observe different propensities of proposing as well as accepting equal splits by weekly and monthly bargainers, which further support the prediction. Hence,

⁶We 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 (2020b) 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.

⁷Negligible bargaining delay implies that only a single dated self of any individual gets to make all decisions. Without any room for dynamic inconsistency in behavior, *a fortiori*, there is no room for naïveté (in the sense of O’Donoghue and Rabin, 1999) or learning about naïveté, which are notoriously challenging to analyze theoretically. In this respect our design is similar to the standard time preference elicitation paradigm where participants choose between variously delayed monetary rewards at a single point in time. We therefore also face no selection issues, whereas with significant bargaining delay, attrition is likely to be systematically related to time preferences (see Sprenger, 2015); e.g., Kim (2020a) indeed finds that patience and present bias measured at the beginning of his experiment were predictive of how long participants would take part in his longitudinal study.

⁸Our main test for this advantage compares the distributions of initial proposals for first-order stochastic dominance; importantly, we obtain similar results when comparing *accepted* initial proposals.

time preferences matter, and—in the broad sense of our manipulation—we confirm that patience is a significant source of bargaining power.

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 asymmetry in bargaining power favoring weekly over monthly bargainers. While differences in propensities to propose or reject an equal split between the types, as well as differences in accepted proposals, are somewhat in line with the predicted advantage of weekly bargainers, 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—also equal-split behavior again further supports the advantage—and it is evidence that participants not only expect but also strategically exploit a present bias in others.

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, we also find suggestive evidence that, beyond present bias, bargainers perceive and respond to diminishing impatience, as in general hyperbolic discounting (Loewenstein and Prelec, 1992).

Since our design does not itself induce incomplete information, we derive and test comparative statics predictions from a complete information theory.⁹ 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 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 was successful in keeping its effects both relatively mild and roughly constant.

Nonetheless, from this more general perspective, our experimental manipulation may be mainly one

⁹Explicitly modeling incomplete information about time preferences to capture their observed heterogeneity seems elusive. See, however, Fanning and Kloosterman (2020) who successfully apply this alternative approach to studying fairness concerns.

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 offers) 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.

The rest of this paper is organized as follows. We first present the general theoretical background for our experimental study in Section 2. This is followed by a description of our experimental design, including the behavioral predictions for the most important classes of time preferences, in Section 3. We then report and discuss the findings from our experiment in Section 4, and subsequently relate our study to the existing literature in Section 5. Section 6 offers concluding remarks. All proofs are relegated to this paper's Appendix. An Online Appendix consists of five parts and provides the following supplemental material: additional figures that complement those in the main body of the paper (part A); the results of alternative statistical tests (part B); experimental instructions and selected screenshots for one exemplary experimental treatment (Treatment *WM*, parts C and D); all details of our additional time preference elicitation and results on how measured time preferences relate to bargaining behavior (part E).

2 Theoretical Background

2.1 The Model

Consider two individuals $i \in \{1, 2\}$ deciding on how to share a fixed monetary amount via indefinite alternating-offers bargaining as in Rubinstein (1982). 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 $j = 3 - i$ (we will use this convention for i and j throughout), 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 proposal in round 1, and the game continues until a proposal is accepted. Denoting by r_n the responding player of round n , $r_n = 2$ for n odd, and $r_n = 1$ for n even.

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) $u_i : [0, 1] \rightarrow [0, 1]$ is continuous and strictly increasing from $u(0) = 0$ to $u(1) = 1$;¹⁰
3. (Intertemporal Utility) There exists $\alpha_i < 1$ such that for all $n \in \mathbb{N}$, and for all $q \in [0, 1)$ and $q' \in (q, 1]$,

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

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](#)) or “immediacy” (see the utility formulation of [Schweighofer-Kodritsch, 2018](#)) to the non-stationary setting studied here, and it is implied by standard assumptions; e.g., u_i concave and $\sup_n \delta_i(n) < 1$.¹¹

¹⁰The assumption that $u(1) = 1$ is a mere normalization and without loss of generality.

¹¹Let 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 $\lim_{n \rightarrow \infty} \delta(n) \rightarrow 1$.

2.2 Equilibrium

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.

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

$$x_n = 1 - u_{r_n}^{-1}(\delta_{r_n}(n) \cdot u_{r_n}(x_{n+1})). \quad (2.1)$$

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 $q \leq x_n$.*

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. In the special case where the model reduces to the benchmark of [Rubinstein \(1982\)](#), the infinite sequence in (2.1) reduces to two equations:

$$\begin{aligned} x_1 &= 1 - u_2^{-1}(\delta_2 \cdot u_2(x_2)), \\ x_2 &= 1 - u_1^{-1}(\delta_1 \cdot u_1(x_1)). \end{aligned}$$

We generate several behavioral predictions from this exponential-discounting benchmark for our concrete experimental treatments, and we employ the general characterization to also derive the behavioral predictions under various alternative forms of discounting (in particular, quasi-hyperbolic discounting capturing a present bias). We present all of these theoretical predictions in Section 3 after defining our specific treatments.

2.3 Remarks

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.¹² Given this, our abstract formulation of preferences in terms of rounds of agreement allows us to capture 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.¹³

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.¹⁴ 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.¹⁵ 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”); thus, we are able to study *general* time preferences, including dynamically inconsistent ones, without actually confronting any issues of dynamic inconsistency in the bargaining itself (importantly including naïveté, see O’Donoghue and Rabin, 1999).

This is the bargaining version we implement in our experiment. It further offers the practical advantage over the usual interpretation of alternating-offers bargaining, according to which bargaining and payoff delay coincide, that there are no selection issues related to time preferences, neither into the experiment nor during the running of the experiment. Moreover, it is worthwhile pointing out that if preferences really are dynamically consistent (or individuals share a common belief in such consistency), it is equivalent to a game with this usual interpretation: Proposition 1 and any behavioral predictions derived from it then directly extend to the setting where bargaining itself takes time. This applies in particular to the special case of our model with exponential discounting and constant payoff delay, which is that of Rubinstein (1982) and will serve as our benchmark.

¹²Formalizing 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.

¹³We 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.

¹⁴See, however, Halevy (2015) for evidence that some violations of exponential discounting may be due to *time variance* rather than dynamic inconsistency of discounting.

¹⁵As long as payoff delay remains significant, the model is not susceptible to the “smallest-units” critique of van Damme, Selten, and Winter (1990).

3 Experimental Design and Behavioral Predictions

3.1 Experimental Design

In line with the theory just developed, our experiment implements indefinitely alternating offers bargaining games with frequent offers and significant payoff delay. The monetary surplus to be divided is fixed and amounts to US\$50. Table 1 presents our experimental design, which consists of three treatments. Each of these treatments corresponds to a particular pairing of “bargainer types.” This type is the exogenously imposed payoff delay profile that an individual faces, according to the effective discounting procedure.

Table 1: Experimental Treatments

Bargainer 1	Bargainer 2		
	Monthly (M)	Monthly with Delay (MD)	Weekly with Delay (WD)
Weekly (W)	<i>WM</i>	N/A	<i>WW1D</i>
Weekly with Delay (WD)	N/A	<i>WM2D</i>	N/A

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

In Treatment *WM*, one bargainer faces one week of delay per round of disagreement, whereas the other faces one month of such delay. Treatment *WM2D* is similar, but both bargainers additionally face a front-end delay of one week so that immediate agreements are about payoffs to be received in one week’s time. In Treatment *WW1D*, both bargainers face identical delays per round of disagreement of one week, but one of them additionally faces a front-end delay of one week.

Every treatment therefore matches different payoff types. The treatment is public, whereby the payoff delay types of any two matched participants are common knowledge. 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.

Weeks and months are both natural and significant time units, so our treatments should be able to create meaningful differences in effective discounting. To credibly implement delayed payments, we relied on the popular mobile payment system *Venmo*.¹⁶ In the rest of the paper, we will call a bargainer type whose payment window is weekly/monthly/delayed a *weekly/monthly/delayed* bargainer.

¹⁶Venmo is a service provided by PayPal that allows account holders to transfer funds to others via a mobile phone app. It handled \$12 billion in transactions during the first quarter of 2018 (<https://en.wikipedia.org/wiki/venmo>). For more information, please visit <https://help.venmo.com/hc/en-us/articles/210413477>. When recruiting our participants, we clearly announced that those without a Venmo account were not eligible to participate in the experiment. 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.

For a concrete illustration of the different payoff delays, consider agreements that are reached in Round 3. In Treatment *WM*, this would mean that the weekly bargainer receives the associated payoff in two weeks from the day of the experiment, and the monthly bargainer in two months; in Treatment *WM2D*, these delays would be three weeks for the delayed weekly and a week plus two months for the delayed monthly bargainer; in Treatment *WW1D*, this would be two weeks for the weekly and three weeks for the delayed weekly bargainer. Appendices C and D provide the instructions and selected screenshots for exemplary Treatment *WM*.

We coupled this 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. Note that keeping bargainers away from possible indifference to delay is also required by Property 3 of our preference assumptions. 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 (assuming expected utility).¹⁷

3.2 Behavioral Predictions

We now employ Proposition 1 to derive the comparative statics predictions related to time preferences that our experiment is designed to test.¹⁸ We begin by establishing the important and influential benchmark predictions from exponential discounting, as in Rubinstein (1982), then derive the differential predictions under its leading alternative, quasi-hyperbolic discounting, and finally also discuss predictions under various other forms of discounting as they appear in the literature on time preferences. This will show that the leading treatment *WM* allows us to test for a general patience advantage, and the remaining two treatments *WM2D* and *WW1D* allow us to investigate the robustness of this hypothesized advantage as well as discriminate between different classes of time preferences based on a revealed preference argument. All formal proofs are in the Appendix.

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).

Recall that the only universal prediction from Proposition 1 is immediate agreement. We formulate the comparative statics predictions in terms of relative bargaining power, as reflected in this immediate

¹⁷With expected utility, a constant probability of breakdown simply proportionally reduces each $\delta_i(n)$ by this fraction.

¹⁸Since we leverage our bargainers' unobserved underlying/natural time preferences, point predictions are unavailable.

agreement. For predictions within a treatment that matches two bargainer types A and B (e.g., weekly and monthly bargainers in Treatment *WM*), call the game where type A is the initial proposer AB-game and the game where type B is the initial proposer BA-game; we then say that **the type A bargainer is stronger than the type B bargainer** if the type A bargainer’s equilibrium share in the AB-game is greater than the type B bargainer’s equilibrium share in the BA-game.¹⁹ If neither type is stronger than the other in the above sense, we say they are equally strong.

For between-treatment predictions, take another treatment that matches types A and C; we then say that **the type A bargainer is stronger against the type B bargainer than against the type C bargainer as initial proposer (resp., respondent)** if type A bargainer’s equilibrium share in the AB-game (resp., BA-game) is greater than type A bargainer’s equilibrium share in the AC-game (resp., CA-game). If this is true both as initial proposer and initial respondent we simply say that the type A bargainer is stronger against the type B bargainer than against the type C bargainer.²⁰

Exponential Discounting (EXD). Since any given bargainer type faces a constant payoff delay, the stationarity property of EXD implies that any such delay is discounted with the same discount factor, irrespective of any front-end delay. Let $\delta \in (0, 1)$ be the (common) discount factor for a weekly delay, and let $\phi\delta$ be the (common) discount factor for a monthly delay, where $0 < \phi < 1$ due to impatience.²¹ Using notation $\phi_i \in \{\phi, 1\}$ with $\phi_i = 1$ if and only if bargainer i is a weekly bargainer, any bargainer i ’s type is fully captured by ϕ_i , such that $U_i(q, n) = (\phi_i\delta)^{n-1} u(q)$ and $\delta_i(n) = \phi_i\delta$ is constant across rounds n . *WM* and *WM2D* both correspond to pairing $\{1, \phi\}$, and *WW1D* corresponds to pairing $\{1, 1\}$.

Prediction 1. *Symmetric EXD implies:*

- (A1) *In Treatment WM, the weekly bargainer is stronger than the monthly bargainer.*
- (A2) *In Treatment WM2D, the weekly bargainer is stronger than the monthly bargainer.*
- (A3) *In Treatment WW1D, the weekly bargainer and the delayed weekly bargainer are equally strong.*
- (B1) *Between Treatments WM and WW1D, the weekly bargainer is stronger against the monthly bargainer than against the delayed weekly bargainer.*
- (B2) *Between Treatments WM2D and WW1D, the delayed weekly bargainer is stronger against the delayed monthly bargainer than against the weekly bargainer.*

¹⁹This compares A and B as the initial proposer only; however, since shares add up to one, A is stronger than B as the initial proposer if and only if A is stronger than B as the initial respondent.

²⁰In comparisons between treatments, the observation in footnote 19 does not apply. The fact that shares add up to one only means that A is stronger against B than against C as the initial proposer if and only if B’s share in the AB-game is smaller than C’s share in the AC-game, but does not imply anything about A’s share in the BA-game versus A’s share in the CA-game.

²¹If we take a month to equal four weeks, then $\phi\delta = \delta^4$ pins down $\phi = \delta^3$.

These predictions are 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.

Quasi-Hyperbolic Discounting (QHD). Present bias, the excessive weight put on immediate rewards relative to delayed rewards, is the most important deviation from EXD. By adding a single parameter $\beta \in (0,1)$, the model of quasi-hyperbolic discounting parsimoniously captures this empirically well-established phenomenon. The bias may here play a role only in the first round because upon failure to agree immediately, all possible payoffs lie in the future. Moreover, it will do so only when the initial respondent faces no front-end delay because the proposer’s discounting of the first round’s delay is irrelevant anyways due to the proposer’s strategic advantage, and a front-end delay for the respondent pushes any immediate-agreement payoff into the future. Keeping the earlier EXD notation and adding $\beta_i \in \{\beta, 1\}$ with $\beta_i = 1$ if and only if bargainer i is delayed, any bargainer i ’s type is fully captured by (ϕ_i, β_i) , such that $U_i(q, n) = \beta_i (\phi_i \delta)^{n-1} u(q)$; now $\delta_i(1) = \beta_i \phi_i \delta$ and, for $n > 1$, $\delta_i(n) = \phi_i \delta \geq \delta_i(1)$. *WM* corresponds to pairing $\{(1, \beta), (\phi, \beta)\}$, *WM2D* corresponds to pairing $\{(1, 1), (\phi, 1)\}$, and *WW1D* corresponds to pairing $\{(1, \beta), (1, 1)\}$. In the predictions under QHD below we highlight those that differ from Prediction 1 under EXD. Table 2 summarizes all predictions.

Prediction 2. *Symmetric QHD implies the same as symmetric EXD except:*

- (A3) *In Treatment WW1D, **the delayed weekly bargainer is stronger** than the weekly bargainer.*
- (B2) *Between Treatments WM2D and WW1D, the delayed weekly bargainer is stronger against the delayed monthly bargainer than against the weekly bargainer as initial respondent, but there is **no general prediction concerning the weekly bargainer as initial proposer.***

The predictions under QHD are straightforward from those under EXD, upon noting 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

Table 2: Summary of Predictions 1 (EXD) and 2 (QHD)

Prediction	Treatments	Initial Role	EXD	QHD
(A1)	WM	Both	$W > M$	
(A2)	$WM2D$	Both	$WD > MD$	
(A3)	$WW1D$	Both	$WD = W$	$WD > W$
(B1)	WM and $WW1D$	Both	W in $WM > W$ in $WW1D$	
(B2)	$WM2D$ and $WW1D$	Prop.	WD in $WM2D$ $> WD$ in $WW1D$	No general pred.
		Resp.	WD in $WM2D > WD$ in $WW1D$	

Notes: “>” or “=” indicates, respectively, greater or equal bargaining power / equil. share; in within-treatment predictions (A1)–(A3), $A > B$ (resp., $A = B$) as initial proposer if and only if $A > B$ (resp., $A = B$) as initial responder, in contrast to between-treatment predictions (B1)–(B2).

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.

Our focus here lies on comparative statics predictions specific to time preferences. These are predicated on immediate agreement equilibrium. We investigate this universal prediction and discuss the interpretation of our findings in the absence of immediate agreement in Section 4.3. There, we also discuss another commonly considered comparative statics prediction due to the alternating-offers protocol’s inherent asymmetry, which is that of a proposer advantage. This prediction is not specific to payoff delay as the cost of disagreement, but it does obtain here as well (under both EXD and QHD). A proposer advantage is a within-treatments prediction, and in the terminology introduced earlier, type

A has a proposer advantage if this very type’s equilibrium share is greater in the AB-game than in the BA-game. Note that this is true if and only if type B has a proposer advantage. This comparison concerns a given type in the two different initial roles, whereas the patience advantage compares the two different types in the same initial role. Neither advantage implies or rules out the other.

Other Forms of Discounting. 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](#)). At the same time, 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](#)). We now discuss the implications of these alternatives.

First, consider diminishing impatience, meaning $\delta_i(n)$ increases in n . 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 (2.1), 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 in what follows, imposes the structure of $d(t) = (1 + \alpha \cdot t)^{-\beta/\alpha}$ (with $\alpha, \beta > 0$), and this implies that a weekly bargainer always remains more patient than a monthly bargainer; notably, this extends also to when *both* are delayed.²² 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 altogether permissive with respect to the comparison between Treatments *WM2D*

²²The 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 $\delta_A(n) = [(1 + \alpha \cdot n)/(1 + \alpha \cdot (n - 1))]^{-\beta/\alpha}$; 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.

and *WWID*.

Finally, consider near-future bias. 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 *WWID*, the weekly bargainer rather than the delayed weekly bargainer would be stronger. Using that $\beta\phi < 1$, which follows from $\beta < 1/\delta$ and $\phi < \delta$, it is straightforward to show that the between-treatment predictions under such near-future bias coincide with those under EXD.

3.3 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. 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 \times 3 treatments). 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.²³ For the selected match, if agreement was reached, the agreed number of tokens for

²³ Azrieli, Chambers, and Healy (2018) offer a theoretical justification and discussion of this payment rule’s incentive compatibility in settings like ours.

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.²⁴

In addition, participants received a show-up payment of \$10. Any amount a participant was due to receive was paid electronically via *Venmo*, including immediate payments. Earnings were \$37.90 on average, and the average duration of a session was approximately 1.5 hours.²⁵

4 Experimental Results

This section presents our experimental results regarding Predictions 1 and 2, as summarized in Table 2. We first consider the predictions within treatments, (A1)–(A3), and then those between treatments, (B1)–(B2). Subsequently, we provide additional results and discussion.

In line with the literature, we conduct our main tests 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).²⁶ We compare these distributions for first-order stochastic dominance, using the test procedure proposed by Barrett and Donald (2003).²⁷ Since demands and offers add up to a constant amount (500 tokens), if type A’s demands first-order stochastically dominate type B’s demands, type A faces an unambiguously more favorable distribution of initial proposals than type B.

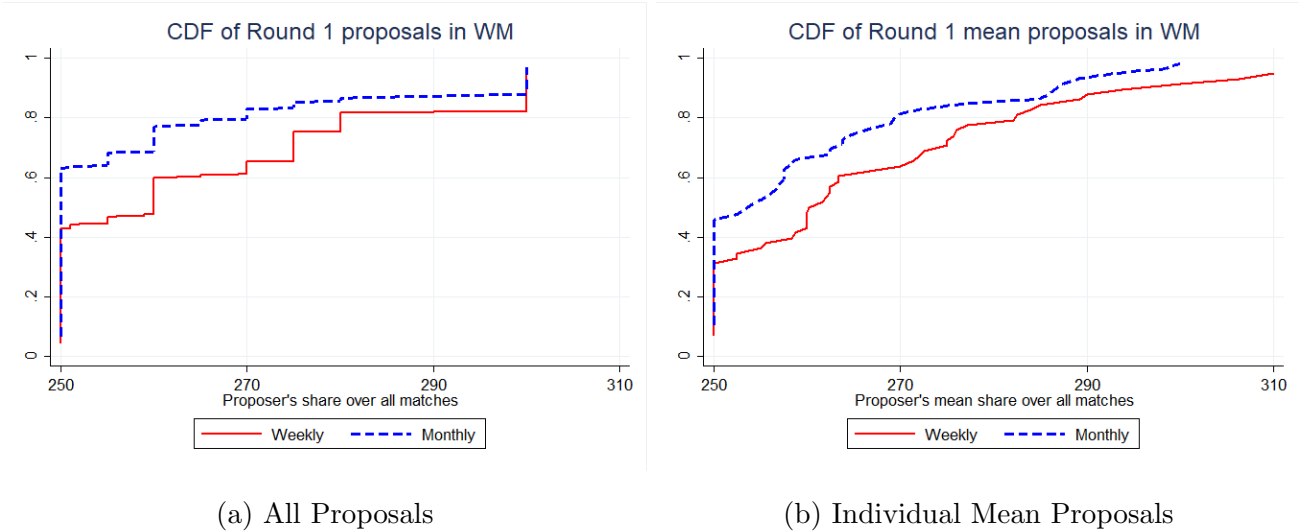


Figure 1: Round-1 Proposals in Treatment *WM*

4.1 Within-Treatment Comparisons

The most general prediction of a patience advantage concerns our leading treatment *WM*. Figure 1(a) 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 concerns appear important, as approximately 50% of proposals are equal splits (corresponding to 250).²⁸ 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 statics prediction, the weekly CDF clearly lies below the monthly CDF, and we have statistically highly significant first-order stochastic dominance.²⁹ 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,

²⁴We implemented the elicitation task in 4 sessions per treatment. This allows us to check whether the random assignment was successfully implemented in terms of participants' underlying time preferences, which is a crucial aspect of our design and which our data confirm. See Online Appendix E for details.

²⁵We conducted 6 sessions in May and June 2018, and 12 sessions in October and December 2018. The longest delay among the matches selected for payment was 7 months, and the corresponding amount was paid on May 17, 2019.

²⁶The CDF figures present the cumulative distributions in the range of [250, 310] for ease of graphical representation. This range contains, on average, more than 95% of the data.

²⁷Let F and G be two distributions/CDFs. The procedure involves testing both Null hypotheses $F \leq G$ and $G \leq F$, with inference that F first-order stochastically dominates G if and only if $F \leq G$ is accepted while $G \leq F$ is rejected. For the p -values in each test, we employ a bootstrap of size 1,000. Online Appendix B additionally presents a basic Kolmogorov–Smirnov test and OLS-based tests, which yield qualitatively similar results.

²⁸This does not necessarily mean proposers are fair-minded themselves; even if they have “selfish” preferences but believe they are facing a fair-minded respondent, they may optimally propose an equal split.

²⁹We cannot reject the Null that weekly demands (weakly) first-order stochastically dominate monthly demands (p -value = 0.961), while we strongly reject the reversed Null (p -value = 0.001).

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 Delay Advantage in Treatment *WM*). *In Treatment WM, weekly bargainers are observed to be significantly stronger than monthly bargainers.*

Online Appendix A’s Figures 8 and 11 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. We also obtain further confirmation from *respondent* behavior when facing an equal split: While such proposals are generally hardly ever rejected, weekly bargainers reject them somewhat more often (7% vs. 4%).

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

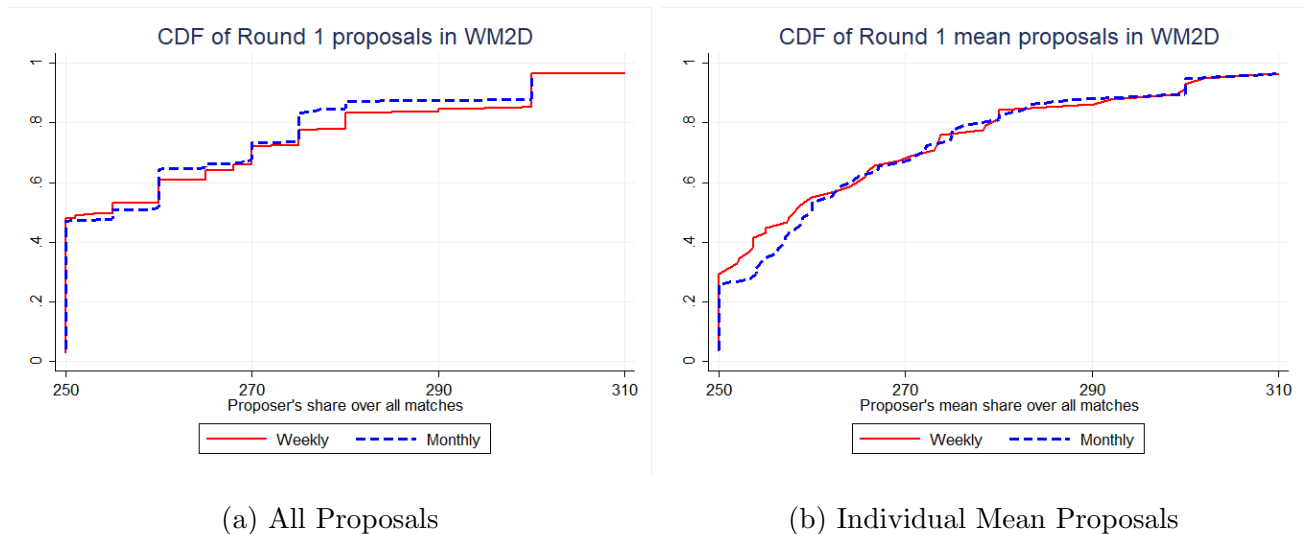


Figure 2: Round-1 Proposals in Treatment *WM2D*

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. 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.³⁰ Figure 2(b) confirms this in terms of individual average demands.

Result 2 (Basic Delay Advantage in Treatment *WM2D*). *In Treatment WM2D, delayed weekly and delayed monthly bargainers are observed to be equally strong.*

³⁰None of the Null hypotheses is rejected (p -values > 0.131), implying that there is no first-order stochastic dominance.

Online Appendix A’s Figures 9 and 12 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 more often than delayed monthly ones (8% vs. 2%)—and also in the later phase of the experiment learn to demand more than their monthly counterparts. This suggests that there is less of a clear common perception of relative bargaining power in this treatment.

We can only speculate that when payoffs are anyways delayed, differences in longer-run effective discounting may be cognitively overwhelmed by the symmetric breakdown risk (or expected to be). Alternatively, time preferences may also exhibit diminishing impatience such that the front-end delay makes both types more patient but more so for the monthly bargainer and thereby level the playing field. In any case, at the very least we find (A2), shared by both Predictions 1 and 2, not confirmed here.

Treatments *WM* and *WM2D* pair two types that differ solely in their basic delay of payoffs per round of disagreement. The third treatment, Treatment *WW1D*, is symmetric in this respect. The only asymmetry between types here is that one is facing a front-end delay whereas the other is not. Under EXD, this “fixed cost” asymmetry is irrelevant, while it confers an advantage under QHD (present bias), as derived in Predictions 1 and 2, respectively (A3).

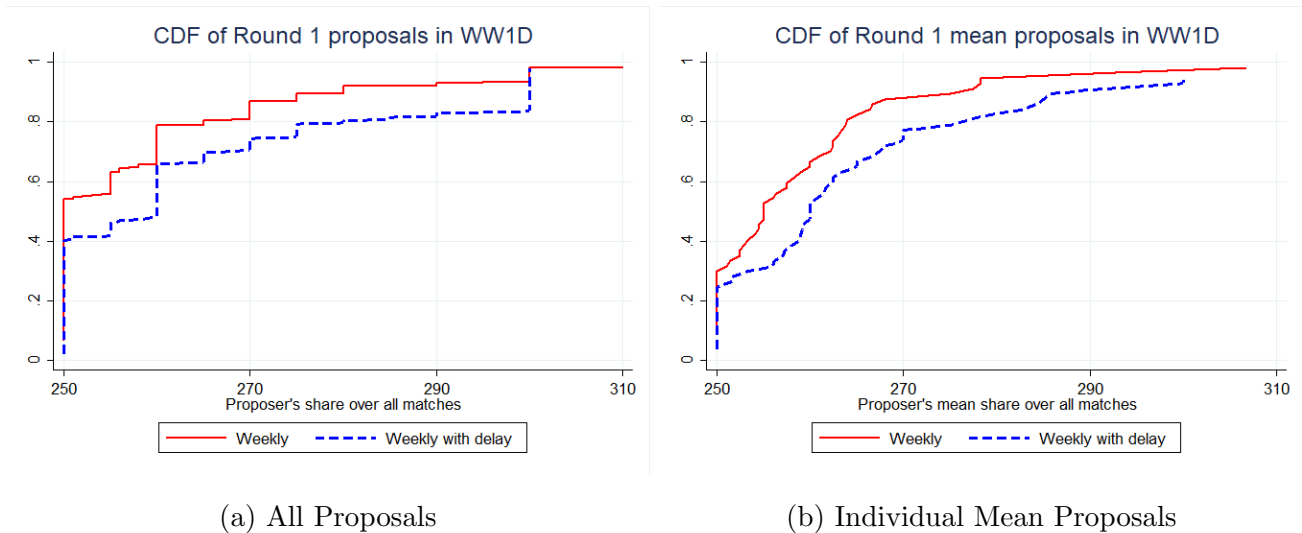


Figure 3: Round-1 Proposals in Treatment *WW1D*

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 (without front-end delay), and the dotted line indicates the CDF of delayed weekly demands. Once again, a large proportion of proposals, between 45% and 50%, 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 first-order stochastically dominate the weekly ones (without front-end delay), as alternatively predicted

under QHD.³¹ We again obtain the same result for individuals' average proposals shown in Figure 3(b).

Result 3 (Front-End Delay Advantage in Treatment *WW1D*). *In Treatment WW1D, delayed weekly bargainers are observed to be significantly stronger than weekly bargainers.*

As for Treatment *WM*, here Online Appendix A's Figures 10 and 13 again provide further confirmation: The result holds up also for accepted proposals, and it obtains immediately after the initial match. Regarding *respondent* behavior, delayed weekly bargainers also reject equal splits at a greater rate (8% vs. 2%).

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 strongly rejected in favor of the alternative prediction under present bias as in QHD (or also HYD). Note also that this result most strongly rejects near-future bias.

Overall, our within-treatment comparisons show that (i) patience is a significant source of bargaining power, (ii) this patience advantage may hinge on the availability of immediate payoffs, and (iii) it exploits present bias.

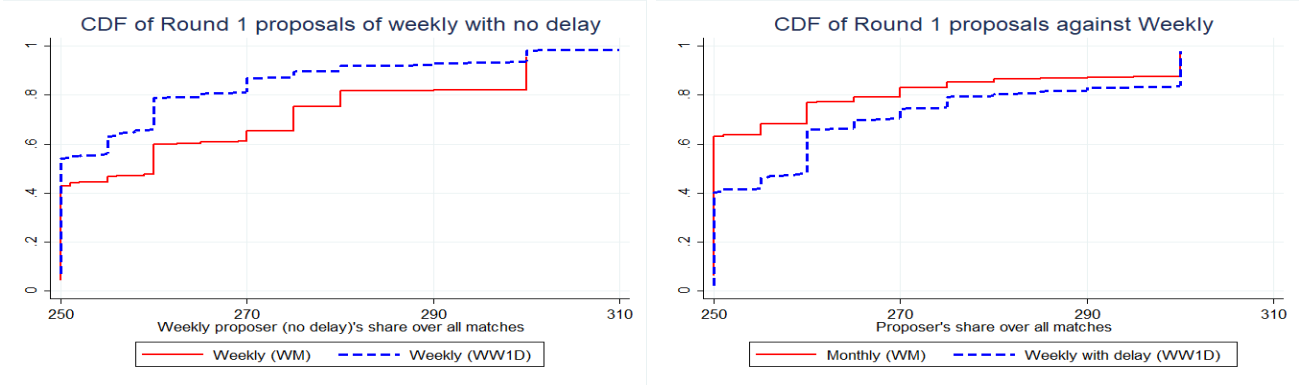
4.2 Between-Treatment Comparisons

Our comparisons between treatments always concern two treatments that have one type in common, matching types A and B vs. matching types A and C. We compare the distributions of initial proposals between the two treatments' games in which this common type has the same role—i.e., AB vs. AC games, and BA vs. CA games—for first-order stochastic dominance. Each comparison now involves two different treatments, hence predictions regarding type A as proposer—asking, for instance, whether type B faces an unambiguously more favorable initial offer distribution from A than type C—and predictions regarding type A as respondent—asking, for instance, whether A faces an unambiguously more favorable initial offer distribution from type B than type C—are tested separately. For simplicity we now only present the results for all Round-1 proposals, since—as seen for within-treatment comparisons—they are similar for individual averages.

First, we compare demands by and offers to weekly bargainers between treatments *WM* and *WW1D*, where both EXD and QHD predict (B1) that the weekly bargainer type is stronger in the former treatment. Figure 4(a) compares 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 clearly first-order stochastically dominate the latter. The difference is also highly statistically significant.³²

³¹We cannot reject the Null that delayed weekly demands (weakly) first-order stochastically dominate weekly demands (p -value = 0.971), while we strongly reject the reversed Null (p -value = 0.001).

³²We cannot reject the Null that weekly demands in Treatment *WM* (weakly) first-order stochastically dominate



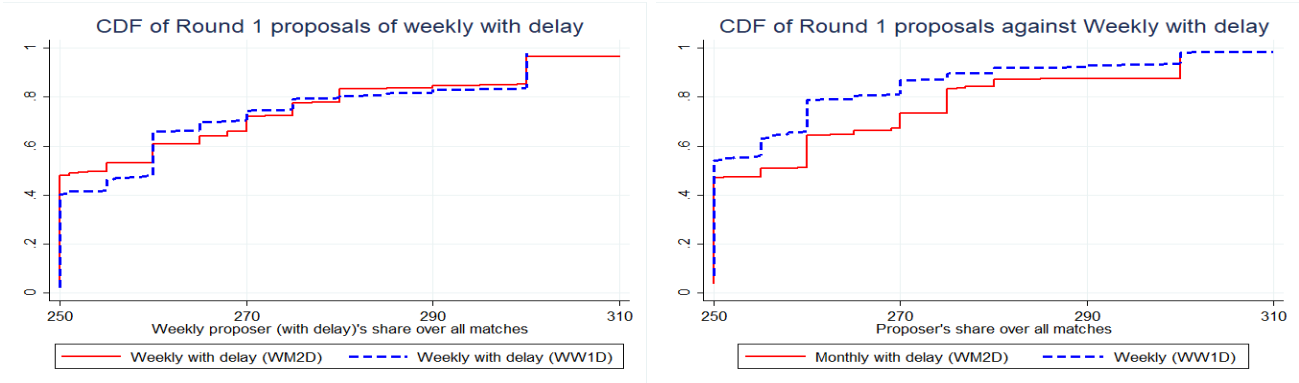
(a) Weekly Proposer in WM vs. $WW1D$

(b) Weekly Respondent in WM vs. $WW1D$

Figure 4: Response to Different Types by Weekly – All Matches

Figure 4(b) compares 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 first-order stochastically dominate the latter, and this difference is highly statistically significant.³³ Hence, we fully and strongly confirm the general prediction (B1) of both EXD and QHD (and also HYD).

Result 4 (Between Treatments WM and $WW1D$). *Weekly bargainers are observed to be significantly stronger against monthly bargainers than against delayed weekly bargainers, both as initial proposer and as initial respondent.*



(a) Del. Weekly Proposer $WM2D$ vs. $WW1D$

(b) Del. Weekly Respondent $WM2D$ vs. $WW1D$

Figure 5: Response to Different Types by Delayed Weekly Bargainers– All Matches

Finally, we turn to prediction (B2) under EXD, and its more permissive qualification under QHD, which compares, respectively, the demands by and the offers to the delayed weekly bargainer between

weekly demands Treatment $WW1D$ (p -value = 0.845), while we strongly reject the reversed Null (p -value = 0.001).

³³We cannot reject the Null that delayed weekly demands in Treatment $WW1D$ (weakly) first-order stochastically dominate monthly demands in Treatment WM (p -value = 0.885), while we strongly reject the reversed Null (p -value = 0.001).

treatments $WM2D$ and $WW1D$. Figure 5(a) compares delayed weekly demands, those 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, since there is neither any qualitative first-order dominance relationship visible between these distributions nor any statistically significant difference between them.³⁴

Figure 5(b) 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 significantly different, they are so with the opposite first-order stochastic dominance relationship.³⁵

Result 5 (Between Treatments $WM2D$ and $WW1D$). *Delayed weekly bargainers are observed to be equally strong against delayed monthly bargainers and weekly bargainers as initial proposer, and to be significantly stronger against weekly bargainers than against delayed monthly bargainers as initial respondent.*

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 (it also rejects near-future bias); the latter finding, however, rejects that part of prediction (B2) under QHD where it agrees with EXD. Though the qualifications and discussion regarding Treatment $WM2D$, following Result 2, apply here as well, 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 3.2). 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.

³⁴If anything, we observe a tendency towards the opposite of what EXD predicts: We cannot reject the Null that delayed weekly demands in Treatment $WW1D$ (weakly) first-order stochastically dominate delayed weekly demands in Treatment $WM2D$ (p -value = 0.239), while we marginally reject the reversed Null (p -value = 0.090).

³⁵We cannot reject the Null that delayed monthly demands in Treatment $WM2D$ (weakly) first-order stochastically dominate weekly demands in Treatment $WW1D$ (p -value = 0.748), while we strongly reject the reversed Null (p -value = 0.001).

4.3 Additional Results and Discussion

4.3.1 Proposer Advantage

A distinct strategic advantage in indefinitely alternating offers bargaining is the so-called proposer advantage. It is due to the asymmetry in the protocol, whereby the respondent carries the burden of delaying agreement and the proposer is able to fully capture the gains from agreeing now rather than later. The proposer advantage is another comparative statics prediction to be tested for within treatments. While it is possible to construct examples of time preferences that violate it, this prediction obtains under EXD, and a present bias as in QHD only reinforces it. Since there exists no prior study of this protocol with actual payoff delay, we also test it here.

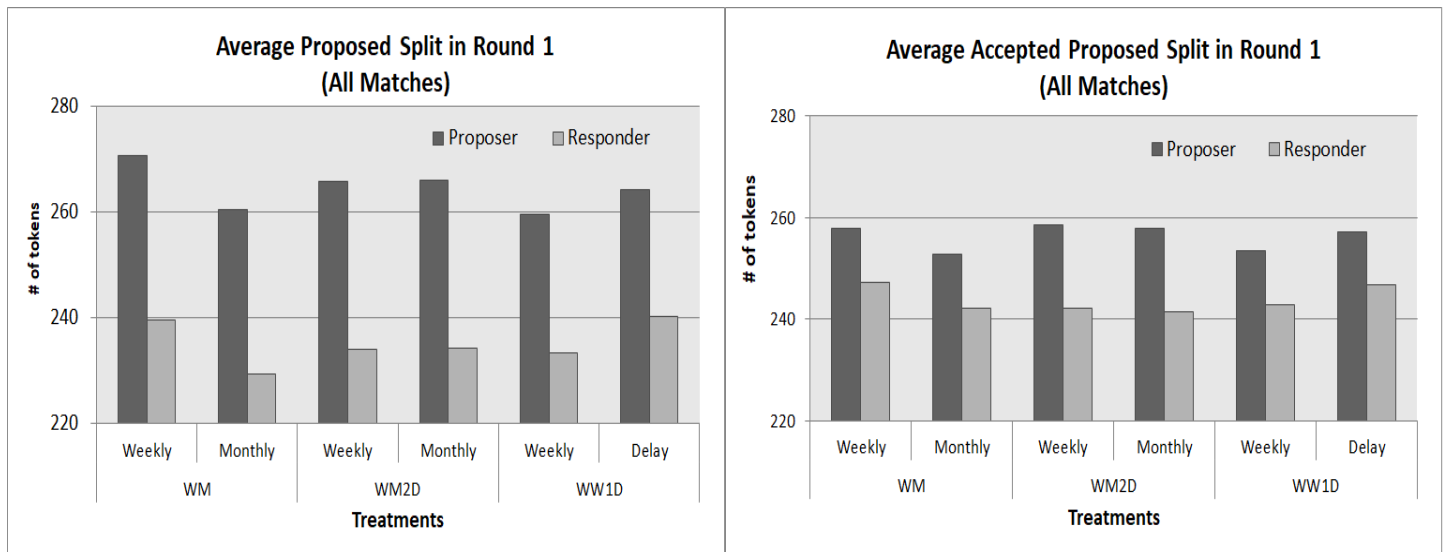


Figure 6: Proposer Advantage – All Matches' Proposals

Again using all initial proposals, Figure 6's left panel reports the average share for proposers and respondents of each type in each treatment over the all matches. For every type, the average share for proposers is clearly larger than that for respondents, and the differences are substantial in magnitude (25–40 tokens). Indeed, in every treatment and for every type, the distribution of shares demanded by initial proposers of this type significantly first-order stochastically dominates that of shares offered to initial respondents of this type.

As already observed, not all proposals are accepted, and we would naturally expect more rejections for proposals that leave less to respondents. We therefore provide a similar comparison for accepted proposals (immediate agreements) in Figure 6's right panel which confirms a proposer advantage. Similar confirmation obtains for final payoffs (see Online Appendix A's Figure 16 regarding all final payoffs, hence including zero payoffs due to exogenous termination, and Figure 17 regarding only final agreement payoffs, hence excluding those under random exogenous termination). Altogether, these findings firmly support the predicted proposer advantage when the cost of disagreement is payoff delay.

4.3.2 Immediate vs. Delayed Agreement, and Incomplete Information

Proposition 1 shows that under minimal preference assumptions, equilibrium predicts immediate agreement. This implausible prediction has been the main criticism against Rubinstein (1982), and it has led to the development of the theory of bargaining under incomplete information to explain delay. While the focus of this paper lies on comparative statics, 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.

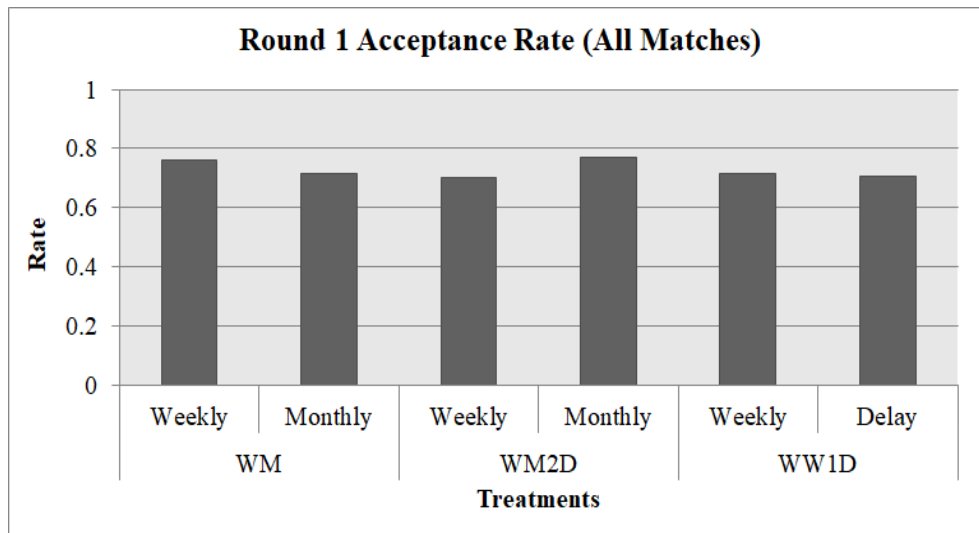


Figure 7: The Proportions of Immediate Agreements – All Matches

Figure 7 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%. 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 results for accepted proposals as for all proposals, so that differences are not driven by rejected proposals.

Online Appendix A's Figures 14 and 15 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.

4.3.3 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 E for details of the task, measurement and analysis). This allows us, on the one hand, to 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. 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.

5 Related Literature

Regarding our basic question of whether time preferences really matter in bargaining, the most closely related work is the experiment by [Manzini \(2001\)](#), which is the only other bargaining study that implements actual time delay of payoffs. [Manzini's](#) design and conclusion are radically different from ours. She first elicits participants' limit prices for avoiding a delay of one or two months, respectively, of a given monetary prize that is otherwise paid the next day, via a variation of the BDM ([Becker, DeGroot, and Marschak, 1964](#)) procedure. The participants are then paired for a single bargaining game with alternating offers over just two rounds, so the second round would be an ultimatum game. Immediate agreement results in payment the subsequent day, whereas delayed agreement results in payment with a month's delay. Providing the bargainers with information on their respective limit prices for a month's delay, these turn out to have no significant correlation with the opening offers.³⁶ Hence, she concludes that time preferences do not matter in bargaining and suggests that the task of bargaining distracts attention completely away from time considerations.

³⁶She also studies two additional treatments implementing shrinking pies in a way that is comparable to the treatment with delayed payments. For both treatments, she finds much higher correlations of opening offers with the opponent's cost of disagreeing.

Our results qualify this negative conclusion. [Manzini](#)'s very careful design quite compellingly shows that time preferences are not *all* that matters. However, in view of the well-established relevance of fairness concerns in bargaining, it is highly unclear how informative the opponent's time preference measure is about their minimum acceptable offer. Apart from the likely noise due to the elicitation method, the measure is taken from individuals' choices over dated *own* payoffs, with no social component, hence relates to a different domain than the distributional one in the bargaining problem. In addition, the fact that this information is "leaked" to the proposer by the experimenter may well induce reluctance by the proposer to exploit it, as well as by the responder to have it exploited. The proposers' opening offers may therefore above all reflect the heterogeneity of proposers' beliefs and tastes, rather than being systematically related to this information about the opponent. By contrast, our effective discounting procedure transparently induces differences in time preferences between groups of participants, holding various other characteristics constant. Thus, we are able to isolate a causal effect of time preferences (or beliefs about time preferences), and we find that it is significant.³⁷

Our review of other related literature focuses on (1) theoretical analyses of time preferences in the canonical bargaining environment with an infinite horizon and alternating offers and (2) experimental studies that investigate this bargaining model. There are large areas of work on bargaining that we do not cover, including the vast experimental literature on ultimatum bargaining and finite-horizon sequential bargaining (though some remarks seem to apply also to finite-horizon experiments), and the theoretical literature that extends the original [Rubinstein \(1982\)](#) model in several other directions, such as asymmetric/incomplete information (about aspects other than time preferences), multilateral bargaining, or endogenous proposer determination. For a recent review of the ultimatum bargaining literature, see [Güth and Kocher \(2014\)](#), and for a survey of sequential bargaining experiments, see [Roth \(1995\)](#). For a comprehensive survey of non-cooperative bargaining theory during its most active period of research, see [Binmore, Osborne, and Rubinstein \(1992\)](#); for a more recent survey focusing on incomplete information, see [Ausubel, Cramton, and Deneckere \(2002\)](#).

Theory. In his seminal paper, [Rubinstein \(1982\)](#) introduces the canonical bargaining model in which two players alternate in making offers to each other on how to divide a given surplus until they reach agreement. Assuming EXD with concave utility and complete information, there is a unique subgame-perfect Nash equilibrium. This equilibrium occurs in stationary strategies that imply immediate agreement in every round. Given impatience and that the burden of delay is with the player responding to an offer, a proposing player enjoys a strategic advantage. Moreover, *ceteris paribus*, the more patient a player is—in particular, the higher her discount factor for given utility—the greater her bargaining power, in the sense of capturing a larger share of the surplus in the equilibrium agreement. With symmetric preferences, as offers become infinitely frequent and players approach perfect patience, the proposer advantage vanishes, and the equilibrium outcome converges to an immediate equal split, as prescribed by the [Nash \(1950\)](#) bargaining solution (see also [Binmore, Rubinstein, and Wolinsky, 1986](#)).

³⁷Note, however, that our [Result 2](#) that the significance of this effect appears to hinge on the presence of immediate payoffs is consistent with [Manzini](#)'s findings, as her design features a common front-end delay.

Motivated by empirical evidence, several theoretical attempts have recently been made to generalize this model in terms of time preferences. Almost all of them have focused on “stable” (time-invariant) preferences to maintain the game’s stationarity property, which makes the game tractable. In this case, any deviation from EXD implies dynamic inconsistency, and [Schweighofer-Kodritsch \(2018\)](#) provides a comprehensive equilibrium characterization under minimal preference assumptions when these preferences are common knowledge (for related work see also [Ok and Masatlioglu, 2007](#); [Noor, 2011](#); [Pan, Webb, and Zank, 2015](#); [Lu, 2016](#)). In particular, he finds that with concave utility, a weak form of present bias is sufficient to obtain a unique equilibrium similar to that under EXD. However, as [Akin \(2007\)](#) and [Haan and Hauck \(2019\)](#) show for QHD, naïveté about present bias may lead to even perpetual disagreement.

We contribute and exploit the observation that under EXD, only payoff delay matters, not bargaining delay. In particular, the [Rubinstein \(1982\)](#) model can be interpreted both as one where payoff delay coincides with bargaining delay and one where there is no bargaining delay but only payoff delay. Under the former interpretation, any disagreement delays the next bargaining round, and the timing of payoffs coincides with that of agreement. Under the latter interpretation, bargaining itself is essentially instantaneous, but payoffs nonetheless are significantly delayed with any disagreement. Based on this latter interpretation, we generalize the model to arbitrary bargaining and payoff delays upon disagreement and general time preferences, under the sole substantial assumption of dynamic consistency. This is similar to bargaining over a time-varying surplus as considered and geometrically analyzed by [Binmore \(1987\)](#), where the variation in surplus stems from non-constant discounting (see also [Coles and Muthoo, 2003](#)). Relative to this prior work, we show that, under very mild assumptions on time preferences, there is a unique equilibrium, and we provide an algebraic proof.

Regarding incomplete information about time preferences, there exists very little theoretical work (and none in recent years). [Rubinstein \(1985\)](#) studied an extension of his seminal complete-information model in which only the initial respondent’s time preferences are not commonly known, such that her (constant) discount factor may take one of two values. Equilibrium becomes subject to a severe multiplicity issue, and while [Rubinstein](#) proposes a selection criterion that delivers a unique prediction, this issue has not been satisfactorily resolved and hindered further progress (see [Binmore et al., 1992](#), for further discussion and references). It seems fair to say that moving towards more realism in modeling incomplete information about time preferences remains elusive until fundamental issues in the theory of games under incomplete information are resolved.

Experiments. [Weg, Rapoport, and Felsenthal \(1990\)](#) and [Rapoport, Weg, and Felsenthal \(1990\)](#) are the first experimental studies of an infinite-horizon, alternating-offers bargaining game. Both implement a within-subjects shrinking-pie design. They compare two conditions, equal and unequal “discount factors,” which correspond to the rates at which the players’ value of the pie shrinks over bargaining rounds.³⁸ To prevent their experiments from lasting too long, they program the computer to

³⁸[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 regarding equilibrium under this cost specification.

terminate the bargaining once the number of rounds exceeds 20 while informing their participants only that a game would be terminated by the experimenters if it lasted “too long.” Based on an analysis of their experimental data on final agreements, initial offers, the number of rounds to reach agreement and the characteristics of counteroffers, they reject the most basic predictions of the [Rubinstein \(1982\)](#) model’s unique equilibrium and argue for the importance of fairness concerns. In particular, they observe neither a significant proposer advantage nor any significant cost advantage.

[Zwick, Rapoport, and Howard \(1992\)](#) experimentally study an environment in which the number of bargaining periods is unlimited and the pie’s value is fixed but bargaining is subject to exogenous random termination. This takes the form of a constant and commonly known breakdown probability. They implement three different probabilities in a between-subjects design. Based on their experimental results, they again reject basic predictions of the [Rubinstein \(1982\)](#) model; e.g., average Round-1 demands are the same under a breakdown probability of 1/10 as under a breakdown probability of 5/6. Furthermore, they also reject the equal split solution.

Like [Weg et al. \(1990\)](#) earlier, [Binmore, Swierzbinski, and Tomlinson \(2007\)](#) employ a shrinking-pie design with unequal discount factors. They adopt a similar forced termination procedure: participants are informed that there will be exogenous termination but not of the exact rule. In fact, the computer intervenes and terminates the game after a randomly drawn number of rounds ranging from 3 to 7. These authors find some behavioral support for the basic predictions of the [Rubinstein \(1982\)](#) model, especially for a proposer advantage. Unlike any of the above studies and ours, however, they have a long and incentivized training/conditioning phase where participants play against a robot programmed to a specific strategy. Moreover, they do not implement the deterministic alternating-offers protocol but instead a random proposer protocol, where the proposer of any round is always randomly chosen from the two players with equal probability, and the pie in their experiment consists of lottery tickets.

Notably, none of these studies features any payoff delay, meaning they cannot speak *directly* to the question of whether or how time preferences matter in bargaining. The domain of outcomes over which preferences are defined is either that of immediate monetary rewards or of lotteries over monetary rewards.³⁹ While shrinking pie designs may mimic discounting, including individually different discounting, this practically takes the form of EXD, unlike most people’s natural discounting. Moreover, the cognitive response to natural time delay may differ from that to computational discounting. This may at least in part explain why our findings are much more favorable towards the theory in every respect. Indeed, we conjecture that the prominent observation of “disadvantageous counteroffers” in finite-horizon alternating-offers bargaining experiments ([Ochs and Roth, 1989](#)) may be an artefact of (mistakes in) computational discounting.

³⁹Somewhat relatedly, [Andreoni and Sprenger \(2012\)](#) conclude that “risk preferences are not time preferences.”

6 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 very 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 two main results are that (i) being (perceived as) more patient increases one’s bargaining power, and (ii) (this perception of) greater patience importantly depends on whether one is subject to a present bias. The first main 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 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.

At the same time, our second main result provides what appears to be the first evidence that people strategically respond to and exploit the present bias of others. This means that the notion of “the” discount factor of a person as capturing her patience under EXD needs to differentiate between the very short run involving immediate rewards and the longer run, as parsimoniously captured by QHD. Hence, our results especially promote theoretical analyses of dynamic strategic interaction incorporating present-biased individuals. With regard to the particular setting of bargaining, our design allows us to abstract from any behavioral implications of dynamic inconsistency (including naïveté), and our findings establish a benchmark for experiments that use a more natural longitudinal design.

We also obtain suggestive evidence for present bias as a feature of general hyperbolic discounting. Our design offers a rare opportunity to investigate the strategic role of time preferences at such a level of detail. While largely unexplored in strategic interaction (though see [Obara and Park, 2017](#), for a notable exception in the context of repeated games), this finding warrants further consideration in both empirical and theoretical work.

Finally, our leading result 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.⁴⁰ We hope our work will help raise awareness of this question's importance and promote future empirical research that quantitatively addresses it.

⁴⁰See 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 Higgins, 1942](#)) and subject to debate among legal scholars (e.g., [Barnhizer, 2005](#)).

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.
- Andreoni, James and Charles Sprenger (2012), “Risk preferences are not time preferences.” *American Economic Review*, 102, 3357–3376.
- 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.
- 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, Martin J. Osborne, and Ariel Rubinstein (1992), “Noncooperative models of bargaining.” In *Handbook of Game Theory with Economic Applications* (Robert Aumann and Sergiu Hart, eds.), volume 1, chapter 7, 179–225, North Holland.
- Binmore, Ken, Joseph Swierzbinski, and Chris Tomlinson (2007), “An experimental test of Rubinstein’s bargaining model.” Mimeo.
- Binmore, Kenneth G., Ariel Rubinstein, and Asher Wolinsky (1986), “The Nash bargaining solution in economic modelling.” *The RAND Journal of Economics*, 17, 176–188.
- 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.
- Chakraborty, Anujit, Yoram Halevy, and Kota Saito (2020), “The relation between behavior under risk and over time.” *American Economic Review: Insights*, 2, 1–16.
- 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.
- 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.
- 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 and Andrew Kloosterman (2020), “An experimental test of the coase conjecture: Fairness in dynamic bargaining.” Mimeo.
- 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.
- Frederick, Shane, George Loewenstein, and Ted O’Donoghue (2002), “Time discounting and time preference: A critical review.” *Journal of Economic Literature*, 40, 351–401.
- Friedenberg, Amanda (2019), “Bargaining under strategic uncertainty: The role of second-order optimism.” *Econometrica*, 87, 1835–1865.
- 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 (2019), “Games with possibly naive hyperbolic discounters.” Mimeo.
- 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.
- Ho, Kate and Robin S. Lee (2017), “Insurer competition in health care markets.” *Econometrica*, 85, 379–417.
- Kim, Jeongbin (2020a), “Discounting, dynamic consistency, and cooperation in an infinitely repeated game experiment.” Mimeo.
- Kim, Jeongbin (2020b), “The effects of time preferences on cooperation: Experimental evidence from infinitely repeated games.” Mimeo.
- 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.
- 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.

- Mischel, W., Y. Shoda, and M. I. Rodriguez (1989), “Delay of gratification in children.” *Science*, 244, 933–938.
- Nash, John F. (1950), “The bargaining problem.” *Econometrica*, 155–162.
- Noor, Jawwad (2011), “Intertemporal choice and the magnitude effect.” *Games and Economic Behavior*, 72, 255–270.
- Obara, Ichiro and Jaeok Park (2017), “Repeated games with general discounting.” *Journal of Economic Theory*, 172, 348–375.
- 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), *The Strategic Approach: A Model of Alternating Offers*, chapter 3, 29–68. 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.
- 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 (2020), “What drives the gender wage gap? Examining the roles of sorting, productivity differences, bargaining and discrimination.” *Review of Economics and Statistics*, (forthcoming).
- 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 (2021), “Patience and comparative development.” *Review of Economic Studies*, (forthcoming).
- 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.
- Vannetelbosch, Vincent J. (1999), “Alternating-offer bargaining and common knowledge of rationality.” *Theory and Decision*, 47, 111–138.
- 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.

Appendix: Proofs

This Appendix provides proofs for all theoretical results in the paper, in the order of their appearance: Lemma 1, Proposition 1, Prediction 1 and Prediction 2.

Lemma 1

Proof. Define, for each player i , the function $f_i : [0, 1] \rightarrow [0, 1]$ 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 (2.1) 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 (2.1) 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 assumption 3 with $\alpha \equiv \max\{\alpha_1, \alpha_2\}$ —observe that

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

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

Proposition 1

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 $u_1(a_1)^N \geq U_1 \geq u_1(b_1^N)$, 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. \square

Prediction 1

Proof. First, define $f(U) \equiv 1 - u^{-1}(U)$ for any $U \in [0, 1]$, so Proposition 1 implies that the unique equilibrium is characterized by

$$x_1^E = f(\phi_2 \delta u(f(\phi_1 \delta u(x_1^E)))) \text{ and } x_2^E = f(\phi_1 \delta u(x_1^E)), \quad (.1)$$

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 $q \in [0, 1]$.⁴¹ 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 .1, this covers all parts except for (A3). The latter follows directly from the irrelevance of front-end delay under EXD.

Finally, regarding a Proposer Advantage, simply observe that $x_i^E > u^{-1}(\phi_i \delta u(x_i^E)) = 1 - f(\phi_i \delta u(x_i^E)) = 1 - x_j^E$. \square

Prediction 2

Proof. The second-round continuation equilibrium is characterized by the shares x_i^E solving the two equations (.1). Backward induction then yields immediate agreement in the first round, with the initial proposer's share given by

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

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.

⁴¹Our 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.

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, $(\phi_2, \beta_2) = (1, \beta)$, $1 - x_1^Q$ is greater against the monthly bargainer, $(\phi_1, \beta_1) = (\phi, \beta)$, than against the delayed weekly bargainer, $(\phi_1, \beta_1) = (1, 1)$. 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.

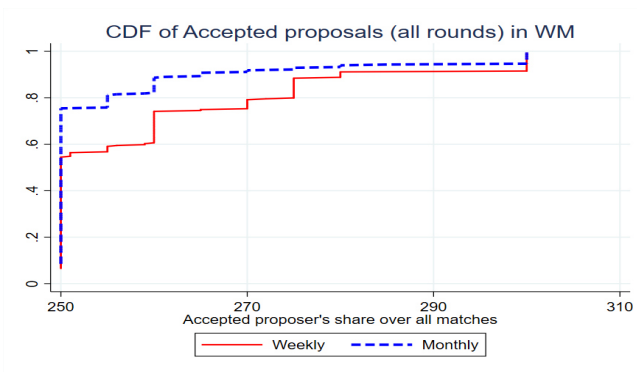
Regarding (B2), first observe that with the initial respondent's type equal to $(\phi_2, \beta_2) = (1, 1)$, her continuation share—hence also $1 - x_1^Q$ —is smaller against the weekly than the monthly bargainer, as under EXD. Second, fixing $(\phi_1, \beta_1) = (1, 1)$, it should be clear from continuity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

The proposer advantage follows straight from the corresponding proof for EXD upon noting that $\beta_2 \leq 1$ implies $x_1^Q \geq x_1^E$, since $x_1^E > 1 - x_2^E = u^{-1}(\phi_1 \delta u(x_1^E)) \geq u^{-1}(\beta_1 \phi_1 \delta u(x_1^E))$. \square

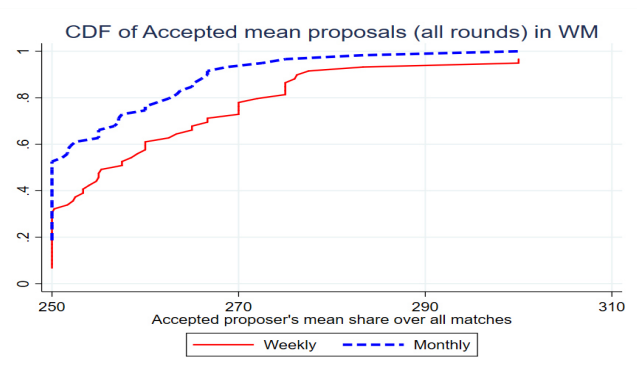
Online Appendix: Supplemental Material

This Online Appendix consists of five parts and provides the following supplemental material: Appendix **A** provides additional figures that complement those provided in the main body of the paper; Appendix **B** shows the results of alternative statistical tests; Appendices **C** and **D** contain experimental instructions and selected screenshots, respectively, for one exemplary experimental treatment (Treatment *WM*); final Appendix **E** presents all details of our additional time preference elicitation and results on how measured time preferences relate to bargaining behavior.

A Additional Figures

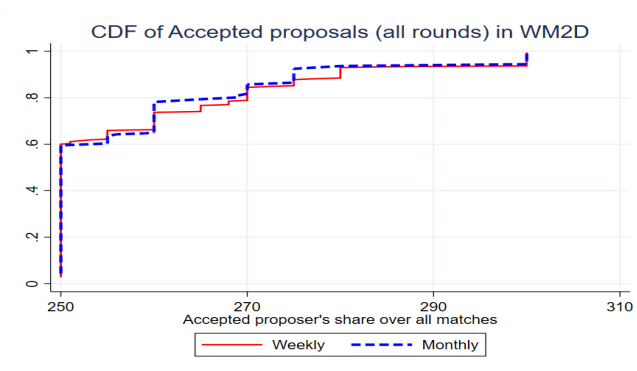


(a) Accepted Proposals

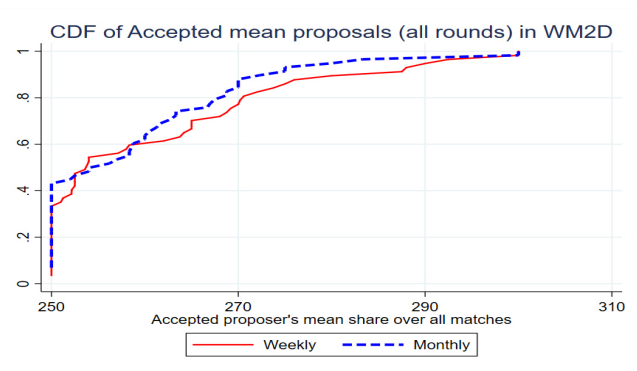


(b) Accepted Mean Proposals

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

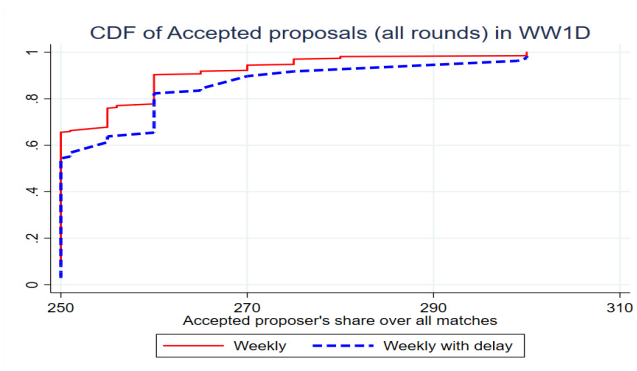


(a) Accepted Proposals

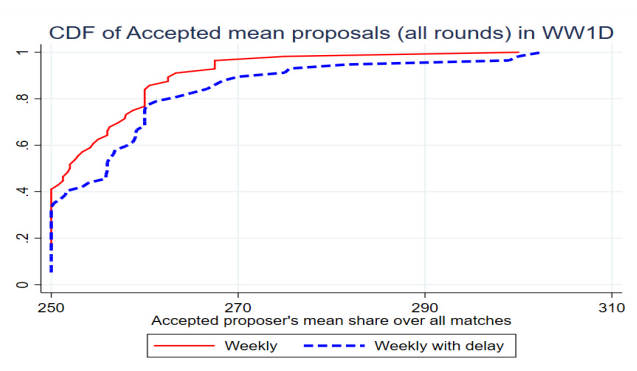


(b) Accepted Mean Proposals

Figure 9: Accepted Proposals over All Matches in Treatment *WM2D*



(a) Accepted Proposals



(b) Accepted Mean Proposals

Figure 10: Accepted Proposals over All Matches in Treatment *WW1D*

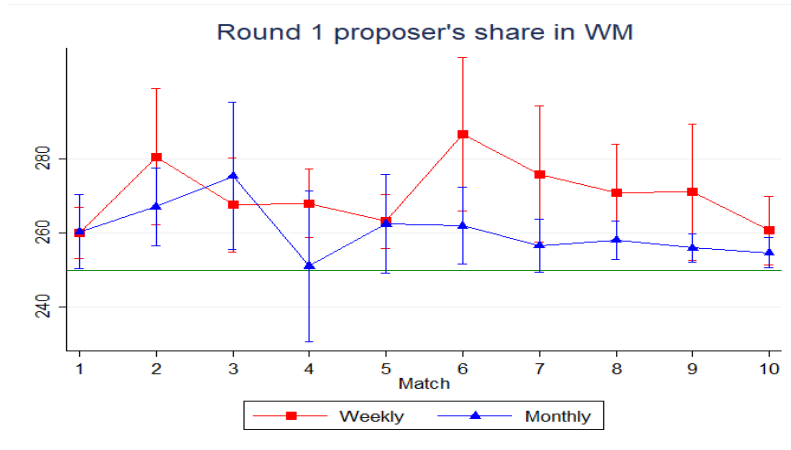


Figure 11: Round-1 Proposals over Matches in Treatment *WM*

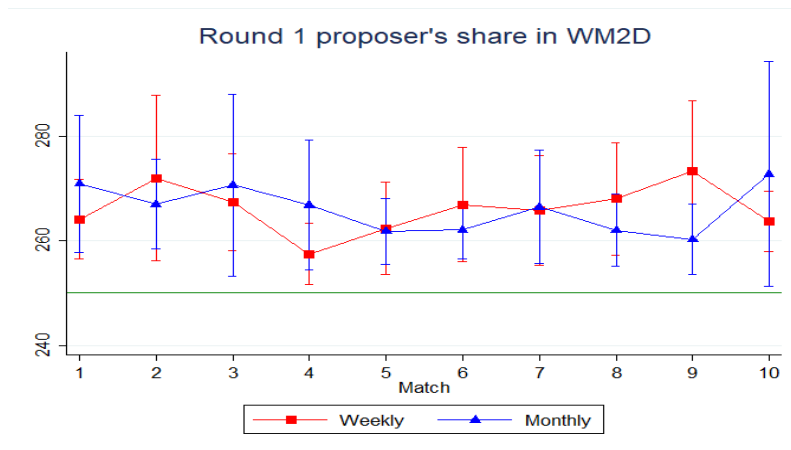


Figure 12: Round-1 Proposals over Matches in Treatment *WM2D*

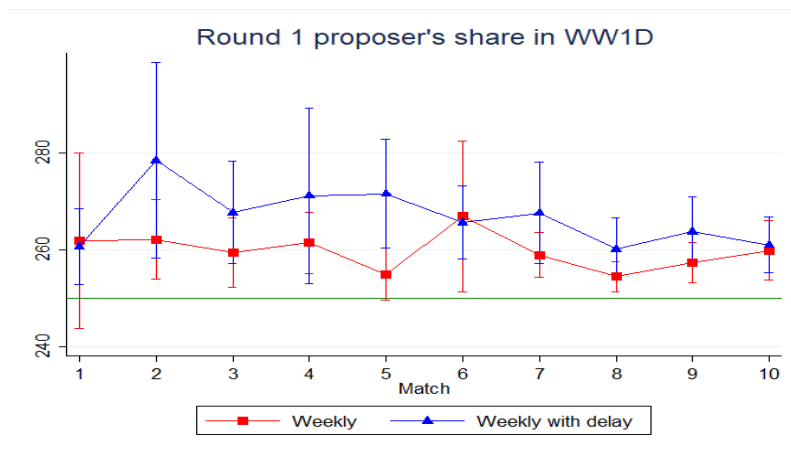


Figure 13: Round-1 Proposals over Matches in Treatment *WW1D*

Immediate vs. Delayed Agreements

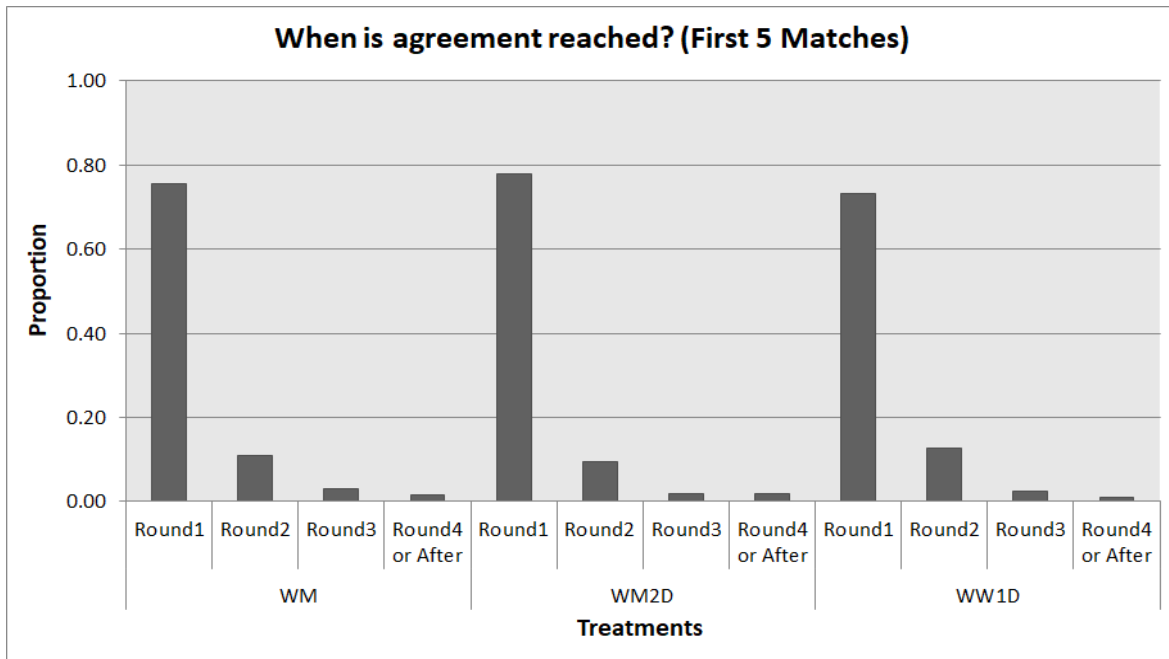


Figure 14: The Proportions of Agreements over Rounds – First 5 Matches

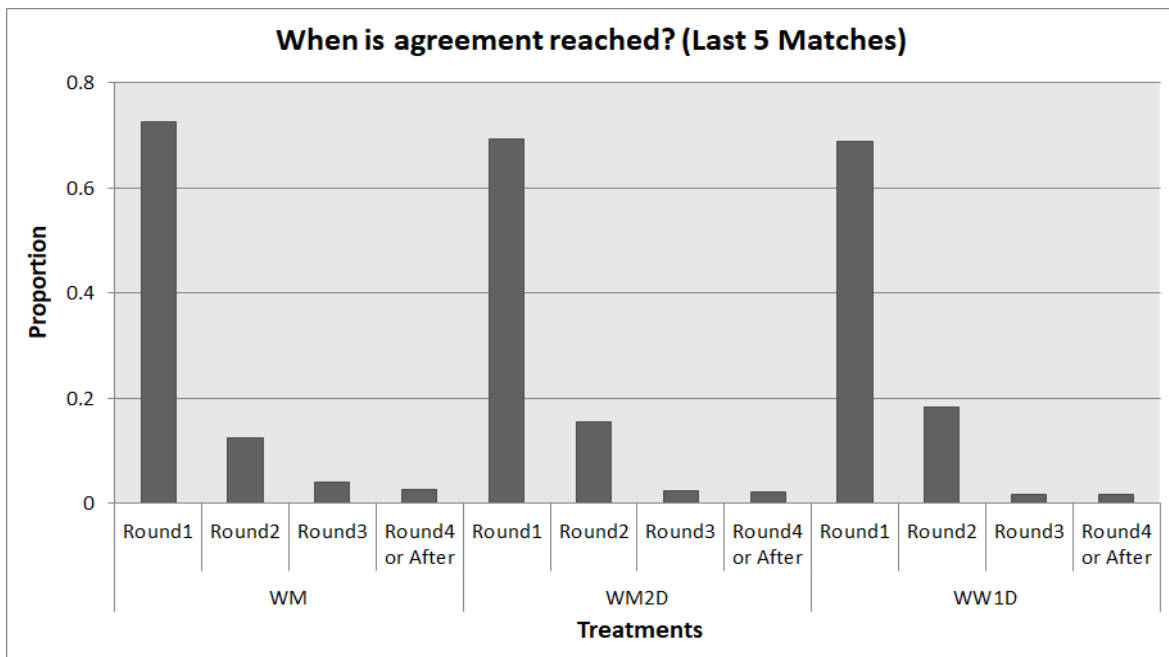


Figure 15: The Proportions of Agreements over Rounds – Last 5 Matches

Proposer Advantage: Final Payoffs incl. Random Termination

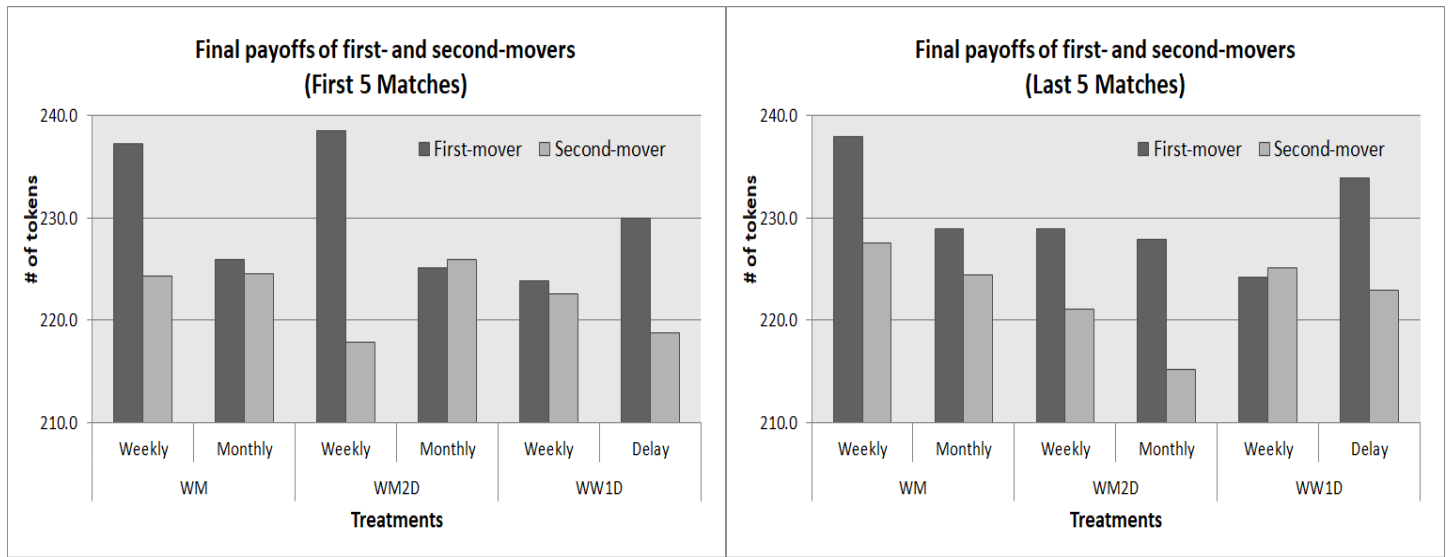


Figure 16: Final Payoffs (All) – First and Last 5 Matches

Proposer Advantage: Final Payoffs excl. Random Termination

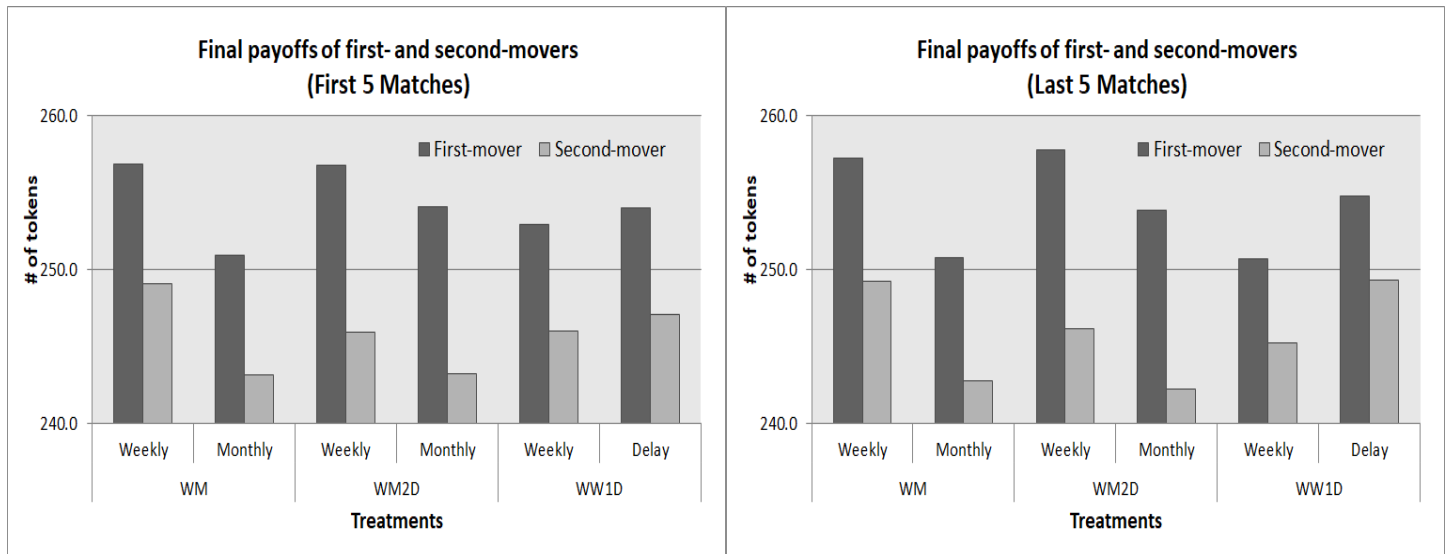


Figure 17: Final Payoffs (excl. Random Terminations) – First and Last 5 Matches

B Alternative Statistical Tests

Table 3: Alternative Statistical Tests

	K-S test		OLS clustered S.E.		OLS clustered S.E. with fixed effects	
	Panel (a)	Panel (b)	Panel (a)	Panel (b)	Panel (a)	Panel (b)
Figure 1	0.218**	0.248**	10.210** (4.958)	8.891 ^b (5.614)	9.701** (4.851)	8.911 ^c (5.594)
Figure 2	0.066	0.138	-0.251 (3.791)	-1.747 (3.826)	-0.748 (3.454)	-7.747 (3.749)
Figure 3	0.183***	0.228 ^a	-7.052** (3.252)	-6.631* (3.558)	-7.018** (3.266)	-6.631* (3.594)
Figure 4	0.218***	0.228***	13.790** (6.375)	-6.423** (3.152)	15.270*** (5.053)	5.028 (7.256)
Figure 5	0.083	0.149***	3.709 (4.140)	5.077 (3.833)	-1.140 (4.694)	15.210 (13.330)

Notes: Kolmogorov–Smirnov test, OLS with standard errors clustered at the individual level, 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) of the OLS, standard errors are not clustered. ***Significant at 1%; **5%; *10%.

a: p-value=0.103, *b:* p-value=0.115, *c:* p-value=0.114.

C Experimental Instructions – 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.

D Selected z-Tree Screenshots

Timing of Payment		
Round	The Other's Type (A)	Your Type (B)
Round 1	Today	Today
Round 2	In 1 week	In 1 month
Round 3	In 2 weeks	In 2 months
...

Previous Round Summary will be provided here.

Match: 1 **Round: 1**

Your Role in this round = **PROPOSER**.

Payment timing for the accepted split:
 Your tokens **Today**.
 The other person's tokens **Today**.

Please propose a split of 500 tokens:

Tokens for yourself:

Tokens for the other person:

Submit

Figure 18: *Proposer's* Screen

Timing of Payment		
Round	Your Type (A)	The Other's Type (B)
Round 1	Today	Today
Round 2	In 1 week	In 1 month
Round 3	In 2 weeks	In 2 months
...

Previous Round Summary will be provided here.

Match: 1 **Round: 1**

Your Role in this round = **RESPONDER**.

Payment timing for the accepted split:
 Your tokens **Today**.
 The other person's tokens **Today**.

The proposed split from the other person:
 250 for you and 250 for the other.

Would you like to accept / reject the proposal?
Please select a column

Accept

Reject

Submit

Figure 19: *Responder's* Screen

E 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, 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 4: Description of the Elicitation Task

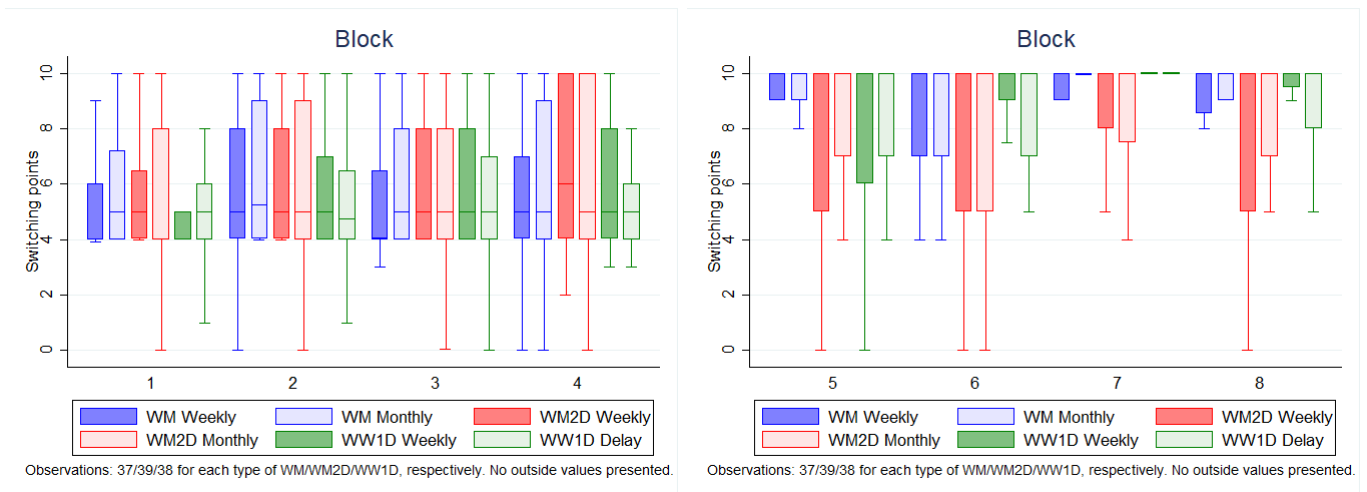
Switching	Sooner \Rightarrow Later			
Block	(1)	(2)	(3)	(4)
Sooner	\$4 Today	\$4 Today	\$4 1 month	\$4 1 month
Later	\$X 1 week	\$X 1 month	\$X 1 month and 1 week	\$X 2 months
Switching	Sooner \Leftarrow Later			
Block	(5)	(6)	(7)	(8)
Sooner	\$X Today	\$X Today	\$X 1 month	\$X 1 month
Later	\$10 1 week	\$10 1 month	\$10 1 month and 1 week	\$10 2 months

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

Table 4 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 $k \in \{1, 2, \dots, 8\}$ refers to the block number, by treatment and bargainer type, to check whether our randomization in terms of underlying time preferences was successful. Figure 20 provides the corresponding box plots. We use the same test as for our bargaining predictions, the Kolmogorov-Smirnov test, to compare 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 carry out analogous tests on the time preference task responses. Comparing, first, the switching points between the two bargainer types within any treatment—e.g., weekly vs. monthly 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).⁴² Overall, we therefore conclude that our randomization into treatments and types in terms of underlying time preferences was successful indeed.



(a) Blocks 1-4

(b) Blocks 5-8

Figure 20: Distribution of Switching Points by Type/Treatment/Block

⁴²We 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 comparison (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.⁴³ 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 5 in terms of averages with standard deviations, and in Figure 21 in terms of box-plots, by types and treatments.

Table 5: 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 5 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 21 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

⁴³Inconsistency 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.

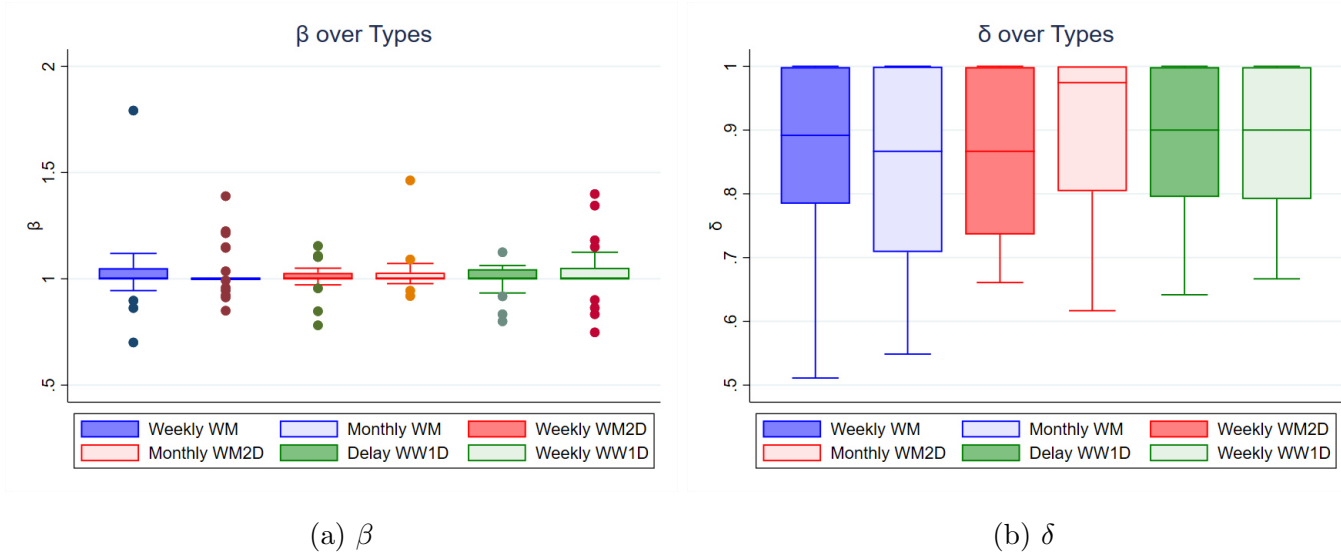


Figure 21: Distribution of β and δ by Type/Treatment

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 6 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.⁴⁴

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 7 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

⁴⁴The number of observations overall is 131, because one of the final 132 participants with estimated discount factors (see Table 5) happened to never be selected as initial proposer.

Table 6: β , δ , and Round-1 Average Proposer behavior (OLS)

	All Matches	First Five	Last Five
	(1)	(2)	(3)
β	14.78 (16.85)	2.00 (11.27)	20.46 (20.36)
δ	26.33 (14.98)	36.11** (15.09)	7.11 (19.54)
Constant	224.5*** (23.66)	230.5*** (18.81)	232.8*** (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 7: β , δ , and Round-1 Average Acceptance of the Equal Splits (OLS)

	All Matches	First Five	Last Five
	(1)	(2)	(3)
β	0.14 (0.16)	-0.06 (0.12)	0.50 (0.29)
δ	-0.12 (0.16)	-0.46** (0.17)	0.16 (0.15)
Constant	0.90*** (0.26)	1.39** (0.21)	0.32 (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.

Instructions

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 Option A to Option B, or from Option B to Option A.

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 22: Elicitation Task Screen-shot Block 1