

Schweighofer-Kodritsch, Sebastian

Working Paper

Time Preferences and Bargaining

Discussion Paper, No. 38

Provided in Cooperation with:

University of Munich (LMU) and Humboldt University Berlin, Collaborative Research Center
Transregio 190: Rationality and Competition

Suggested Citation: Schweighofer-Kodritsch, Sebastian (2017) : Time Preferences and Bargaining, Discussion Paper, No. 38, Ludwig-Maximilians-Universität München und Humboldt-Universität zu Berlin, Collaborative Research Center Transregio 190 - Rationality and Competition, München und Berlin

This Version is available at:

<https://hdl.handle.net/10419/185708>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

Time Preferences and Bargaining

Sebastian Schweighofer-Kodritsch (Humboldt University Berlin and WZB Berlin)

Discussion Paper No. 38

June 18, 2017

Time Preferences and Bargaining*

Sebastian Schweighofer-Kodritsch[†]

June 17, 2017

Abstract

This paper presents an analysis of general time preferences in the canonical Rubinstein (1982) model of bargaining, allowing for arbitrarily history-dependent strategies. I derive a simple sufficient structure for optimal punishments and thereby fully characterize (i) the set of equilibrium outcomes for any given preference profile, and (ii) the set of preference profiles for which equilibrium is unique. Based on this characterization, I establish that a weak notion of present bias—implied, e.g., by any hyperbolic or quasi-hyperbolic discounting—is sufficient for equilibrium to be unique, stationary and efficient. Conversely, I demonstrate how certain violations of present bias give rise to multiple (non-stationary) equilibria that feature delayed agreement under gradually increasing offers.

Keywords: time preferences, dynamic inconsistency, alternating offers, bargaining, optimal punishments, delay

JEL classification: C78, D03, D74

*This paper generalizes my earlier work “On Time-Inconsistency in Bargaining”, which was itself based on my dissertation submitted to the LSE. I am deeply grateful to Erik Eyster, Balázs Szentes and Yves Breitmoser for their support, guidance and inspiration during various stages of this project. I also thank the co-editor Itzhak Gilboa and three anonymous referees for insightful comments on a prior draft, Can Çeliktumur, Leonardo Felli, Paul Heidhues, Alex Imas, Kristóf Madarász, Francesco Nava, Motty Perry, Michele Piccione, Matthew Rabin, Ariel Rubinstein, Roland Strausz, Justin Valasek, Roel van Veldhuizen and Georg Weizsäcker for illuminating discussions, and my audiences at the LSE, STICERD, RES PhD Presentation Meeting 2013, DIW Berlin, University of Bonn, University of Cologne, Goethe-University Frankfurt, Florida State University, University of Oslo, NOeG 2013 Conference, University of Innsbruck, NAG 2014 Conference, 6th Israeli Game Theory Conference, CEA 2014 Conference, FUR 2014 Conference, 25th International Conference on Game Theory, 13th Viennese Workshop on Optimal Control and Dynamic Games, ESEM 2016, University of Mannheim, WZB Berlin and UECE Lisbon Meetings 2016 for helpful remarks. Financial support by the German Science Foundation (CRC TRR 190) is gratefully acknowledged.

[†]Humboldt-University Berlin and WZB Berlin Social Science Center. Mail: Wirtschaftswissenschaftliche Fakultät, Humboldt-Universität zu Berlin, Spandauer Str. 1, D-10099 Berlin, Germany. Phone: +49 (0)30 2093 5998. Email: sebastian.kodritsch@gmail.com.

1 Introduction

In the absence of irrevocable commitments, time is the prime variable of bargaining agreements: the parties may agree not only now or never, but also sooner or later. The question of how the parties’ time preferences govern bargaining outcomes lies at the heart of modern bargaining theory (Ståhl, 1972; Rubinstein, 1982). Under the traditional assumption of exponential discounting (ED), their impatience drives the parties towards a sharply predictable immediate agreement, which is efficient. Beyond this special case, a full understanding of the fundamental role of time preferences in bargaining has remained elusive, however. Since any violation of ED implies dynamic inconsistency, standard techniques for characterizing equilibrium fail to be applicable. While bargaining theory has turned towards informational frictions to explain inefficient delay, it has remained an open question whether time preferences alone might already impose a friction on the parties’ ability to reach agreement immediately, when they are dynamically inconsistent.¹

This paper provides a general analytic framework for the canonical Rubinstein (1982) bargaining model: I derive a simple sufficient off-path “punishment” structure, supporting all equilibrium behavior, that renders arbitrarily history-dependent strategies analytically tractable under only minimal restrictions on time preferences and dynamic inconsistency. The resulting equilibrium characterization puts the aforementioned influential conclusions from ED on a solid basis: it shows that they extend to all time preferences satisfying a weak notion of *present bias*, which covers all established models of dynamic inconsistency. The characterization further reveals that, more generally, the extension to dynamic inconsistency is non-trivial, and changes in relative impatience can matter in equilibrium: a novel kind of equilibrium delay emerges when at least one party exhibits instead a *near-future bias*.

I consider any profile of time preferences such that a party i evaluates delayed agreements with a continuous utility function $U_i(x_i, t)$ and assume only that she prefers a greater surplus share x_i , holding the delay t constant, and a shorter delay t , holding her surplus share $x_i > 0$ constant. This covers all existing models of time preferences, with ED as the only special case for which preferences are dynamically consistent (Halevy, 2015). The standard solution technique for characterizing equilibrium under ED exploits the game’s stationarity via recursions on the parties’ extreme continuation values (see Shaked and Sutton, 1984). Under dynamic inconsistency, however, the possibility of multiple and delayed equilibrium agreements means that a party may not rank these agreements the same way across different points in time, hence continuation values do not encode sufficient information to determine

¹Abreu, Pearce, and Stacchetti (2015) give a recent account of the literature on bargaining under incomplete information; see, however, also Yildiz (2011). Another line of research has examined variations of the bargaining protocol (e.g., Muthoo, 1990), or even endogenized it (see Perry and Reny, 1993; Sákovics, 1993).

present values.² As a consequence, the recursion breaks down, and it is not known what additional information regarding the underlying set of continuation agreements would be required or how to obtain it.

To overcome this problem I directly analyze the off-path punishments (continuation equilibria, upon rejection) that support all equilibrium play, i.e. *optimal penal codes* (cf. [Abreu, 1988](#)). I show it is sufficient to consider *simple* penal codes described by four “extreme” outcomes. These, jointly, define four punishments such that the exact same punishment is used to deter *any* deviation by a given player in a given role (hence four), independent of the deviation’s history; e.g., any deviation by player 1 as the proposer triggers the exact same continuation equilibrium (upon rejection), on as well as off the path.³ This simplified structure renders equilibrium analysis tractable for general time preferences; e.g., it reveals that the crucial piece of information required for recursion on a player’s extreme (continuation) values is the extreme/maximal (continuation) delay, which is itself jointly determined with the players’ extreme values by the optimal punishments. Moreover, it allows me to exploit a fixed-point property of any optimal simple penal code—that each of its four outcomes is extreme among all those they jointly support—to arrive at the paper’s core results: a full characterization of both (i) the set of equilibrium outcomes for any given preference profile, and (ii) the set of preference profiles that imply a unique equilibrium.

While the strategic implications of dynamic inconsistency in bargaining can be subtle, for the purposes of applied work the characterization bears the good news that the conclusions from ED are confirmed. Under standard concavity assumptions on preferences concerning the surplus share, a weak notion of *present bias* (to be satisfied by both parties) turns out sufficient for equilibrium to be unique. In this case equilibrium is also stationary and implies an immediate agreement that depends only on the parties’ attitudes to a single (first) period of delay.⁴ The notion of present bias identified here means that a decision maker finds a given delay most costly when it concerns an otherwise immediate reward, as opposed to further delaying a future reward. This property is readily testable empirically and easily checked for any given model of time preferences; in particular, any discounting that is hyperbolic ([Chung and Herrnstein, 1967](#); [Ainslie, 1975](#)) or quasi-hyperbolic ([Phelps and Pollak, 1968](#);

²As solution concept I use the natural two-player version of multiple-selves equilibrium, which is equivalent to subgame perfection under ED (cf. [Chade, Prokopovych, and Smith, 2008](#)).

³[Mailath, Nocke, and White \(2015\)](#) present related examples of repeated *sequential* games where no simple penal code is optimal due to incentive trade-offs between within-round and continuation punishment. By contrast, here a single round’s play determines all payoffs.

⁴The curvatures of the parties’ utilities in the surplus share govern *stationary* equilibrium, and they are essentially orthogonal to dynamic (in-)consistency. Even under ED there are multiple stationary equilibria if utilities are sufficiently convex in the surplus share. Under non-separability there is no atemporal utility from surplus—e.g., see the magnitude-effects model advanced by [Noor \(2011\)](#)—and this curvature may become more convex for delayed shares, which may also produce such multiplicity (section 5.2 has details).

Laibson, 1997) satisfies it. Thus this paper equips applied work with a strategically founded bargaining solution for these commonly considered preferences. The solution is then not only sharp and simple, with straightforward comparative statics, but—by virtue of the more general sufficiency result given present bias—also robust to misspecification of the parties’ attitudes to delay beyond a single period.

Conversely, a theoretically novel kind of equilibrium delay arises when at least one of the two parties has instead a *near-future bias*. Such a decision maker finds delaying a near-future reward by a given amount of time more costly than delaying an immediate one. For instance, any discounting function that is initially concave, hence falling steepest not at zero but at some positive delay, implies this property (e.g., Ebert and Prelec, 2007; Bleichrodt, Rohde, and Wakker, 2009). In contrast to present bias, a party with a near-future bias does not exert immediate control over the delay she finds most costly. The price that her current self is willing to pay to avoid a near-future delay is excessive to her near-future self, to whom that same delay is immediate; i.e., she subsequently becomes more patient. This makes delay *self-enforcing*: a delay off the path—as a threat that commands a “self-control premium” for immediate agreement—supports itself on the path.⁵ Moreover, any such delay is supported as a gradual agreement, where the parties gradually increase their offers over the course of the bargaining. This delay-result under near-future bias clarifies the importance of present bias for reaching immediate agreement, and it informs future theoretical work with dynamically inconsistent preferences by showing how changes in relative impatience can matter in bargaining.

Related Literature. There exists little prior work on bargaining that analyzes dynamically inconsistent time preferences: Burgos, Grant, and Kajii (2002a); Akin (2007); Ok and Masatlioglu (2007); Noor (2011).⁶ All of these papers restrict attention to stationary strategies, however, thus severely limiting the potential for dynamic inconsistency to matter.⁷ This paper studies a general class of preferences that covers all of those studied previously and at the same time generalizes the analysis to arbitrarily history-dependent strategies.

Other closely related work investigates non-stationary time preferences that are, however, dynamically consistent (Binmore, 1987; Rusinowska, 2004; Pan, Webb, and Zank, 2015); e.g., a player may apply different discount rates to June 30, 2016, and July 1, 2016, but

⁵Rather than relying on stationary equilibrium off the path, as in prior constructions (Avery and Zemsky, 1994), such delay equilibria are non-stationary in *every* subgame.

⁶Burgos et al. (2002a) study bargaining with breakdown risk for certain non-expected-utility preferences; Akin (2007) also investigates naïveté and learning by quasi-hyperbolic discounters.

⁷The sole exception is Lu (2016) who studies bargaining by sophisticated quasi-hyperbolic discounters. In his model bargaining is over an infinite stream of cakes rather than a single one, however, so agreements are infinite consumption commitments.

independent of the delay to these dates.⁸ I abstract from such exogenous effects of time on the players’ preferences, which would also appear negligible under frequent offers; instead, the discount rate for any given period may depend only on the delay to this period, not on its absolute time. Moreover, I maintain that preferences are history-independent; i.e., unlike in [Fershtman and Seidmann \(1993\)](#), and [Li \(2007\)](#), where the best past offer acts as a “reference point”, the parties are consequentialist, caring only about the eventual surplus division and its delay, not *how* agreement is reached.

Regarding the power of history-dependent strategies in generating delay, also the work that endogenizes the timing of offers, starting with [Perry and Reny \(1993\)](#) and [Sákovics \(1993\)](#), as well as that on “negotiation” by [Busch and Wen \(1995\)](#) where, as long as parties fail to agree, they repeatedly play a disagreement game, share similarities. The underlying reason for why history-dependent strategies are powerful here—namely, dynamic inconsistency—is fundamentally different, however.

Finally, this paper contributes to the wider literature that explores the bargaining implications of relaxing certain hitherto standard but “unrealistic” (or extreme) assumptions about the players. Whereas this model’s only non-standard feature is dynamically inconsistent *preferences*, relaxing ED, most of the recent literature has been concerned with non-standard *beliefs*, relaxing common knowledge of the bargaining protocol or of players’ rationality (e.g., [Yildiz, 2011](#); [Friedenberg, 2016](#)).

Outline. After introducing the formal model in section 2, section 3 already describes the main results of this paper for the special case where players maximize their discounted share of the surplus for arbitrary discounting; this generalizes the most widely used version of the [Rubinstein \(1982\)](#) model. Section 4 then contains the full-fledged equilibrium characterization, and I further investigate equilibrium uniqueness and multiplicity/delay in section 5. Finally, I offer some concluding remarks in section 6. All formal proofs (as well as additional notation) are found in appendix A; appendix B contains supplementary material.

2 Bargaining and Time Preferences

2.1 Bargaining Protocol, Histories and Strategies

I follow [Rubinstein \(1982\)](#) exactly with regards to the bargaining protocol of (possibly indefinitely) alternating offers. There are two players $\{1, 2\} \equiv I$, who bargain over a perfectly

⁸This is similar to time-varying surplus as in [Coles and Muthoo \(2003\)](#); see also [Merlo and Wilson \(1995\)](#) and [Cripps \(1998\)](#), who investigate Markovian surplus processes. All of these models maintain dynamic consistency of preferences; indeed, delay typically occurs only when efficient.

divisible surplus of (normalized) size one. Throughout the paper, whenever $i \in I$ denotes one player, $j \equiv 3 - i$ denotes the other. In round $n \in \mathbb{N}$, player $P(n)$ proposes a surplus division $x \in \{(x_1, x_2) \in \mathbb{R}_+^2 \mid x_1 + x_2 = 1\} \equiv X$ to her opponent $R(n)$ (equivalently, $P(n)$ offers $R(n)$ share $x_{R(n)}$), who then responds by either accepting or rejecting the proposal. If it is accepted, the game ends with agreement on x ; otherwise, one period of time elapses until the subsequent round $n + 1$, where the roles of proposer and respondent are reversed, so $P(n + 1) = R(n)$. This process of alternating offers begins with player 1's proposal, i.e. $P(1) = 1$, and continues until there is agreement, possibly without ever terminating.

A *history* of play to the beginning of round $n \in \mathbb{N}$ is a sequence of $n - 1$ rejected proposals $h^{n-1} \in X^{n-1}$, where $X^0 \equiv \{\emptyset\}$; throughout, “history” always refers to such a beginning-of-round history. A *strategy* σ_i of a player i assigns to every possible such history h^{n-1} an available action: if $i = P(n)$, then $\sigma_i(h^{n-1})$ specifies a proposal $x \in X$, and if $i = R(n)$, then it specifies for every possible proposal whether she accepts or rejects it; without loss of generality, I identify this response rule $\sigma_{R(n)}(h^{n-1})$ with the set of accepted proposals $Y \in \mathcal{P}(X)$. If i 's response rule Y has $x \in Y \Leftrightarrow x_i \geq q$, I say that i *accepts with threshold* q . A strategy σ_i is *stationary* if it specifies the same proposal x and response rule Y , irrespective of history. Finally, a *strategy profile* σ is a pair of strategies $(\sigma_{P(1)}, \sigma_{R(1)})$, and its prescribed *play* after history h^{n-1} is $\sigma(h^{n-1}) \equiv (\sigma_{P(n)}(h^{n-1}), \sigma_{R(n)}(h^{n-1}))$.

2.2 Outcomes and (Time) Preferences

If the players agree on division x with a delay of t periods, I call the outcome (x, t) , and if they perpetually fail to agree, I call it $((0, 0), \infty)$. Thus defined in terms of relative time (delay), the set of possible outcomes is the same after any history. A player i 's preferences are formulated over the set $A_i \equiv [0, 1] \times T$, for $T \equiv \mathbb{N}_0 \cup \{\infty\}$, of i 's *personal outcomes* that are her own share and the delay of agreement.

Assumption 1. *In any round n , a player i 's preferences over personal outcomes are represented by the same utility function $U_i : A_i \rightarrow \mathbb{R}$, satisfying the following properties:*

1. *Continuity:* $\{a \in A_i \mid U_i(a) \geq k\}$ and $\{a \in A_i \mid U_i(a) \leq k\}$ are closed for all $k \in \mathbb{R}$,⁹
2. *Desirability:* $q < q'$ implies $U_i(q, t) < U_i(q', t)$ for all t ;¹⁰

⁹Closedness refers to the product topology on A_i , where $[0, 1]$ and T are endowed with the relative standard and discrete topologies, respectively.

¹⁰Absent separability, desirability cannot be formulated independent of the time dimension; specifically, (2.) rules out that a player be entirely indifferent regarding her share once delay gets “too long”. A slight generalization can accommodate such preferences as well, however, without requiring a single change in the results or proofs presented: replace property (2.) with “for any $t \in T$, either U_i is constant on $[0, 1] \times \{t' \in T \mid t' \geq t\}$ or $q < q'$ implies $U_i(q, t) < U_i(q', t)$.”

3. Impatience:

- (a) $t > t'$ implies $U_i(q, t) \leq U_i(q, t')$ for all q ,
- (b) $q > 0$ implies $U_i(q, 0) > U_i(q, 1)$, and
- (c) $\lim_{t \rightarrow \infty} U_i(1, t) \leq U_i(0, 0)$ or there exists a finite \hat{t} such that $U_i(q, t) = U_i(q, \hat{t})$ for all q and all $t \geq \hat{t}$.

Continuity (1.) is a standard technical assumption, and desirability (2.) defines the conflict of interest in the bargaining problem. Property (3.) corresponds to a general notion of impatience regarding agreement: for any given division of the surplus, players do not prefer later over sooner agreement (3.a), if a division yields them a positive share they prefer immediate agreement over delayed agreement (3.b), and whenever they do not become “overwhelmingly” impatient for delay approaching infinity (the standard case guaranteeing “continuity at infinity”), they must be impatient only regarding a finite horizon (3.c). In what follows, by “impatience” I refer only to the two properties (3.ab). The role of property (3.c) is technical: together with continuity, it guarantees existence of a “worst” equilibrium, and I point out explicitly where it is used.

Assumption 1 covers all models of time preferences with impatience put forward in the literature (see [Manzini and Mariotti, 2009](#)).¹¹ It generalizes the most widely studied class of separable time preferences (i.e., discounted utility) axiomatized by [Fishburn and Rubinstein \(1982, thm. 1\)](#), where $U_i(q, t) = d(t) \cdot u(q)$ with $d(\cdot)$ a decreasing “discounting” function, to also cover various non-separable time preferences such as those proposed by [Benhabib, Bisin, and Schotter \(2010\)](#) or [Noor \(2011\)](#).¹² An *instantaneous utility function* can nonetheless be defined by $u_i(q) \equiv U_i(q, 0)$, and it is continuous and increasing.

[Halevy \(2015, prop. 4\)](#) shows that a player’s preferences satisfying assumption 1 are dynamically consistent if and only if they satisfy the stationarity axiom. The latter requires that the preference over two outcomes (q, t) and (q', t') depend only on their relative delay: $U_i(q, t) \geq U_i(q', t')$ if and only if $U_i(q, t + \tau) \geq U_i(q', t' + \tau)$ for any $\tau \in T$; this would here

¹¹The focus of this paper is on time preferences in the usual broad sense of preferences over delayed rewards, which have been extensively researched empirically. However, assumption 1 can also (alternatively or additionally) accommodate costs that are proper to the bargaining activity; e.g., with $U_i(q, t) = q - c(t)$ for $c(\cdot)$ increasing, party i would rather quit bargaining altogether if she expected it to take some time but eventually result only in a very small payoff (e.g., consider $q = 0$).

¹²[Ok and Masatlioglu \(2007\)](#) propose a theory of *relative* discounting that relaxes transitivity for comparisons across three different delays, thus capturing also sub-additive discounting ([Read, 2001](#)) and similarity-based choice ([Rubinstein, 2003](#)). Within the simplified structure of equilibria established below, these failures of transitivity play no role, however. Hence, the characterization of equilibrium *outcomes* also covers these “preferences” (formally, in their notation, let $d(t) \equiv \eta(0, t)$).

yield ED, where $U_i(q, t) = \delta^t \cdot u(q)$ (Fishburn and Rubinstein, 1982, thm. 2). With the exception of ED, all time preferences studied here are therefore dynamically inconsistent.

2.3 Equilibrium Concept

I abstract from informational frictions by assuming that the players' preferences are common knowledge. In the terminology coined by O'Donoghue and Rabin (1999), players are then fully “sophisticated” about their own as well as their opponent's dynamic inconsistency. The equilibrium concept has to incorporate how intertemporal conflict within a player's own preferences is resolved. In single-person decision problems, the standard solution concept for such sophisticated decision makers is that of Strotz-Pollak equilibrium (Strotz, 1956; Pollak, 1968), also known as multiple-selves equilibrium (Piccione and Rubinstein, 1997); it is the subgame-perfect Nash equilibrium (SPNE) of an auxiliary game in which the decision-maker at any point in time is a distinct non-cooperative player. Technically, one then looks for strategy profiles that are robust to one-stage deviations, and this formalizes the presumption that a decision-maker cannot internally commit to future behavior.

The equilibrium notion employed here is the natural extension of this concept to strategic interaction by multiple decision-makers (cf. Chade et al., 2008). To facilitate its definition, let $z_i^{h^{n-1}}(x, Y|\sigma)$ denote the personal outcome of player i , as of round n , that obtains if, following history h^{n-1} , $P(n)$ proposes x , $R(n)$ uses response rule Y , and in case there is no agreement, i.e. $x \notin Y$, both players subsequently adhere to strategy profile σ ; e.g., if $\sigma_{P(n+1)}(h^{n-1}, x) = x' \in \sigma_{R(n+1)}(h^{n-1}, x)$, then $z_i^{h^{n-1}}(x, Y|\sigma)$ equals $(x_i, 0)$ whenever $x \in Y$, and $(x'_i, 1)$ otherwise; accordingly, $z_i^{h^{n-1}, x}(\sigma(h^{n-1}, x)|\sigma) = (x'_i, 0)$.

Definition 1. A strategy profile σ is a **multiple-selves equilibrium (“equilibrium”)** if, for any round n , history h^{n-1} , division x and response rule Y ,

$$\begin{aligned} U_{P(n)}\left(z_{P(n)}^{h^{n-1}}\left(\sigma\left(h^{n-1}\right)\middle|\sigma\right)\right) &\geq U_{P(n)}\left(z_{P(n)}^{h^{n-1}}\left(x, \sigma_{R(n)}\left(h^{n-1}\right)\middle|\sigma\right)\right); \\ U_{R(n)}\left(z_{R(n)}^{h^{n-1}}\left(x, \sigma_{R(n)}\left(h^{n-1}\right)\middle|\sigma\right)\right) &\geq U_{R(n)}\left(z_{R(n)}^{h^{n-1}}\left(x, Y\middle|\sigma\right)\right). \end{aligned}$$

Observe that this indeed defines the SPNE of the auxiliary game where the set of players is taken to be $I \times \mathbb{N}$. The well-known one-stage deviation principle (e.g., Fudenberg and Tirole, 1991, thm. 4.2) says that it coincides with SPNE of the actual game played by I whenever both players' preferences satisfy ED; hence this paper's model contains that of Rubinstein (1982) as a special case.¹³

¹³As in Rubinstein (1982), I consider only pure strategies—a common restriction in this literature, even in models with inherent risk (e.g., Binmore, Rubinstein, and Wolinsky, 1986; Merlo and Wilson, 1995).

2.4 Preliminaries

A central property for the analysis of this game is its stationarity: conditional on failure to agree, the game repeats itself every two rounds. Hence, ignoring history, all subgames beginning with the very same player i 's proposal are identical and, in particular, have the same equilibria; denote this game by G_i . The above defines G_1 ; the sole modification of specifying player 2 as the initial proposer, $P(1) = 2$, defines game G_2 . To distinguish absolute and relative time, throughout, I use n for rounds of a given bargaining game (absolute time) and t for delays to a given agreement (relative time).

Let then $A_i^* \subseteq A_i$ be the set of player i 's personal outcomes that are equilibrium outcomes in G_i . The equilibrium characterization will center on a player i 's *minimal proposer value* v_i^* and *minimal rejection value* w_i^* , as well as the *supremal delay* t_i^* in G_i , given by:

$$\begin{aligned} v_i^* &\equiv \min \{U_i(q, t) \mid (q, t) \in A_i^*\} \\ w_i^* &\equiv \min \{U_i(q, t + 1) \mid (q, t) \in A_i^*\} \\ t_i^* &\equiv \sup \{t \in T \mid \exists q \in [0, 1], (q, t) \in A_i^*\}. \end{aligned}$$

3 The Case of Discounted Shares

The most widely used version of the [Rubinstein \(1982\)](#) model has the bargainers maximize their exponentially discounted surplus share. To make the key results of this paper quickly accessible, this section illustrates them for the generalization of this case only in terms of *discounting*; in fact, under the following common strengthening of assumption [1](#) it summarizes all information necessary to apply the results in either theoretical or empirical work.

Assumption 2. *In any round n , a player i 's preferences over personal outcomes are represented by the same utility function $U_i : A_i \rightarrow \mathbb{R}$ such that*

$$U_i(q, t) = \left(\prod_{s=1}^t \delta_i(s) \right) \cdot q,$$

where (i) $0 < \delta_i(s) < 1$ for any positive s , and (ii) $\lim_{t \rightarrow \infty} \prod_{s=1}^t \delta_i(s) = 0$.¹⁴

A player i 's total discount factor for a delay of t periods, denoted $d_i(t)$, is the product $\prod_{s=1}^t \delta_i(s)$ of the intermittent per-period discount factors. Indifference between two outcomes

Permitting randomization devices, while unlikely to enlarge the set of equilibrium outcomes (cf. [Binmore, 1987](#)), would come at the cost of augmenting the domain of preferences by risk, however, adding a layer of cardinality.

¹⁴I follow the convention that the empty product for $t = 0$ equals one.

$(q, t - 1)$ and (q', t) means that $q = \delta_i(t) \cdot q'$, so unless discounting is constant— $\delta_i(\cdot) = \delta_i$, i.e. ED—it is dynamically inconsistent.

The burden of deciding about delay in bargaining is ultimately on the player responding to an offer. Regardless of its exact form, the respondent’s impatience bestows a strategic advantage upon the proposing player, guaranteeing the latter a minimal rent. In particular, perpetual disagreement can therefore not be an equilibrium outcome.

Given the bargainers eventually agree, it is straightforward to characterize *stationary* equilibrium and establish equilibrium existence: starting from agreement on division x when player i makes an offer, two rounds of backwards induction must lead to the same agreement. Under assumption 2 there exists a unique such agreement, hence a unique stationary equilibrium: i always proposes the same division x and accepts with the same threshold y_i —equal to j ’s offer—such that

$$x_i = 1 - \delta_j(1) \cdot (1 - y_i) \text{ and } y_i = \delta_i(1) \cdot x_i. \tag{1}$$

Stationary equilibrium assumes that bargainers are unresponsive to their opponent’s past behavior (as well as their own). Under this restriction, each party’s decision problem boils down to a two-period consideration, hence dynamic inconsistency cannot unfold and there is immediate agreement after any history as under ED. Consider then the following example, in which parties may condition their bargaining on history.

Example 1. Od (player 1) and Eve (player 2) bargain over how to “split a dollar”. Their preferences satisfy assumption 2, where it is only specified that both discount a first period of delay with common factor $\delta_i(1) = \delta$, and that Od discounts a second period of delay with factor $\delta_1(2) = \gamma\delta$ for $\gamma < 1$. (It is instructive to think first of $\delta \approx 1$ and $\gamma \approx 0$.) Since $\delta_1(2) < \delta_1(1)$, Od is dynamically inconsistent with a “near-future bias”: e.g., facing the prospect of agreement on x in two periods, he would prefer agreeing instead next period for any share q with $\gamma\delta x_1 < q$, but in this next period reverse his preference if also $q < \delta x_1$. (ED would require $\gamma = 1$, hence $\delta_1(2) = \delta_1(1)$.)

Figure 1 describes equilibrium strategies for (once) delayed agreement on a given continuation equilibrium division z (filled green); for concreteness, take $z = \left(\frac{\delta}{1+\delta}, \frac{1}{1+\delta}\right)$, as under continuation according to the unique stationary equilibrium. Delay requires a supporting off-path threat that prevents Od from exploiting his proposer advantage. (If the second-round had agreement on z regardless of first-round play, Od could simply offer Eve her (then unique) rejection value δz_2 —which she had no reason to reject—and thus appropriate the full efficiency gain from immediate rather than delayed agreement.) This threat is alternative second-round agreement y (shaded green), which is more favorable to Eve than z and




Figure 1: Delay equilibrium in example 1 (assuming $\delta y_2 \geq 1 - \delta z_1$). The equilibrium path uses solid lines/borders, and dashed ones indicate supporting off-path behavior.

played in case Od initially offered Eve a share in excess of δz_2 . Hence, Eve initially accepts with threshold δy_2 , and for $1 - \delta y_2 \leq \delta z_1$ initial proposer Od prefers the delayed z over any available immediate agreement; he therefore chooses his initial offer x_2 so low (e.g., zero) that Eve in turn prefers the delayed z over acceptance, $x_2 \leq \delta z_2$.

Of course, threat y such that $1 - \delta y_2 \leq \delta z_1$ (implying $y_2 > z_2$) must be credible. It is Od's near-future bias that lends credibility to it: the strategies in figure 1 specify that *any* failure to agree when Eve makes her offer off-path (shaded green) leads to continuation play identical to that from round 1, with once delayed agreement on z . Od's rejection would therefore always entail *two* periods of delay and have value $\gamma \delta^2 z_1$, enabling proposer Eve to appropriate the full efficiency gain from immediate agreement, with her share equal to $y_2 = 1 - \gamma \delta^2 z_1$. For a sufficiently strong bias of Od (sufficiently low γ), y satisfies the equilibrium condition $1 - \delta y_2 \leq \delta z_1$, and the delayed agreement on z produces its own supporting threat. Such values of γ exist for any z with $z_1 \geq \frac{1-\delta}{\delta}$; as $\delta \rightarrow 1$, this means *any* z (in particular the stationary continuation equilibrium). Moreover, regardless of how small Od's bias is (γ close to one), the strategies then form an equilibrium for sufficiently frequent offers (δ large enough).

Two points are worth emphasizing about this example. (It is readily extended to exhibit also longer delays; see example 3 below.) First, for delay to occur it suffices that the proposer (Od) makes an unacceptably low offer. Though inefficient, he may well do so if he expects any attempt at compromise (Pareto-improvement) to be rejected as well. The intuitive difference between an “unambiguously” low and a compromise offer is, however, that the latter's rejection would allow the respondent (Eve) to credibly adopt an uncompromising

stance. It is this off-path belief that rationalizes the low offer that eschews the respondent’s such opportunity.

Second, near-future bias provides a foundation for this belief, hence delay. In contrast to prior explanations, which depended on multiple *stationary* equilibria (Avery and Zemsky, 1994), this dynamic inconsistency means delay can be “self-enforcing”: any delay at the proposer stage comes with the *threat* of one additional (future) delay at the respondent stage (see Od in round 2 off-path, shaded green); under near-future bias this additional delay can be so costly (γ low enough) as to rationalize an agreement that in turn supports unacceptable offers—hence delay—at the proposer stage. To outweigh the proposer advantage, which ensures a minimal rent to the proposer over her worst threat, the bias needs to be sufficiently strong. As offers become frequent, this rent vanishes, however, and delay equilibria arise for arbitrarily small such biases.

Given the possibility of equilibrium delay, standard recursive arguments fail in characterizing equilibrium. When preferences are dynamically inconsistent, knowledge of a player’s *continuation value* is insufficient to determine her *rejection value*, which is the strategically relevant one. In particular, the relationship $w_i^* = \delta_i(1) \cdot v_i^*$ between i ’s minimal (continuation) value v_i^* as proposer and i ’s minimal (rejection) value w_i^* as respondent generally holds true only when no equilibrium of G_i has delay (cf. Shaked and Sutton, 1984).

To circumvent this problem, I directly investigate the structure of optimal punishments delivering the minimal values v_i^* and w_i^* . The main insight towards characterizing equilibrium is that, given any equilibrium delay t , a proposing player is indifferent between her least preferred immediate equilibrium agreement and her least preferred equilibrium agreement with that delay t ; both yield proposer i her minimal value v_i^* . This indifference property allows to solve for player i ’s minimal values (v_i^*, w_i^*) given the maximal delay t_i^* in game G_i : letting $\Delta_i(t) \equiv \inf \{\delta_i(s) \mid s \in T, 0 < s \leq t\}$ denote player i ’s minimal per-period discount factor over horizon t ,

$$v_i^* = 1 - \delta_j(1) \cdot (1 - w_i^*) \text{ and } w_i^* = \Delta_i(t_i^* + 1) \cdot v_i^*. \quad (2)$$

Proposer i cannot do worse than by making the smallest offer that respondent j would never refuse, j ’s maximal rejection value. This value obtains when j would subsequently receive her maximal share $1 - w_i^*$ with least delay—i.e., immediately following rejection—and equals $\delta_j(1) \cdot (1 - w_i^*)$. For the second equation in (2) suppose an equilibrium of game G_i with delay t . From the indifference property, initial proposer i ’s worst such equilibrium has her share equal to $\frac{1}{d_i(t)} \cdot v_i^*$, and this implies rejection value $\frac{d_i(t+1)}{d_i(t)} \cdot v_i^* \equiv \delta_i(t+1) \cdot v_i^*$ for i as the respondent prior to G_i . This rejection value is minimal whenever $\delta_i(t+1)$ is so, meaning

that the one additional delay—the $t + 1$ -th period—that i 's rejection would entail is most costly (over $t \leq t_i^*$).

Conversely, the maximal delay t_i^* in game G_i is uniquely determined by the minimal proposer values v_i^* and v_j^* , as they capture the players' incentives, as proposer, to make an unacceptable offer rather than settle for the worst immediate agreement:

$$t_i^* = \sup \left\{ t \in T \mid \kappa_i(t, v_i^*, v_j^*) \leq 1 \right\}, \text{ for } \kappa_i(t, v_i, v_j) \equiv \begin{cases} 0 & t = 0 \\ \frac{v_i}{d_i(t)} + \frac{v_j}{d_j(t-1)} & t > 0 \end{cases}. \quad (3)$$

The function $\kappa_i(t, v_i, v_j)$ measures the incentive cost of delay t in game G_i : if initial proposer i could obtain up to value v_i by making an accepted offer rather than incurring delay t , she requires at least the share $\frac{v_i}{d_i(t)}$ with this delay in order not to do so; similarly, player j 's share must be at least $\frac{v_j}{d_j(t-1)}$, since when she gets to propose the first time along the path, the delay would be $t - 1$. As the delay shrinks, these shares become smaller, so the above two incentive constraints are not only necessary but sufficient. They can be satisfied under *some* feasible division if and only if $\kappa_i(t, v_i, v_j) \leq 1$. When the proposer values are minimal, so is the incentive cost, and an equilibrium with delay t exists as long as this minimal incentive cost does not exceed the total available surplus (3).

The values $(v_i^*, w_i^*, t_i^*)_{i \in I}$ are jointly determined by the system of six equations in (2) and (3), to which they are the unique *extreme* solution: if $(v_i, w_i, t_i)_{i \in I}$ is any solution, then $v_i^* \leq v_i$, $w_i^* \leq w_i$ and $t_i^* \geq t_i$ for both i . They fully characterize equilibrium: agreement on division x with delay t is an equilibrium outcome of game G_i if and only if

$$t \leq t_i^* \text{ and } \frac{v_i^*}{d_i(t)} \leq x_i \leq \begin{cases} 1 - w_j^* & t = 0 \\ 1 - \frac{v_j^*}{d_j(t-1)} & t > 0 \end{cases}.$$

The set of divisions that players might agree upon is monotonically shrinking with the delay, where the bounds trace the players' time preferences according to the aforementioned indifference property.

The characterization yields several further insights. First, equilibrium is unique if and only if there is a unique solution to the system of equations. Indeed, the unique stationary equilibrium values in (1), together with $t_1^* = t_2^* = 0$, always form a solution. It is then immediate from (2) that a weak manifestation of present bias, namely $\delta_i(1) \leq \delta_i(s)$ for all $s \geq 1$, is sufficient for uniqueness (then $\Delta_i(\infty) = \delta_i(1)$, and hence $w_i^* = \delta_i(1) \cdot v_i^*$): if both parties find the first period of delay that rejection always entails most costly, then the proposer advantage is only reinforced and delay cannot be self-enforcing. Thus the

uniqueness under ED extends to any form of present bias, in particular any quasi-hyperbolic or hyperbolic discounting.

A future bias of at least one of the bargainers is therefore necessary for equilibrium delay. When this bias concerns the relatively *near future*—relative referring to the players’ overall impatience that drives the incentive cost in (3)—then it is sufficient (e.g., under frequent offers). The resulting equilibrium set has two noteworthy features in this case: i) gradual agreement, and ii) immediate equal division under symmetry.

First, any delayed agreement is reached through gradual agreement, where, as bargaining unfolds, each party’s “concessions” (offers as proposer, and maximum accepted/conceded opponent shares as respondent) increase towards that of the eventual agreement (see section 5.2.1 for a formal definition). The closer in time is the agreement, the smaller is the set of Pareto-improvements, hence ever higher concessions are consistent with delay. For instance, in example 1’s delay equilibrium, Od’s concessions are x_2 and z_2 , and Eve’s are $1 - \delta y_2$ and z_1 ; both sequences are increasing.

Second, if both players’ preferences are symmetric, existence of a delay equilibrium always implies a credible threat such that the minimal proposer value/share is less than one half; $\kappa_i(1, v^*, v^*) \leq 1$ implies $v^* < \frac{1}{2}$. At the same time, due to the proposer advantage, there is then also an equilibrium in which the proposer obtains a value/share greater than one half (e.g., the symmetric stationary equilibrium), hence this threat supports an immediate equal split.¹⁵ In example 1 (which permits symmetry) the equilibrium condition for delayed agreement when z is the stationary equilibrium division implies $1 - \delta y_2 < \frac{1}{2}$, and immediate agreement on an equal division can be supported by only slightly modified threats: if round 2 is reached following an offer of less than *one half*, they agree on y , otherwise on z .

4 Equilibrium for General Time Preferences

The previous section has outlined the fundamental strategic considerations that may emerge in bargaining when the parties discount their shares in a dynamically inconsistent manner. I now turn to the rigorous analysis of the general model, which will clarify how the above intuition is established for general time preferences. As already indicated, I allow for arbitrarily history-dependent strategies to provide a complete account of the strategic considerations that arise under dynamic inconsistency. The otherwise common assumption of stationary strategies would conflict with this objective, because it strongly restricts the parties’ beliefs *a priori*: however systematically player i has deviated from a given stationary

¹⁵More generally, for any $t < t^*$, an equal division with t periods of delay is an equilibrium outcome; also, whenever an equal split is an equilibrium agreement for some delay t , so is an *immediate* equal split.

strategy in the past, it restricts the other to still believing that i will comply with it (see Rubinstein, 1991, p. 912). This point is of special importance here due to the additional presence of *intra*-personal conflict (dynamic inconsistency). First, a player’s beliefs about her own future behavior are as central as those regarding the opponent, as she may have reason to “doubt herself”. Second, the potential of stationary strategies for creating/exploiting dynamic preference reversals is severely limited.

The combination of dynamically inconsistent preferences with the possibility of multiple equilibria and delay (through history-dependent strategies) poses an analytical challenge, however.¹⁶ Standard recursive techniques (see Shaked and Sutton, 1984) fail to be applicable, because a player’s *continuation value* alone provides insufficient information to pin down a unique *rejection value*; yet, this is the strategically relevant value one round earlier, hence required for recursion.

To illustrate, consider a (β, δ) -discounter with linear instantaneous utility, say player 1. Immediate agreement on x and once delayed agreement on y with the same (continuation) value $U_1 = x_1 = \beta\delta y_1$ imply the different rejection values $\beta\delta x_1 = \beta\delta U_1$ and $\beta\delta^2 y_1 = \delta U_1$, respectively. Without further knowledge regarding the underlying equilibrium outcomes, a player i ’s minimal proposer value v_i^* (which is i ’s minimal continuation value when responding) is hence insufficient to determine her minimal rejection value w_i^* .

The approach proposed in this paper directly analyzes the off-path “punishments” (continuation equilibria) that support all equilibrium play and underlie the minimal values (v_i^*, w_i^*) . Its basic idea is that the game’s stationarity property will nonetheless entail a tractable structure for such punishments, since only two types of round need to be distinguished in terms of deviations: any round in which the same party $i \in \{1, 2\}$ gets to make an offer has the same sets of both equilibrium plays and continuation equilibria. If a particular “optimal” assignment of the latter as punishments deters deviations from *any* equilibrium play, it achieves this at any such stage, also off-path, independent of history. How much tractability is thus gained then depends on how “simple” this optimal assignment can be made. In the next section I show what optimality of punishment means, and how four appropriately chosen equilibrium outcomes suffice to describe all off-path play.

The following two reservation shares of a player i (subject to feasibility) will feature prominently in the analysis. (Under the stronger assumption 2 this extra notation could easily be dispensed with.) First, her (immediate) *reservation share* for a given rejection value $U \in U_i(A_i)$ is

$$\pi_i(U) \equiv \min \{q \in [0, 1] \mid u_i(q) \geq U\};$$

¹⁶It is straightforward to show that stationary equilibrium implies immediate agreement after any history (see appendix A.5).

player i then accepts any offer above $\pi_i(U)$ whose rejection would yield value U . Second, her *delayed reservation share* for delay t and immediate value (instantaneous utility) $u \in u_i([0, 1])$ is

$$\phi_i(u, t) \equiv \max \{q \in [0, 1] \mid u \geq U_i(q, t)\};$$

player i then rejects offer q with value $u = u(q)$ for any promised agreement with delay t that has her share greater than $\phi_i(u, t)$.¹⁷

4.1 Optimal Simple Penal Codes and Simple Play

Due to the conceptual similarity, I adopt the terminology introduced by [Abreu \(1988\)](#) for infinitely repeated games.¹⁸ The major difference as well as innovation is that, due to the sequential nature of moves (see below), I base the analysis on sequences of play—for short “plays”—rather than paths; such a play extends paths to include the entire response *rules* used along the path rather than just the on-path responses.¹⁹ I then call an assignment of punishments supporting all equilibrium play (of both G_1 and G_2) an *optimal penal code* (OPC), and I call it an *optimal simple penal code* (OSPC) if punishment is history-independent, with a single punishment per player per role (proposer or respondent) in which this player may deviate.

The sequential nature of moves within a round complicates the analysis relative to repeated games because an OPC cannot simply assign a deviant player’s worst continuation equilibrium. The proposer’s punishment for a deviant offer is constrained by the respondent’s incentives after such a deviation, which affords the proposer a strategic advantage; e.g., a worse continuation equilibrium for the proposer may at the same time weaken the respondent’s current bargaining position and thus make deviant offers more attractive. Indeed, [Mailath et al. \(2015\)](#) present related examples of infinitely repeated *sequential-move* games in which the second mover’s “incentive constraint” forces any OPC to fine-tune punishment to the first mover’s particular deviation, so that no OSPC exists.

Optimal Simple Punishment. The trade-off between providing incentives within-round and under continuation is, however, less complicated here: the respondent’s acceptance ends the game, and the agreement round’s actions determine all payoffs. Punishment therefore takes place only after deviations that result in a rejection, and for a given punishment

¹⁷Since T contains infinity, for completeness, set $\phi_i(u, \infty) = 1$ for any $u \in u_i([0, 1])$.

¹⁸I am deeply grateful to my former colleague Can Çeliktemur for pointing out this similarity to me at an early stage of this project.

¹⁹Against the background of [Abreu’s](#) influential work, I define various concepts of this section only verbally; the full-fledged formalism can be found in appendix [A](#).

the offer that led to it is inconsequential. Call then (i) any deviant rejection of an offer a *respondent deviation*, and (ii) any deviant offer that the respondent may reject without deviating herself a *proposer deviation*. These two types exhaust all (one-stage) deviations that lead to punishment: e.g., given a strategy profile prescribes proposal x and response rule Y , if a proposal $x' \in Y$ is rejected, this constitutes a respondent deviation, and if a proposal $x' \notin Y \setminus \{x\}$ is rejected, this constitutes a proposer deviation. The following result shows that optimality of punishments is a property of their rejection values and optimal punishments can always be made simple. (Existence of an OPC will be established constructively, as part of the equilibrium characterization in theorem 1.)

Lemma 1. *Any OPC's punishments (i) minimize the respondent's rejection value after respondent deviations, and (ii) maximize the respondent's rejection value after proposer deviations. Whenever an OPC exists, there exists an OSPC.*

The first property, regarding a responding player's *deviant rejection*, is straightforward: if rejection of some offer cannot be deterred by her least preferred continuation equilibrium (i.e., one with minimal rejection value) then there cannot be an equilibrium in which she accepts this offer; conversely, if it can be deterred by some continuation equilibrium then *a fortiori* by her least preferred one. Hence any outcome $(x^{R,i}, t^{R,i})$ of a player i 's optimal respondent punishment—an equilibrium outcome of game G_i —satisfies $w_i^* = U_i(x_i^{R,i}, t^{R,i} + 1)$.

The second property is driven by the proposer advantage. A proposer can always deviate to an offer that the respondent will accept and thus evade punishment. In particular, a responding player accepts any offer whose value exceeds her maximal rejection value, in any equilibrium. This guarantees a minimal rent to the proposer, equal to the full efficiency gain from immediate agreement over the respondent's most preferred rejection outcome (which is inefficient due to the delay). Given (ii), any deviant offer that the respondent compliantly rejects would dissipate this rent, as the respondent obtains the same value—her maximal rejection value—but in this case inefficiently. Hence, a proposer can never do better by deviating than by making the lowest accepted offer. However, a play where at some stage the proposing player would gain by deviating to an accepted offer could not be supported by *any* specification of punishments.²⁰

Note the following immediate consequence: letting $(x^{P,i}, t^{P,i})$ be any outcome of player i 's optimal proposer punishment—i.e., an equilibrium outcome of game G_j such that respondent j 's rejection value $U_j(x_j^{P,i}, t^{P,i} + 1)$ is maximal—it must be that i 's minimal proposer value satisfies $v_i^* = u_i(1 - \pi_j(U_j(x_j^{P,i}, t^{P,i} + 1)))$. Not only could proposer i always obtain at

²⁰Recall that we are concerned with one-stage deviations only; hence, whether such a deviation exists can be determined from play alone. Allowing for any punishments, there may also be a deviation to a rejected offer that is even more attractive, but it would be a profitable deviation from prescribed play in any case.

least this value by making an accepted offer, but immediate agreement on the division x with $x_j = \pi_j \left(U_j \left(x_j^{P,i}, t^{P,i} + 1 \right) \right)$ is itself clearly also an equilibrium outcome of game G_i (take $(x^{P,i}, t^{P,i})$ as “unconditional” continuation outcome). Because she may always make an offer that the respondent would never refuse, there cannot be a delay equilibrium that is worse for the proposer than her least preferred immediate-agreement equilibrium.

The first part of lemma 1 shows that it is without loss of generality to restrict OPCs to four optimal punishments, one per player per type of deviation, with the respective properties (i) and (ii); these then support any equilibrium play, of both G_1 and G_2 . Given how it identifies the perpetrator, an OPC is then simple in the sense that punishment need not fit the crime. However, so far this simplicity concerns only first deviations from prescribed play; the punishments themselves may still be rather complex.

The second part of lemma 1 extends the simplicity of an OPC to its own punishments, thus creating an OSPC. It is based on the observation that any OPC supports, in particular, the play of its own constituent punishments. Intuitively, we can therefore iteratively apply the same optimal punishments also to deviations from first punishment play (second deviations), and then also to deviations from second punishment play (third deviations) etc. Thus we create an OPC in which player i 's proposer and respondent deviations are followed by the same respective punishment, entirely independent of their history, i.e. an OSPC; e.g., a proposer deviation by player 1 from its own punishment's play then simply “restarts” this very punishment play. It is therefore without loss of generality to restrict OPCs to OSPCs, and these are fully described by four optimal punishment *plays*.

Simple Play. Consequentialist parties care only about outcomes of play, not play itself; making an offer that is commonly known to be rejected is therefore tantamount to not offering anything at all. The final simplification result removes such redundancy regarding equivalent types of equilibrium play (in particular, optimal punishment play).

Call a play that ends in agreement on division x in round n (perpetual disagreement means $x = (0, 0)$ and $n = \infty$) a *simple play* if (i) all rejected offers are *minimal offers* (i.e., zero offers), and (ii) all response rules specify *maximal acceptance thresholds*, equal to the respective respondent's reservation share for her maximal rejection value in a disagreement round $m < n$, and to $x_{R(n)}$ in the terminal agreement round n . Note that, given the players' maximal rejection values, simple play is fully determined by its ultimate outcome, here (x, t) for $t = n - 1$.²¹ For the purpose of characterizing equilibrium outcomes, with optimal punishments, this is indeed without loss of generality.

²¹As defined here, simple play exists for every equilibrium outcome, but not necessarily for every *possible* outcome; e.g., if player 2's maximal rejection value implies a zero reservation share, then there is no simple play of G_1 with delayed agreement.

Lemma 2. *Whenever an OPC exists and (x, t) is an equilibrium outcome of game G_i , the simple play of this outcome is an equilibrium play of G_i .*

In conclusion, all strategic complexity off the equilibrium path can be summarized by merely four optimal punishment *outcomes* $\left(\left(x^{P,i}, t^{P,i}\right), \left(x^{R,i}, t^{R,i}\right)\right)_{i \in I}$; these define four simple plays that form an OSPC supporting all equilibrium play, of both (sub-) games G_1 and G_2 . Moreover, to check whether an outcome is an equilibrium outcome it suffices to check only for one-stage deviations from its simple play, which is straightforward. These insights afford a greatly simplified structure for equilibrium analysis.

4.2 Equilibrium Characterization

The equilibrium characterization exploits a “fixed-point property” of any quadruple of optimal punishment outcomes: by means of their implied OSPC they support themselves as the most extreme outcomes—in terms of their rejection values (lemma 1)—among all the outcomes that they support. Since there may be multiple OSPCs, I first map this fixed-point property into a system of equations that the unique associated punishment values $(v_i^*, w_i^*, t_i^*)_{i \in I}$ necessarily solve. These equations, in general, have multiple solutions, and the values $(v_i^*, w_i^*, t_i^*)_{i \in I}$ are found as their unique *extreme* solution, whose existence follows from the continuity assumptions on preferences. These values then characterize the set of OSPCs, and thus also the set of equilibrium outcomes. This is the central result of this paper.

Define first the function $\kappa_i : T \times U_i(A_i) \times U_j(A_j) \rightarrow \mathbb{R}_+$ such that

$$\kappa_i(t, v_i, v_j) \equiv \begin{cases} 0 & t = 0 \\ \phi_i(v_i, t) + \max\{\phi_j(v_j, t-1), \phi_j(u_j(0), t)\} & t > 0 \end{cases},$$

which measures the *surplus-cost of delay* t in G_i given proposer values v_i and v_j , and which is non-decreasing in each of its arguments. Its significance derives from the fact that, given the minimal proposer values v_i^* and v_j^* from optimal punishment, game G_i has an equilibrium outcome with (positive) delay t if and only if $\kappa_i(t, v_i^*, v_j^*) \leq 1$. The restriction to simple play allows to reduce the necessary and sufficient incentive constraints for agreement on x with this delay to $x_i \geq \phi_i(v_i^*, t)$ and $1 - x_i \equiv x_j \geq \max\{\phi_j(v_j^*, t-1), \phi_j(u_j(0), t)\}$; κ_i therefore measures the incentive cost of delay t as the minimal amount of surplus, so that both players can be promised a large enough share with this delay.

Let then $E \subseteq \prod_{i \in I} (u_i([0, 1]) \times U_i(A_i) \times T)$ be the set of sextuples $(v_i, w_i, t_i)_{i \in I}$ such

that, for each $i \in I$,

$$v_i = u_i(1 - \pi_j(U_j(1 - \pi_i(w_i), 1))) \quad (4)$$

$$w_i = \inf \{U_i(\phi_i(v_i, t), t + 1) \mid t \in T, t \leq t_i\} \quad (5)$$

$$t_i = \sup \{t \in T \mid \kappa_i(t, v_i, v_j) \leq 1\} \quad (6)$$

Lemma 6 in appendix A.3 shows how each element $(v_i, w_i, t_i)_{i \in I}$ of E corresponds to a quadruple of punishment outcomes that are “constrained” optimal in the following sense: by means of a construction similar to an OSPC, they support a *subset* of equilibrium outcomes that includes themselves (so they are indeed equilibrium outcomes), and on which they are optimal; i.e., constrained to this subset, they yield the minimal punishment values $(v_i, w_i)_{i \in I}$ and supremal delays $(t_i)_{i \in I}$.

If optimal punishments, and thus an OSPC, exist, the associated values $(v_i^*, w_i^*, t_i^*)_{i \in I}$ are necessarily in E . However, in general, due to the interdependency of punishments—harsher punishments permit longer delays, and longer delays permit harsher punishments—there may be (other) constrained OSPCs. In fact, the set E always contains an element $(v_i, w_i, t_i)_{i \in I}$ with $t_1 = t_2 = 0$ that corresponds to a “trivial” constrained OSPC: irrespective of who deviated in a given round, it specifies the same punishment; thus this OSPC reduces to a single stationary equilibrium in which player i always offers $1 - \phi_i(v_i, 0) = \pi_j(w_j)$ and always accepts with threshold $\pi_i(w_i) = \pi_i(U_i(\phi_i(v_i, 0), 1))$, so there is immediate agreement after any history.

In view of potential multiplicity in E , the actual values $(v_i^*, w_i^*, t_i^*)_{i \in I}$ must then be its unique *extreme element*; i.e., any other element $(v_i, w_i, t_i)_{i \in I}$ satisfies $v_i^* \leq v_i$, $w_i^* \leq w_i$ and $t_i^* \geq t_i$ for both i .

Theorem 1. *The values $(v_i^*, w_i^*, t_i^*)_{i \in I}$ exist, and they are equal to the unique extreme element of the set E . For each $i \in I$, $(x^{P,i}, t^{P,i})$ and $(x^{R,i}, t^{R,i})$ are outcomes of player i ’s optimal proposer and respondent punishment, respectively, if and only if*

$$\left\{ \begin{array}{l} t^{P,i} = 0 \\ x_i^{P,i} = \pi_i(w_i^*) \end{array} \right\} \quad \text{and} \quad \left\{ \begin{array}{l} t^{R,i} \in \arg \min \{U_i(\phi_i(v_i^*, t), t + 1) \mid t \in T, t \leq t_i^*\} \\ x_i^{R,i} = \phi_i(v_i^*, t^{R,i}) \end{array} \right\},$$

and the set A_i^* of player i ’s personal equilibrium outcomes in game G_i equals

$$\left\{ (q, t) \in A_i \mid \phi_i(v_i^*, t) \leq q \leq \begin{cases} 1 - \pi_j(w_j^*) & t = 0 \\ 1 - \max \{ \phi_j(v_j^*, t - 1), \phi_j(u_j(0), t) \} & t > 0 \end{cases} \right\}.$$

A few features of optimal punishments are noteworthy in view of the strategic advantage

enjoyed by a proposing player. First, a player’s optimal proposer punishment is unique and involves no delay: given her impatience, the respondent’s rejection value is maximized by the maximal credible share with least delay following rejection (4). Second, an initially proposing player i ’s least preferred equilibrium outcomes for various delays are necessarily indifferent, all yielding her the same minimal value v_i^* , and this allows to pin down optimal respondent punishment (5). Finally, whether and how long agreement may be delayed is fully determined by the players’ incentives as proposer (6); this drives the aforementioned indifference property (see also the characterization of A_i^* in theorem 1).

Example 1 shows that the equilibrium characterization neither reduces to uniqueness nor to stationarity of equilibrium, nor to stationarity of optimal punishments. It is never an “anything goes”-type result, however, as the players’ impatience imposes a certain structure on equilibrium through the proposer advantage: as a function of delay, the set of equilibrium divisions monotonically shrinks (since $\phi_i(u, \cdot)$ is increasing, the upper and lower bounds on each player’s share converge), and perpetual disagreement is never an equilibrium outcome (note that $v_i^* > u_i(0) \geq U_i(0, \infty)$). In section 5, I present further detail, examples and discussion regarding the structure of equilibria for various preferences.

Theorem 1 is partly reminiscent of Merlo and Wilson (1995, thms 7 and 8), who assume ED and analyze bargaining by multiple players under a Markovian process governing the protocol as well as the size of the cake. They also characterize the set of equilibrium values by means of an extremal fixed point, but its nature differs significantly. ED implies that there is a stationary equilibrium outcome that maximizes one player’s value at the same time as it minimizes all other players’ values. In the two-player case this simple relationship between punishment and reward implies that optimal punishments are efficient and, without loss of generality, also stationary. Only in the case of more than two players, one player’s optimal punishment might necessitate some punishment of another player and some inefficiency, thus complicating the incentive structure (cf. Burgos et al., 2002b).

By contrast, here such a complication arises already with two players, and from a very different source: the dynamic inconsistency of a player’s time preferences. Optimal punishment might *require* delay, in which case it is both inefficient and non-stationary. The extreme equilibria are then “truly” non-stationary in the sense that their continuation is non-stationary after any history. Equilibrium delay does not necessitate multiple stationary equilibria; indeed, it does not even depend on the existence of a stationary equilibrium.

This distinguishes the delay obtained here from that obtained in other extensions of the original Rubinstein (1982) model that maintain a stationary game structure and ED, all of which rely on multiple stationary equilibria to support delay (Haller and Holden, 1990; Muthoo, 1990; van Damme, Selten, and Winter, 1990; Fernandez and Glazer, 1991; Myerson,

1991; Avery and Zemsky, 1994). The sole exception I am aware of is that of Busch and Wen (1995).²² Their model of negotiation enriches bargaining by a disagreement game, which is a fixed simultaneous-move game played after any rejected offer and determines a stream of payoffs before agreement. The truly non-stationary equilibria they construct exploit the resultingly richer preference domain through non-stationary play of the disagreement game similar to folk theorems for repeated games, but constrained by the parties' incentives to reach agreement.

Existence of an OSPC is equivalent to the existence of minimum values v_i^* and w_i^* (as argued, a “constrained” OSPC and hence an equilibrium always exist, however). This is non-trivial here, as the set of equilibrium outcomes need not be closed.²³ The generality of assumption 1 means that the length of equilibrium delay might have no upper bound, despite the fact that perpetual disagreement is never an equilibrium outcome due to the proposer advantage (see appendix B.2 for an example). While existence of a minimal value v_i^* follows from standard continuity even with unbounded delay, the (only) role played by impatience property (3.c) is to ensure that the minimal value w_i^* also exists in this case, because the delay of agreement that is required for *optimal punishment* is then bounded.

5 Uniqueness v. Multiplicity, and Delay

For economic applications, where bargaining arises naturally in various contexts (household decision-making, wage setting, international trade agreements etc.), uniqueness of the bargaining prediction is an important concern. Any uncertainty about this one aspect of a model feeds through all of the conclusions drawn from it. The following characterization of those preference profiles (within the general class defined by assumption 1) for which equilibrium is indeed unique is immediate from theorem 1.

Corollary 1. *Equilibrium is unique if and only if the set E is a singleton. Whenever unique, equilibrium is stationary and has immediate agreement after any history: player i always offers the share $1 - \phi_i(v_i^*, 0) = \pi_j(w_j^*)$ and always accepts with the threshold $\pi_i(w_i^*) = \pi_i(U_i(\phi_i(v_i^*, 0), 1))$, $i \in I$.*

These necessary and sufficient conditions for uniqueness do not isolate preference properties at the individual level: fixing one party's preferences, whether equilibrium is unique or displays multiplicity generally depends on those of the opponent. For the purposes of applied

²²I am indebted to Paola Manzini for drawing my attention to these authors' work.

²³Although the equilibrium concept introduced in definition 1 is equivalent to a version of subgame-perfect Nash equilibrium, existing results based on the upper hemi-continuity of its equilibrium correspondence (e.g., Börgers, 1991) cannot be applied here, because they assume finitely many players.

work, this is hardly useful. Below, I therefore investigate what broad qualitative properties of preferences at the individual level imply uniqueness on the one hand, and multiplicity and delay on the other. For the latter case I also highlight general properties of the equilibrium set.

5.1 Uniqueness

Already *stationary* equilibrium need not always be unique, and this is so even under ED (see Rubinstein, 1982). However, the set of stationary equilibria is fully determined by the curvature properties of the parties' preferences regarding their *surplus share*, which are essentially orthogonal to their dynamic (in-)consistency.²⁴ Indeed, the same axioms that have been postulated in order to guarantee uniqueness of stationary equilibrium under ED (e.g., Binmore et al., 1986; Hoel, 1986; Osborne and Rubinstein, 1990) also do so within the much more general class of preferences analyzed here. For instance, consider the following property.²⁵

Definition 2. Player i 's preferences exhibit **immediacy** if, for any two shares q and q' , and any positive ϵ ,

$$u_i(q) = U_i(q', 1) \Rightarrow u_i(q + \epsilon) > U_i(q' + \epsilon, 1).$$

Starting from indifference between an immediate and a once delayed agreement, immediacy says that an increase in one's surplus share is more valuable when immediate. With impatience, indifference requires that the delayed share exceed the immediate one, so immediacy extends a basic property of any discounted concave utility to non-separable preferences. Because it is concerned with comparisons of only immediate and once delayed agreements, it does not restrict whether or how preferences are dynamically inconsistent.

Lemma 3. *If both players' preferences exhibit immediacy, stationary equilibrium is unique.*

Immediacy ensures that the proposer's surplus rent in immediate rather than once delayed (history-independent) agreement is monotonically increasing in the share that the respondent would obtain by rejecting; e.g., if any offer's rejection would subsequently result in immediate agreement on division x , then the proposing player i 's such surplus rent equals $(1 - \pi_j(U_j(x_j, 1))) - (1 - x_j) = x_j - \pi_j(U_j(x_j, 1))$. Its increasingness implies that the backwards-induction dynamics are well behaved: starting from any (history-independent) agreement, backwards induction produces a unique limit, i.e. a unique stationary point.²⁶

²⁴Appendix A.5 provides a full characterization of stationary equilibrium, for the general case.

²⁵This is essentially a reformulation in utility terms of the "increasing loss to delay" axiom of Osborne and Rubinstein, 1990, pp. 35-36.

²⁶If the rent were non-monotonic, the limit may depend on the starting division, yielding multiple stationary points.

A unique stationary equilibrium is the only equilibrium with immediate agreement after any history. This equilibrium is unique overall whenever delay is not self-enforcing in the sense that it enlarges the scope for punishment so much that it effectively supports itself. Consider then the following preference property.

Definition 3. Player i 's preferences exhibit a **weak present bias** if, for any two shares q and q' , and any delay t ,

$$u_i(q) = U_i(q', t) \Rightarrow U_i(q, 1) \leq U_i(q', t + 1). \quad (7)$$

Present bias means that a party becomes more patient when an immediate and an indifferent delayed reward are pushed into the future. Hence, if a present-biased individual, in a period's time, would be indifferent between receiving a reward q immediately and receiving a reward q' with t periods of delay, she currently prefers the larger later reward.

Recall now that, due to the proposer advantage, delay cannot hurt a proposing party beyond her least preferred immediate agreement. Under weak present bias, delay cannot hurt this party as the respondent either: rejection necessarily entails a minimal delay of one period, but beyond this "critical" period she is more patient. Hence, subject to indifference as the proposer, she cannot be made worse off as the respondent; delay cannot be self-enforcing.

Proposition 1. *If, in addition to immediacy, both players' preferences exhibit a weak present bias, then equilibrium is unique.*

Together with immediacy, weak present bias provides a simple set of sufficient conditions for uniqueness. Both properties are readily checked for any given preferences, and both are readily testable empirically.

The interpretation of property (7) as weak present bias is most straightforward for discounted utility, where $U(q, t) = d(t) \cdot u(q)$. Letting $d(t) \equiv \prod_{s=1}^t \delta(s)$, weak present bias then reduces to $\delta(1) \leq \delta(t)$, saying that no future period of delay is discounted more heavily than the first one from the immediate present.²⁷ Any hyperbolic or quasi-hyperbolic discounting exhibits this property, with an actual bias: the (β, δ) -model of quasi-hyperbolic discounting has $\delta(1) = \beta\delta < \delta = \delta(t)$ for any $t > 1$, and hyperbolic discounting has $\delta(\cdot)$ increasing.²⁸

Proposition 1 establishes the robustness of the bargaining wisdom received from the study of ED to various forms of present bias: equilibrium is unique as well as efficient, it is easily

²⁷Halevy (2008) introduces a strict version of this discounting property, which he calls "diminishing impatience", and relates it to non-linear probability weighting of consumption risk. The weak formulation of property (7) means it also covers ED as the limiting case where $\delta(\cdot)$ is constant.

²⁸The non-separable models of Benhabib et al. (2010) and Noor (2011) were both designed to capture the very same pattern of preference reversals that hyperbolic and quasi-hyperbolic discounting explain, and it can easily be verified that they, too, exhibit a weak present bias.

computed on the basis of only the players’ attitudes to a single (the first) period of delay and has familiar comparative statics. If one believes in the essence of present bias but finds the evidence inconclusive as to what exact functional form it assumes, it is comforting to learn that equilibrium is robust to any mis-specification of higher-order delay attitudes. Moreover, the finding that the historically main mode of surplus sharing is efficient under present bias is good news for its evolutionary explanations (e.g., [Dasgupta and Maskin, 2005](#); [Netzer, 2009](#)): otherwise, communities without a present bias would have had an evolutionary advantage, making its survival hard to understand.

Most importantly, [proposition 1](#) expands the scope of applied work, which shows strong interest in the study of present-biased time preferences—in particular (β, δ) -discounting—but has hitherto lacked a strategically founded bargaining solution. Its application requires some caution, however, as the following example indicates.

Example 2. Let the two parties’ preferences be given by $U_i(q, t) = d_i(t) \cdot q$ with $d_i(0) = 1 > d_i(t) = \beta_i \delta_i^t$ for all $t > 0$, $(\beta_i, \delta_i) \in (0, 1)^2$. The unique equilibrium of the game in which player 1 makes the initial offer has immediate agreement on division x such that

$$x_1 = \frac{1 - \beta_2 \delta_2}{1 - \beta_1 \delta_1 \beta_2 \delta_2}.$$

For a given positive period-length, this prediction is indistinguishable from that under ED where each player i has preferences $U_i(q, t) = \tilde{\delta}_i^t q$ with $\tilde{\delta}_i \equiv \beta_i \delta_i$ (cf. [Bernheim and Rangel, 2009](#), pp. 69-71).

Whichever continuous-time version of (β, δ) -discounting is adopted (cf. [Harris and Laibson, 2013](#); [Pan et al., 2015](#)), the limiting case of very frequent offers that is commonly focused on in applications becomes problematic. Either a player’s bias is taken to discontinuously differentiate instantaneous from delayed gratification (let $t \in \mathbb{R}_+$ above), in which case $x_1 \rightarrow \frac{1 - \beta_2}{1 - \beta_1 \beta_2}$ as $\delta_i \rightarrow 1$ (regardless of relative speeds of convergence), and the bargaining outcome is fully determined by the players’ very short-run impatience; the initial proposer’s advantage then prevails for arbitrarily frequent offers, and—failing to generate an equal split—the model is at odds with the Nash bargaining solution.²⁹

Or an extended notion of the “present” of length $\tau_i > 0$ is adopted, such as $d_i(t)$ equal to δ_i^t whenever $t \leq \tau_i$ and $\beta_i \delta_i^t$ otherwise. Then, however, as the length of a bargaining period falls below some player’s τ_i , the model exhibits multiple equilibria and delay, of the type presented in [example 1](#) (there $1 \leq \tau_1 < 2$).

A related conceptual issue arises concerning the possibly distinct times of agreement and

²⁹Notice that any bias $\beta_i < 1$, however small, means that in the limit this player obtains none of the surplus in bargaining against an exponential discounter.

consumption feasibility. If there is an exogenous lag $\hat{\tau}$ between agreement and consumption, exceeding the length of time for which there is a “present bias”, the unique equilibrium has immediate agreement with player 1’s share equal to $x_1 = \frac{1-\delta_2^{\hat{\tau}+1}}{1-\delta_1^{\hat{\tau}+1}\delta_2^{\hat{\tau}+1}}$; only the “long-run” discounting matters, because each player i discounts even immediate *agreements* with extra factor β_i .

Taking a broad perspective on what is being consumed, it could also be a bargainer’s relevant others’ esteem, proportional to the surplus she fetches (e.g., when a union leader negotiates on behalf of her union). The agreement reached might then differ drastically, depending on whether the bargaining is done behind closed doors (there is a lag between agreement and consumption, and only long-run discounting matters) or in the presence of such relevant others (when the timing of agreement and consumption coincide, and the degrees of present bias are the main determinant of the division).³⁰

5.2 Multiplicity and Delay

In view of the sufficient conditions for uniqueness in proposition 1, multiplicity of equilibrium can arise from two conceptually distinct sources: (i) violations of immediacy, and (ii) violations of weak present bias. The former relate to the curvature of utility in the surplus share and entail multiple *stationary* equilibria, which may also support delay when used as history-dependent (non-stationary) punishments. The latter relate to the particular form of dynamic inconsistency and allow delay to be *self-enforcing* (rather than relying on stationary equilibria to support it). This section first highlights a few general structural properties of the equilibrium set whenever there exist delay equilibria, regardless of their source. Then it goes on to separately discuss each of the two potential sources.

5.2.1 Gradual Agreement and Equal Split

Delay can only arise in a non-stationary equilibrium: if there were a unique, history-independent continuation equilibrium, the proposing party could appropriate any efficiency gains from immediate agreement with an accepted offer. Equilibrium disagreement requires a “punishment” for any such attempt, favoring the responding party, to rationalize the following strategic reasoning: although Pareto-improvements are available, the proposing party believes that by offering one she would induce the opponent to expect an even superior (non-Pareto-improving) agreement and, accordingly, reject the proposal. This belief supports an

³⁰I thank Erik Eyster and David Cooper for independently pointing out the following: any (common) lag between time of agreement and time of consumption does not affect the unique bargaining outcome under ED (this can be seen from the functions π_i), but under (β, δ) -discounting would shift bargaining power toward the player who is more patient in the long-run.

offer that is unfavorable vis-à-vis the delayed outcome for the respondent, hence the delay.

Given supporting punishments exist, the eventual agreement determines all restrictions on possible equilibrium play during any disagreement round: (i) the proposer’s (rejected) offer is no better for the respondent than the eventual agreement, and (ii) the respondent rejects all offers that are better for the proposer than the eventual agreement. Observe now that, since the parties are impatient, the value of the eventual agreement increases *across subsequent rounds* along the equilibrium path, as the remaining delay gets shorter. Hence, the set of Pareto-improvements shrinks, and the parties may make ever greater “concessions” that nonetheless result in disagreement; thus they may always agree gradually.

Formally, for any equilibrium play $(x^n, Y^n)_{n=1}^{t+1}$, define party i ’s *concession* in round n , denoted b_i^n , as her offer if i is the proposer, i.e. $b_i^n = x_j^n$ if $i = P(n)$, and as the supremal opponent share that she would accept if i is the respondent, i.e. $b_i^n = \sup \{x_j \in [0, 1] \mid x \in Y^n\}$ if $i = R(n)$. Call an equilibrium with outcome (x, t) a *gradual-agreement equilibrium* if its play has both players’ concessions b_i^n increasing in n , i.e. $b_i^{n+1} > b_i^n$ for both i and all $n \leq t$. Any such equilibrium has the intuitive property that both parties become more and more conciliatory over the course of bargaining as they keep failing to reach agreement. Gradual agreement meaningfully applies only to equilibria with delay, of course; then, however, its requirement is rather strong, as it treats a player’s offers and response rules symmetrically in terms of concessions (it clearly implies increasing offers by each player). Nonetheless, gradual agreement is without loss of generality.

Proposition 2. *If both parties $i \in I$ are uniformly impatient, so that for any positive share q , $t < t'$ implies $U_i(q, t) > U_i(q, t')$, then every equilibrium outcome is the outcome of a gradual-agreement equilibrium.*³¹

Under gradual agreement, a player’s concession has the interpretation of the credible promise that she will subsequently always be willing to give up at least this share, as long as the other player keeps to her promise. The fact that this promise has no material counterpart—rejected offers enter neither payoffs nor preferences directly, only strategically—makes it distinct from the commitment mechanisms in related work explaining such “gradualism” (Admati and Perry, 1991; Compte and Jehiel, 2004).³²

The final result of this section relates equilibrium to the influential axiomatic bargaining solution proposed by Nash (1950), which imposes the intuitive property that symmetric bargaining problems should yield a symmetric, i.e. equal, division. Under (symmetric) ED,

³¹If the requirement for gradual agreement were weakened to *non-decreasing* concessions, this proposition would hold true for any preference profile.

³²In these papers the value of a player’s outside option increases in the opponent’s past concessions. Li (2007) obtains a similar effect with history-dependent preferences.

given immediacy, this is here also the *limiting* outcome of the unique equilibrium as offers become arbitrarily frequent (Binmore et al., 1986, prop. 4). For more general time preferences, if delay can be supported then the following symmetry result obtains.

Proposition 3. *If the two bargaining parties' preferences are symmetric, then an immediate equal split is an equilibrium outcome whenever there exists an equilibrium with delayed agreement. More generally, an equal split with delay $t - 1$ is an equilibrium outcome whenever there exists an equilibrium in which agreement is delayed by t periods.*

In reasonably symmetric bargaining situations, the possibility of delay implies that the parties may instead quickly agree on an equal split. This holds true here without recourse to a limiting argument, hence even for non-negligible costs of disagreement; as offers become more frequent, the required delay equilibria are, however, more likely to exist (see example 3 below).

5.2.2 Non-Immediacy, Multiple Stationary Equilibria, and Delay

Almost any model of time preferences assumes separability in reward and delay; i.e., an atemporal utility function on rewards can be defined that is being discounted for delay. Concavity, or even weaker strict log-concavity, of these utilities then implies immediacy, hence a unique stationary equilibrium, and this is the unique equilibrium overall under weak present bias. Conversely, if at least one party's utility exhibits sufficiently strong convexity, multiple stationary equilibria arise, and these may then also support delay, irrespective of the dynamic (in-)consistency of discounting. Already Rubinstein (1982) presents an example of this possibility under ED, when both parties have symmetric preferences represented by $U(q, t) = \delta^t \cdot \exp(q)$.³³

When time preferences are not separable, however, the curvature of utility from the reward can depend on its delay. Concavity of utility from *immediate* rewards then ceases to be sufficient for a unique stationary equilibrium, and the preference property of immediacy imposes a restriction on how the curvatures for immediate and once-delayed rewards are related: at indifference, marginal utility should be greater immediately.

An interesting model that is well-suited to illustrate this point, and also how immediacy might fail, is that of magnitude-dependent discounting proposed by Noor (2011). It has time

³³This is the best-known example of multiplicity under ED. Originally, it uses representation $q - c \cdot t$, but this equals $\ln(\delta^t \cdot \exp(q))$ for $c = -\ln(\delta)$ and is therefore equivalent. (Recall that there is no uncertainty.) While he does not fully characterize equilibrium outcomes under multiplicity, in particular concerning the possible delays, appendix B.1 shows how theorem 1 applies in a straightforward manner to deliver this characterization.

preferences represented by $U(q, t) = \delta(q)^t \cdot u(q)$, where the discount factor $\delta(\cdot)$ is an increasing function of the reward, and reward-utility $u(\cdot)$ is concave. Thus the model captures empirically observed magnitude effects, where larger rewards are discounted less than smaller ones, and it behaviorally subsumes the (separable) hyperbolic discounting model. Indeed, it is straightforward to show that such preferences satisfy weak present bias.³⁴ Yet, due to the reward-dependence of discounting, immediacy may fail despite the concavity of u : supposing indifference $u(q) = \delta(q') \cdot u(q')$, immediacy would require $u(q + \epsilon) > \delta(q' + \epsilon) \cdot u(q' + \epsilon)$, which can be rewritten as

$$u(q + \epsilon) > \delta(q') \cdot u(q' + \epsilon) + (\delta(q' + \epsilon) - \delta(q')) \cdot u(q' + \epsilon).$$

While concavity of u implies that $u(q + \epsilon) > \delta(q') \cdot u(q' + \epsilon)$, since discount factor $\delta(\cdot)$ increases in the size of the reward, the second term on the right-hand side is positive and may well outweigh the concavity. Hence multiple stationary equilibria may arise, and also delay equilibria. This is true even when both $u(\cdot)$ and $\delta(\cdot) \cdot u(\cdot)$ are concave. Since the basic construction of delay equilibria in this case is familiar from the literature (Avery and Zemsky, 1994), I present a simple numeric example only in the appendix B.1; note here, however, that the results of section 5.2.1 apply.

5.2.3 Near-Future Bias, and Self-Enforcing Delay

Given immediacy, a violation of weak present bias is necessary for the emergence of delay equilibria. When it concerns the relatively near future, it is sufficient; i.e., in contradiction to (7), for some relatively *small* t , an indifference $u_i(q) = U_i(q', t)$ is broken in favor of the sooner agreement once both outcomes lie in the future, $U_i(q, 1) > U_i(q', t + 1)$.³⁵

Under discounting, using decomposition $d_i(t) \equiv \prod_{s=1}^t \delta_i(s)$, a near-future bias means that $\delta_i(s) < \delta_i(1)$ for $s > 1$ not too large; i.e., a near-future period of delay is discounted more heavily than the first one. Whereas under weak present bias the minimal per-period discount factor $\Delta_i(t) \equiv \inf \{\delta_i(s) | s \in T, 0 < s \leq t\}$ is independent of the horizon t and constant at $\Delta_i(\infty) = \delta_i(1)$, under near-future bias it initially decreases as the horizon is extended: $\Delta_i(s) < \Delta_i(1)$ for $s > 1$ not too large. Ebert and Prelec (2007), Bleichrodt et al. (2009), Takeuchi (2011) and Pan et al. (2015) have advanced functional forms for near-future-biased discounting; in graphical terms, all of these discounting functions are initially

³⁴Take any indifference $u(q) = \delta(q')^t \cdot u(q')$ and note that $q' \geq q$ by impatience, which implies that $\delta(q) \cdot u(q) \leq \delta(q') \cdot u(q) = \delta(q')^{t+1} \cdot u(q')$.

³⁵Of course, offers must not take too much time for the “bias horizon” to be relevant; e.g., if a counter-offer would take forever, the first offer is an ultimatum, and there is a unique equilibrium in which the initial proposer obtains the entire surplus without delay.

concave, so their decline is steepest at some positive delay rather than at zero.

For a near-future biased bargainer a further period of delay in the near future is more critical than the first, initial period of delay. To avoid a costly future delay, she has to rely on her future self. However, to her future self the same delay, in absolute time, will not be as critical any more, in relative time. Put succinctly, a given future delay is more painful now than it will be later—she will subsequently become more patient and, accordingly, tougher in bargaining than she would initially want herself to be.³⁶

This type of dynamic inconsistency makes delay self-enforcing, because any delay on path automatically implies the threat of an additional delay off-path, in the event of a rejection: assuming the additional delay would be particularly costly to her, such a bargainer may accept so bad a deal as the respondent now, that—in terms of a threat—this supports her unacceptable offers as the proposer later, when she will be more patient. Although her proposer advantage limits the power of this threat, as offers become frequent and this advantage vanishes, delay equilibria emerge for an arbitrarily small such bias.

The following final example of a near-future bias extends example 1 to illustrate the usefulness of theorem 1 for fully characterizing a rich equilibrium set, to demonstrate its general properties highlighted in section 5.2.1, and to assess the potential costs of delay.

Example 3. Let the two parties’ preferences be symmetrically given by $U_i(q, t) = d(t) \cdot q$ with

$$d(t) = \begin{cases} \delta^t & t \leq \tau \\ \gamma\delta^t & t > \tau \end{cases}, \text{ for } (\delta, \gamma) \in (0, 1)^2 \text{ and } \tau > 0.$$

First, note that the $\tau + 1$ -th period of delay is discounted most heavily: whereas the per-period discount factors are $\delta(t) = \delta$ for all $t \neq \tau + 1$, for that period it is $\delta(\tau + 1) = \gamma\delta$. Since $\tau > 0$, weak present bias is violated, and there is instead a bias toward not experiencing more than τ periods of delay. (Immediacy is clearly satisfied.) Hence $\Delta(t)$ equals δ for all $t \leq \tau$ and $\gamma\delta$ for all $t > \tau$; given Δ determines whether non-stationary delay equilibria emerge, this minimal deviation from ED is made only for convenience, to keep the number of parameters down to a mere three, $\{\delta, \gamma, \tau\}$. Due to preference symmetry, the player subscript is omitted throughout this example.

Suppose there is an equilibrium in which agreement is delayed by τ periods: then $v^* = \frac{1-\delta}{1-\gamma\delta^2}$ and $w^* = \gamma\delta v^*$ (see (2) in section 3); delay $\tau > 0$ is then “self-enforcing” if and only if

³⁶As an extreme but instructive example imagine someone who—at *any point in time*—does not mind bargaining for, say, 5 rounds, but is extremely averse to bargaining any longer; such a shifting personal “deadline” (in *relative* time) is dynamically inconsistent, since as soon as the first round is over this player will already not mind delaying agreement until round 6.




Figure 2: Graphs regarding equilibrium delay in example 3. The panel on the left shows the parametric regions (δ, γ) such that delay equilibria exist for three given values of τ , which are 1 (blue, orange and green), 25 (brown and green) and 1000 (green). The panel on the right plots $\hat{\tau}(\delta, \gamma)$ as a function of δ for three given values of γ , which are 0.5 (blue), 0.75 (orange) and 0.99 (green).

$1 \geq \kappa(\tau, v^*, v^*) = \frac{v^*}{\delta^\tau} + \frac{v^*}{\delta^{\tau-1}}$, which reduces to

$$\delta^\tau \geq (1 + \delta) \cdot \frac{1 - \delta}{1 - \gamma\delta^2} \quad (8)$$

after substituting for v^* . The left-hand side is the present value of the surplus, and the right-hand side is the present value of the incentive cost of a delay of τ periods: each proposer requires $v^* = \frac{1-\delta}{1-\gamma\delta^2}$, and the factor $(1 + \delta)$ is due to the fact that the initial proposer does so immediately whereas the other player does so only next round. Observe that, for any given $\tau > 0$ and $\gamma < 1$, there exist large enough values of δ such that inequality (8) is satisfied (the left-hand side limits to one whereas the right-hand side limits to zero as $\delta \rightarrow 1$); generally, as δ increases, the set of parameters γ and τ for which delay equilibria exist expands, as the left-hand-side panel of figure 2 illustrates. Whenever such delay equilibria exist, the minimal proposer and rejection values are obtained only by means of a “truly” non-stationary delay equilibrium, using optimal punishments.

Notice also that inequality (8) implies $w^* < v^* < \frac{v^*}{\delta^{\tau-1}} \leq \frac{1}{2}$, and an equal split with any delay up to $\tau - 1$ periods is then an equilibrium outcome (in particular under immediate agreement). It may also be reached gradually, say with delay \hat{t} , $0 < \hat{t} < \tau$: define a sequence $(b^n)_{n=1}^{\hat{t}+1}$ of concessions such that $b^1 \equiv 0$ and $b^n \equiv \frac{1}{2} \left(b^{n-1} + \delta^{\hat{t}+1-n} \cdot \frac{1}{2} \right)$, noting that the

sequence is increasing, and that b^n falls short of a player's present value of agreeing on an equal split with the delay $\hat{t} + 1 - n$ that remains as of the n -th round, which is $\delta^{\hat{t}+1-n} \cdot \frac{1}{2}$. It is straightforward to verify that the following describes equilibrium play with gradual agreement: in any (disagreement) round $n < \hat{t} + 1$ the proposing player $P(n)$ offers the share b^n , and the responding player $R(n)$ accepts with threshold $1 - b^{n+1}$ ($b^n < 1 - b^{n+1}$ follows from $b^n < b^{n+1} < \frac{1}{2}$); in the (agreement) round $n = \hat{t} + 1$ the proposing player $P(\hat{t} + 1)$ offers the share $\frac{1}{2}$, and the responding player $R(\hat{t} + 1)$ accepts with threshold $\frac{1}{2}$.

Solving for τ , inequality (8) becomes

$$\tau \leq \frac{\ln(1 - \delta^2) - \ln(1 - \gamma\delta^2)}{\ln(\delta)} \equiv \hat{\tau}(\delta, \gamma),$$

and if it is satisfied, the maximal delay t^* equals $\lfloor \hat{\tau}(\delta, \gamma) \rfloor$, i.e. the greatest integer not exceeding $\hat{\tau}(\delta, \gamma)$. For any $\gamma < 1$, this maximal delay approaches infinity as $\delta \rightarrow 1$, showing how small deviations from ED result in the emergence of delay equilibria as offers become very frequent; e.g., $\lfloor \hat{\tau}(\delta, \gamma) \rfloor = 404$ in case $\delta = \gamma = 0.999$. The right-hand-side panel of figure 2 illustrates this numerically.

The resulting delays can be very costly. The present value of the surplus in an equilibrium where agreement is maximally delayed equals $\gamma\delta^{t^*}$ whenever $\tau \leq \hat{\tau}(\delta, \gamma)$. As $\delta \rightarrow 1$, for any given $\gamma < 1$, not only is $\tau \leq \hat{\tau}(\delta, \gamma)$ going to be satisfied, but the entire surplus vanishes. For instance, while in the case of $\delta = \gamma = 0.99$ the maximal surplus loss amounts to roughly one third of the total, for values of γ that fall short of δ , the loss can be dramatic: up to 99.8% of the surplus can be lost through delay when $\delta = 0.99999$ and $\gamma = 0.99$.

When players discount the future only up to a finite number of delays, equilibrium delay can even be unbounded. Example 6 in appendix B.2 demonstrates this point, by only slightly modifying the example given here.

6 Concluding Remarks

The reason why two bargaining parties will reach agreement is that delay is costly. A basic cost of delay stems from impatience, as modeled in economics by time preferences. When information is perfect, time preferences are the sole driving force of strategic interaction, and this paper has examined their full implications in alternating-offers bargaining by two strategically sophisticated parties.

Based on a novel analytical approach that renders the game tractable under minimal assumptions on time preferences (dynamic inconsistency), its main insights are that any present bias pushes the parties towards immediate and efficient agreement, while a near-

future bias allows inefficient delay to be self-enforcing. With respect to the received literature, the notion of impatience covered here is comprehensive: it requires only that, *ceteris paribus*, a party prefers sooner over later, as well as more over less. I have, however, also maintained two standard assumptions: that time preferences are defined over sure outcomes, and that they are a fundamental stable trait of how an individual trades off delay and reward. I now discuss these in turn.

First, any meaningful uncertainty over outcomes would have necessitated extra assumptions regarding how the parties evaluate uncertainty that is distributed over time. Yet, in reduced form, the model also captures bargaining under the shadow of breakdown risk, with non-linear probability weighting of this risk as the source of the parties’ dynamic inconsistency. Given a (constant) probability $1-p$ that bargaining breaks down and yields nothing to both parties whenever they fail to agree, consider preferences $U_i(q, t) = g_i(p^t) \cdot u_i(q)$, where g_i is a general probability-weighting function applied to the “survival” probability for delay t , and where $u_i(0) = 0$ (see [Halevy, 2008](#); [Saito, 2015](#)). These preferences are dynamically consistent if and only if g_i is the identity, in which case i maximizes expected utility. Recasting $g_i(p^t)$ as a discounting function, all results of this paper apply in a straightforward manner. Since the preference domain then naturally includes risk, however, this paper’s restriction to pure strategies warrants reconsideration. This issue, and especially the relationship between the non-cooperative bargaining solution thus obtained and the axiomatic Nash solution (see [Rubinstein, Safra, and Thomson, 1992](#); [Grant and Kajii, 1995](#)) are most notable extensions for future research.³⁷

Second, the theoretical notion of time preferences is that of a fundamental stable individual trait in almost any economic analysis. Recent research suggests, however, that the choices between various delayed monetary rewards that are usually studied to infer time preferences additionally reflect transitory financial circumstances, as confounds of “pure” time preferences: e.g., [Ambrus, Ásgeirsdóttir, Noor, and Sándor \(2015\)](#), [Carvalho, Meier, and Wang \(2016\)](#) and [Dean and Sautmann \(2016\)](#) argue theoretically and show empirically how such choices systematically respond to liquidity (see also [Noor, 2009](#), for a related point). This recent attention to rigorously dealing with confounds in time preference research conveys optimism that, ultimately, it will be able to reliably identify pure time preferences empirically, and to also explain the otherwise puzzling amount of heterogeneity of “unconditionally” measured time preferences found in any study.³⁸ In view of the minimal preference

³⁷Other issues addressed in extensions or variations of the [Rubinstein \(1982\)](#) model under ED may also warrant reconsideration under dynamic inconsistency. While details of optimal penal codes may vary with the particular model, the basic structural properties established in section 4.1 appear robust—they essentially only depend on the proposer’s strategic advantage in bargaining—and accordingly useful for deriving the optimal punishments explicitly.

³⁸The vast heterogeneity concerns not only the broad quantitative impatience measures implied by in-

assumptions of this paper, it is likely to cover whatever preference model will emerge from this research, and hence its basic bargaining implications. Moreover, insights into how exogenous observables affect attitudes to delay (in addition to pure time preferences) open a wide range of interesting applications and extensions of this paper’s model; e.g., to study how bargaining behavior is influenced by liquidity.

Notwithstanding the above empirical issues regarding traditional choice experiments, a present bias is psychologically intuitive for hedonic utility, and this intuition is supported by both neurological evidence (e.g., [McClure, Laibson, Loewenstein, and Cohen, 2004](#)) and evolutionary arguments (e.g., [Netzer, 2009](#)). Neither currently exists for near-future bias as a stable preference trait. Moreover, this paper’s theoretical findings could also be interpreted as lending further support for a (weak) present bias over a near-future bias, if one is willing to assume that evolution has favored preferences that promote predictably efficient agreements. Concluding from this perspective, when the parties are strategically sophisticated, dynamic inconsistency of time preferences is unlikely to be a major reason for delay in bargaining.

dividuals’ choices but also their basic qualitative classification into no bias, present bias or future bias. Considering the importance of time preferences in human decision-making, and especially from an evolutionary perspective, such heterogeneity would be rather surprising for pure time preferences. Note that the finding of a future bias is essentially one of a *near*-future bias, which makes it indistinguishable from a present bias for delay horizons beyond a few weeks. It has also been called “reverse time-inconsistency” ([Sayman and Öncüler, 2009](#)), “increasing impatience” ([Attema, Bleichrodt, Rohde, and Wakker, 2010](#)), “hypobolic discounting” ([Eil, 2012](#)) or “patient shifts” ([Read, Frederick, and Airoldi, 2012](#)).

References

- Abreu, D. (1988). On the theory of infinitely repeated games with discounting. *Econometrica* 56(2), 383–396. 3, 16
- Abreu, D., D. Pearce, and E. Stacchetti (2015). One-sided uncertainty and delay in reputational bargaining. *Theoretical Economics* 10(3), 719–773. 2
- Admati, A. R. and M. Perry (1991). Joint projects without commitment. *The Review of Economic Studies* 58(2), 259–276. 27
- Ainslie, G. (1975). Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin* 82(4), 463–496. 3
- Akin, Z. (2007). Time inconsistency and learning in bargaining games. *International Journal of Game Theory* 36(2), 275 – 299. 4
- Ambrus, A., T. L. Ásgeirsdóttir, J. Noor, and L. Sándor (2015, March). Compensated discount functions - an experiment on the influence of expected income on time preferences. 33
- Attema, A. E., H. Bleichrodt, K. I. M. Rohde, and P. P. Wakker (2010). Time-tradeoff sequences for analyzing discounting and time inconsistency. *Management Science* 56(11), 2015–2030. 34
- Avery, C. and P. B. Zemsky (1994). Money burning and multiple equilibria in bargaining. *Games and Economic Behavior* 7(2), 154–168. 4, 12, 22, 29
- Benhabib, J., A. Bisin, and A. Schotter (2010). Present-bias, quasi-hyperbolic discounting, and fixed costs. *Games and Economic Behavior* 69(2), 205–223. 7, 24
- Bernheim, B. D. and A. Rangel (2009). Beyond revealed preference: Choice-theoretic foundations for behavioral welfare economics. *The Quarterly Journal of Economics* 124(1), 51–104. 25
- Binmore, K. G. (1987). Perfect equilibria in bargaining models. In K. G. Binmore and P. Dasgupta (Eds.), *The Economics of Bargaining*, pp. 77–105. Oxford: Basil Blackwell. 4, 9
- Binmore, K. G., A. Rubinstein, and A. Wolinsky (1986). The Nash bargaining solution in economic modelling. *The RAND Journal of Economics* 17(2), 176–188. 8, 23, 28
- Bleichrodt, H., K. I. M. Rohde, and P. P. Wakker (2009). Non-hyperbolic time inconsistency. *Games and Economic Behavior* 66(1), 27–38. 4, 29
- Börger, T. (1991). Upper hemicontinuity of the correspondence of subgame-perfect equilibrium outcomes. *Journal of Mathematical Economics* 20(1), 89–106. 22
- Burgos, A., S. Grant, and A. Kajii (2002a). Bargaining and boldness. *Games and Economic Behavior* 38(1), 28–51. 4
- Burgos, A., S. Grant, and A. Kajii (2002b). Corrigendum to “Bargaining and boldness”. *Games and Economic Behavior* 41(1), 165–168. 21
- Busch, L.-A. and Q. Wen (1995). Perfect equilibria in a negotiation model. *Econometrica* 63(3), 545–565. 5, 22

- Carvalho, L. S., S. Meier, and S. W. Wang (2016). Poverty and economic decision-making: Evidence from changes in financial resources at payday. *The American Economic Review* 106(2), 260–284. 33
- Chade, H., P. Prokopovych, and L. Smith (2008). Repeated games with present-biased preferences. *Journal of Economic Theory* 139(1), 157–175. 3, 8
- Chung, S.-H. and R. J. Herrnstein (1967). Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior* 10(1), 67–74. 3
- Coles, M. G. and A. Muthoo (2003). Bargaining in a non-stationary environment. *Journal of Economic Theory* 109(1), 70–89. 5
- Compte, O. and P. Jehiel (2004). Gradualism in bargaining and contribution games. *The Review of Economic Studies* 71(4), 975–1000. 27
- Cripps, M. W. (1998). Markov bargaining games. *Journal of Economic Dynamics and Control* 22(3), 341–355. 5
- Dasgupta, P. and E. Maskin (2005). Uncertainty and hyperbolic discounting. *The American Economic Review* 95(4), 1290–1299. 25
- Dean, M. and A. Sautmann (2016, January). Credit constraints and the measurement of time preferences. 33
- Ebert, J. E. J. and D. Prelec (2007). The fragility of time: Time-insensitivity and valuation of the near and far future. *Management Science* 53(9), 1423–1438. 4, 29
- Eil, D. (2012, June). Hypobolic discounting and willingness-to-wait. GMU Working Paper in Economics No. 12-28. 34
- Fernandez, R. and J. Glazer (1991). Striking for a bargain between two completely informed agents. *The American Economic Review* 81(1), 240–252. 21
- Fershtman, C. and D. J. Seidmann (1993). Deadline effects and inefficient delay in bargaining with endogenous commitment. *Journal of Economic Theory* 60(2), 306–321. 5
- Fishburn, P. C. and A. Rubinstein (1982). Time preference. *International Economic Review* 23(3), 677–694. 7, 8
- Friedenberg, A. (2016, July). Inefficient agreements: The role of second-order optimism. 5
- Fudenberg, D. and J. Tirole (1991). *Game Theory*. The MIT Press. 8
- Grant, S. and A. Kajii (1995). A cardinal characterization of the Rubinstein-Safra-Thomson axiomatic bargaining theory. *Econometrica* 63(5), 1241–1249. 33
- Halevy, Y. (2008). Strotz meets Allais: Diminishing impatience and the certainty effect. *The American Economic Review* 98(3), 1145–1162. 24, 33
- Halevy, Y. (2015). Time consistency: Stationarity and time invariance. *Econometrica* 83(1), 335–352. 2, 7
- Haller, H. and S. Holden (1990). A letter to the editor on wage bargaining. *Journal of Economic Theory* 52(1), 232–236. 21

- Harris, C. and D. Laibson (2013). Instantaneous gratification. *The Quarterly Journal of Economics* 128(1), 205–248. [25](#)
- Hoel, M. (1986). Perfect equilibria in sequential bargaining games with nonlinear utility functions. *The Scandinavian Journal of Economics* 88(2), 383–400. [23](#)
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics* 112(2), 443–478. [4](#)
- Li, D. (2007). Bargaining with history-dependent preferences. *Journal of Economic Theory* 136(1), 695–708. [5](#), [27](#)
- Lu, S. E. (2016). Self-control and bargaining. *Journal of Economic Theory* 165, 390–413. [4](#)
- Mailath, G. J., V. Nocke, and L. White (2015, February). When and how the punishment must fit the crime. University of Mannheim Working Paper 15-04. [3](#), [16](#)
- Manzini, P. and M. Mariotti (2009). Choice over time. In P. Anand, P. Pattanaik, and C. Puppe (Eds.), *The Handbook of Rational and Social Choice*, pp. 239–270. Oxford University Press. [7](#)
- McClure, S. M., D. I. Laibson, G. Loewenstein, and J. D. Cohen (2004). Separate neural systems value immediate and delayed monetary rewards. *Science* 306, 503–507. [34](#)
- Merlo, A. and C. Wilson (1995). A stochastic model of sequential bargaining with complete information. *Econometrica* 63(2), 371–399. [5](#), [8](#), [21](#)
- Muthoo, A. (1990). Bargaining without commitment. *Games and Economic Behavior* 2(3), 291–297. [2](#), [21](#)
- Myerson, R. B. (1991). *Game Theory: Analysis of Conflict*. Harvard University Press. [21](#)
- Nash, J. F. (1950). The bargaining problem. *Econometrica* 18(2), 155–162. [27](#)
- Netzer, N. (2009). Evolution of time preferences and attitudes toward risk. *The American Economic Review* 99(3), 937–955. [25](#), [34](#)
- Noor, J. (2009). Hyperbolic discounting and the standard model: Eliciting discount functions. *Journal of Economic Theory* 144(5), 2077–2083. [33](#)
- Noor, J. (2011). Intertemporal choice and the magnitude effect. *Games and Economic Behavior* 72(1), 255–270. [3](#), [4](#), [7](#), [24](#), [28](#), [A-20](#)
- O’Donoghue, T. and M. Rabin (1999). Doing it now or later. *The American Economic Review* 89(1), 103–124. [8](#)
- Ok, E. A. and Y. Masatlioglu (2007). A theory of (relative) discounting. *Journal of Economic Theory* 137(1), 214–245. [4](#), [7](#)
- Osborne, M. J. and A. Rubinstein (1990). *Bargaining and Markets*. Academic Press, Inc. [23](#)
- Pan, J., C. S. Webb, and H. Zank (2015). An extension of qasi-hyperbolic discounting to continuous time. *Games and Economic Behavior* 89, 43–55. [4](#), [25](#), [29](#)

- Perry, M. and P. J. Reny (1993). A non-cooperative bargaining model with strategically timed offers. *Journal of Economic Theory* 59(1), 50–77. [2](#), [5](#)
- Phelps, E. S. and R. A. Pollak (1968). On second-best national saving and game-equilibrium growth. *The Review of Economic Studies* 35(2), 185–199. [3](#)
- Piccione, M. and A. Rubinstein (1997). On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior* 20, 3–24. [8](#)
- Pollak, R. A. (1968). Consistent planning. *The Review of Economic Studies* 35(2), 201–208. [8](#)
- Read, D. (2001). Is time-discounting hyperbolic or subadditive? *Journal of Risk and Uncertainty* 23(1), 5–32. [7](#)
- Read, D., S. Frederick, and M. Airoldi (2012). Four days later in Cincinnati: Longitudinal tests of hyperbolic discounting. *Acta Psychologica* 140(2), 177–185. [34](#)
- Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica* 50(1), 97–109. [2](#), [5](#), [8](#), [9](#), [21](#), [23](#), [28](#), [33](#), [A-19](#)
- Rubinstein, A. (1991). Comments on the interpretation of game theory. *Econometrica* 59(4), 909–924. [15](#)
- Rubinstein, A. (2003). “Economics and psychology”? The case of hyperbolic discounting. *International Economic Review* 44(4), 1207–1216. [7](#)
- Rubinstein, A., Z. Safra, and W. Thomson (1992). On the interpretation of the Nash bargaining solution and its extension to non-expected utility preferences. *Econometrica* 60(5), 1171–1186. [33](#)
- Rusinowska, A. (2004). Bargaining model with sequences of discount rates and bargaining costs. *International Game Theory Review* 6(2), 268–280. [4](#)
- Saito, K. (2015, April). A relationship between risk and time preferences. [33](#)
- Sákovics, J. (1993). Delay in bargaining games with complete information. *Journal of Economic Theory* 59(1), 78–95. [2](#), [5](#)
- Sayman, S. and A. Öncüler (2009). An investigation of time inconsistency. *Management Science* 55(3), 470–482. [34](#)
- Shaked, A. and J. Sutton (1984). Involuntary unemployment as a perfect equilibrium in a bargaining model. *Econometrica* 52(6), 1351–1364. [2](#), [12](#), [15](#)
- Ståhl, I. (1972). *Bargaining Theory*. EFI The Economics Research Institute, Stockholm. [2](#)
- Strotz, R. H. (1955-1956). Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies* 23(3), 165–180. [8](#)
- Takeuchi, K. (2011). Non-parametric test of time consistency: Present bias and future bias. *Games and Economic Behavior* 71(2), 456–478. [29](#)
- van Damme, E., R. Selten, and E. Winter (1990). Alternating bid bargaining with a smallest money unit. *Games and Economic Behavior* 2(2), 188–201. [21](#)
- Yildiz, M. (2011). Bargaining with optimism. *Annual Review of Economics* 3(4), 451–478. [2](#), [5](#)

Appendix

A Proofs

A.1 Additional Notation

The set $\mathcal{A} \equiv X \times \mathcal{P}(X)$ defines the possible pairs of proposals and response rules. The stationary strategy σ_i that specifies “always propose x ” and “always respond using rule Y ” is identified with the pair $(x, Y) \in \mathcal{A}$. The particular division that has player i ’s share equal to one (player j ’s share is zero) is denoted by $e^{(i)}$, and a player i ’s response rule “accept if and only if your share is at least q ” is denoted by $X_{i,q}$.

Take any strategy profile σ , and suppose that if both players act according to σ the outcome is division x in round m (hence with delay $m - 1$), where $x = (0, 0)$ and $m = \infty$ in case of perpetual disagreement. For any $n \leq m$, let then $h^{n-1}(\sigma) \in X^{n-1}$ be the round- n history σ induces, and let $(h^{n-1}(\sigma), x) \in X^m$ be its induced (terminal) path. I formally define σ ’s *play* to be the sequence $\langle \sigma \rangle \equiv (\langle \sigma \rangle_n)_{n=1}^m \in \mathcal{A}^m$ of offers and response *rules* it prescribes along its induced path, i.e. $\langle \sigma \rangle_n \equiv \sigma(h^{n-1}(\sigma))$ for any $n \leq m$.

To isolate plays from strategy profiles, call any sequence $(x^n, Y^n)_{n=1}^m \in \mathcal{A}^m$, for $m \in \mathbb{N}$, a play of game G_i if there exists a strategy profile σ in this game such that $\langle \sigma \rangle = (x^n, Y^n)_{n=1}^m$; this holds true if and only if $x^n \in Y^n \Leftrightarrow n = m$ (the condition is identical for both games G_1 and G_2), and for a given game, a play defines an equivalence class of strategy profiles.

Next, consider the following mapping that produces “simple” strategy profiles. Given any quadruple of plays $S \equiv (\langle \sigma^{P,i} \rangle, \langle \sigma^{R,i} \rangle)_{i \in I}$, define, for each $i \in I$, a mapping $\sigma^{S,i}(\cdot)$ that assigns to any play $\langle \hat{\sigma} \rangle$ a strategy profile in game G_i as follows: interpreting any play $\langle \sigma \rangle \in \{\langle \hat{\sigma} \rangle\} \cup \{\langle \sigma^{P,i} \rangle, \langle \sigma^{R,i} \rangle\}_{i \in I}$ as a sequence of “states”, say a strategy profile is in “state” $\langle \sigma \rangle_n$ if it prescribes play $\langle \sigma \rangle_n$ after a given history, and then define $\sigma^{S,i}(\langle \hat{\sigma} \rangle)$ by the rule that

(1) in round 1 $\sigma^{S,i}(\langle \hat{\sigma} \rangle)$ is in state $\langle \hat{\sigma} \rangle_1$, and

(2) if in round m it is in state $\langle \sigma \rangle_m = (x, Y)$, and proposal x' is rejected, then in round $m + 1$ it is in state

$$\tau(\langle \sigma \rangle_n, x') = \begin{cases} \langle \sigma \rangle_{n+1} & x' = x \notin Y \\ \langle \sigma^{P,P(n)} \rangle_1 & x' \neq x \notin Y \\ \langle \sigma^{R,R(n)} \rangle_1 & x' \neq x \in Y \end{cases}.$$

This is a well-defined strategy profile with the property that it distinguishes only four types of deviations from a given prescribed play—one per player per role—and always specifies the same continuation play after the same type of deviation. It is thus simple in the sense of minimal history-dependence.

Finally, given any pair of reservation shares $Q \equiv (q_1, q_2)$, define, for each $i \in I$, the mapping $\alpha^{Q,i}(\cdot)$ that assigns to any outcome (\hat{x}, \hat{t}) the sequence $(x^n, Y^n)_{n=1}^{\hat{t}+1} \in \mathcal{A}^{\hat{t}+1}$ such that

$$(x^n, Y^n) = \begin{cases} (e^{(P(n))}, X_{R(n), q_{R(n)}}) & n < \hat{t} + 1 \\ (\hat{x}, X_{R(n), \hat{x}_{R(n)}}) & n = \hat{t} + 1 \end{cases} \text{ for } (P(n), R(n)) \equiv \begin{cases} (i, j) & n \text{ odd} \\ (j, i) & n \text{ even} \end{cases}.$$

Note that $\alpha^{Q,i}(\hat{x}, \hat{t})$ is a play of game G_i if and only if

$$\left\{ \begin{array}{l} \hat{t} = 1 \Rightarrow q_j > 0 \\ \hat{t} > 1 \Rightarrow q_1 \cdot q_2 > 0 \end{array} \right\}.$$

A.2 Lemmas 1 and 2

Take any strategy profile σ and any round- n history h^{n-1} : first, let $\sigma|_{h^{n-1}}$ denote the restriction of σ to continuation histories of h^{n-1} , i.e. histories of the form (h^{n-1}, h^{m-1}) where $h^{m-1} \in X^{m-1}$ for $m \in \mathbb{N}$, and second, let $\sigma|^{h^{n-1}}$ denote the strategy profile in game $G_{P(n)}$ that is obtained from $\sigma|_{h^{n-1}}$ upon replacing h^{n-1} by the initial history h^0 . (Observe that, given h^{n-1} , $\sigma|^{h^{n-1}}$ completely characterizes $\sigma|_{h^{n-1}}$.) Fixing any quadruple of strategy profiles $(\sigma^{P,i}, \sigma^{R,i})_{i \in I}$ such that, for each $i \in I$, $\sigma^{P,i}$ is a strategy profile in game G_j and $\sigma^{R,i}$ is a strategy profile in game G_i , define, for each $i \in I$, the mapping $\sigma^{*,i}(\cdot | (\sigma^{P,i}, \sigma^{R,i})_{i \in I})$ as follows: for any strategy profile σ in game G_i , it is the unique strategy profile $\sigma^{*,i}$ in this game such that $\langle \sigma^{*,i} \rangle = \langle \sigma \rangle$ and

$$\sigma^{*,i}|_{(h^{n-1}(\sigma), x)} = \begin{cases} \sigma^{P,P(n)} & x \notin \sigma_{R(n)}(h^{n-1}(\sigma)) \setminus \{\sigma_{P(n)}(h^{n-1}(\sigma))\} \\ \sigma^{R,R(n)} & x \in \sigma_{R(n)}(h^{n-1}(\sigma)) \end{cases}.$$

Using this definition, lemmas 1 and 2 are formally summarized in the proposition below; part (i) establishes the defining property of optimal punishment, part (ii) shows that it is without loss of generality for optimality to restrict attention to simple punishment, and part (iii) shows it is without loss of generality for equilibrium to restrict attention to simple play.

Proposition 4. *Let the quadruple of outcomes $((x^{P,i}, t^{P,i}), (x^{R,i}, t^{R,i}))_{i \in I}$ be such that, for each $i \in I$,*

$$(x_j^{P,i}, t^{P,i}) \in \arg \max_{(q,t) \in A_j^*} U_j(q, t+1) \quad \text{and} \quad (x_i^{R,i}, t^{R,i}) \in \arg \min_{(q,t) \in A_i^*} U_i(q, t+1). \quad (9)$$

(i) Fix a quadruple of equilibria $(\sigma^{P,i}, \sigma^{R,i})_{i \in I}$ such that, for each $i \in I$, $\sigma^{P,i}$ is an equilibrium of game G_j supporting outcome $(x^{P,i}, t^{P,i})$ and $\sigma^{R,i}$ is an equilibrium of game G_i supporting outcome $(x^{R,i}, t^{R,i})$. Then, for any $k \in I$ and strategy profile $\hat{\sigma}$ in game G_k , $\langle \hat{\sigma} \rangle$ is an equilibrium play of G_k if and only if $\sigma^{*,k} \left(\hat{\sigma} \left| (\sigma^{P,i}, \sigma^{R,i})_{i \in I} \right. \right)$ is an equilibrium of G_k .

(ii) The quadruple of equilibria $(\sigma^{P,i}, \sigma^{R,i})_{i \in I}$ in (i) can be chosen such that

$$\sigma^{P,i} = \sigma^{*,j} \left(\sigma^{P,i} \left| (\sigma^{P,i}, \sigma^{R,i})_{i \in I} \right. \right) \quad \text{and} \quad \sigma^{R,i} = \sigma^{*,i} \left(\sigma^{R,i} \left| (\sigma^{P,i}, \sigma^{R,i})_{i \in I} \right. \right). \quad (10)$$

(iii) For any $k \in I$, (\hat{x}, \hat{t}) is an equilibrium outcome of game G_k if and only if $\alpha^{Q^*,k}(\hat{x}, \hat{t})$, with $Q^* = (\pi_1(U_1(x_1^{P,2}, t^{P,2} + 1)), \pi_2(U_2(x_2^{P,1}, t^{P,1} + 1)))$, is an equilibrium play of G_k .

Proof. Part (i). Sufficiency is immediate, since $\langle \sigma^{*,k} \rangle = \langle \hat{\sigma} \rangle$.

For necessity, let $\langle \hat{\sigma} \rangle$ be an equilibrium play of G_k with outcome (\hat{x}, \hat{t}) , where it is without loss of generality to assume $\hat{\sigma}$ is itself an equilibrium of G_k , and also let $\sigma^* = \sigma^* \left(\hat{\sigma} \left| (\sigma^{P,i}, \sigma^{R,i})_{i \in I} \right. \right)$. By construction, $\langle \sigma^* \rangle = \langle \hat{\sigma} \rangle$, and continuation play under σ^* following any deviation from its path is an equilibrium of the resulting subgame. In order to verify that σ^* is an equilibrium it therefore suffices to verify that there are no profitable one-stage deviations at the histories $h^{n-1}(\sigma^*)$ along its path.

Take then any such history $h = h^{n-1}(\sigma^*)$, where player P makes an offer to player R , and $\sigma^*(h) = \hat{\sigma}(h) = (\tilde{x}, \tilde{Y})$. Consider any proposal $x' \in \tilde{Y}$; $\hat{\sigma}$'s being an equilibrium and the construction of σ^* imply that

$$u_R(x'_R) \geq U_R(z_R^h(x', \emptyset | \hat{\sigma})) \geq \min \{U_R(x_R, t+1) | (x_R, t) \in A_R^*\} = U_R(z_R^h(x', \emptyset | \sigma^*)),$$

whereby acceptance is optimal for R under σ^* .

Next, consider any proposal $x' \notin \tilde{Y} \setminus \{\tilde{x}\}$; $\hat{\sigma}$'s being an equilibrium and the construction of σ^* imply that

$$u_R(x'_R) \leq U_R(z_R^h(x', \emptyset | \hat{\sigma})) \leq \max \{U_R(x_R, t+1) | (x_R, t) \in A_R^*\} = U_R(z_R^h(x', \emptyset | \sigma^*)),$$

whereby rejection is optimal for R under σ^* .

The only remaining case at the responding stage is that of proposal \tilde{x} such that $\tilde{x} \notin \tilde{Y}$; this implies that $n < \hat{t} + 1$, and then $\hat{\sigma}$'s being an equilibrium play and the construction of σ^* imply that

$$u_R(\tilde{x}_R) \leq U_R(z_R^h(\tilde{x}, \emptyset | \hat{\sigma})) = U_R(\hat{x}_R, \hat{t} + 1 - n) = U_R(z_R^h(\tilde{x}, \emptyset | \sigma^*)),$$

whereby rejection is optimal for R under σ^* .

Finally, consider the proposing player P 's incentive to propose $x' \neq \tilde{x}$: if $x' \in \tilde{Y}$, then $u_P(x'_P) \leq U_P(z_P^h(\tilde{x}, \tilde{Y} | \hat{\sigma}))$ by $\hat{\sigma}$'s being an equilibrium, and because of $z_P^h(\tilde{x}, \tilde{Y} | \sigma^*) = z_P^h(\tilde{x}, \tilde{Y} | \hat{\sigma}) = (\hat{x}_P, \hat{t} + 1 - n)$ such deviations are not profitable to P under σ^* .

Letting $q_R^* = \pi_R(U_R(x_R^{P,P}, t^{P,P} + 1))$, it follows from $\hat{\sigma}$'s being an equilibrium that $\{x \in X | x_R > q_R^*\} \subseteq \tilde{Y}$ and $u_P(1 - q_R^*) \leq U_P(\hat{x}_P, \hat{t} + 1 - n)$: R must accept any offer which exceeds her maximal credible reservation share, and if $u_P(1 - q_R^*) > U_P(\hat{x}_P, \hat{t} + 1 - n)$ were true, then, because $u_P(\cdot)$ is continuously increasing and $q_R^* < 1$ due to R 's impatience, there would exist $\epsilon > 0$ such that P 's offering the accepted share $q_R^* + \epsilon$ would be a profitable deviation under $\hat{\sigma}$. Under σ^* any deviant proposal $x' \notin \tilde{Y}$ yields utility $U_P(x_P^{P,P}, t^{P,P} + 1)$; using the fact that $\pi_P(U_P(x_P^{P,P}, t^{P,P} + 1)) + \pi_R(U_R(x_R^{P,P}, t^{P,P} + 1)) < 1$ by impatience,

$$U_P(x_P^{P,P}, t^{P,P} + 1) \leq u_P(\pi_P(U_P(x_P^{P,P}, t^{P,P} + 1))) < u_P(1 - \pi_R(U_R(x_R^{P,P}, t^{P,P} + 1)));$$

hence no such deviation is profitable for P , concluding the proof.

Part (ii). If $(\sigma^{P,i}, \sigma^{R,i})_{i \in I}$ is a quadruple of equilibria as in part (i), then $S = (\langle \sigma^{P,i} \rangle, \langle \sigma^{R,i} \rangle)_{i \in I}$ is a quadruple of plays, so the quadruple of strategy profiles $(\sigma^{S,j}(\langle \sigma^{P,i} \rangle), \sigma^{S,i}(\langle \sigma^{R,i} \rangle))_{i \in I}$ is well-defined. When used as punishments in mapping $\sigma^{*,i}$ this quadruple supports the same set of plays in game G_i , $i \in I$, as does $(\sigma^{P,i}, \sigma^{R,i})_{i \in I}$, since the punishments for various deviations from initial play are outcome equivalent. In particular, $(\sigma^{S,j}(\langle \sigma^{P,i} \rangle), \sigma^{S,i}(\langle \sigma^{R,i} \rangle))_{i \in I}$ therefore supports its own constituent (equilibrium) plays in S , so at no point is there a profitable deviation from any of these strategy profiles; it is therefore a quadruple of equilibria as in part (i).

Finally, by construction, any of them specifies the same punishment after any deviation by the same player in the same role, irrespective of history: if proposing player i makes a deviant offer that is compliantly rejected, this is $\sigma^{S,j}(\langle \sigma^{P,i} \rangle)$, and if responding player i deviantly rejects an offer, this is $\sigma^{S,i}(\langle \sigma^{R,i} \rangle)$. Hence it satisfies (10).

Part (iii). Sufficiency is immediate. Suppose then that agreement on \hat{x} with delay \hat{t} is an equilibrium outcome of G_k , and let $(\sigma^{P,i}, \sigma^{R,i})_{i \in I}$ be a quadruple of equilibria as in part (i). Define also each player i 's shares $q_i^* \equiv \pi_i(U_i(x_i^{P,j}, t^{P,j} + 1))$ and $q_i^{**} \equiv \pi_i(U_i(x_i^{R,i}, t^{R,i} + 1))$.

The first step is to show that $\alpha^{Q^*,k}(\hat{x}, \hat{t})$ is a play. This is immediate only for $\hat{t} = 0$; for $\hat{t} = 1$, it is necessary and sufficient that $q_{3-k}^* > 0$, and for $\hat{t} > 1$, it is necessary and sufficient that both $q_2^* > 0$ and $q_1^* > 0$. Suppose then that $q_i^* = 0$ and note that any equilibrium must then have respondent i accept any offer. While immediate for any positive offer, there

cannot be an equilibrium in which respondent i rejects a zero offer by proposer j , because $u_j(1 - \epsilon) > U_j(1, 1)$ for small enough positive and hence accepted offers ϵ ; i 's rejecting a zero offer would therefore imply that such offers constitute profitable deviations by proposer j . Hence, $\hat{t} = 1$ implies $q_{3-k}^* > 0$, and $\hat{t} > 1$ implies both $q_2^* > 0$ and $q_1^* > 0$.

The second step is to show that, whenever σ' is a strategy profile in game G_k whose play equals $\alpha^{Q^*,k}(\hat{x}, \hat{t})$, then $\sigma \equiv \sigma^{*,k} \left(\sigma' \left| \left(\sigma^{P,i}, \sigma^{R,i} \right)_{i \in I} \right. \right)$ is an equilibrium of G_k . It suffices to verify that there are no profitable one-stage deviations at the histories $h^{n-1}(\sigma)$ for $n \leq \hat{t} + 1$, since the continuation strategy profiles $(\sigma^{P,i}, \sigma^{R,i})_{i \in I}$ are all equilibria of their respective subgames. Consider then any such history $h = h^{n-1}(\sigma)$, where player P makes an offer to player R and $\sigma(h) = (\tilde{x}, X_{R,\tilde{q}})$. Observe the following inequalities:

$$q_R^{**} \leq \tilde{q} \leq q_R^*. \quad (11)$$

While (11) holds by construction if $n < \hat{t} + 1$, in the case of $n = \hat{t} + 1$ it means that $q_R^{**} \leq \hat{x}_R \leq q_R^*$; however, $\hat{x}_R < q_R^{**}$ would imply that there could not be an equilibrium in which R accepts an offer as low as \hat{x}_R , and $\hat{x}_R > q_R^*$ would imply that there could not be an equilibrium in which P offers as much as \hat{x}_R .

R 's rejection of any *deviant* offer $q \neq \tilde{x}_R$ such that $q < \tilde{q}$ is optimal: by (11), such offers exist only if $q_R^* > 0$, in which case their rejection value $U_R(x_R^{P,P}, t^{P,P} + 1)$ equals $u_R(q_R^*)$, and this exceeds that of acceptance, $u(q)$, since $q_R^* \geq \tilde{q} > q$. Moreover, R 's impatience implies that $x_R^{P,P} > q_R^*$, and combined with (11) this yields $U_P(x_P^{P,P}, t^{P,P} + 1) < u_P(1 - q_R^*) \leq u_P(1 - \tilde{q})$, showing that P has no profitable deviation to rejected offers $q < \tilde{q}$.

Also, R 's acceptance of any offer $q \geq \tilde{q}$ is optimal, because it yields a value of at least $u_R(\tilde{q})$, whereas rejection yields no more than $u_R(q_R^{**})$, where $u_R(\tilde{q}) \geq u_R(q_R^{**})$ by (11). Among these offers, \tilde{q} is clearly the best accepted offer for P .

For $n = \hat{t} + 1$, we can already conclude that there is no profitable deviation for either player, since all offers $q < \tilde{q}$ are deviant. Consider then the remaining case of deviations in a round $n < \hat{t} + 1$: if R 's rejection of the minimal possible, i.e. the zero offer failed to be optimal, then $u_R(0) > U_R(\hat{x}_R, \hat{t} + 1 - n)$, so there is no offer that R could optimally reject in favor of agreement on \hat{x} after $\hat{t} + 1 - n$ more rounds—in contradiction to this outcome's equilibrium property; to a similar effect, if P 's compliant zero offer were worse than the lowest accepted offer $\tilde{q} = q_R^*$, then $u_P(1 - q_R^*) > U_P(\hat{x}_P, \hat{t} - n + 1)$, so there is no rejected offer that P could optimally make in return for agreement on \hat{x} after $\hat{t} + 1 - n$ more rounds. \square

A.3 Theorem 1

In what follows, let

$$\begin{aligned}\tilde{v}_i &\equiv \inf \{U_i(q, t) \mid (q, t) \in A_i^*\} \\ \tilde{w}_i &\equiv \inf \{U_i(q, t+1) \mid (q, t) \in A_i^*\}\end{aligned}$$

denote each player i 's infimal punishment values. The theorem is proven via a series of lemmas. The first one, lemma 4, shows that the set E is non-empty. Lemma 5 then shows that for every element $(v_i, w_i, t_i)_{i \in I}$ of E there exists a quadruple of outcomes that deliver the values $(v_i, w_i)_{i \in I}$ when used as punishment outcomes. (This is the only result that uses impatience property (3.c), and it will imply that optimal punishments exist.) Lemma 6 goes on to establish that any such quadruple of outcomes in fact defines a “constrained” OSPC: as punishment outcomes they support a subset of equilibrium outcomes that includes them, and constrained to which they are optimal (see equation (9)). This means, in particular, that for any element $(v_i, w_i, t_i)_{i \in I}$ of E , $\tilde{v}_i \leq v_i$ and $\tilde{w}_i \leq w_i$ for each i . The final two lemmas show that E also contains an element $(v_i, w_i, t_i)_{i \in I}$ such that $v_i \leq \tilde{v}_i$ and $w_i \leq \tilde{w}_i$ for each i . Thus we conclude that E has an extreme element, which is $(v_i^*, w_i^*, t_i^*)_{i \in I}$. (Lemma 6 then implies the characterization of equilibrium outcomes based on the associated OSPC from lemma 5.)

Lemma 4. *The set E is non-empty.*

Proof. Consider the following functions $f_i : [0, 1] \rightarrow [0, 1]$ for each i :

$$f_i(q) \equiv 1 - \pi_j(U_j(1 - \pi_i(U_i(q, 1)), 1)). \quad (12)$$

f_i is continuous, and it is non-decreasing, with $0 < f_i(0) \leq f_i(1) \leq 1$. Hence it possesses a fixed point that is positive. Take any $\hat{q}_1 = f_1(\hat{q}_1)$ and define $\hat{q}_2 \equiv 1 - \pi_1(U_1(\hat{q}_1, 1))$; note that then also $\hat{q}_1 = 1 - \pi_2(U_2(\hat{q}_2, 1))$ and

$$\begin{aligned}\hat{q}_2 &= 1 - \pi_1(U_1(1 - \pi_2(U_2(\hat{q}_2, 1)), 1)) \\ &\equiv f_2(\hat{q}_2).\end{aligned}$$

I will prove that E contains the values $(v_i, w_i, t_i)_{i \in I} = (u_i(\hat{q}_i), U_i(\hat{q}_i, 1), 0)_{i \in I}$.

Given $t_i = 0$, the identity $\phi_i(u_i(\hat{q}_i), 0) \equiv \hat{q}_i$ immediately yields that the chosen values satisfy equations (4) and (5), for each i . At the same time, again for each i , whenever t is

positive,

$$\begin{aligned}
\kappa_i(t, u_i(\hat{q}_i), u_j(\hat{q}_j)) &\geq \kappa_i(1, u_i(\hat{q}_i), u_j(\hat{q}_j)) \\
&\geq \hat{q}_i + \hat{q}_j \\
&= \hat{q}_i + 1 - \pi_i(U_i(\hat{q}_i, 1)) \\
&> 1,
\end{aligned}$$

where the last inequality uses that $\hat{q}_i > 0$ implies $\hat{q}_i > \pi_i(U_i(\hat{q}_i, 1))$. This shows that the chosen values also satisfy equation (6), for each i . \square

Lemma 5. *For every element $(v_i, w_i, t_i)_{i \in I}$ of the set E , there exists a quadruple of outcomes $((y^{(i)}, 0), (x^{(i)}, t^{(i)}))_{i \in I}$ such that, for each $i \in I$,*

$$v_i = u_i\left(1 - \pi_j\left(U_j\left(1 - y_i^{(i)}, 1\right)\right)\right) \quad (13)$$

$$w_i = U_i\left(x_i^{(i)}, t^{(i)} + 1\right). \quad (14)$$

Proof. Let $(v_i, w_i, t_i)_{i \in I} \in E$ and define a quadruple of outcomes $((y^{(i)}, 0), (x^{(i)}, t^{(i)}))_{i \in I}$ such that, for each $i \in I$,

$$y_i^{(i)} = \pi_i(w_i) \text{ and } \left\{ \begin{array}{l} t^{(i)} \in \arg \min \{U_i(\phi_i(v_i, t), t + 1) \mid t \in T, t \leq t_i\} \\ x_i^{(i)} = \phi_i(v_i, t^{(i)}) \end{array} \right\}. \quad (15)$$

Recalling equations (4) and (5), it only remains to show that such values $t^{(i)}$ exist, so that the quadruple is well-defined. This is clearly true when each t_i is finite, and the following three steps prove it also for the case that $t_i = \infty$ (for some i).

Step 1: For any t , $\phi_i(v_i, t) > 0$. From equation (4) it follows that $v_i \geq u_i(1 - \pi_j(U_j(1, 1))) > u_i(0)$, since $\pi_j(U_j(q, t + 1)) \leq \pi_j(U_j(1, 1)) < 1$ for all $(q, t) \in A_j$ due to j 's impatience. Using identity $v_i \equiv u_i(\phi_i(v_i, 0))$, $v_i > u_i(0)$ is equivalent to $\phi_i(v_i, 0) > 0$, and the claim follows from the non-decreasingness of $\phi_i(u, \cdot)$ for any $u \in u_i([0, 1])$.

Step 2: For any $t \leq t_i$, $U_i(\phi_i(v_i, t), t) = v_i$. Since this holds true for $t = 0$ by definition, consider it for $0 < t \leq t_i$ and note that it suffices to show that $\phi_i(v_i, t) < 1$ (recall the definition of ϕ_i): from equation (6), $\kappa_i(t, v_i, v_j) \leq 1$, and using that $\phi_j(v_j, t - 1) > 0$ from step 1, this implies $\phi_i(v_i, t) < 1$.

Step 3: There exists a finite \bar{t}_i such that $w_i = \min \{U_i(\phi_i(v_i, t), t + 1) \mid t \in T, t \leq \bar{t}_i\}$. Since we can simply set $\bar{t}_i = t_i$ if t_i is finite, consider the case of $t_i = \infty$ and distinguish the two possible cases according to impatience property (3.c). Suppose first that player i 's preferences satisfy $\lim_{t \rightarrow \infty} U_i(1, t) \leq u_i(0)$. Since $v_i > u_i(0)$ from step 1, there then exists

a finite delay \hat{t} such that $t \geq \hat{t}$ implies $U_i(1, t) < v_i$, and hence $U_i(\phi_i(v_i, t), t) < v_i$, which contradicts step 2. The alternative case is that there exists a finite delay \hat{t} such that $t \geq \hat{t}$ implies $U_i(q, t) = U_i(q, \hat{t})$ for all q ; hence $U_1(\phi_1(v_1, t), t+1) = U_1(\phi_1(v_1, \hat{t}), \hat{t}+1)$ for all such t , which proves the claim upon setting $\bar{t}_1 = \hat{t}$. \square

Statement and proof of the next lemma use the following definition: for any values $(v_k, w_k)_{k \in I} \in \times_{k \in I} (u_k([0, 1]) \times U_k(A_k))$ and any player i , $A_i(v_1, w_1, v_2, w_2)$ is the set

$$\left\{ (q, t) \in A_i \left| \phi_i(v_i, t) \leq q \leq \begin{cases} 1 - \pi_j(w_j) & t = 0 \\ 1 - \max\{\phi_j(v_j, t-1), \phi_j(u_j(0), t)\} & t > 0 \end{cases} \right. \right\}.$$

Lemma 6. *Take any element $(v_i, w_i, t_i)_{i \in I}$ of the set E and associated quadruple of outcomes $((y^{(i)}, 0), (x^{(i)}, t^{(i)}))_{i \in I}$ satisfying (15). Then, for each $i \in I$,*

$$\left\{ (1 - y_j^{(j)}, 0), (x_i^{(i)}, t^{(i)}) \right\} \subseteq A_i(v_1, w_1, v_2, w_2) \subseteq A_i^*,$$

and the following equalities hold true:

$$\begin{aligned} v_i &= \min \{U_i(q, t) \mid (q, t) \in A_i(v_1, w_1, v_2, w_2)\} \\ w_i &= \min \{U_i(q, t+1) \mid (q, t) \in A_i(v_1, w_1, v_2, w_2)\} \\ t_i &= \sup \{t \in T \mid \exists q \in [0, 1], (q, t) \in A_i(v_1, w_1, v_2, w_2)\}. \end{aligned}$$

Proof. The following observation, for each i , will be helpful:

$$v_i > \max \{u_i(0), w_i\}. \quad (16)$$

Since $v_i > u_i(0)$ was established in step 1 of lemma 5, it only remains to prove that $v_i > w_i$: this follows from equation (5), implying $w_i \leq U_i(\phi_i(v_i, 0), 1)$, because $\phi_i(v_i, 0) > 0$ and i is impatient.

Let then, for each i , $\hat{q}_i \equiv \pi_i(U_i(1 - y_j^{(j)}, 1))$ and note that equation (13) implies that

$$\hat{q}_i = 1 - \phi_j(v_j, 0). \quad (17)$$

First, I will show that, given $Q = (\hat{q}_1, \hat{q}_2)$, $\alpha^{Q,i}(x^{(i)}, t^{(i)})$ is a play of game G_i , for each i . To simplify notation, let $i = 1$, which is without loss of generality. There is nothing to check if $t^{(1)} = 0$, so consider the case of $t^{(1)} > 0$. This implies that $t_1 > 0$ and hence $\kappa_1(1, v_1, v_2) \leq 1$; using that $\phi_2(v_2, 0) > 0$ by (16), we obtain $\phi_1(v_1, 1) < 1$, which implies $\phi_1(v_1, 0) < 1$, and hence, via equation (4) (for $i = 1$), $\hat{q}_2 > 0$. While necessary for any $t^{(1)} > 0$, this is

sufficient to prove the claim for $t^{(1)} = 1$. Suppose then $t^{(1)} > 1$; this implies $t_1 > 1$ and hence $\kappa_1(2, v_1, v_2) \leq 1$. Using $\phi_1(v_1, 2) > 0$ from combining (16) with the non-decreasingness of $\phi_1(u, \cdot)$, this in turn implies that $\phi_2(v_2, 0) < 1$, from which $\hat{q}_1 > 0$ follows via equation (4) (for $i = 2$).

Since any immediate-agreement outcome defines a play, it immediately follows from the previous argument that $S \equiv \left(\alpha^{Q,j} \left(y^{(j)}, 0 \right), \alpha^{Q,i} \left(x^{(i)}, t^{(i)} \right) \right)_{i \in I}$ is a quadruple of plays. I will now show that, for any outcome (\hat{x}, \hat{t}) and each i , $\alpha^{Q,i}(\hat{x}, \hat{t})$ is a play of G_i such that $\sigma^{S,i}(\alpha^{Q,i}(\hat{x}, \hat{t}))$ is an equilibrium of G_i if and only if $(\hat{x}_i, \hat{t}) \in A_i(v_1, w_1, v_2, w_2)$. Since $\left\{ (1 - y_j^{(j)}, 0), (x_i^{(i)}, t^{(i)}) \right\} \subseteq A_i(v_1, w_1, v_2, w_2)$, it is sufficient to prove that $\alpha^{Q,i}(\hat{x}, \hat{t})$ is a play of G_i such that there are no profitable deviations from this play under the strategy profile $\sigma^{S,i}(\alpha^{Q,i}(\hat{x}, \hat{t}))$ if and only if $(\hat{x}_i, \hat{t}) \in A_i(v_1, w_1, v_2, w_2)$. Again, only to simplify notation, I prove this claim for $i = 1$; also, I let $\hat{\sigma} \equiv \sigma^{S,1}(\alpha^{Q,1}(\hat{x}, \hat{t}))$ and $\hat{A}_1 \equiv A_1(v_1, w_1, v_2, w_2)$.

First, consider **immediate-agreement outcomes** $(\hat{x}, 0)$; $\alpha^{Q,1}(\hat{x}, 0)$ is a play for any division \hat{x} , and it remains to show that there is no profitable deviation from this play under $\hat{\sigma}$ if and only if $(\hat{x}_1, 0) \in \hat{A}_1$. Player 2's accepting all offers $q \geq \hat{x}_2$ is optimal if and only if $\hat{x}_2 \geq \pi_2(w_2)$, because deviantly rejecting such an offer would trigger her respondent punishment, which has continuation outcome $(x^{(2)}, t^{(2)})$ and associated rejection value w_2 ; her rejecting all other offers is optimal if and only if $\hat{x}_2 \leq \hat{q}_2$ because non-deviantly rejecting such a deviant offer would trigger player 1's proposer punishment, which has continuation outcome $(y^{(1)}, 0)$ and associated rejection value $U_2(1 - \pi_1(w_1), 1)$; using equation (17), $\hat{x}_2 \leq \hat{q}_2$ is equivalent to $\phi_1(v_1, 0) \leq \hat{x}_1$. To summarize, in terms of player 1's share in \hat{x} , player 2's response rule is optimal if and only if $\phi_1(v_1, 0) \leq \hat{x}_1 \leq 1 - \pi_2(w_2)$; this is equivalent to $(\hat{x}_1, 0) \in \hat{A}_1$.

Given player 2 optimally accepts with threshold \hat{x}_2 , this is the lowest immediately accepted offer, and there is no profitable deviation for player 1 if and only if $u_1(\hat{x}_1) \geq U_1(\pi_1(w_1), 1)$, because any deviation to a rejected offer triggers her proposer punishment which has continuation outcome $(y^{(1)}, 0)$ and associated rejection value $U_1(\pi_1(w_1), 1)$; inequality (16) implies $\phi_1(v_1, 0) > \pi_1(w_1)$, whereby $v_1 \geq U_1(\pi_1(w_1), 1)$ from player 1's impatience, and there is no profitable deviation for proposing player 1 whenever there is none for responding player 2. Hence, there is no profitable deviation from $\alpha^{Q,1}(\hat{x}, 0)$ if and only if $(\hat{x}_1, 0) \in \hat{A}_1$.

Next, consider **once delayed agreement outcomes** $(\hat{x}, 1)$; $\alpha^{Q,1}(\hat{x}, 1)$ is a play if and only if $\hat{q}_2 > 0$. Observe that $\hat{q}_2 = 0$ is equivalent to $\phi_1(v_1, 0) = 1$, by equation (17), and jointly with inequality (16) (for $i = 2$), this would indeed mean that \hat{A}_1 contains no delayed agreements at all. Hence it remains to establish the claim for this case under the assumption that $\hat{q}_2 > 0$.

Regarding the second round on the path, the above finding for the case of immediate-agreement outcomes—by mere relabeling—shows that there are then no profitable one-stage deviations if and only if $\phi_2(v_2, 0) \leq \hat{x}_2 \leq 1 - \pi_1(w_1)$. In terms of player 1's share this is equivalent to

$$\pi_1(w_1) \leq \hat{x}_1 \leq 1 - \phi_2(v_2, 0).$$

In the first round $\hat{\sigma}$ specifies that player 2 respond to offers by accepting with threshold \hat{q}_2 . Accepting offers $q \geq \hat{q}_2$ is optimal if and only if accepting offer \hat{q}_2 is optimal, i.e. if $u_2(\hat{q}_2) \geq U_2(x_2^{(2)}, t^{(2)} + 1)$, since the (deviant) rejection of any such offer is followed by continuation outcome $(x^{(2)}, t^{(2)})$. Note that $U_2(x_2^{(2)}, t^{(2)} + 1) = w_2$ from equation (14), and $w_2 \leq U_2(\phi_2(v_2, 0), 1)$ from equation (5); recalling equation (17), if acceptance were not optimal, then $u_2(1 - \phi_1(v_1, 0)) < U_2(\phi_2(v_2, 0), 1)$, which would imply that $\phi_2(v_2, 0) + \phi_1(v_1, 0) > 1$ and there would be no delayed agreement in \hat{A}_1 .

Rejection of all (deviant) offers q such that $0 < q < \hat{q}_2$ is followed by continuation outcome $(y^{(1)}, 0)$ and is optimal by construction, since $\hat{q}_2 > 0$ implies that $u_2(\hat{q}_2) = U_2(1 - \pi_1(w_1), 1)$ is the associated rejection value. Rejecting the zero offer specified for the proposer in this round is optimal if and only if $u_2(0) \leq U_2(\hat{x}_2, 1)$; either $u_2(0) \leq U_2(1, 1)$, in which case $u_2(0) \leq U_2(\hat{x}_2, 1)$ is equivalent to $\hat{x}_1 \leq 1 - \phi_2(u_2(0), 1)$, or $u_2(0) > U_2(1, 1)$, in which case $\phi_2(u_2(0), 1) = 1$ together with inequality (16) (for $i = 1$) implies that \hat{A}_1 contains no delayed agreements.

By equation (17), the initial proposer 1 can obtain at most the value v_1 from making a deviant accepted offer $q \geq \hat{q}_2$; making a deviant rejected offer $q < \hat{q}_2$ yields value $U_1(\pi_1(w_1), 1)$, which is no greater than v_1 due to inequality (16); hence making her supposed (rejected) offer of a zero share is optimal if and only if $v_1 \leq U_1(\hat{x}_1, 1)$. This is equivalent to $\hat{x}_1 \geq \phi_1(v_1, 1)$ unless $v_1 > U_1(\phi_1(v_1, 1), 1)$; however, the latter would imply $\phi_1(v_1, 1) = 1$ and together with inequality (16) (for $i = 2$) would yield that \hat{A}_1 contains no delayed-agreement outcomes. In summary of this case for $\hat{q}_2 > 0$, using that $\pi_1(w_1) < \phi_1(v_1, 1)$ from inequality (16), and noting that $\min\{1 - \phi_2(v_2, 0), 1 - \phi_2(u_2(0), 1)\}$ equals $1 - \max\{\phi_2(v_2, 0), \phi_2(u_2(0), 1)\}$, we obtain there is no profitable deviation if and only if $(\hat{x}_1, 1) \in \hat{A}_1$.

Finally, consider **further delayed agreement outcomes** (\hat{x}, \hat{t}) such that $\hat{t} > 1$; $\alpha^{Q,1}(\hat{x}, \hat{t})$ is a play if and only if $\hat{q}_1 \cdot \hat{q}_2 > 0$. From the previous case we know that if $\hat{q}_2 = 0$ then \hat{A}_1 would not contain any delayed agreement; now note that $\hat{q}_1 = 0$ is equivalent to $\phi_2(v_2, 0) = 1$, by equation (17), and in combination with inequality (16) (for $i = 1$) would imply that \hat{A}_1 contains no agreements delayed by more than one period. Hence it remains to establish the claim for this case under the assumption that $\hat{q}_1 \cdot \hat{q}_2 > 0$.

In the last round of play $\alpha^{Q,1}(\hat{x}, \hat{t})$, which is round $\hat{t} + 1$, we can use the previous findings

to conclude that there is no profitable deviation if and only if

$$\begin{cases} \pi_1(w_1) \leq \hat{x}_1 \leq 1 - \phi_2(v_2, 0) & \hat{t} \text{ odd} \\ \phi_1(v_1, 0) \leq \hat{x}_1 \leq 1 - \pi_2(w_2) & \hat{t} \text{ even} \end{cases}.$$

Consider then play $\alpha^{Q,1}(\hat{x}, \hat{t})$ for any round $n < \hat{t} + 1$, in which player P makes an offer to player R . Optimality of R 's response rule is characterized in a manner similar to optimality of initial respondent 2's response rule when we considered agreement-outcomes with one round of delay; it is therefore characterized by $u_R(0) \leq U_R(\hat{x}_R, \hat{t} + 1 - n)$. Since $U_R(\hat{x}_R, \hat{t} + 1 - n)$ is non-decreasing in n , this yields only two restrictions, namely those for the first two rounds' respondent stages, which are $u_2(0) \leq U_2(\hat{x}_2, \hat{t})$ and $u_1(0) \leq U_1(\hat{x}_1, \hat{t} - 1)$, respectively. These two inequalities are equivalent to

$$\phi_1(u_1(0), \hat{t} - 1) \leq \hat{x}_1 \leq 1 - \phi_2(u_2(0), \hat{t})$$

whenever both $u_2(0) \leq U_2(1, \hat{t})$ and $u_1(0) \leq U_1(1, \hat{t} - 1)$ hold true; otherwise, however, \hat{A}_1 contains no outcome that has agreement delayed by \hat{t} periods.

Again, similar to optimality for initial proposer 1 when we considered one round of delay, proposer P 's zero offer is here optimal if and only if $v_P \leq U_P(\hat{x}_P, \hat{t} + 1 - n)$. Since $U_P(\hat{x}_P, \hat{t} + 1 - n)$ is non-decreasing in n , this yields only two restrictions, namely those for the first two rounds' proposer stages, which are $v_1 \leq U_1(\hat{x}_1, \hat{t})$ and $v_2 \leq U_2(\hat{x}_2, \hat{t} - 1)$, respectively. These two inequalities are equivalent to

$$\phi_1(v_1, \hat{t}) \leq \hat{x}_1 \leq 1 - \phi_2(v_2, \hat{t} - 1)$$

whenever both $v_1 \leq U_1(1, \hat{t})$ and $v_2 \leq U_2(1, \hat{t} - 1)$ hold true; otherwise, however, \hat{A}_1 contains no outcome that has agreement delayed by \hat{t} periods. Now observe that $\phi_1(v_1, \hat{t})$ is at least as large as any of $\pi_1(w_1)$, $\phi_1(v_1, 0)$ or $\phi_1(u_1(0), \hat{t} - 1)$, due to 1's impatience and inequality (16); moreover, also $\phi_2(v_2, \hat{t} - 1)$ is at least as large as both $\phi_2(v_2, 0)$ and $\pi_2(w_2)$ due to 2's impatience and inequality (16). Hence we can summarize this case for $\hat{q}_1 \cdot \hat{q}_2 > 0$ by the condition that (\hat{x}, \hat{t}) is such that

$$\phi_1(v_1, \hat{t}) \leq \hat{x}_1 \leq 1 - \max\{\phi_2(v_2, \hat{t} - 1), \phi_2(u_2(0), \hat{t})\},$$

which is again equivalent to $(\hat{x}_1, \hat{t}) \in \hat{A}_1$.

A similar proof applies to the case of $i = 2$, hence $\{(1 - y_j^{(j)}, 0), (x_i^{(i)}, t^{(i)})\} \subseteq A_i(v_1, w_1, v_2, w_2) \subseteq$

A_i^* , and the lemma's claimed equations are easily verified. \square

Lemma 7. *The following relationships hold true for each $i \in I$:*

$$\tilde{v}_i = u_i (1 - \pi_j (U_j (1 - \pi_i (\tilde{w}_i), 1))) \quad (18)$$

$$\tilde{w}_i \geq \inf \{U_i (\phi_i (\tilde{v}_i, t), t + 1) \mid t \in T, t \leq t_i^*\} \quad (19)$$

$$t_i^* \leq \sup \{t \in T \mid \kappa_i (t, \tilde{v}_i, \tilde{v}_j) \leq 1\} \quad (20)$$

Proof. First, observe that, for each i ,

$$(q, t) \in A_i^* \Rightarrow (1 - \pi_i (U_i (q, t + 1)), 0) \in A_j^*. \quad (21)$$

Let σ be an equilibrium of game G_i which supports i 's personal outcome (q, t) , denote the share $1 - \pi_i (U_i (q, t + 1))$ by \hat{q} and the division such that j 's share equals \hat{q} by \hat{x} . The strategy profile $\hat{\sigma}$ in game G_j such that $\hat{\sigma} (h^0) = (\hat{x}, X_{i,\hat{q}})$ and $\hat{\sigma} (x, h) = \sigma (h)$ for any division x and history h , is an equilibrium supporting j 's personal outcome $(1 - \hat{q}, 0)$: following any initial rejection, $\hat{\sigma}$ specifies equilibrium σ , which induces personal outcome (q, t) for player i and thus implies that the initial response rule of accepting with threshold \hat{q} is optimal for i ; the initial proposer j best-responds by offering this share, because this is the lowest accepted offer and, moreover, satisfies $u_j (1 - \hat{q}) \geq U_j (1 - q, t + 1)$, due to $\hat{q} \leq q$, which follows from i 's impatience, together with the desirability and impatience properties of j 's preferences.

Using this observation, I will now prove all three conditions (18)-(20) for the case of $i = 1$; mere relabeling yields them for $i = 2$.

To show that the pair $(\tilde{v}_1, \tilde{w}_1)$ satisfies equation (18), combine (21) (for $i = 2$) with the fact that any equilibrium of game G_1 must have the initial respondent 2 accept all offers greater than $\sup \{\pi_2 (U_2 (q, t + 1)) \mid (q, t) \in A_2^*\}$, to obtain

$$\tilde{v}_1 = u_1 (1 - \sup \{\pi_2 (U_2 (q, t + 1)) \mid (q, t) \in A_2^*\}).$$

It then remains to prove that $\pi_2 (U_2 (1 - \pi_1 (\tilde{w}_1), 1)) = \sup \{\pi_2 (U_2 (q, t + 1)) \mid (q, t) \in A_2^*\}$. For this, also combine (21) (now for $i = 1$) with the fact that any equilibrium of G_2 must have the initial respondent 1 reject all offers less than $\pi_1 (\tilde{w}_1)$, which yields that

$$1 - \pi_1 (\tilde{w}_1) = \sup \{q \in [0, 1] \mid (q, 0) \in A_2^*\}.$$

Now observe that any $(q, t) \in A_2^*$ with $t > 0$ satisfies $U_1 (1 - q, t) \geq \tilde{w}_1$, which implies $1 - q \geq \pi_1 (U_1 (1 - q, t)) \geq \pi_1 (\tilde{w}_1)$ by 1's impatience and the non-decreasingness of π_1 , and

therefore

$$\pi_2(U_2(q, t+1)) \leq \pi_2(U_2(1 - \pi_1(\tilde{w}_1), t+1)) \leq \pi_2(U_2(1 - \pi_1(\tilde{w}_1), 1))$$

by the desirability and impatience properties of 2's preferences, together with the non-decreasingness of π_2 .

Regarding the proof that $(\tilde{v}_1, \tilde{w}_1, t_1^*)$ satisfies inequality (19), simply note that $(q, t) \in A_1^*$ implies $U_1(q, t) \geq \tilde{v}_1$ by the definition of \tilde{v}_1 , and thus $q \geq \phi_1(\tilde{v}_1, t)$; the claim then follows from the desirability property of 1's preferences.³⁹

Inequality (20) certainly holds true if $t_1^* = 0$; for the case of $t_1^* > 0$, note that $(q, t) \in A_1^*$ implies both $U_1(q, t) \geq \tilde{v}_1$ and $U_2(1 - q, t) \geq u_2(0)$. These two inequalities imply, respectively, that $q \geq \phi_1(\tilde{v}_1, t)$ and $1 - q \geq \phi_2(u_2(0), t)$. Moreover, if $(q, t) \in A_1^*$ with $t > 0$ also implies that $(q, t - 1) \in A_2^*$, hence $U_2(1 - q, t - 1) \geq \tilde{v}_2$, and thus $1 - q \geq \phi_2(\tilde{v}_2, t - 1)$. Altogether, for any $t > 0$ there exists a share q such that $(q, t) \in A_1^*$ only if $\kappa_1(t, \tilde{v}_1, \tilde{v}_2) \leq 1$, concluding the proof. \square

Lemma 8. *There exist values $(v_i, w_i, t_i)_{i \in I} \in E$ such that $v_i \leq \tilde{v}_i$, $w_i \leq \tilde{w}_i$ and $t_i \geq t_i^*$ for both $i \in I$.*

Proof. Consider the following sequence $(v_i^n, w_i^n, t_i^n)_{i \in I}$: $(w_1^1, w_2^1) \equiv (\tilde{w}_1, \tilde{w}_2)$ and, for any $n \in \mathbb{N}$ and each i ,

$$\begin{aligned} v_i^n &\equiv u_i(1 - \pi_j(U_j(1 - \pi_i(w_i^n), 1))) \\ t_i^n &\equiv \sup \{t \in T \mid \kappa_i(t, v_i^n, v_j^n) \leq 1\} \\ w_i^{n+1} &\equiv \inf \{U_i(\phi_i(v_i^n, t), t+1) \mid t \in T, t \leq t_i^n\}. \end{aligned}$$

Note that $v_i^1 = \tilde{v}_i$ and $t_i^1 \geq t_i^*$, by lemma 7. It is straightforward that $w_i^{n+1} \leq w_i^n$, $v_i^{n+1} \leq v_i^n$ and $t_i^{n+1} \geq t_i^n$. I will establish the claim by proving that the sequence $(v_i^n, w_i^n, t_i^n)_{i \in I}$ possesses a limit in E .

The first step is to prove that the sequence (w_1^n, w_2^n) converges: since each component sequence w_i^n is non-increasing and bounded from below by $U_i(0, \infty) \in \mathbb{R}$, it converges. Denoting this limit by (\hat{w}_1, \hat{w}_2) , the continuity properties of the functions involved imply the

³⁹Under the weakening of desirability suggested in fn. 10, the observation $\tilde{v}_1 > u_1(0)$ from (16) means that no equilibrium delay t can be such that player 1 does not care about her share: otherwise, there would exist $(q, t) \in A_1^*$ with $U_1(q, t) = U_1(0, t)$, but $U_1(0, t) \leq u_1(0)$ by impatience; hence $U_1(q, t) < \tilde{v}_1$, a contradiction.

following convergence properties of the sequences v_i^n and t_i^n , for each i :

$$\begin{aligned} v_i^n &\rightarrow u_i(1 - \pi_j(U_j(1 - \pi_i(\hat{w}_i), 1))) \equiv \hat{v}_i \\ t_i^n &\rightarrow \sup \{t \in T \mid \kappa_i(t, \hat{v}_i, \hat{v}_j) \leq 1\} \equiv \hat{t}_i; \end{aligned}$$

i.e., $(\hat{v}_i, \hat{w}_i, \hat{t}_i)_{i \in I} \in E$. □

A.4 Corollary 1

Proof. Lemma 6 implies that E 's being a singleton is necessary for equilibrium uniqueness.

Concerning its sufficiency, the proof of lemma 4 shows that whenever E is a singleton, its unique element $(v_i^*, w_i^*, t_i^*)_{i \in I}$ equals $(u_i(\hat{q}_i), U_i(\hat{q}_i, 1), 0)_{i \in I}$, for \hat{q}_1 the unique fixed point of f_1 and $\hat{q}_2 \equiv 1 - \pi_1(U_1(\hat{q}_1, 1))$. Characterization theorem 1 then implies that each A_i^* equals the singleton $\{(\hat{q}_i, 0)\}$. Consider then any round in which player P makes an offer to responding player R : since any equilibrium has the outcome that offer \hat{q}_P is accepted, it must be that P indeed offers \hat{q}_P , and that R accepts this offer. Since any equilibrium has the same continuation outcome with R 's associated rejection value equal to $U_R(\hat{q}_R, 1)$, any optimal response rule must have R accept any offer $q > \pi_R(U_R(\hat{q}_R, 1))$ as well as reject any offer $q < \pi_R(U_R(\hat{q}_R, 1))$. This pins down a unique equilibrium that is, moreover, stationary. □

A.5 Lemma 3

This lemma will be proven based on the following characterization of stationary equilibrium, which establishes a one-to-one relationship between stationary equilibria and fixed points of f_1 (defined by equation (12) to prove lemma 4, as part of theorem 1). Note that in terms of the players' impatience (3.) the characterization of stationary equilibrium relies only on property (3.b), players' attitudes to delay beyond a single (first) period are irrelevant.

Lemma 9. *The profile of stationary strategies $(x^{(i)}, Y^{(i)})_{i \in I}$ is an equilibrium if and only if*

$$\left\{ \begin{array}{l} x_1^{(1)} = f_1(x_1^{(1)}) \\ x_2^{(2)} = 1 - \pi_1(U_1(x_1^{(1)}, 1)) \end{array} \right\}, \text{ and for each } i \in I, Y^{(i)} = X_{i, x_i^{(j)}}.$$

A stationary equilibrium exists, and it is unique if and only if f_1 has a unique fixed point.

Proof. First, note that any equilibrium, hence any stationary equilibrium, has agreement, since $v_i^* > u_i(0) \geq U_i(0, \infty)$ from inequality (16). Consider then a stationary equilibrium $(x^{(i)}, Y^{(i)})_{i \in I}$. If $x^{(1)} \notin Y^{(2)}$ then its outcome in G_1 must be $(x^{(2)}, 1)$. Because this outcome

obtains irrespective of play in the initial round of G_1 , responding player 2 must accept any proposal x with $x_2 > \pi_2(U_2(x_2^{(2)}, 1))$. Player 2's impatience property (3.b) implies that either (i) $\pi_2(U_2(x_2^{(2)}, 1)) < x_2^{(2)}$ or (ii) $\pi_2(U_2(x_2^{(2)}, 1)) = x_2^{(2)} = 0$. In case of (i) there exist values $\epsilon > 0$ such that $\epsilon < x_2^{(2)} - \pi_2(U_2(x_2^{(2)}, 1))$, and any of them satisfy

$$u_1(1 - \pi_2(U_2(x_2^{(2)}, 1)) - \epsilon) > U_1(1 - \pi_2(U_2(x_2^{(2)}, 1)) - \epsilon, 1) \geq U_1(1 - x_2^{(2)}, 1)$$

by impatience property (3.b) and desirability of player 1's preferences, applied in this sequence. In case of (ii), impatience property (3.b) together with continuity of player 1's preferences imply existence of $\epsilon > 0$ such that $u_1(1 - \epsilon) > U_1(1, 1)$. In any case player 1 can therefore propose immediately accepted divisions that yield a value greater than that from proposing $x^{(1)}$, contradicting equilibrium. After a symmetric argument, it is then proven that $x^{(i)} \in Y^{(j)}$ for both $i \in I$.

Given this immediate-agreement property of stationary equilibrium, by desirability, (i) a responding player j must accept any offered share $x_j > \pi_j(U_j(x_j^{(j)}, 1))$ as well as reject any $x_j < \pi_j(U_j(x_j^{(j)}, 1))$, and (ii) there cannot exist a proposal x by player i with $x_i > x_i^{(i)}$ such that $x \in Y^{(j)}$, whereby

$$x_i^{(i)} = 1 - \pi_j(U_j(x_j^{(j)}, 1)) \text{ and } Y^{(i)} = X_{i, x_i^{(i)}},$$

and substituting the expression for $x_2^{(2)}$ into that for $x_1^{(1)}$ yields $x_1^{(1)} = f_1(x_1^{(1)})$, establishing necessity. Sufficiency is easily verified, and its proof omitted here.

Existence of a fixed point of f_1 and hence stationary equilibrium is established by the proof of lemma 4, and the characterization shows that there are as many distinct stationary equilibria as there are fixed points of f_1 . \square

Lemma 3 follows from combining the above characterization with the next result.

Lemma 10. *If both players' preferences exhibit immediacy, then f_1 has a unique fixed point.*

Proof. Suppose that player i 's preferences exhibit immediacy, take any share q and any $\epsilon > 0$ such that $q + \epsilon \leq 1$, and consider various possible cases to establish that $l_i(q) \equiv q - \pi_i(U_i(q, 1))$ is increasing. First, if $U_i(q + \epsilon, 1) \leq u_i(0)$, then also $U_i(q, 1) \leq u_i(0)$ and $l_i(q) = q < q + \epsilon = l_i(q + \epsilon)$. Second, if $U_i(q, 1) \leq u_i(0) < U_i(q + \epsilon, 1)$, then continuity and impatience imply existence of a share $q' \in [q, q + \epsilon)$ such that $U_i(q', 1) = u_i(0)$; letting $\epsilon' \equiv q + \epsilon - q'$, immediacy implies $u_i(\epsilon') > U_i(q' + \epsilon', 1) \equiv U_i(q + \epsilon, 1)$, and hence $l_i(q + \epsilon) > q + \epsilon - \epsilon' \geq q = l_i(q)$. Finally, if $u_i(0) < U_i(q, 1)$, then continuity and impatience imply

existence of a share $q' \in (0, q)$ such that $u_i(q') = U_i(q, 1)$; immediacy implies $u_i(q' + \epsilon) > U_i(q + \epsilon, 1)$, and hence $l_i(q + \epsilon) > q + \epsilon - (q' + \epsilon) = l_i(q)$.

Consider then the following difference:

$$\begin{aligned} q - f_1(q) &= q - 1 + \pi_2(U_2(1 - \pi_1(U_1(q, 1)), 1)) \\ &= [q - \pi_1(U_1(q, 1))] - [(1 - \pi_1(U_1(q, 1))) - \pi_2(U_2(1 - \pi_1(U_1(q, 1)), 1))]. \\ &\equiv l_1(q) - l_2(1 - \pi_1(U_1(q, 1))). \end{aligned}$$

If l_i is increasing for both i , then l_1 is increasing in q and l_2 is increasing in $1 - \pi_1(U_1(q, 1))$. Since $1 - \pi_1(U_1(q, 1))$ is non-increasing in q , overall the two terms' difference is increasing in q , and $q - f_1(q)$ has at most one root; by existence of a fixed point, established earlier, it has exactly one. \square

A.6 Proposition 1

Proof. As a first step, I will show the following: if $w_i^* = U_i(\phi_i(v_i^*, 0), 1)$ for both $i \in I$, then equilibrium is unique if and only if stationary equilibrium is unique. Theorem 1 implies that the outcome $(x^{R,i}, 0)$ such that $x_i^{R,i} = \phi_i(v_i^*, 0)$ is an optimal respondent punishment outcome for player i , and that her optimal proposer punishment therefore has outcome $(x^{P,i}, 0)$ such that $x_i^{P,i} = \pi_i(U_i(\phi_i(v_i^*, 0), 1))$. Using equation (4),

$$\begin{aligned} \phi_i(v_i^*, 0) &= 1 - \pi_j(U_j(1 - \pi_i(U_i(\phi_i(v_i^*, 0), 1)), 1)) \\ &= f_i(\phi_i(v_i^*, 0)), \end{aligned}$$

which, by lemma 9, reveals that $x_i^{R,i} = f_i(x_i^{R,i})$ as well as $x_j^{P,i} \equiv 1 - x_i^{P,i} = 1 - \pi_i(U_i(x_i^{R,i}, 1))$ are the two players' respective proposer shares in one particular stationary equilibrium. If there is a unique stationary equilibrium, then $(x^{R,1}, 0) = (x^{P,2}, 0)$ and $(x^{P,1}, 0) = (x^{R,2}, 0)$ such that $x_1^{R,1} = x_1^{P,2} = \hat{q}_1$ and $x_2^{P,1} = x_2^{R,2} = 1 - \pi_1(U_1(\hat{q}_1, 1))$ for $\hat{q}_1 = \phi_1(v_1^*, 0)$ the unique fixed point of f_1 . Letting $\hat{q}_2 \equiv 1 - \pi_1(U_1(\hat{q}_1, 1))$, theorem 1 then says that $(v_i^*, w_i^*, t_i^*)_{i \in I} = (u_i(\hat{q}_i), U_i(\hat{q}_i, 1), 0)_{i \in I}$, and $A_i^* = \{(\hat{q}_i, 0)\}$, so uniqueness of equilibrium follows from the argument in the proof of corollary 1. This proves sufficiency. Necessity holds trivially.

The second step shows that $w_i^* = U_i(\phi_i(v_i^*, 0), 1)$ follows whenever a player i 's preferences exhibit a weak present bias. This establishes the proposition, because under immediacy stationary equilibrium is indeed unique. The proof of lemma 5 and theorem 1 imply a finite delay \bar{t}_i such that

$$w_i^* = \min \left\{ U_i(\phi_i(v_i^*, t), t + 1) \mid t \in T, t \leq \bar{t}_i \right\},$$

where $U_i(\phi_i(v_i^*, t), t) = v_i^*$ holds true for any $t \leq \bar{t}_i$. A weak present bias then implies that $U_i(\phi_i(v_i^*, 0), 1) \leq U_i(\phi_i(v_i^*, t), t+1)$ for all such t , and hence $w_i^* = U_i(\phi_i(v_i^*, 0), 1)$, proving the claim. \square

A.7 Proposition 2

Proof. The proposition holds trivially for immediate-agreement equilibria. Suppose therefore that (\hat{x}, \hat{t}) with $\hat{t} > 0$ is an equilibrium outcome of game G_1 ; the case of game G_2 follows from mere relabeling. Theorem 1 implies that \hat{x} is an interior division, since

$$0 < \phi_1(v_1^*, \hat{t}) \leq \hat{x}_1 \leq \max\{\phi_2(v_2^*, \hat{t}-1), \phi_2(u_2(0), \hat{t})\} < 1. \quad (22)$$

For every round $n \leq \hat{t} + 1$, define each player i 's reservation share for the rejection value corresponding to agreement on \hat{x} with remaining delay $\hat{t} + 1 - n$: $\pi_i^n \equiv \pi_i(U_i(\hat{x}_i, \hat{t} + 1 - n))$. The inequalities in (22) imply $u_i(\pi_i^n) = U_i(\hat{x}_i, \hat{t} + 1 - n)$ because of $U_i(\hat{x}_i, \hat{t} + 1 - n) \geq u_i(0)$, and the stronger impatience property assumed in the proposition yields that π_i^n is increasing, since $\hat{x}_i > 0$.

Define a play as follows: in round 1, player 1 offers a share of $b_1^1 = 0$, and player 2 accepts with threshold $1 - b_2^1$ such that $b_2^1 = \phi_1(v_1^*, 0)$; in round n such that $1 < n < \hat{t} + 1$, player $P(n)$ offers a share of $b_{P(n)}^n = \frac{1}{2}(b_{P(n)}^{n-1} + \pi_{R(n)}^n)$ and player $R(n)$ accepts with threshold $1 - b_{R(n)}^n$ such that $b_{R(n)}^n = \frac{1}{2}(b_{R(n)}^{n-1} + \pi_{P(n)}^n)$, with the sole exception that $b_1^2 = \phi_2(v_2^*, 0)$; in round $n = \hat{t} + 1$, player $P(n)$ offers a share $b_{P(n)}^n = \hat{x}_{P(n)}$ and player $R(n)$ accepts with threshold $1 - b_{R(n)}^n$ such that $b_{R(n)}^n = \hat{x}_{R(n)}$.

First, verify that each sequence $(b_i^n)_{n=1}^{\hat{t}+1}$ is increasing since $b_i^{n-1} < \pi_j^n$: this is true for $n - 1 = 1$, because $b_i^1 \leq \pi_j^1 < \pi_j^2$, and if it is true for $n - 1 \geq 1$ such that $n < \hat{t} + 1$, it is true for n , because $b_i^n = \frac{1}{2}(b_i^{n-1} + \pi_j^n) < \pi_j^n < \pi_j^{n+1}$. Second, observe that $b_{P(n)}^n < 1 - b_{R(n)}^n$ for all $n < \hat{t} + 1$: since $\pi_1^n + \pi_2^n < 1$ for all such n , this follows from $b_i^n \leq \pi_j^n$; hence this indeed defines a play with outcome (\hat{x}, \hat{t}) .

The final step is to show that this defines equilibrium play. Taken then any strategy profile σ of game G_1 such that $\langle \sigma \rangle$ equals the above play (clearly, one exists) and define the strategy profile $\hat{\sigma} \equiv \sigma^* \left(\sigma \left| \left(\sigma^{P,i}, \sigma^{R,i} \right)_{i \in I} \right. \right)$, where $(\sigma^{P,i}, \sigma^{R,i})_{i \in I}$ is an OPC, as in proposition 4, part (i). Hence $\langle \hat{\sigma} \rangle = \langle \sigma \rangle$ and $\hat{\sigma}$ is an equilibrium if and only if there are no profitable one-stage deviations from its play $\langle \hat{\sigma} \rangle$.

Consider then any round $n \leq \hat{t} + 1$ of play $\langle \hat{\sigma} \rangle$. Rejecting an offer $q \geq 1 - b_{R(n)}^n$ is no better than accepting it for $R(n)$, since it yields the minimal credible rejection value $w_{R(n)}^*$

due to optimal punishment, but

$$w_{R(n)}^* \leq U_{R(n)}(\hat{x}_{R(n)}, \hat{t} + 1 - n) = u_{R(n)}(\pi_{R(n)}^n) \leq u_{R(n)}(1 - \pi_{P(n)}^n) \leq u_{R(n)}(1 - b_{R(n)}^n),$$

using that $(\hat{x}, \hat{t} - n)$ is a continuation equilibrium outcome (by assumption), that $\pi_1^n + \pi_2^n \leq 1$ and that $b_{R(n)}^n \leq \pi_{P(n)}^n$; accepting an offer $q < 1 - b_{R(n)}^n$ such that $q \neq b_{P(n)}^n$ is no better than rejecting it, since

$$u_{R(n)}(1 - b_{R(n)}^n) \leq u_{R(n)}(1 - \phi_{P(n)}(v_{P(n)}^*, 0)) = U_{R(n)}(1 - \pi_{P(n)}(w_{P(n)}^*), 1),$$

using that any *responding* player i 's concession is at least $\phi_j(v_j^*, 0)$, by construction, and theorem 1, which shows that continuation with optimal punishment of a proposing player i has rejection value $U_j(1 - \pi_i(w_i^*), 1)$ for respondent j , and that this is equal to $u_j(1 - \phi_i(v_i^*, 0))$; finally, accepting offer $q = b_{P(n)}^n < 1 - b_{R(n)}^n$, which can only be the case for $n < \hat{t} + 1$, is no better than rejecting it, since

$$u_{R(n)}(b_{P(n)}^n) \leq u_{R(n)}(\pi_{R(n)}^n) = U_{R(n)}(\hat{x}_{R(n)}, \hat{t} + 1 - n).$$

Consider then the proposer's incentives, given the respondent's behavior and punishments for deviations: the minimal offer which the respondent accepts equals $b_{R(n)}^n$, which is no greater than $\pi_{P(n)}^n$, whereby

$$u_{P(n)}(b_{R(n)}^n) \leq u_{P(n)}(\pi_{P(n)}^n) = U_{P(n)}(\hat{x}_{P(n)}, \hat{t} + 1 - n),$$

so there is no profitable deviation to any (alternative) accepted offer; any other deviant offer has (rejection) value $U_{P(n)}(\pi_{P(n)}(w_{P(n)}^*), 1)$ which is no greater than $v_{P(n)}^*$ by theorem 1, and since $U_{P(n)}(\hat{x}_{P(n)}, \hat{t} + 1 - n) \geq v_{P(n)}^*$, because (\hat{x}, \hat{t}) is an equilibrium outcome, there is no profitable deviation to a rejected offer either. \square

A.8 Proposition 3

Proof. Omitting player indices due to symmetry, by theorem 1, if there exists an equilibrium with agreement delayed $t > 0$ periods, then $\kappa(t, v^*, v^*) \leq 1$. This implies that $\phi(v^*, t') \leq \frac{1}{2}$

for all $t' < t$, since

$$\begin{aligned} \kappa(t, v^*, v^*) \leq 1 &\Leftrightarrow \underbrace{\phi(v^*, t)}_{\geq \phi(v^*, t-1)} + \underbrace{\max\{\phi(v^*, t-1), \phi(u(0), t)\}}_{\geq \phi(v^*, t-1)} \leq 1 \\ &\Rightarrow \phi(v^*, t-1) \leq \frac{1}{2}, \end{aligned}$$

and $\phi(v^*, \cdot)$ is non-decreasing. Using again theorem 1, recalling also that $\pi(w^*) \leq \phi(v^*, 0)$, $(\frac{1}{2}, t') \in A^*$ follows for all $t' < t$. \square

B Supplementary Material

B.1 Multiplicity and Delay under Weak Present Bias

Supplementing section 5.2.2, I here present two examples of how violations of immediacy result in multiplicity and, possibly, also delay. The first is one of dynamically consistent preferences (ED) and was presented already by Rubinstein (1982, concl. I). To the best of my knowledge, its set of equilibria has not yet been explicitly characterized, however.

Example 4. Let the two parties' preferences be given by $U_i(q, t) = q - ct$, for $c \in (0, 1)$. Due to preference symmetry, player indices are omitted in what follows. The preferences are covered by assumption 1 once $U(0, \infty) \equiv -\infty$ is specified; in particular, impatience property (3.c) is satisfied: $U(1, t)$ tends to minus infinity, whereas $u(0) = 0$.⁴⁰ In the assumed absence of uncertainty, they actually satisfy ED, albeit with “strongly” convex instantaneous utility: $U(q, t) = \ln(\delta^t u(q))$ for $\delta \equiv \exp(-c)$ and $u(q) \equiv \exp(q)$. Hence they exhibit a weak present bias but violate immediacy (increasing shares by the same amount leaves indifferent).⁴¹

This results in a multiplicity of stationary equilibrium: any $q \in [c, 1]$ is a proposer's equilibrium share in some stationary equilibrium (with immediate agreement, of course). Applying the characterization of theorem 1, $v^* = c$ and $w^* = 0$, where both of these minimal proposer and rejection values correspond to a player's least preferred stationary equilibrium. Using these two least preferred stationary equilibria as optimal punishments, non-stationary delay equilibria can be constructed, and equation (6) offers a formula to compute the maximal

⁴⁰ U violates the requirement of assumption 1 that $U(0, \infty) \in \mathbb{R}$, but the positive monotonic transformation $\exp(U)$ represents the same preferences and satisfies also this property.

⁴¹One may interpret such preferences as there being a cost to bargaining. To justify the non-negativity of each player's share in any proposal, assume then that players have an “outside option” of leaving the bargaining table forever, which is equivalent to obtaining a zero share immediately.

such delay for any $c \in (0, 1)$:

$$\begin{aligned}
\kappa(t, c, c) &= \min\{c + ct, 1\} + \min\{ct, 1\} \\
&= \begin{cases} (2t + 1)c & t \leq \frac{1-c}{c} \\ 1 + ct & \frac{1-c}{c} \leq t \leq \frac{1}{c} \\ 2 & \frac{1}{c} \leq t \end{cases} \\
\Rightarrow t^* &= \sup\{t \in T \mid \kappa(t, c, c) \leq 1\} \\
&= \max\left\{t \in T \mid t \leq \frac{1}{2} \cdot \frac{1-c}{c}\right\} \\
&= \left\lfloor \frac{1}{2} \cdot \frac{1-c}{c} \right\rfloor.
\end{aligned}$$

For instance, if $c = \frac{1}{100}$, so that the cost per bargaining round equals one percent of the surplus per player, then the maximal equilibrium delay is 49 periods, with an associated efficiency loss of 98 percent of the surplus. To determine the values of c for which delayed agreement is an equilibrium outcome, simply solve $\kappa(1, c, c) \leq 1$ for c , yielding $c \leq \frac{1}{3}$. The set of equilibrium divisions with a given delay $t \leq t^*$ in game G_1 equals $\{x \in X \mid c + ct \leq x_1 \leq 1 - ct\}$ and is monotonically shrinking in t .

The second example is one of dynamically inconsistent preferences (with an actual present bias) that are non-separable, following the magnitude-effects model of [Noor \(2011\)](#).

Example 5. Let the two parties' preferences be symmetrically given by $U_i(q, t) = \delta(q)^t \cdot u(q)$ with $\delta(q) = 0.5 + 0.49 \cdot q^{0.5}$ and $u(q) = q^{0.5}$.

While both $U_i(q, 0) = q^{0.5}$ and $U_i(q, 1) = 0.5 \cdot q^{0.5} + 0.49 \cdot q$ are concave, these preferences violate immediacy; e.g., once delayed share $q' = 0.75$ is indifferent to immediate share $q \approx 0.64$, but upon increasing both by $\epsilon = 0.05$ the delayed one is preferred. Equations (4) and (5) for $t_i = 0$ have here three solutions, all of which correspond to a (symmetric) stationary equilibrium, with respective proposer shares 0.04, 0.57 and 0.98. (All numbers are rounded.) These different stationary equilibria can be used as (non-stationary) threats to support further equilibrium outcomes.

Indeed, given weak present bias (see footnote 34), the extreme *stationary* equilibria deliver the extreme equilibrium values; hence, they constitute optimal punishments supporting *all* equilibrium outcomes. Here the smallest stationary-equilibrium proposer share equals 0.04, and any *immediate* division with the initial proposer's share between this smallest amount and the largest stationary share of 0.98 can be supported. For any such division x , it can easily be verified that the following is an equilibrium: player 1 initially proposes division x , player 2 accepts with threshold x_2 , and in case of a rejection, (i) if the initial offer was

less than x_2 , the players continue with the stationary equilibrium in which player 2, as the proposer of round 2, receives the largest credible share of 0.98, and (ii) if the initial offer was at least x_2 , the players continue with the stationary equilibrium in which player 2, as the proposer of round 2, receives the smallest credible share of 0.04.

Computing all other equilibrium outcomes is straightforward using the indifference property (due to preference symmetry, player indices are omitted in what follows): for a single period of delay, the delayed share indifferent to the smallest immediate share of 0.04 equals 0.10, and the surplus cost κ of this delay therefore equals $0.04 + 0.10 = 0.14$, which is feasible. Hence, any once-delayed division with the initial proposer's share between 0.10 and $1 - 0.04 = 0.96$ can be supported. Let player 2 be the initial proposer and take any such division x ; it can easily be verified that the following is an equilibrium: player 2 initially demands the entire surplus (offers zero), player 1 accepts with threshold 0.96, and in case of a rejection, (i) if the initial offer was zero, then the players continue with the immediate-agreement equilibrium described above for division x , and (ii) if the initial offer was positive, then the players continue with the stationary equilibrium in which player 1, as the proposer of round 2, receives the largest credible share of 0.98.

Continuing this way until the surplus cost of the delay becomes infeasible—i.e., $\kappa > 1$ —we can describe the set of equilibrium divisions for any feasible delay. The maximal delay t^* equals seven rounds, and the set of equilibrium divisions with this delay equals that of all divisions with the initial proposer's share between 0.48 and $1 - 0.43 = 0.57$.

B.2 Unbounded Equilibrium Delay

The following example slightly modifies example 3 to exhibit unbounded equilibrium delay.

Example 6. Let the two players' preferences be symmetrically given by $U_i(q, t) = d(t) \cdot q$ with

$$d(t) = \begin{cases} \delta^t & t \leq \tau \\ \gamma\delta^{\tau+1} & t > \tau \end{cases}, \quad (\delta, \gamma) \in (0, 1)^2 \text{ and } \tau > 0.$$

Due to preference symmetry, the player subscript is again omitted in what follows.

The difference to example 3 is that delays beyond horizon $\tau + 1$ are not discounted. Observe, however, that $\Delta(t)$ equals δ for all $t \leq \tau$ and $\gamma\delta$ for all $t > \tau$, exactly as in example 3. Hence, whenever there is an equilibrium in which agreement is delayed by τ periods, $v^* = \frac{1-\delta}{1-\gamma\delta^2}$ and $w^* = \gamma\delta v^*$, as was found there.

The absence of discounting beyond a delay of $\tau + 1$ periods implies that equilibrium delay




Figure 3: Graphs regarding unbounded equilibrium delay in example 3. The panel on the left shows the parametric regions (δ, γ) such that equilibrium delay is unbounded for three given values of τ , which are 1 (blue, brown and green), 25 (brown and green) and 1000 (green). The panel on the right illustrates how the respective parametric regions for existence of delay equilibria (superset, bounded by solid line) and unbounded equilibrium delay (subset, bounded by dashed line) are related for the case of $\tau = 50$.

is unbounded if and only if $1 \geq \kappa(\tau + 2, v^*, v^*) = 2 \frac{v^*}{\gamma \delta^{\tau+1}}$, which reduces to

$$\delta^\tau \geq \frac{2}{\gamma \delta} \cdot \frac{1 - \delta}{1 - \gamma \delta^2} \quad (23)$$

after substituting for v^* . Notice that this inequality is more stringent than example 3's inequality (8), which shows when delay equilibria exist; in particular, $\gamma > 0$ is here required. Indeed, γ might be too low: despite existence of an equilibrium with delay τ , which fully determines the optimal punishments, proposing players would then require too large a compensation for longer delays, as those would involve additional discounting through γ . Nonetheless, for any given $\tau > 0$ and $\gamma < 1$, there again exist large enough values of δ such that also inequality (23) is satisfied, with the set of parameters γ and τ such that equilibrium delay is unbounded expanding as δ increases. Figure 3 illustrates this.