On Time-Inconsistency in Bargaining*
JOB-MARKET PAPER

Sebastian Kodritsch†

November 20, 2012

Abstract

This paper analyses dynamically inconsistent time preferences in Rubinstein's (1982) seminal model of bargaining. When sophisticated bargainers have time preferences that exhibit a form of "present bias"—satisfied by the hyperbolic and quasi-hyperbolic time preferences increasingly common in the economics literature—equilibrium is unique and lacks delay. However, when one bargainer is more patient about a single period's delay from the present than one that occurs in the near future, the game permits a novel form of equilibrium multiplicity and delay. Time preferences with this property have most recently been empirically documented; they also arise when parties who weight probabilities non-linearly bargain under the shadow of exogenous breakdown risk, as well as in settings of intergenerational bargaining with imperfect altruism. The paper's main contributions are (i) a complete characterisation of the set of equilibrium outcomes and payoffs for separable time preferences, and (ii) present bias as a readily interpretable sufficient condition for uniqueness at the level of individual preferences.

Keywords: bargaining, time preference, dynamic inconsistency, delay
JEL classification: C78, D03, D74

---

*First and foremost, I thank my supervisor Erik Eyster; I feel priviledged for the numerous stimulating conversations with him and his enormous academic support. This work has also greatly benefitted from the benign challenges that Balázs Szentes has posed to me, for which I am deeply grateful. I acknowledge illuminating discussions with Francesco Nava, Michele Piccione and Matthew Rabin as well as helpful remarks by Patrick Blanchenay, Can Celiktemur, Leonardo Felli, Gilat Levy, Matthew Levy, John Moore, and Ronny Razin.

†Department of Economics, and Suntory and Toyota International Centres for Economics and Related Disciplines (STICERD), London School of Economics and Political Science (LSE). Email: s.kodritsch@lse.ac.uk.

1 Introduction

As a mechanism to distribute economic surplus, bargaining is pervasive in real economic exchange and accordingly fundamental to the economic analysis of contracts. In the absence of irrevocable commitments, time becomes a significant variable of bargaining agreements; parties may not only agree now or never but also sooner or later. At the heart of economists' understanding of how the bargaining parties' "time preferences" shape the agreement they will reach, lies the so-called strategic approach to bargaining, which was pioneered by Stahl (1972) and firmly established in economics by the seminal work of Rubinstein (1982). Building on Stahl's disciplined formal description of the bargaining process as one where parties alternate in making and answering proposals, and extending it to a situation without an exogenous deadline, Rubinstein (1982) reaches surprisingly sharp conclusions about how two completely informed and impatient parties share an economic surplus: there is a unique subgame-perfect Nash equilibrium with the properties that (i) agreement is reached immediately, (ii) a player's "bargaining power" is her tolerance of a round's delay, and (iii) the initial proposer enjoys a strategic advantage. Moreover, this equilibrium has a simple—"stationary"—structure: whenever it is her turn in the respective role, a player always makes the very same offer and follows the very same acceptance rule, and in any round the offer of the proposer equals the smallest share the respondent accepts, given that upon a rejection the roles are reversed and the same property holds true.

Rubinstein derives these results for players whose time preferences satisfy exponential discounting.[1] Within the past fifteen years, however, a large body of evidence challenging this assumption has received attention in economics. In numerous empirical studies, surveyed by Frederick et al. (2002), psychologists have measured periodic discount rates which are declining in delay, a finding termed "decreasing impatience" or hyperbolic discounting. Based on this evidence, Laibson (1997) introduced the (β, δ)-model of quasi-hyperbolic discounting, which emphasises a distinct time preference for the very short run, governed by parameter β, for intertemporal trade-offs involving the immediate present; the long-run time preference over prospects which are in the future is assumed to satisfy exponential discounting with parameter δ. Most recently, and in response to this model's success, experimental economists, studying the particular domain of single-dated monetary rewards, have produced both (i) defence of exponential discounting, e.g. Andreoni and Sprenger (2012), and (ii) further qualification of its violations for short delays, where increasing impatience has been observed, e.g. Takeuchi (2011).

In view of this evidence, one is naturally led to wonder whether the aforementioned results, upon which economists' understanding of bargaining is based, remain valid once time preferences

[1] Although Rubinstein (1982) neither derives nor employs a utility representation, Fishburn and Rubinstein (1982) show that the axioms imposed on the players' preferences imply exponential discounting (see also Osborne and Rubinstein (1990, Section 3.3)).

take "richer" forms than exponential discounting. More specifically, when is there still a *unique prediction*, and *which discount factor* matters in this case? This question is all the more important given the increased interest in economic applications using non-exponential time preferences. Or may players' dynamic inconsistency, which results once exponential discounting is violated, invite *multiplicity and non-stationary equilibria*? Is there a meaningful notion of *"bargaining power"* more generally? This paper addresses all of these questions.

I revisit the Rubinstein (1982) model for general separable and time-invariant preferences: for each of the two players $i \in \{1, 2\}$, there exist a decreasing *discount function* d_i and a continuous increasing *utility function* u_i such that, *at any time* during the process of bargaining, i's preferences over divisions $x = (x_1, x_2)$ agreed upon with a (further) delay of t (bargaining) periods have a *separable* utility representation

$$U_i(x,t) = d_i(t)\, u_i(x_i) \tag{1}$$

While the main part of the paper, including this introduction, presents and discusses results for the special case of linear utility, $u_i(x_i) = x_i$, the appendix analogously deals with the above preferences in general.[2] A player i's time-invariant preferences with representation 1 are time-consistent if and only if $d_i(t) = \delta_i^t$ for some constant $\delta_i \in (0, 1)$, i.e. they satisfy exponential discounting. The standard equilibrium concept for games played by time-inconsistent players is *Strotz-Pollak equilibrium (SPE)*. It assumes that a player cannot commit to future actions and, accordingly, requires robustness against one-shot deviations only; as is important for comparability, it coincides with subgame perfect Nash equilibrium in the case of time-consistent players. Not a single departure from exponential discounting has been analysed without the restriction to stationary equilibrium (see section 1.2). However, as I argue in section 3, the *assumption* of stationarity is particularly problematic under time-inconsistency: a stationary bargaining strategy is incapable of even creating preference reversals for the opponent. The analysis throughout this paper allows for general strategies.

In the space of preferences defined above, "patience" is a more complex category than under exponential discounting: e.g. in the context of the (β, δ)-model, for a given utility function, having inferred δ from choices over long-term prospects does not permit conclusions about how trade-offs between the immediate present and future prospects are resolved, because these are governed separately by β. Nonetheless, for the present context a player's discount factor—her patience—for rewards delayed by $t + 1$ periods, $d_i(t + 1)$, can be usefully decomposed as the product of $d_i(t)$

[2] The reasons for this choice of exposition are that the linear case has received by far the greatest attention in the literature and permits a significant reduction of notation. If anything, the restriction is beneficial to the transparency of my results, because linear utility implies a unique equilibrium in the exponential discounting case and has the discount factor capture all aspects of time preference; in particular, this parameter measures bargaining power. Focusing on this case therefore allows to cleanly disentangle the new source of multiplicity and delay identified here from that of "curved utility", which—in "extreme" cases—can imply multiplicity of stationary equilibrium, and even delay, also under exponential discounting; see Rubinstein (1982).

and what I call the *marginal patience* at delay t, denoted $P_i(t) \equiv d_i(t+1)/d_i(t)$. My first main result identifies the following *simple condition as sufficient for equilibrium uniqueness* whenever it is satisfied by both players i: for any t,

$$P_i(0) \leq P_i(t)$$

In this case, equilibrium indeed takes the simple stationary form described above, and attitudes to delay beyond one single period of bargaining are irrelevant. The above property of preferences can be interpreted as a weak form of *present bias*: it says that, for any given reward, a "marginal" delay is most costly when it is one from the immediate present. This is satisfied by quasi-hyperbolic and hyperbolic as well as exponential discounting preferences; in fact, the property of a constant marginal patience defines exponential discounting, where $P_i(t) = \delta_i$ for all t. Since present bias is a restriction on individual preferences it also lends itself to empirical testing.

This uniqueness result may be highly useful for economic applications that feature both a self-control problem of "over-consumption", e.g. to generate demand for commitment savings products, and bargaining, e.g. intra-household bargaining: it guarantees that there is a unique prediction, which is moreover simple to compute and has clear as well as familiar comparative statics properties. Furthermore, if one believes in the essence of present bias identified here, but finds the evidence inconclusive as to which particular functional form it assumes, then my result is comforting: since the details of such preferences beyond the first period of delay do not matter, the analysis is robust to such mis-specification. Care should then, however, be taken when calibrating or interpreting tje model on the basis of empirical estimates of discount factors: since it is the very short-run discount factors that determine the bargaining split, imputing values from choices with longer-term trade-offs entails the risk of implicitly using the wrong model.

The second main result generalises the analysis to incorporate the possibility that present bias may fail to hold for one of the players. I obtain a *full characterisation of equilibrium outcomes for general separable and time-invariant preferences.* This implies a *characterisation of those pairs of preferences for which equilibrium is unique*—generalising the sufficiency of present bias—and reveals a *novel form of equilibrium multiplicity and inefficient delay.* When a bargaining party's marginal patience is decreasing relative to $P_i(0)$ over the short run, then the anticipation of future delay creates scope for additional threats by the opponent, which are more severe than any threat that is based on subsequent immediate agreement and thus can support delay in a self-confirming manner via non-stationary strategies. Equilibrium multiplicity, delay and non-stationarity are tightly linked: the unique equilibrium with immediate agreement at any stage is the unique stationary equilibrium. A more general notion of bargaining power which emerges is the *minimal marginal patience* over a sufficiently long horizon from the present.[3]

Interestingly, the property of preferences which is conducive to such delay has recently been documented by several studies in the domain of single monetary rewards for a majority of participants, e.g. by Attema et al. (2010) and Takeuchi (2011). While it is at present too early to gauge the reliability of these findings, such preferences may also arise from the specific context of bargaining; I sketch two such instances to which my results can be immediately applied: the first is one where bargaining takes place under the shadow of exogenous breakdown risk, and parties weight probabilities non-linearly. Such preferences have been proposed to simultaneously account for violations of expected utility and time-inconsistent discounting. The second instance I propose is intergenerational bargaining, where each generation is finitely-lived and displays imperfect altruism, so a party's marginal patience is then lowest for the final period of life.

The remainder of this introduction illustrates the main results by means of contrasting two simple examples, discusses the most closely related literature which studies Rubinstein's model for dynamically inconsistent preferences under the assumption of stationary strategies and ends with a brief outline.

1.1 Two Examples

Consider two players, Od and Eve, labelled $i \in \{1, 2\}$ with $i = 1$ for Od and $i = 2$ for Eve, who bargain over how to split a dollar. They alternate in making and answering proposals which are elements of $X = \{(x_1, x_2) \in [0, 1]^2 | x_1 + x_2 = 1\}$ until one is accepted. The first proposal is made by Od in round (period) 1. In any potential round $t \in \mathbb{N}$, a player cares only about the relative delay and the size of her share in a prospective agreement. Specifically, assume their preferences over delayed agreements $(x, t) \in X \times T$, $T = \mathbb{N}_0$, in any period satisfy the representation of equation 1 as follows, where $(\alpha, \beta) \in (0, 1)^2$ and $k \in \{0, 1\}$:

$$\begin{aligned} U_1(x, t) &= \alpha^t x_1 \\ U_{2|k}(x, t) &= \begin{cases} x_2 & t = 0 \\ \beta x_2 & t = 1 \\ k\beta x_2 & t \in T \setminus \{0, 1\} \end{cases} \end{aligned}$$

Eve's preferences are extreme, but in ways which differ strongly over a short horizon of two rounds of delay for the two possible values of k; the illustration of how this contrast translates into possible equilibrium behaviour should serve as a caricature for the general points of this paper.

Assume that both players' preferences be common knowledge—in particular, both players fully

[3] The extent of this horizon is the maximal equilibrium delay which depends on the opponent's preferences.

understand Eve's time preferences—and that players cannot commit to future actions, i.e. use Strotz-Pollak equilibrium (SPE). It is straightforward to show that this game has a unique stationary equilibrium, which is independent of k (of higher-order discounting generally) and which I will refer to as Rubinstein equilibrium (RubE): Od always offers Eve a share of x_2^* and Eve always offers Od a share of y_1^*, with each offer equal to the smallest share the respective respondent is willing to accept when anticipating that the subsequent offer is accepted, i.e.

$$\begin{aligned} x_2^* &= \beta y_2^* \\ y_1^* &= \alpha x_1^* \end{aligned}$$

These two equations have a unique solution: these offers are

$$\begin{aligned} x_2^* &= 1 - \frac{1-\beta}{1-\alpha\beta} \\ y_1^* &= 1 - \frac{1-\alpha}{1-\alpha\beta} \end{aligned}$$

Most textbooks' proofs that this particular equilibrium is the unique one in the case where Eve is also time-consistent, i.e. where instead $U_2(x,t) = \beta^t x_2$, owe to Shaked and Sutton (1984). Their insight is that, despite the history-dependence that any particular equilibrium may display, one may still use backwards induction on the payoff extrema—taken over all equilibria—for each player. This is true because the worst payoff to a proposer occurs when her opponent threatens with her own best subsequent proposer payoff, and "vice versa". After two rounds of backwards induction from the best proposer payoff of a player the resulting payoff must then again equal this best payoff, and similarly for the worst payoff. The resulting system of four equations for these equilibrium payoff extrema has a unique solution revealing payoff uniqueness and efficiency, whence equilibrium uniqueness follows.

When studying dynamically inconsistent time preferences, it is, however, not clear how to use backwards induction: unless equilibrium delay can be ruled out, a player's rankings of equilibrium outcomes of the game where she makes the first proposal may disagree when comparing her two perspectives of (i) the actual initial proposer who evaluates equilibrium outcomes and (ii) the respondent, who evaluates continuation equilibrium outcomes to determine her threat point, because from the latter perspective all equilibrium outcomes are delayed by one additional round. Hence, the relationship between a player i's best/worst proposer payoff and best/worst respondent payoff is more complicated: the lattter need not equal the former times $d_i(1)$. Adding this distinction is my first innovation over Shaked and Sutton (1984). It turns out that, while the two perspectives of a player, generally, agree on what is best—due to the proposer's natural advantage and the player's impatience this is the best immediate agreement—the challenging part is the relationship between a player i's worst threat as a respondent (the worst continuation equilibrium payoff) and i's worst equilibrium payoff as a proposer; it depends on the particular type of dynamic inconsistency as illustrated below by contrasting $k = 1$ and $k = 0$.

Present Bias Consider the case of $k = 1$. Eve is then indifferent to the timing of agreements that all occur in the future but is impatient about postponing agreement from the immediate present; intuitively, Eve's preferences display a form of *present bias*.[4] Her dynamic inconsistency takes the following form: whereas Eve is indifferent between receiving the entire surplus after one or two more rounds because both such prospects have a present value of β, once she finds herself in the next round she will be *more impatient* and prefer the earlier agreement because it is immediate; the comparison is then $1 > \beta$.

In order for this dynamic inconsistency to matter for equilibrium there must be a delay at some stage, possibly only off the equilibrium path. Clearly, not both players can benefit over the efficient RubE from a delayed agreement, and we might reasonably suspect that Eve will lose as soon as her inconsistency is made to bear on the equilibrium outcome. Now suppose v_2 is her worst payoff among all those that may obtain in an equilibrium of the (sub-) game that begins with her proposal and, moreover, suppose it is obtained in an agreement on x which has some delay $t > 0$, i.e. $v_2 = \beta x_2$. This cannot be less than her worst immediate-agreement payoff because she can always choose to satisfy Od's most severe threat immediately: there is an immediate agreement x' with $x'_2 = v_2$, where $1 - v_2 = x'_1 = \alpha V_1$ and V_1 is Od's best subsequent proposer payoff. Since, when responding, Eve further discounts only subsequent immediate agreements, her weakest threat against Od is $\beta x'_2 = \beta v_2$, whence $V_1 = 1 - \beta v_2$. Combining the two equations, we find that

$$v_2 = \frac{1 - \alpha}{1 - \alpha\beta} = y_2^*$$

Because Eve is most impatient about immediate agreement, she cannot be made to lose further from delay; this could only make her stronger as the respondent. But the same argument goes through for Od, and, letting v_1 denote his analogous worst proposer payoff, implies

$$v_1 = \frac{1 - \beta}{1 - \alpha\beta} = 1 - \beta y_2^* = V_1 = x_1^*$$

Then also $v_2 = V_2$ holds true, from which uniqueness and the characterisation as the above RubE follow.

As theorem 1 shows, this argument establishes uniqueness whenever both players' preferences satisfy present bias, which requires that marginal patience is minimal for a delay from the immediate

[4]This case corresponds to the limiting case of quasi-hyperbolic discounting (β, δ)-preferences, where $\delta = 1$.

present, i.e. each player i has $P_i(t) \equiv d_i(t+1)/d_i(t)$ minimal at $t=0$; this is true e.g. for (β,δ)-discounting,

$$P_i(t) = \begin{cases} \beta\delta & t=0 \\ \delta & t>0 \end{cases}$$

Eve's preferences in this example are a limiting case of such preferences where $\delta = 1$. Present bias ensures that a responding player i's worst threat is her worst subsequent immediate agreement, which is worth $d_i(1)\, v_i$. This is the property that allows to exploit the backwards induction approach of Shaked and Sutton (1984) and establish uniqueness more generally.

Violation of Present Bias If $k=0$ then Eve also discounts the first round of delay with a factor β. But she is now willing to accept even the smallest offer in return for not experiencing a delay of more than one round. Note the different nature of her time-inconsistency compared with the previous case: while she is indifferent between receiving the entire dollar with a delay of two rounds and receiving nothing with a delay of one round, at the beginning of the next round she will prefer receiving the entire dollar with one further round's delay over any (at this stage) immediate share less than $\beta > 0$; she will be *more patient* once the sooner option is immediate.

Now suppose $\alpha = \frac{90}{99}$ and $\beta = \frac{99}{100}$, so the RubE has Eve expecting an offer of $x_2^* = \frac{90}{100}$ in the initial round. Yet, Od opens bargaining with a bold move, claiming the entire dollar, and Eve accepts. The following (non-stationary) strategies indeed implement this extreme immediate agreement as an equilibrium outcome:

- Round 1: Od demands the entire dollar and Eve accepts. Upon rejection bargaining progresses to
- Round 2: Eve demands the entire dollar and Od rejects. Upon rejection bargaining continues through
- Round 3:
  - if, previously, Eve demanded the entire dollar then play continues with the stationary equilibrium; there is agreement on x^*.
  - otherwise, play continues as from round 1, leading to immediate agreement on $(1,0)$.

At round-3 histories there is nothing to check: x^* is an equilibrium outcome, and the other continuation strategies' equilibrium property needs to be checked as of round 1. Given the history-dependent continuation, in round 2, Od is willing to accept only proposals x such that $x_1 \geq \alpha = \frac{10}{11}$, and Eve prefers continuation agreement x^*, which has a present value of $\beta x_2^* = \frac{891}{1000}$, over any such proposal because this would yield at most $1-\alpha=\frac{1}{11}$. Anticipating this further delay, which ensues in case she rejects Eve is willing to agree to any division, which Od then exploits by demanding the entire dollar.

The novel phenomenon in this case is how the anticipation of a delay—exploiting the sharp drop to zero in Eve's patience about a further delay from one period in the future relative to her patience about such a delay from the immediate present which is $\beta=\frac{99}{100}$—creates the extreme split in favour of Od as a threat vis-à-vis the RubE, which is powerful enough to "rationalise" itself as an equilibrium outcome, thus resulting in multiplicity. This cannot happen under present bias, where, starting from any payoff less than the RubE's payoff, two steps of backwards induction which involve only the single-period discount factors result in a decrease towards the RubE payoff; it can therefore not rationalise itself as in this example. Indeed, repetition of this step leads to convergence towards the RubE payoff.

Note Eve's intra-personal conflict: as a best reply against Od's strategy, from the point of view of the initial round, Eve would like to reject and subsequently offer a share of α for a present value of $\beta(1-\alpha)=\frac{9}{100}$. However, once round 2 comes around, Eve is not willing to be as generous and prefers forcing a rejection. Restricting Od to a stationary strategy would deprive him of the ability to exploit Eve's such preference reversal. While this equilibrium demonstrates multiplicity it features delay only off the equilibrium path; however, to observe delay on the equilibrium path, simply consider the variant where Eve makes the first proposal and modify strategies accordingly.

The key to characterising the set of equilibrium outcomes beyond present bias is the general insight that a player i's worst equilibrium payoff v_i in the (sub-) game starting with i's proposal, denoted G_i, is constant across all possible equilibrium delays; this is proven in lemma 1. Intuitively, whenever there is delayed agreement, say x with delay t, in equilibrium, the maximal threats must be severe enough to deter players from making too generous an offer when proposing. Since the incentives to do so are strongest for a proposer when the envisaged agreement on x lies furthest ahead in the future, it is sufficient to deter the player(s) from doing so in the earliest round of proposing; for player i in G_i, this can be done up to the point of indifference between yielding to the maximal threat, giving v_i, and obtaining the delayed outcome with a present value of $d_i(t)\,x_i$. If $t_i<\infty$ is the maximal equilibrium delay in G_i and $Z_i^*$is the set of equilibrium outcomes of G_i, then the worst threat of player i when responding (considering all continuation equilibrium payoffs), denoted w_i, is therefore the following function of v_i and t_i:

$$w_i \equiv \inf_{(x,t)\in Z_i^*}\{d_i(t+1)\,x_i\} = \inf_{t\le t_i}\left\{d_i(t+1)\cdot\frac{v_i}{d_i(t)}\right\} = \min_{t\le t_i}\{P_i(t)\}\cdot v_i \tag{2}$$

Note that this reveals the minimal marginal patience over a horizon equal to the maximal equilibrium delay as the determinant of a player's worst threat and a generalised notion of bargaining power.

This argument, however, introduces a further unknown: t_i. By the previous argument, however, t_i is obtained from tracing the set of outcomes that can be implemented via the most severe threats to the proposers, which yield v_1 and v_2, respectively: if $t_i > 0$ then it is the maximal delay $t > 0$ such that the cost of the threats does not exceed the available cake, i.e. $(v_i/d_i(t)) + (v_j/d_j(t-1)) \leq 1$. Building on these results, a system of equations is obtained which theorem 2 studies to establish uniqueness of a solution to this system as both necessary and sufficient for uniqueness of equilibrium, generalising uniqueness under present bias. Theorem 3 further generalises this result, producing a characterisation of equilibrium outcomes and payoffs when the system of equations may have multiple solutions. Appendix B illustrates the general characterisation result via the example studied here.

Neither of the two extreme preferences of player Eve studied here fall into the class dealt with in the main part of this paper; they are, however, limiting cases. The remainder serves to demonstrate that the findings illustrated here generalise.

1.2 Literature on Bargaining with Non-Exponential Discounting

There exist a number of contributions to decision theory that use the classic Rubinstein (1982) model of bargaining as an application. Ok and Masatlioglu (2007) axiomatise preferences that are more general than the ones I consider here in that they allow for non-transitivity; while there is separability, discounting is then relative to a particular comparison instead of absolute as part of a present-value calculation. In the context of the present work, their proposition 2 claims that when the players' utility functions are strictly concave then there is a unique "time-consistent" SPE (p. 230), which is the familiar stationary equilibrium.[5] However, they do not define what they mean by "time-consistent" when used to qualify SPE nor provide a proof, only indicating that the arguments of Shaked and Sutton (1984) apply.[6] My theorem 2 proves that, without a refinement of SPE, their proposition fails to hold in general, because their class of preferences covers also those that are shown to imply multiplicity of SPE.

Noor (2011) generalises the exponential discounting model to allow for the discount factor to depend on the size of the reward, relaxing separability, which also induces preference reversals of the type rationalised by hyperbolic discounting. This is done in order to additionally accomodate another empirical phenomenon called "magnitude effect" where, for a given delay, smaller rewards appear to be discounted more heavily than larger rewards. In applying these preferences to bargaining, he simplifies them to linear ones (in the share) and focuses on stationary equilibrium with immediate agreement; he finds the possibility of multiplicity and of a more patient initial proposer obtaining a smaller share than her opponent. For the kind of equilibria he studies, which involve

[5]Due to the strict concavity of the utility functions there is a unique such stationary equilibrium.

[6]They mention that it is "possibly a refinement" of SPE (see their footnote 15).

only attitudes to delay of a single period, however, those preferences are indistinguishable from standard exponential preferences with non-linear utility.[7]

Akin (2007) studies bargaining by two players with linear utility and quasi-hyperbolic preferences. His focus is on naïveté about future preferences and learning through delay in bargaining. Assuming stationary equilibrium conditional on beliefs (the equilibrium concept is of course more involved than SPE due to naïveté and learning assumptions), he finds that delay may arise due to a naïve player's learning from a sophisticated player who has an incentive to forgo earlier agreements in exchange for such learning of the opponent and accordingly better later splits. Theorem 1 lends some comfort to this analysis by showing that under sophistication there is indeed a unique stationary SPE for such preferences.

In the context of an exogenous risk of breakdown, results similar to those in Rubinstein (1982) have been obtained (see Binmore et al. (1986)). Burgos et al. (2002) study risk preferences which allow for non-separability and time-inconsistency, where their equilibrium concept permits full commitment to future actions. The authors provide assumptions which yield a unique stationary equilibrium; this equilibrium is then further analysed. Volij (2002) shows that when these preferences are restricted to being time-consistent the model becomes equivalent to that of Rubinstein (1982).

All of the works discussed here assume stationary strategies in one way or another, so even when they include the preferences I am considering, they rule out the type of delay equilibria which theorem 3 reveals, and which necessarily rely on non-stationary strategies. As argued elsewhere, e.g. Osborne and Rubinstein (1990, p. 39), stationarity of strategies is problematic as an assumption in particular in bargaining. This is even more so in the presence of time-inconsistency because the restriction to a stationary strategy deprives a player of the ability to even create, let alone exploit, preference reversals of a time-inconsistent opponent.

1.3 Outline

Section 2 defines the bargaining game, including the general class of preferences considered in this paper; its last subsection highlights a stationarity property of the game and, on this basis, defines various concepts which the proofs for the subsequent sections will be based upon. Section 3 studies stationary equilibrium and finishes by arguing that stationarity, as an assumption on strategies, is particularly problematic in the analysis of time-inconsistent preferences. Section 4 presents the first main result, which is the sufficiency of present bias for equilibrium uniqueness. This is generalised in section 5 where a characterisation of those preferences for which equilibrium is unique as well as a general characterisation of equilibrium outcomes and payoffs are provided. Section 6 sketches

[7] Without restrictions on the curvature, the latter permit the same kind of multiplicity. In fact, there may also arise delay out of the multiplicity of stationary equilibria (see e.g. Osborne and Rubinstein (1990, Section 3.9.2) which actually refers to an example in Rubinstein (1982, pp. 107-108)).

two "applications" of these results where dynamic inconsistency is motivated from specific aspects of the bargaining environment, and section 7 concludes.

While the exposition of the paper concentrates on the case of linear utility functions, appendix A provides analogous versions of results for the general case; most proofs in the main paper will be made by referring to the more general results there.

2 Model and Definitions

2.1 Protocol, Histories and Strategies

Two players $I = \{1, 2\}$ bargain over how to share a perfectly divisible surplus of (normalised) size one. In each round $t \in \mathbb{N}$, player $\rho(t) \in I$ proposes a split $x \in X = \{(x_1, x_2) \in \mathbb{R}_+ | x_1 + x_2 = 1\}$ to opponent player $3 - \rho(t)$, equivalently, offers the opponent a share $x_{3-\rho(t)}$, who responds by choosing $a \in \{0, 1\} = A$, either accepting the proposal, $a = 1$, or rejecting it, $a = 0$. Upon the first acceptance, bargaining terminates with the agreed split x being implemented, and upon rejection players move to the next round in $t+1$. Bargaining begins in round $t = 1$ with a proposal by player 1 and has the players alternate in their roles of proposer and respondent, i. e. $\rho(t+1) = 3 - \rho(t)$.

Histories of such a game at the beginning of a round $t \in \mathbb{N}$ are sequences of proposals and responses: $h^{t-1} = (x^s, a^s)_{s \leq t-1} \in (X \times A)^{t-1}$. Since bargaining concludes following the first accepted proposal, such non-terminal histories are elements of $H^{t-1} = (X \times \{0\})^{t-1}$, and a terminal history ending in round t is an element of $H^{t-1} \times (X \times \{1\}) = \mathcal{H}^t$; for completeness, let $H^0 = \{h^0\}$. H^∞ denotes the set of non-terminal histories of infinite length.

A strategy of a player i is a mapping σ_i such that, for any $t \in \mathbb{N}$, $h^{t-1} \in H^{t-1}$ and $x \in X$,

$$\begin{aligned} i = \rho(t) \quad &\Rightarrow \quad \sigma_i\left(h^{t-1}\right) \in X \\ i = \rho(t+1) \quad &\Rightarrow \quad \sigma_i\left(h^{t-1}, x\right) \in A \end{aligned}$$

Let the space of all such functions of player i be denoted by Σ_i.[8] Any pair of strategies $\sigma = (\sigma_1, \sigma_2)$ generates either a terminal history in $\cup_{t \in \mathbb{N}} \mathcal{H}^t$ or an infinite non-terminal history in H^∞ in an obvious way: the first-round actions are $(\sigma_1(h^0), \sigma_2(h^0, \sigma_1(h^0))) \equiv h^1_\sigma$ so if $\sigma_2(h^0, \sigma_1(h^0)) = 1$ then $h^1_\sigma \in \mathcal{H}^1$ and the game ends after the first round, otherwise add the second-round actions to generate a history $(h^1_\sigma, \sigma_1(h^1_\sigma), \sigma_2(h^1_\sigma, \sigma_1(h^1_\sigma))) \equiv h^2_\sigma$ etc. Call a terminal history that is thus obtained h^t_σ if it is in $\mathcal{H}^t$ for $t \in \mathbb{N}$; if none exists then call the corresponding infinite non-terminal

[8] A player's strategy must specify her action for every contingency, including all those that the play of this strategy actually rules out. For instance, although a strategy by player 2 may specify acceptance of every possible first round proposal, it must also specify what she would propose in round 2 following a rejection; see Rubinstein (1991) on how to interpret strategies in extensive form games.

history h_σ^∞.

This can in fact be done starting from any $h \in H^t \cup (H^t \times X)$ in the very same way, in which case the history obtained is the continuation history of h under σ; if it yields a terminal history after s more rounds then it is some $h_\sigma^s \in H^s \cup (A \times H^{s-1})$ such that $(h, h_\sigma^s) \in \mathcal{H}^{t+s}$, and otherwise it is an element of $H^\infty \cup (A \times H^\infty)$. Note that for any two histories $h^t \in H^t$ and $h^s \in H^s$ with $\rho(t) = \rho(s)$ the sets of possible continuation histories are identical; therefore this holds true also for (h^t, x) and (h^s, x) for any $x \in X$, and in this sense the protocol is stationary. In particular there exist stationary strategies.

Definition 1. A bargaining strategy $\sigma_i \in \Sigma_i$ of player i is a *stationary strategy* if there exist $\hat{x} \in X$ and $\hat{a} : X \to \{0, 1\}$ such that, for any $t \in \rho^{-1}(i)$, $h^{t-1} \in H^{t-1}$ and $(h^t, x) \in H^t \times X$,

$$\begin{aligned} \sigma_i \left(h^{t-1}\right) &= \hat{x} \\ \sigma_i \left(h^t, x\right) &= \hat{a}(x) \end{aligned}$$

Given any equilibrium concept (see section 2.3), an equilibrium (σ_1^*, σ_2^*) is a *stationary equilibrium* if each σ_i^*, $i \in I$, is stationary.

A stationary strategy does not respond to history; indeed, if σ is a pair of stationary strategies then, for any non-terminal histories $h^t \in H^t$ and $h^s \in H^s$ with $\rho(t) = \rho(s)$ the same continuation history is obtained.

A further note on the generality of this model may be warranted at this stage: it makes two normalisations, that of the size of the surplus and that of the amount of time elapsing between rounds of bargaining. Unless one is interested in comparative statics involving these, by defining players preferences relative to these parameters there is no loss of generality; indeed, this is how the assumptions on preferences below are to be understood.

Another restriction implicit in the protocol is that proposals are non-wasteful (players' shares add up to one). This is without loss of generality, which will become clear from the properties of the preferences assumed: players only care about their own share which they want to maximise and obtain sooner rather than later, and they can always choose to claim the entire cake. Hence, a proposer who wants an offer accepted will not waste anything of what is left for herself, and, by claiming the entire cake, a proposer makes the least attractive feasible offer anyways.

2.2 Preferences

When discussing players' time preferences in the discrete time setting considered here I will refer to "dates", which correspond to bargaining periods or "rounds" in the context of the bargaining protocol considered here. At the most basic level it is assumed that players care intrinsically only about the size and the timing of their own share in any agreement, and not about how a particular agreement is obtained nor about the details of disagreement. Because I want to study dynamically inconsistent time-preferences, where a player's preference over two *dated* future rewards may change depending on the *date* at which she makes the choice, the primitive of a player i's preferences is a sequence of *dated* preference orderings $\{\succeq_{(i,t)}\}_{t\in\mathbb{N}}$.[9] Each element $\succeq_{(i,t)}$ is defined on the set of feasible outcomes at t, which is $Z_t = (X \times \mathbb{N}_t) \cup \{D\}$, where $\mathbb{N}_t = \{t' \in \mathbb{N} | t' \geq t\}$ and D is disagreement.[10]

Now let $T = \mathbb{N}_0$ denote the possible delay of an agreement. I assume that, at any date t, $\succeq_{(i,t)}$ has a separable utility representation, where $d_i : T \to [0,1]$ is continuous and decreasing, with $d_i(0) = 1$ and $\lim_{t\to\infty} d_i(t) = 0$, and $u_i : [0,1] \to [0,1]$ is continuous and increasing, with $u_i(1) = 1$:[11]

$$U_{(i,t)}(z) = \begin{cases} d_i(s)\, u_i(x_i) & z = (x, t+s) \in X \times \mathbb{N}_t \\ 0 & z = D \end{cases}$$

Note that this representation ignores t and involves time only in relative terms, i. e. as delay from t. In terms of feasible *delayed* agreements the domain at any t is identically equal to $(X \times T) \cup \{D\} = Z$, which I will refer to as the set of *outcomes*. I assume that the player has in fact identical preferences at any date over Z: for any $z \in Z$, and any (i,t) and (i,t'),

$$U_{(i,t)}(z) = U_{(i,t')}(z) = \begin{cases} d_i(t)\, u_i(x_i) & z = (x,t) \in X \times T \\ 0 & z = D \end{cases} = U_i(z)$$

The image of $[0,1]$ under u_i will be denoted $\mathscr{U}_i \equiv [u_i(0), 1]$. The following assumption summarises.

Assumption 1. *For each player $i \in I$, there exist a continuous decreasing function $d_i : T \to [0,1]$ with $d_i(0) = 1$ and $\lim_{t\to\infty} d_i(t) = 0$, and a continuous increasing function $u_i : [0,1] \to [0,1]$ with $u_i(1) = 1$, such that, for any date $t \in \mathbb{N}$, preference $\succeq_{(i,t)}$ over continuation outcomes Z are represented by*

$$U_i(z) = \begin{cases} d_i(s)\, u_i(x_i) & z = (x,s) \in X \times T \\ 0 & z = D \end{cases} \tag{3}$$

Although the various dated "selves" of a player "look the same" in terms of their relative preferences—*the* individual's time preference—dynamic inconsistency arises whenever the "marginal"

[9]Interestingly, Rubinstein mentions such a generalisation in several remarks of his original bargaining article, see Rubinstein (1982, remarks on pages 101 and 103).

[10]Formally, each such outcome is an equivalence class of (continuation) histories.

[11]An axiomatisation of such separable time preferences for discrete time is provided in Fishburn and Rubinstein (1982). Due to the discreteness of time in this model, continuity of d_i is without loss of generality; axiomatisations for continuous time with this property are available in Fishburn and Rubinstein (1982) and Ok and Masatlioglu (2007).

patience about an additional period of delay from a given delay of t rounds, denoted $P_i(t)$, is not constant, where

$$\frac{d_i(t+1)}{d_i(t)} \equiv P_i(t)$$

Constancy of $P_i(t)$ is in fact the defining property of exponential discounting which is therefore the only dynamically consistent form of discounting in the class delineated by assumption 1. Just as $d_i(t+1)$ is interpreted as a measure of patience about a delay of $t+1$ periods, $P_i(t)$ can be interpreted as measuring the marginal patience at delay t: one util with delay $t+1$ is worth $P_i(t)$ utils with delay t.

I now define a property that refines the class of preferences above, which—at least in related forms—has appeared elsewhere in the literature and turns out to be of great interest in the bargaining context.[12]

Definition 2. A player i's preferences satisfy *present bias* if, for any $t \in T$, $P_i(0) \leq P_i(t)$. They satisfy *strict present bias* if this inequality holds strictly for every $t > 0$.

The significance of the property is clear: an individual with present bias considers a one-period delay most costly when it involves a delay from the immediate present. Observe that present bias is actually equivalent to saying that any fixed delay is (weakly) more costly when it occurs from the present than when it occurs from the subsequent period because, for any $t \in T$, from cross-multiplication,[13]

$$P_i(0) \leq P_i(t) \Leftrightarrow \frac{d_i(t)}{d_i(0)} \leq \frac{d_i(1+t)}{d_i(1)} \tag{4}$$

Exponential discounting, i. e. $d_i(t) = \delta^t$ for some $\delta \in (0,1)$, satisfies present bias in its weak form: $P_i(t) = \delta$ for all $t \in T$; in other words, marginal patience is independent of delay and measured by a single parameter. Quasi-hyperbolic discounting, where, for $t > 0$, $d_i(t) = \beta\delta^t$ with $\beta \in (0,1)$, satisfies strict present bias; see the initial example in section 1.1.

To simplify the exposition of my results and improve their transparency, I strengthen assumption 1 to a representation with linear utility throughout the main part of the paper. This is the best known and most widely used class of utility functions in bargaining. The general case is dealt with in appendix A.

[12]The definitions of "present bias" and "strong present bias" in Ok and Masatlioglu (2007, p. 225) are closely related, but stronger, ordinal versions of the two I provide here. Halevy (2008, Definition 1) has a definition identical to mine of present bias but calls the property "diminishing impatience"; he is, however, interested in how different degrees of "mortality risk" translate into properties of discounting under non-linear probability weighting.

[13]It does not, however, imply the stronger property that, for any $\{s,t\} \subseteq T$,

$$\frac{d_i(t)}{d_i(0)} \leq \frac{d_i(s+t)}{d_i(s)}$$

Assumption 2. *Each player i's preferences satisfy assumption 1 with a linear representation*

$$U_i(x,t) = d_i(t)\, x_i$$

2.3 Equilibrium Concepts

In this section I will introduce two equilibrium concepts for games with time-inconsistent players which are both adaptations of subgame-perfect Nash equilibrium (SPNE) and discuss them.[14] Underlying these definitions is the assumption that both the protocol and preferences are common knowledge. This implies that the game has perfect information, and for the case of dynamically inconsistent preferences it also means that players are "sophisticated" about their own as well as their opponent's future preferences; see O'Donoghue and Rabin (1999, 2001).

In the definitions below, denote by $z_h(\sigma) \in Z$ the continuation outcome of a history $h \in H^t \cup (H^t \times X)$, $t \in \mathbb{N}$, that obtains under the two parties' playing according to strategy profile σ. If a terminal continuation history h_σ obtains such that $(h, h_\sigma) \in \mathcal{H}^{t+s}$ for some $s \in T$ then $z_h(\sigma) = (x, s)$, where x is the last (accepted) proposal; otherwise $z_h(\sigma) = D$.

2.3.1 Strotz-Pollak Equilibrium

When a player's preferences over actions may change over the course of a game, a theory is required for how this intrapersonal conflict is resolved. It has become standard to consider each player i's dated self (i,t) as a distinct non-cooperative player and derive individual behaviour from SPNE of this game; for the origins of this concept see Strotz (1956) and in particular Pollak (1968).[15] Game-theoretically, thus the intrapersonal conflict is dealt with in exactly the same manner as interpersonal conflict. Specifically, this means that at any history of round t at which (i,t) is to move, this self of player i takes as given (the beliefs about) not only the behaviour of the oppent (or opponent's selves) but also the behaviour of all other selves of player i; in other words, changing (i,t)'s strategy is equivalent to one-shot deviations. Adapting this idea to the present context results in the following definition.

Definition 3. A strategy profile σ^* is a *Strotz-Pollak equilibrium (SPE)* if, for any $t \in \mathbb{N}$, $h^{t-1} \in H^{t-1}$, $x \in X$ and $a \in \{0,1\}$, the following holds:

$$\rho(t) = i \;\Rightarrow\; U_i\left(z_{h^{t-1}}(\sigma^*)\right) \geq U_i\left(z_{(h^{t-1},x)}(\sigma^*)\right)$$
$$\rho(t+1) = i \;\Rightarrow\; U_i\left(z_{(h^{t-1},x)}(\sigma^*)\right) \geq U_i\left(z_{(h^{t-1},x,a)}(\sigma^*)\right)$$

[14] For an introduction to SPNE see a textbook on game theory, e.g. Osborne and Rubinstein (1994, Part II).
[15] Further developments, in particular with regard to existence of SPE, can be found in Peleg and Yaari (1973) and Goldman (1980).

This definition is really just an application of SPNE to the game when the set of players is taken to be $I \times \mathbb{N}$. The well-known one-shot deviation principle guarantees that SPE coincides with SPNE whenever players have time-consistent preferences (e.g. Fudenberg and Tirole (1991, Theorem 4.1), where continuity at infinity holds for every self's preference because of $\lim_{t\to\infty} d_i(t) = 0$). SPE is the main concept I will use in this work: when referring to "equilibrium" I will mean SPE.

2.3.2 Perfect Commitment Equilibrium

At the other extreme lies the assumption that every self (i,t) can perfectly control i's (future) behaviour, which the following solution concept is based upon.

Definition 4. A strategy profile σ^* is a *Perfect Commitment Equilibrium (PCE)* if, for any $t \in \mathbb{N}$, $h^{t-1} \in H^{t-1}$, $x \in X$, and any $\sigma \in \Sigma$ such that $\sigma_{3-i} = \sigma^*_{3-i}$, the following holds:

$$\begin{aligned} \rho(t) = i \quad &\Rightarrow \quad U_i\left(z_{h^{t-1}}(\sigma^*)\right) \geq U_i\left(z_{h^{t-1}}(\sigma)\right) \\ \rho(t+1) = i \quad &\Rightarrow \quad U_i\left(z_{(h^{t-1},x)}(\sigma^*)\right) \geq U_i\left(z_{(h^{t-1},x)}(\sigma)\right) \end{aligned}$$

This definition applies SPNE in the standard sense of robustness to "full-strategy deviations", disregarding any commitment problems, whence PCE and SPNE also coincide under time-consistency of all players. Clearly, in any such equilibrium, *conditional on the opponent's strategy*, there is in fact no intrapersonal conflict. The condition is emphasised because the opponent's strategy determines the set of feasible outcomes of a player, and there may be ways to limit a time-inconsistent player's choice set such that there is no conflict (as a trivial case consider an opponent who is able to dictate an outcome).

Proposition 1. *Any PCE is a SPE.*

Proof. Let σ^* be a PCE and restrict σ_i in definition 4 to coinciding with σ^*_i except for the immediate action which is $\sigma^*_i(h^{t-1})$ if $\rho(t) = i$ and $\sigma^*_i(h^{t-1}, x)$ if $\rho(t+1) = i$, respectively. □

PCE as a refinement of SPE will be of interest here only in so far as it provides information about SPE along the lines of definition 5 below. That, for time-inconsistent players, PCE is stronger than SPE will be demonstrated below (contrast proposition 3 and theorem 2).[16] The observation that the two concepts lie at two opposite extremes in terms of "self-control" or internal coordination motivates the following terminology.

[16]Their connection is tighter in the bargaining game analysed here than in the single-person context which has been the main focus of the respective literature. For an example that proves the point, see Asheim (1997, Example 1). The main issue arising in refinements of SPE based on notions of cooperation is indifference of subsequent selves, as elaborated in Kodritsch (2012). For a dramatic example in a single-person decision problem to argue the weakness of SPE (it exploits a violation of continuity at infinity), see Asheim (1997, Example 2).

Definition 5. Any SPE which is not a PCE is said to exhibit *intrapersonal conflict.*

Weaker refinements of SPE have been proposed, all of them departing from the premise that, notwithstanding the presence of commitment problems, the existence of a single individual to whom these selves "belong" should imply a conceptualisation that does not treat them as entirely distinct non-cooperative players. In various ways these refinements capture different degrees of intrapersonal coordination, but there is yet to emerge a consensus on a viable alternative to SPE.[17]

The natural benchmark for the analysis in this paper is the uniqueness of an SPE which is stationary found in Rubinstein (1982). To simplify terminology, I define stationary equilibrium separately here.

Definition 6. A *Rubinstein equilibrium (RubE)* is any SPE which is in stationary strategies.

2.4 A Stationarity Property and Useful Definitions

In payoff-relevant terms, for any two rounds t and s with respective histories $h^{t-1} \in H^{t-1}$ and $h^{s-1} \in H^{s-1}$ the sets of feasible outcomes are identically equal to Z. Part of assumption 1 is that each player has a single preference over Z. Because of alternating offers, therefore, the subgames starting after these histories are identical if and only if $\rho(t) = \rho(s) = i$.

Denote the subgame starting at a history $h \in H^{t-1}$ for $t \in \rho^{-1}(i)$ by G_i. The set of SPE outcomes of G_i, measured in relative terms as elements of Z, will be referred to as Z_i^*. Proposition 5 below will ensure that both Z_1^* and Z_2^* are non-empty. Based on this set, define the following payoff extrema, where the restriction to $Z_i^* \cap (X \times T)$, exluding disagreement, will also be justified below by corollary 3: for $i \in I$, define

$$\begin{aligned} V_i &= \sup_{(x,t) \in Z_i^*} \{U_i(x,t)\} \\ W_i &= \sup_{(x,t) \in Z_i^*} \{U_i(x,1+t)\} \end{aligned}$$

V_i is the lowest upper bound on the SPE payoff of player i as the initial proposer in G_i, and W_i is the lowest upper bound on the SPE payoff of player i as the respondent conditional on rejection, i.e. the supremum continuation SPE payoff; informally, it is player i's "best threat" when responding. Let the corresponding player-indexed lowercase letters, i.e. v_i and w_i, denote the respective infima, and, moreover, for each of these bounds, let an additional superscript of 0 indicate the restriction

[17] See Laibson (1994, Chapter 1), Ferreira et al. (1995), Kocherlakota (1996) and Asheim (1997). While for single-person problems there is a tendency towards favouring stationary equilibria—see Kocherlakota (1996), Laibson (1997) or Sorger (2004)—the case in strategic settings has not received much attention; however, in the bargaining context, such a restriction appears problematic as argue Osborne and Rubinstein (1990, p. 39); see also section 3 below.

to immediate-agreement SPE outcomes, e.g.

$$w_i^0 = \inf_{(x,0)\in Z_i^*} \{U_i(x,1)\}$$

Moreover, I will introduce another lowest upper bound, the supremum SPE delay in G_i: for each $i \in I$, define

$$t_i = \sup_{(x,t)\in Z_i^*} \{t\}$$

The significance of these functions of Z_i^* will become clear from the proofs below but the idea, going back to Shaked and Sutton (1984), is that while equilibria may in principle display complex history-dependence, arguments akin to backwards induction can be used to relate and determine these variables. This is known for the case of exponential discounters where only $(v_i, V_i)_{i\in I}$ are defined because the equalities $W_i = d_i(1)V_i$ and $w_i = d_i(1)v_i$ which define a player i's best and worst threats, respectively, do not require further arguments. When allowing for time-inconsistent discounting preferences, this is neither obviously nor generally true, however, for the second equality.

3 Stationary Equilibrium

This section will establish the robustness of stationary equilibrium in terms of existence. Thinking about the textbook case of a finite horizon, which is not covered here, one can see that backwards induction results in a unique equilibrium, where at each stage the proposer offers the opponent the present value of the unique continuation agreement and the opponent agrees, i.e. there is immediate agreement at any stage and only one-period discounting enters payoffs. Taking the limit of the respective offers and acceptance rules as the horizon becomes infinite, they become independent of time and the resulting stationary strategies preserve the equilibrium property, which establishes existence of a stationary equilibrium.

I prove that any such strategy profile is in fact a PCE, i.e. it is robust against full commitment deviations. Finally, based on this result, I argue that *assuming* stationary strategies when players have dynamically inconsistent time preferences is even more problematic than it is under time-consistency; if one is genuinely interested in the possible distinctive implications of time-inconsistent (higher-order) discounting then one must allow for non-stationary strategies.

Proposition 2. *Under assumption 2, the bargaining game has a unique RubE which is given by* *the following strategy profile σ^R: for any $t \in \rho^{-1}(i)$, $h \in H^{t-1}$ and $x \in X$,*

$$\begin{aligned} \sigma_i^R(h) &= x^{R,i};\ x_i^{R,i} = \frac{1-d_j(1)}{1-d_i(1)\,d_j(1)} \\ \sigma_j^R(h,x) &= \begin{cases} 1 & x_j \geq 1-x_i^{R,i} \\ 0 & x_j < 1-x_i^{R,i} \end{cases} \end{aligned}$$

This SPE exhibits no intrapersonal conflict.

Proof. See appendix A.1: under assumption 2 there exists a unique pair of proposals—one for each player—with the property that each player, as a respondent, is indifferent between accepting the other player's proposal and a rejection that results in a subsequent immediate agreement on that player's proposal. This pair is given by $\left(x^{R,1}, x^{R,2}\right)$. Proposition 4 shows that any RubE must take the form of the strategy profile in the above proposition which is based on this pair, and proposition 5 establishes that any such strategy profile is a PCE. □

Because of this result, saying that a RubE always exists and is unique, I will refer to "the RubE" in what follows.

The remainder of this section is devoted to a discussion of the RubE's property of being a PCE. Existence of a PCE is remarkable since in such an equilibrium each player's strategy creates a situation for the opponent in which the latter's time-inconsistency is "neutralised", and this despite the availability of full commitment deviations. If one were to assume stationary strategies then there would obtain a unique and simple prediction with this property. It has been argued elsewhere, and in the context of time-consistent preferences, that such an assumption is problematic in bargaining from the point of view of strategic reasoning; see Osborne and Rubinstein (1990, p. 39). I want to add to this point a further observation which suggests it is even more problematic when players are dynamically inconsistent, namely that constraining a player to a stationary strategy robs her of the ability to create—and potentially exploit—preference reversals of a time-inconsistent opponent.

Such a preference reversal must take the form that, subject to the feasible outcomes under the opponent's strategy, a player at some stage most prefers some delayed outcome over the best immediate one—which she can implement herself—while later taking actions that induce a worse outcome than the envisaged delayed outcome. While such preference reversals may take complex forms for general strategies of the opponent, they are easily examined for a stationary strategy. Without loss of generality, consider player 2's problem when facing an opponent player 1 who behaves according to some stationary strategy which means that she always proposes some $\hat{x} \in X$ and follows acceptance rule a_1 such that $\hat{y}$ is the most preferred split for player 2 that she accepts.[18]

[18]Because u_2 is increasing, this is all that matters about player 1's acceptance rule. Strictly speaking, however,

Since disagreement is worst, at any stage, player 2's *favourite* feasible outcome subject to this strategy by player 1 is then

- either $(\hat{y}, 0)$ or $(\hat{x}, 1)$ when proposing
- either $(\hat{x}, 0)$ or $(\hat{y}, 1)$ when responding to player 1's proposal of $\hat{x}$

Note that in order to have any preference reversal there must be one over such most preferred feasible outcomes, i.e. while (as a proposer) player 2 prefers $(\hat{x}, 1)$ over $(\hat{y}, 0)$, (as a respondent) she prefers $(\hat{y}, 1)$ over $(\hat{x}, 0)$, with at least one preference being strict:

$$\begin{aligned} d_2(1)\,\hat{x}_2 &> (\geq) \quad \hat{y}_2 \\ d_2(1)\,\hat{y}_2 &\geq (>) \quad \hat{x}_2 \end{aligned}$$

Yet, this is clearly impossible: by mere impatience, if a player prefers some delayed reward over an immediate reward then this preference for the former reward must intensify when it becomes immediate and the latter is delayed instead. The scope for dynamic inconsistency to matter for equilibrium is accordingly eliminated when stationarity of strategies is imposed on the players' behaviour.

4 Present Bias and Uniqueness

In applied work which uses strategic bargaining, e.g. wage-setting through negotiations by unions and firms or intra-household bargaining over how to share common resources, it is important to have reliable predictions. Under multiplicity of equilibrium, the uncertainty about this one aspect of the model feeds through all conclusions. Therefore it is of great interest to understand when uniqueness obtains in order to gauge whether the assumptions required for it are reasonable within the context of the application. In addition, robustness is certainly desirable: since the parametrisations of preferences, technologies etc. which economic applications employ are only approximations, to have confidence in the conclusions they should remain themselves approximately true once the approximation is not exact.

Ideally, uniqueness can be guaranteed from properties of *individual* preferences which are more readily interpretable as well as testable. This section therefore investigates the question of which individual preferences yield a unique equilibrium once players cannot be assumed to satisfy exponential discounting; this class turns out to be large, including all of the most familiar alternatives to exponential discounting.

there may not exist a minimum of the set $\{y_1 \in [0,1] | a_1(y) = 1\}$. But under assumption 1, in particular continuity, the argument goes through with the modification of player 2's offering player 1 a share of $\hat{y}_1 + \epsilon$ for some $\epsilon > 0$, which then exists.

To simplify notation, whenever $i \in I$, I will let $j = 3 - i$. The main result of this section is the following theorem.

Theorem 1. *Under assumption 2, if each player i's preferences satisfy present bias as in definition 2 then the RubE is the unique SPE.*

Proof. Appendix A.2 proves a more general theorem 4, which implies this claim. □

Present bias ensures that a respondent, by rejecting, cannot obtain a worse payoff under a continuation equilibrium that has itself delay than the worst continuation payoff under subsequent immediate agreement. This is because, due to present bias, as next round's proposer, this player will be at least as impatient about a delay as the current round's respondent. As a consequence, the proposing opponent could not exploit a present-biased player's time-inconsistency by means of delay; in other words, the respondent is weakest—in terms of available threats—under subsequent immediate agreement, or $w_i = w_i^0$.

This is the main insight involving time-inconsistency under present bias for establishing the result. Due to the natural advantage conferred to a proposer by the protocol, which places the onus of delay fully onto the responding player, a player's equilibrium payoffs as the initial proposer of a subgame can easily be seen to be spanned by immediate agreement equilibria, i.e. $V_i = V_i^0$ and $v_i = v_i^0$: whatever the commonly known continuation play, a proposer can exploit the impatience of the respondent; therefore, the proposer can extract the full rent over the worst continuation of the opponent immediately to obtain the best payoff, and also satisfy the most demanding opponent immediately to obtain the worst payoff. Moreover, mere impatience then implies that the strongest threat a player can entertain as the respondent is based on the best immediate agreement payoff this player could obtain as the subsequent proposer, or $W_i = W_i^0$. When both players' preferences satisfy present bias, therefore, all equilibrium as well as continuation payoffs are spanned by immediate and subsequent immediate agreement equilibrium payoffs, respectively; the familiar approach of Shaked and Sutton (1984), using two steps of backwards induction together with the stationarity of the game, can then be applied to yield the following equations

$$V_i^0 = \underbrace{1 - d_j(1)\left(1 - \underbrace{d_i(1)\, V_i^0}_{=W_i^0}\right)}_{=w_j^0}$$

$$v_i^0 = \underbrace{1 - d_j(1)\left(1 - \underbrace{d_i(1)\, v_i^0}_{=w_i^0}\right)}_{=W_j^0}$$

Hence $V_i = V_i^0 = v_i^0 = v_i$ from which uniqueness of payoffs, their coincidence with the RubE payoffs as well as their efficiency follow readily, which then immediately establishes the conclusion.

The following result sheds further light on the RubE, and indirectly also on present bias in the bargaining context: the RubE is the only SPE which exhibits no intrapersonal conflict.

Proposition 3. *Under assumption 2, the RubE is the unique PCE.*

Proof. See proof of proposition 6 in appendix A.2. □

When read in conjunction with theorem 1 this proposition can be interpreted as saying that, under present bias, not only can both players fully exploit each other's time inconsistency through immediate agreement at any stage, but they can do so by means of stationary strategies, which does not induce preference reversals. Not only would it be misleading to conclude that present bias is a form of time-inconsistency that cannot be exploited by the opponent, but the opposite is true, namely that present bias is a form of time-inconsistency which is certain to be exploited by the opponent in some equilibrium of this game, because it is in the RubE, which always exists.

Note, moreover, the implication of this proposition that any SPE other than the RubE, in particular any non-stationary SPE as well as any SPE which features delay (on or off the equilibrium path) exhibits intrapersonal conflict: some player at some stage would then prefer to change her own future actions. In this sense therefore such SPE could arise purely from dynamically inconsistent time preferences; of course, by 1, these must violate present bias for at least one player.

5 General Characterisation Results

First, note that present bias is unnecessarily strong because the maximal delay can be bounded by a simple rationality argument. By mere impatience there always exist proposals that would be immediately accepted by a rational respondent, so even if a proposing player i expected to obtain the entire surplus, there is a finite delay after which i would rather make an offer that entices the most demanding rational respondent to immediately agree. Formally, note that a rational respondent j accepts any proposal x such that $x_j > d_j(1)$, whence a rational proposer's worst immediate agreement payoff is no less than $1 - d_j(1)$. Eventually $d_i(t)$ falls below this number because its limit is zero, which yields the following bound:

$$\bar{t}_i = \max\{t \in T \,|\, d_i(t) \geq 1 - d_j(1)\} \tag{5}$$

Clearly, $t_i \leq \bar{t}_i < \infty$ and $P_i(0) \leq P_i(t)$ for all $t \leq \bar{t}_i$, for both $i \in I$, is a weaker sufficient condition for uniqueness.

Note that even this simple argument involves relating the two player's preferences, however. From the previous section's analysis, the relationship of payoff bounds which depends on the details of a player's preferences is that between the worst threat of a player i as a respondent w_i and i's worst equilibrium payoff v_i. While $w_i \leq d_i(1) v_i$ holds because $v_i = v_i^0$, the main issue is when a player's worst continuation payoff might fall below the present value of the worst subsequent immediate agreement. The key to relating w_i to v_i is the introduction of the maximal delay t_i as an additional variable, because not only will t_i be determined by the maximal threats to the players when proposing, v_1 and v_2, but, when combined with the argument that v_i is the worst payoff for any given possible delay $t \leq t_i$, it also generates an equation relating w_i to v_i through t_i; thus, by expanding the number of unknown characteristics of the set of equilibrium outcomes by the maximal delays t_1 and t_2, one can generate two more restrictions each, which "closes" the system of equations.

While appendix A.3 treats the general case, I provide simplified separate proofs for the case of linear-utility representations. Define for each player i the *minimal marginal patience* within the "equilibrium horizon" t_i of G_i as

$$\delta_i(t_i) = \min \{P_i(t) \,| t \in T,\ t \leq t_i\}$$

Note that for a present-biased player i this equals $P_i(0)$ irrespective of t_i. Also define the *minimal cost of a delay* by t periods in G_i as

$$c_i(t|v_i, v_j) = \begin{cases} 0 & t = 0 \\ \frac{v_i}{d_i(t)} + \frac{v_j}{d_j(t-1)} & t > 0 \end{cases}$$

The aforementioned key step in obtaining a characterisation of uniqueness, payoffs and outcomes is the following lemma, introducing t_i as an additional variable.

Lemma 1. *Under assumption 2, for any $i \in I$ and $t \in \{t' \in T | 0 < t' \leq t_i\}$,*

$$(x,t) \in Z_i^* \Leftrightarrow \frac{v_i}{d_i(t)} \leq x_i \leq 1 - \frac{v_j}{d_j(t-1)}$$

Moreover, $w_i = \delta_i(t_i) v_i$ and $t_i = \max\{t \in T | c_i(t|v_i, v_j) \leq 1\}$.

Proof. Consider G_i and take any $t \in \{t' \in T | 0 < t' \leq t_i\}$. If $(x,t) \in Z_i^*$ then the first inequality follows straight from the fact that $d_i(t) x_i \geq v_i$ by definition of v_i and the fact that i makes the initial proposal; since $(x,t) \in Z_i^*$ necessitates $(x, t-1) \in Z_j^*$, the second inequality follows from the same argument.

Now take any x which satisfies the two inequalities and consider strategies as follows, where, for simplicity, assume here that the payoff bounds are indeed obtained in some equilibrium (this issue is dealt with in the general proof of the appendix): at any round $t' < t$, the respective proposer, say $i' \in I$, offers the respondent, say j', a zero share and upon rejection of a positive offer the respondent obtains his best payoff $W_{j'}$, which satisfies $W_{j'} = 1 - v_{i'}$ (this is part of lemma 7). Upon rejection of a zero share, if $t' + 1 < t$ the same is true with roles reversed, and if $t' + 1 = t$ then the proposer, say $k \in I$, proposes x; upon a rejection by the respondent, say l, of a proposal x' this player's continuation payoff is $d_l\,(1)\, v_l$ if $x'_l \geq x_l$, and it is $W_l = 1 - v_k$ if $x'_l < x_l$. The inequalities ensure that at every on-path stage the respective proposer has no strict incentive to deviate; since the respective respondent's threats are defined via equilibrium payoffs in terms of v_1 and v_2 there is nothing to check except for the on-path round-t history where the inequalities must imply that the respective respondent l's continuation payoff $d_l\,(1)\, v_l$ does not exceed x_l; obviously, they do, however, imply the stronger property that $x_l \geq v_l$.

This implies that for any $t \in T$ with $t \leq t_i$ (now also allowing $t = 0$),

$$\begin{aligned}
\inf\left\{d_i\,(t)\,x_i\,|\exists x \in X,\ (x,t) \in Z_i^*\right\} &= d_i\,(t) \cdot \underbrace{\inf\left\{x_i\,|\exists x \in X,\ (x,t) \in Z_i^*\right\}}_{=\frac{v_i}{d_i(t)}} \\
&= v_i \\
\Rightarrow w_i &\equiv \inf\left\{d_i\,(1+t)\,x_i\,|(x,t) \in Z_i^*\right\} \\
&= \inf\left\{d_i\,(1+t) \cdot \frac{v_i}{d_i\,(t)}\middle| t \in T,\ t \leq t_i\right\} \\
&= \psi_i\,(t_i)\,v_i
\end{aligned}$$

Finally, $t_i = \max\{t \in T | c_i(t|v_i, v_j) \leq 1\}$ is an immediate consequence of the first part. □

Note that while w_i is determined by v_i at the "cost" of introducing the maximal delay t_i as an additional unknown, the maximal threats to the proposers, v_1 and v_2, in turn pin down t_i; one can easily verify that $|t_1 - t_2| \in \{0, 1\}$, as it must be the case. Once $(v_i, t_i)_{i \in I}$ is known, the sets of equilibrium outcomes of G_1 and G_2, respectively, can be characterised. Using the familiar approach of Shaked and Sutton (1984), which is to employ backwards induction on the payoff bounds, provides the following equation for v_i in terms of t_i because starting from i's weakest threat w_i two rounds of backwards induction must yield v_i (for details see lemmata 5-9 in appendix A.2):

$$v_i = 1 - \underbrace{d_j\,(1)\left(1 - \underbrace{\delta_i\,(t_i)\,v_i}_{=w_i}\right)}_{=W_j} \Leftrightarrow v_i = \frac{1 - d_j\,(1)}{1 - \delta_i\,(t_i)\,d_j\,(1)}$$

The expression for v_i is similar to that of the RubE except that player i's worst equilibrium payoff may be lower than i's RubE payoff if there is delay and i violates present bias within the equilibrium horizon.

Summarising this section's results so far, $(v_i, t_i)_{i\in I}$ must be a solution to the following system of four equations in four unknowns $(\tilde{v}_i, \tilde{t}_i)_{i\in I}$:

$$\tilde{v}_1 = \frac{1 - d_2(1)}{1 - \delta_1\left(\tilde{t}_1\right) d_2(1)} \tag{6}$$

$$\tilde{t}_1 = \max\left\{t \in T | c_1\left(t|\tilde{v}_1, \tilde{v}_2\right) \leq 1\right\} \tag{7}$$

$$\tilde{v}_2 = \frac{1 - d_1(1)}{1 - \delta_2\left(\tilde{t}_2\right) d_1(1)} \tag{8}$$

$$\tilde{t}_2 = \max\left\{t \in T | c_2\left(t|\tilde{v}_2, \tilde{v}_1\right) \leq 1\right\} \tag{9}$$

Existence of a solution to this system is guaranteed because the RubE payoffs together with zero delays, i.e. $(\tilde{v}_i, \tilde{t}_i)_{i\in I} = (x^{R,i}, 0)_{i\in I}$ solves it, as is quickly verified. This must be the case, of course, since the treatment here generalises the standard case. However, this observation suggests the following theorem characterising those pairs of players' preferences which yield a unique equilibrium, and whose proof illuminates the significance of a solution to the above system of equations as the existence of mutually self-confirming payoff-delay outcomes, of which the RubE outcomes $(x^{R,i}, 0)_{i\in I}$ are a special case, where j's threat offering i the present value of continuation outcome $(x^{R,i}, 0)$ coincides with j's RubE outcome $(x^{R,j}, 0)$. The necessity part of the theorem rules other such solutions out.

Theorem 2. *Under assumption 2, the RubE is the unique equilibrium if and only if the system of equations 6-9 has* $\left(\tilde{v}_i, \tilde{t}_i\right)_{i\in I} = \left(x^{R,i}, 0\right)_{i\in I}$ *as the unique solution.*

Proof. Because both $\left(v_i, t_i\right)_{i\in I}$ and $\left(x^{R,i}, 0\right)_{i\in I}$ solve this system, if there is a unique solution then they coincide, whence sufficiency follows.

For necessity, first note that any solution $\left(\tilde{v}_i, \tilde{t}_i\right)_{i\in I}$ other than $\left(x^{R,i}, 0\right)_{i\in I}$ has $\tilde{t}_i > 0$ as well as $\delta_i\left(\tilde{t}_i\right) < \delta_i(0)$ for some $i \in I$. Take such a solution $\left(\tilde{v}_i, \tilde{t}_i\right)_{i\in I}$, and, without loss of generality, let $\delta_1\left(\tilde{t}_1\right) = \delta_1\left(\hat{t}_1\right) < \delta_1(0)$ for $\hat{t}_1$ with $0 < \hat{t}_1 \leq \tilde{t}_1$; similarly, let $\hat{t}_2 \leq \tilde{t}_2$ be such that $\delta_2\left(\tilde{t}_2\right) = \delta_2\left(\hat{t}_2\right)$. Now consider the outcomes $\left(x, \hat{t}_1\right)$, with $x_1 = \tilde{v}_1/d_1\left(\hat{t}_1\right)$, and $\left(y, \hat{t}_2\right)$, with $y_2 = \tilde{v}_2/d_2\left(\hat{t}_2\right)$; these will be shown to be mutually self-confirming along the lines of the proof of the first part of lemma 1. This proof considers G_1 and delay $\hat{t}_1$ even only, and establishes $\left(x, \hat{t}_1\right)$ as an equilibrium outcome if both $\left(x, \hat{t}_1\right)$ and $\left(y, \hat{t}_2\right)$ can be used as threats; for the general construction see the proof of theorem 7 in appendix A.3.

For each $t < \hat{t}_1$, the respective proposer, say i, offers the respective respondent, say j, a share of zero, and upon a rejection of a positive offer when roles are reversed in the subsequent round, j offers i the a share equal to the present value of a continuation with $\left(x, \hat{t}_1\right)$ if $i = 1$, and with $\left(y, \hat{t}_2\right)$

if $i = 2$; if these are indeed anticipated as continuation values then the respondent is indifferent, so specify acceptance. Note that, for each $i \in I$, this present value equals $\delta_i \left(\hat{t}_i\right) \tilde{v}_i$ whence proposer i in t could obtain at most $\tilde{v}_i$ by deviating, ensuring no strict incentive to deviate from a zero offer. After $\hat{t}_1 - 1$ such rounds proposing player 1 offers player 2 a share of x_2, which is the lowest share this player accepts, because the two outcomes $\left(x, \hat{t}_1\right)$ and $\left(y, \hat{t}_2\right)$ are specified as continuation outcomes as follows: first, upon rejection of a proposal x' with $x'_2 \geq x_2$ the game continues with $\left(y, \hat{t}_2\right)$ which player 2 does not prefer over x_2 because

$$x_2 = 1 - \frac{\tilde{v}_1}{d_1 \left(\hat{t}_1\right)} \geq \delta_2 \left(\hat{t}_2\right) \tilde{v}_2$$

Second, upon rejection of an offer $x'_2 < x_2$, the game continues with player 2's offering a share of $\delta_1 \left(\hat{t}_1\right) \tilde{v}_1$ which is accepted at indifference, because another rejection is followed by $\left(x, \hat{t}_1\right)$; player 2 does not prefer acceptance of any such offer x'_2 over rejection because

$$1 - \delta_1 \left(\hat{t}_1\right) \tilde{v}_1 \geq x_2 = 1 - \frac{\tilde{v}_1}{d_1 \left(\hat{t}_1\right)}$$

Clearly, player 1 cannot do better than indeed proposing x which is accepted, establishing $\left(x, \hat{t}_1\right)$ as a self-confirming equilibrium outcome, given that $\left(y, \hat{t}_2\right)$ is an equilibrium outcome. Similar constructions can be made for the remaining three cases ($\hat{t}_1$ odd, and the two cases of G_2), proving that also $\left(y, \hat{t}_2\right)$ is self-confirming as an equilibrium outcome when $\left(x, \hat{t}_1\right)$ is an equilibrium outcome. In this sense the two outcomes are "mutually self-confirming". Because $\hat{t}_1 > 0$, this proves the necessity part. □

Recall lemma 1 in view of the construction in the proof of the above theorem for any solution $(\tilde{v}_i, \tilde{t}_i)_{i \in I}$ which is not the RubE: this construction not only establishes mutually self-confirming payoff-delay outcomes but also the associated payoffs $\tilde{v}_1$ and $\tilde{v}_2$; these, as threats, can be used to support the respective delays $\tilde{t}_1$ and $\tilde{t}_2$. This insight is useful for answering the question of which solution to the system of equations 6-9 is $(v_i, t_i)_{i \in I}$ in the general case of multiplicity, and thus for obtaining a characterisation of equilibrium outcomes. Define t^*_1 as the maximum over all $\tilde{t}_1$ such that $(\tilde{v}_i, \tilde{t}_i)_{i \in I}$ solves equations 6-9, and similarly t^*_2; these exist because delay is finite, as shown at the outset of this section. Let v^*_i be the associated solutions, respectively, to equations 6 and 8; note that v^*_1 is then the minimum of all $\tilde{v}_1$ such that $(\tilde{v}_i, \tilde{t}_i)_{i \in I}$ solves equations 6-9, and similarly for v^*_2. From examination of the functions c_i it is, however, clear that $(v^*_i, t^*_i)_{i \in I}$ solves equations 6-9, and by the initial argument of this paragraph, $(v_i, t_i)_{i \in I} = (v^*_i, t^*_i)_{i \in I}$.

Theorem 3. *Under assumption 2, $\left(v_i, t_i\right)_{i \in I} = \left(v^*_i, t^*_i\right)_{i \in I}$, and the sets of equilibrium outcomes and* *players' payoffs in G_i are given by*

$$\begin{aligned}
Z_i^* &= \left\{(x,0)\left|v_i^* \le x_i \le 1-\phi_j\left(t_j^*\right)v_j^*\right.\right\} \\
&\quad \cup \left\{(x,t)\in X\times T\setminus\{0\}\left|t\le t_i^*, \frac{v_i^*}{d_i(t)} \le x_i \le 1-\frac{v_j^*}{d_j(t-1)}\right.\right\} \\
U_i^* &= \left[v_i^*, 1-\psi_j\left(t_j^*\right)v_j^*\right] \\
U_j^* &= \left[\psi_j\left(t_j^*\right)v_j^*, d_j(1)\left(1-\psi_i\left(t_i^*\right)v_i^*\right)\right]
\end{aligned}$$

Proof. See the argument in the paragraph preceding the theorem's statement for the first part. The sets of equilibrium outcomes and player payoffs follow from lemma 1 together with the relationships established in lemmata 5-9 of appendix A.2. □

Note that, in general, the multiplicity obtained here does not require existence of a stationary SPE, because if both $t_1^* > 0$ and $t_2^* > 0$ then all outcomes and payoffs are spanned without the RubE; indeed, the RubE is just a special instance of the general property that any solution to system 6-9 has, which is that of providing mutually self-confirming outcomes.

In the case of exponential discounting two well-known results have been derived: (i) a greater discount factor yields greater bargaining power, meaning a greater equilibrium payoff, and (ii) moving first confers an advantage; see Osborne and Rubinstein (1990, Sections 3.10.2 and 3.10.3). Once there is multiplicity and delay, however, both of these conclusions need to be qualified.

First, the worst payoff a proposer may decrease when her discounting function is uniformly increased in a way that decreases her minimum marginal patience over the SPE-relevant horizon; in this case the payoff sets expand. Observe, however, that within the general class of preferences studied here, despite the multiplicity, a meaningful notion of bargaining power emerges as a generalisation of what the discount factor is under exponential discounting: a strong player is one with a high *minimal* marginal patience for a long horizon of delays.

Second, it may occur that the worst payoff of player 1 in G_1 is less than her best payoff in G_2. In this sense, it is not unambiguously clear that player 1 prefers to make the first proposal but, depending on which equilibrium is played, may prefer to be the respondent of the first round. This is true for player 1 if and only if the analogous property holds for player 2.

Both of these claims can be seen in the characterisation of appendix B for the initial example of section 1.1; while this suffices to prove the two points, general statements are formulated as corollaries 5 and 6, respectively, in appendix A.3.

6 Foundations of Time-Inconsistency and Delay

This section investigates instances of bargaining in which time-inconsistent preferences may arise from the specific environment. The previous results can be readily applied to study how such environmental aspects may inform bargaining outcomes.

First, and based on a recent theoretical literature that relates time-inconsistent discounting to non-linear probability weighthing in the presence of exogenous risk, I translate the basic bargaining game into an environment with a constant exogenous probability of bargaining breakdown. This is straightforward but also permits to investigate what shapes of probability weighting functions may cause delay.

Second, I consider yet another foundation for time-inconsistent preferences which is imperfect altruism—or, more generally, misaligned incentives—across different generations of delegates to a bargaining problem. Two communities bargain over how to share a common resource: each round they nominate a new delegate to the bargaining table where a delegate is biased toward agreements that take place within the horizon of her lifetime.

6.1 Breakdown Risk and Non-linear Probability Weighting

One motive for impatience in the sense of discounting future payoffs is uncertainty, such as mortality risk. Most recently, dynamically inconsistent discounting has been derived from violations of expected utility—specifically, the independence axiom—in an environment with non-consumption risk; see e.g. Halevy (2008) and Saito (2011). This literature seeks to simultaneously explain evidence on risk preferences such as the Allais paradox and evidence on time preferences such as decreasing impatience. In a manner analogous to how Binmore et al. (1986) translate the basic Rubinstein (1982) model into one where bargaining takes place under the shadow of a constant breakdown risk for expected-utility maximisers, I sketch here how the results of this paper can be used to study such a model where the bargaining parties violate expected utility. Bulding on Halevy (2008), suppose that, after each round, there is a constant probability of $1-r \in (0,1)$ that bargaining breaks down, leaving players without any surplus, and that a player i's preferences over agreements $x \in X$ with delay $t \in T$ have the following representation, which—for the sake of simplicity—has a linear utility function and involves breakdown risk as the sole source of discounting:

$$U_i\left(x,t\right) = g_i\left(r^t\right) x_i \tag{10}$$

The function $g_i : [0,1] \to [0,1]$ is continuous and increasing from $g_i(0) = 0$ to $g_i(1) = 1$; it is a so-called probability-weighting function, and such a decision-maker i is time-consistent if and only if g_i is the identity so i maximises expected (linear) utility. Redefining, for a given survival rate r, $g_i(r^t) \equiv d_i(t)$, all previous results can be applied. In particular, one can import theories of risk preferences suggesting non-linear probability weighting such as rank-dependent expected utility (Quiggin (1982)) or cumulative prospect theory (Tversky and Kahneman (1992)) into the basic bargaining model and study their implications.[19]

A qualitative feature of probability weighting that appears widely accepted in the context of cumulative prospect theory is overweighting of small and underweighting of large probabilities; graphically speaking, the probability weighting function has an inverse s-shape, e.g. as the following single-parameter weighting function proposed by Tversky and Kahneman (1992) with $\gamma_i \in (0,1]$:

$$g_i(\pi) = \frac{\pi^\gamma}{(\pi^\gamma + (1-\pi)^\gamma)^{\frac{1}{\gamma}}}$$

If both players' preferences have a representation as in equation 10 then it can easily verified that they satisfy present bias for $g_i(r^t) = d_i(t)$, whence theorem 1 implies that the RubE, where $d_i(1) = g_i(r)$, is the unique SPE. Since increasing γ_i means less underweighting of large probabilities, and more overweighting of small ones, the effect of this parameter on a party's bargaining power depends on the size of the breakdown risk.

The behaviour of the probability weighting function near the extreme points of zero probability and certainty is, however, difficult to assess. Kahneman and Tversky (1979, pp. 282-283) point out that the function is unlikely to be well-behaved there, and that it is both conceivable that there exist discontinuities at the extremes and that small differences are ignored. Proposed parametric forms, however, preserve smoothness with increasing steepness as probabilities approach 0 or 1. While a rigorous analysis of this issue is beyond the scope of this paper, theorem 2 suggests that the following properties of g_i may permit delay while retaining the qualitative property of an inverse s-hape in most of the interior: first, probability underweighting of large probabilities only up to a probability strictly less than one when combined with a sufficiently large survival rate, and, second, sufficient steepness for (strongly) overweighted small probabilities in the presence of a very low survival rate; both would cause present bias to fail within a short horizon of bargaining periods.

6.2 Imperfect Altruism in Intergenerational Bargaining

Suppose there are several communities with access to a productive resource. They decide over how to share it by means of bargaining. As long as these rights have not been settled, some surplus is forgone. Upon failure to agree the communities nominate a new delegate to engage in the bargaining on their behalf. I now sketch a simple version of this general problem.

[19] Of course, these theories are much richer than what the simple preferences I am using here can capture. For instance, in terms of cumulative prospect theory, I assume here that every agreement is perceived as a "gain".

Let there be two communities $i \in I$, each of which has a population of two members in any period $t \in T$: an old member (i, o) and a young member (i, y). Each member lives for two periods, where in the first half of her life a member is called young, and in the second half it is called old, and each young member reproduces so that its synchronous old member is replaced by a young one following disappearance. Assume that the surplus forgone until agreement is constant and preferences over delayed rewards feature imperfect altruism: at any point in time, for any split $x \in X$ of the resource with a delay of $t \in T$ rounds, where community i's share is equal to x_i,

$$
\begin{aligned}
U_{i,o}(x,t) &= \begin{cases} x_i & t = 0 \\ \gamma_i \delta_i^t x_i & t \in T \setminus \{0\} \end{cases} \\
U_{i,y}(x,t) &= \begin{cases} \delta_i^t x_i & t \in \{0,1\} \\ \gamma_i \delta_i^t x_i & t \in T \setminus \{0,1\} \end{cases}
\end{aligned}
$$

The two parameters δ_i and γ_i are both assumed to be lie in the interior of the unit interval. These preferences are supposed to capture that each member is imperfectly altruistic: while they do care somewhat about what happens to their community in their afterlife, they do so to an extent that is less than they care about their own lived future. Note that an old member's preferences are quasi-hyperbolic and satisfy present bias while a young member's preferences violate present bias because they can still look forward to a second period of lifetime.

Suppose then that in each round t a new member of each community is nominated to the bargaining table and contrast two different generational delegation schemes of community i, where each is given a potential rationale:

- i always sends the young member to the bargaining table because the young ones have less to lose which makes them stronger—call such a community Y_i
- i always sends the old member to the bargaining tables, the rationale for this being that the old ones have more to lose which makes them wiser—call such a community O_i

There are four possible games which may arise under such generational discrimination in delegation by each community: the set of player pairs is $\times_{i \in I} \{Y_i, O_i\}$. Note that each of these cases forms a stationary game which fits into the general class of games analysed in this paper, because the preferences over feasible outcomes of the two delegates engaging in bargaining are identical in any round.

To focus on one single community's fate against a given opponent depending on her delegation scheme, I will let community j's preferences be general and contrast Y_i with O_i. In any case, there is a unique RubE; against a given community j's scheme, in this RubE, community Y_i's payoff exceeds that of community O_i because $\gamma_i < 1$ implies a greater proposer payoff (and therefore also respondent payoff):

$$\frac{1-d_j(1)}{1-\delta_i d_j(1)} > \frac{1-d_j(1)}{1-\gamma_i\delta_i d_j(1)}$$

This underlies the rationale which posits that the young ones are stronger in bargaining.

While O_i is present biased and the RubE payoffs the worst possible equilibrium payoffs, Y_i violates present bias, giving rise to the possibility of delay. Instead of providing a full analysis of the respective system of equations 13-14, I propose a simple equilibrium construction similar to that in the first example of section 1.1. Let σ^* be the RubE with Y_i's respondent payoff equal to $\hat{x}_i = \delta_i\left(\frac{1-d_j(1)}{1-\delta_i d_j(1)}\right)$ and consider the following strategies of the (sub-) game in which j is the initial proposer:

- Round 1: j offers Y_i a share of $\gamma_i\delta_i^2\hat{x}_i$ which equals the smallest share which Y_i accepts; if, however, Y_i were to reject, the game moves into
- Round 2: Y_i demands the entire resource while the smallest share that j accepts is $d_j(1)(1-\gamma_i\delta_i^2\hat{x}_i)$; upon a rejection, bargaining continues in
- Round 3:
  - if in the previous round Y_i offered j nothing then the players follow strategies σ^* so there is immediate agreement with Y_i's share equal to $\hat{x}_i$;
  - otherwise, players continue as from round 1.

The crucial stage to check for optimal behaviour is when Y_i makes a proposal in round 2. Given j's strategy, comparing the two available agreements' respective values, these strategies indeed form an equilibrium if and only if

$$\delta_i\hat{x}_i \geq 1-d_j(1)\left(1-\gamma_i\delta_i^2\hat{x}_i\right) \Leftrightarrow \hat{x}_i \geq \frac{1-d_j(1)}{\delta_i\left(1-\gamma_i\delta_i d_j(1)\right)}$$

This is satisfied if both Y_i and j are rather patient about a delay of one period, and community Y_i is sufficiently impatient about a delay of two periods. Now call this equilibrium $\hat{\sigma}$ and repeat the construction where $\hat{\sigma}$ is used instead of σ^* and $\hat{x}_i$ is replaced by $\tilde{x}_i = \gamma_i\delta_i^2\hat{x}_i$. This will result in an equilibrium if and only if

$$\tilde{x}_i \geq \frac{1-d_j(1)}{\delta_i\left(1-\gamma_i\delta_i d_j(1)\right)} \Leftrightarrow 1 \geq \frac{1-\delta_i d_j(1)}{\gamma_i\delta_i^4\left(1-\gamma_i\delta_i d_j(1)\right)}$$

This construction may be further repeated and, depending on parameters, yield an equilibrium or not; for any given $\gamma_i \in (0,1)$, there will be large enough values of δ_i and $d_j(1)$ so equilibrium obtains.

An old community O_i may be considered wise, because when the game is played by (O_1, O_2), the RubE, which is efficient, is the unique equilibrium whereas the presence of a young community may cause delay and thus inefficiency.

7 Conclusion

This paper provides the first analysis of Rubinstein's (1982) seminal bargaining model for dynamically inconsistent time preferences without the restrictive assumption of stationary strategies. It produces a characterisation of equilibrium outcomes for general separable time preferences, theorem 3, from which all other results could be derived. Reflecting both the genesis of this paper and my anticipation of how the various implications would be received, I presented it as two main results. The first main result, theorem 1, establishes that if both players are most impatient about a single period's delay when this means forgoing an immediate agreement, then equilibrium is unique and in stationary strategies. The sufficient property has a clear interpretation as a form of present bias and all time-preferences commonly used in applications satisfy it, in particular quasi-hyperbolic, hyperbolic and exponential discounting preferences. Applied researchers interested in models which feature such preferences in the basic bargaining model may rely on this result: it disposes of the need to argue in favour of selecting the simple stationary equilibrium and thus of the uncertainty previously surrounding predictions based on it. Moreover, once present bias is accepted as a property of preferences, the details of time preferences going beyond the first period of delay from the immediate present are irrelevant to equilibrium; since the empirical evidence is arguably inconclusive about such detail, this robustness is also useful for further work.

In contrast, the second main implication of the general characterisation may, at this stage at least, be mostly of theoretical interest: if some player is more patient about a single round's delay which occurs from the very present than about one from a near future round, then, in general, there may be multiplicity and delay, both based on such a player's preference reversals. While most recent evidence in the domain of money rewards, when controlling for utility curvature and eliciting willingnesses to wait, has documented such violations of present bias—termed increasing impatience—it is too early to confidently judge the validity of this finding which may also be discovered to be an artefact of novel methodology.[20] Should it receive confirmation, however, this paper will constitute a first theoretical investigation of such preferences, and the equilibrium delay obtained may deserve greater as well as wider interest.

[20] For a list of studies which document increasing impatience, see the survey of Attema (2012).

In any event, this paper provides a first step towards the study of psychologically richer preferences in the basic model of bargaining; the short section 6 was designed to hint at this, e.g. suggesting how, in the presence of exogenous breakdown risk, the implications of rank-dependent expected utility and even cumulative prospect theory preferences may be investigated. More generally, the basic generalisation of the approach of Shaked and Sutton (1984) to proving uniqueness in the Rubinstein model, which allowed me to obtain the characterisation, may be useful in further theoretical research on bargaining with non-standard time preferences, or even for the study of general stochastic games with time-inconsistent discounting.

Several extensions of the present analysis beyond such "applications" are easily envisaged: since the quasi-hyperbolic (β, δ)-model is particularly popular in applied modelling, one may explore how robust the uniqueness and basic properties of equilibrium in the game studied here are to variations of the bargaining protocol. As I argued in section 4, it is the fact that such strict present bias causes a player's weakest point to fully enter immediate agreement equilibrium; section 3 might suggest, however, that the particular simplicity of such equilibrium under the alternating-offers protocol may prevent preference reversals from playing a role for equilibrium.

Moreover, the assumption of full sophistication seems unrealistic. One may therefore ask how predictions change once players may be naïve, at least partially.[21] This introduces the potential of learning (and teaching) through delay, which may take different forms for present-biased and non-present-biased players.[22]

Finally, since it has been argued that Strotz-Pollak equilibrium is too extreme in its assumption of fully non-cooperative selves and should be refined, in particular in view of proposition 3, it is an interesting question whether for plausible such refinements the multiplicity results disappears. In other words, how much and what kind of intrapersonal coordination is necessary to restore uniqueness more generally?

References

Akin, Z. (2007). Time inconsistency and learning in bargaining games. *International Journal of Game Theory 36*(2), 275 – 299. 11, 34

Andreoni, J. and C. Sprenger (2012). Estimating time preferences from convex budgets. *The American Economic Review (forthcoming)*. 2

Asheim, G. B. (1997). Individual and collective time-consistency. *The Review of Economic Studies 64*(3), 427–443. 17, 18

[21] O'Donoghue and Rabin (2001) develop such a concept in the context of the quasi-hyperbolic (β, δ)-model.

[22] Akin (2007) studies this aspect for the quasi-hyperbolic (β, δ)-model.

Attema, A. E. (2012). Developments in time preference and their implications for medical decision making. *Journal of the Operational Research Society 63*(10), 1388–1399. 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. 5

Binmore, K., A. Rubinstein, and A. Wolinsky (1986). The nash bargaining solution in economic modelling. *The RAND Journal of Economics 17*(2), 176–188. 11, 29

Burgos, A., S. Grant, and A. Kajii (2002). Bargaining and boldness. *Games and Economic Behavior 38*(1), 28–51. 11

Ferreira, J. L., I. Gilboa, and M. Maschler (1995). Credible equilibria in games with utilities changing during the play. *Games and Economic Behavior 10*(2), 284–317. 18

Fishburn, P. C. and A. Rubinstein (1982). Time preference. *International Economic Review 23*(3), 677–694. 2, 14

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

Fudenberg, D. and J. Tirole (1991). *Game Theory*. The MIT Press. 17

Goldman, S. M. (1980). Consistent plans. *The Review of Economic Studies 47*(3), 533–537. 16

Halevy, Y. (2008). Strotz meets allais: Diminishing impatience and the certainty effect. *The American Economic Review 98*(3), 1145–1162. 15, 29

Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica 47*(2), 263–292. 30

Kocherlakota, N. R. (1996). Reconsideration-proofness: A refinement for infinite horizon time inconsistency. *Games and Economic Behavior 15*(1), 33–54. 18

Kodritsch, S. (2012, April). A note on refinements of strotz-pollak equilibrium. 17

Laibson, D. I. (1994). *Essays in Hyperbolic Discounting*. Ph. D. thesis, MIT. 18

Laibson, D. I. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics 112*(2), 443–478. 2, 18

Noor, J. (2011). Intertemporal choice and the magnitude effect. *Games and Economic Behavior 72*(1), 255–270. 10

O'Donoghue, T. and M. Rabin (1999). Doing it now or later. *The American Economic Review 89*(1), 103–124. 16

O'Donoghue, T. and M. Rabin (2001). Choice and procrastination. *The Quarterly Journal of Economics 116*(1), 121–160. 16, 34

Ok, E. A. and Y. Masatlioglu (2007). A theory of (relative) discounting. *Journal of Economic Theory 137*(1), 214–245. 10, 14, 15

Osborne, M. J. and A. Rubinstein (1990). *Bargaining and Markets*. Academic Press, Inc. 2, 11, 18, 20, 28, 38, 40

Osborne, M. J. and A. Rubinstein (1994). *A Course in Game Theory*. The MIT Press. 16

Peleg, B. and M. E. Yaari (1973). On the existence of a consistent course of action when tastes are changing. *The Review of Economic Studies 40*(3), 391–401. 16

Pollak, R. A. (1968). Consistent planning. *The Review of Economic Studies 35*(2), 201–208. 16

Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior and Organization 3*(4), 323–343. 30

Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica 50*(1), 97–109. 2, 3, 10, 11, 14, 18, 29, 40, 45

Rubinstein, A. (1991). Comments on the interpretation of game theory. *Econometrica 59*(4), 909–924. 12

Saito, K. (2011, February). A relationship between risk and time preferences. 29

Shaked, A. and J. Sutton (1984). Involuntary unemployment as a perfect equilibrium in a bargaining model. *Econometrica 52*(6), 1351–1364. 6, 8, 10, 19, 22, 25, 34

Sorger, G. (2004). Consistent planning under quasi-geometric discounting. *Journal of Economic Theory 118*(1), 118–129. 18

Stahl, 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. 16

Takeuchi, K. (2011). Non-parametric test of time consistency: Present bias and future bias. *Games and Economic Behavior 71*(2), 456–478. 2, 5

Tversky, A. and D. Kahneman (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty 5*(4), 297–323. 30

Volij, O. (2002). A remark on bargaining and non-expected utility. *Mathematical Social Sciences 44*(1), 17 – 24. 11

A Statements and Proofs for the General Case

A.1 Stationary Equilibrium

The presentation of my results requires some further definitions. First, define for each player $i \in I$ a function $f_i : [0,1] \times T \to [0,1]$ which associates with every possible delayed share the minimal share which is sufficient compensation when received immediately, i. e. its present value in share terms:

$$f_i(x_i, t) = u_i^{-1}\left(\max\left\{u_i(0), d_i(t)\, u_i(x_i)\right\}\right)$$

These functions are well-defined and continuous in the first argument because for any $(x,t) \in X \times T$, $d_i(t)\, u_i(x_i) \in [0,1]$ and thus in the domain of u_i, and because each u_i is increasing and continuous, guaranteeing an increasing and continuous inverse function u_i^{-1}. For a fixed delay $t \in T$, each function $f_i(\cdot, t)$ is constant at zero for $x_i \in \left[0, u_i^{-1}\left(u_i(0)/d_i(t)\right)\right]$ (possibly the singleton $\{0\}$), and increasing on this interval's complement in $[0,1]$. Moreover, $t > 0$ implies $f_i(1,t) < 1$ (impatience).

Definition 7. A *Rubinstein pair* is any $(x^*, y^*) \in X \times X$ such that

$$\begin{aligned} y_1^* &= f_1(x_1^*, 1) \\ x_2^* &= f_2(y_2^*, 1) \end{aligned}$$

A Rubinstein pair is a pair of surplus divisions with the following property: facing proposal x^*, player 2 is indifferent between accepting and rejecting when she expects agreement on y^* in the subsequent round, and, similarly, player 1 is indifferent between accepting y^* and rejecting when she expects proposal x^* to be accepted subsequently. Note that

$$\begin{aligned} y_1^* &= f_1(x_1^*, 1) \\ &= u_1^{-1}\left(\max\left\{u_1(0), d_1(1)\, u_1(x_1^*)\right\}\right) \\ &\leq x_1^* \end{aligned} \tag{11}$$

Similarly, also $x_2^* \leq y_2^*$.

A Rubinstein pair only depends on the players' attitudes to single round's delays from the present to the next period. Its definition can be reformulated as a fixed point problem, which can be shown to have a solution on the basis of the properties of the functions $(f_i)_{i\in I}$:

$$\begin{aligned} x_1^* &= 1 - f_2\left(1 - f_1\left(x_1^*, 1\right), 1\right) \\ y_2^* &= 1 - f_1\left(1 - f_2\left(y_2^*, 1\right), 1\right) \end{aligned}$$

Lemma 2. *A Rubinstein pair exists.*

Proof. See (the first part of) the proof of Osborne and Rubinstein (1990, Lemma 3.2). □

The next definition constructs a pair of simple strategies on the basis of any Rubinstein pair; in its statement, $\mathbb{I}$ denotes the indicator function that evaluates to one if its argument is true and to zero otherwise. The definition's terminology will be justified by proposition 5 below.

Definition 8. For any Rubinstein pair (x^*, y^*), a *Rubinstein equilibrium (RubE)* is a strategy profile σ^R that satisfies the following: for any $t \in \rho^{-1}(1)$, $x \in X$, $h \in H^{t-1}$ and $h' \in H^t$,

$$\begin{aligned} \sigma_1^R(h) &= x^* \\ \sigma_2^R(h, x) &= \mathbb{I}\left(x_2 \geq x_2^*\right) \\ \sigma_2^R(h') &= y^* \\ \sigma_1^R(h', x) &= \mathbb{I}\left(y_1 \geq y_1^*\right) \end{aligned}$$

In fact, the following is true.

Proposition 4. *Any stationary SPE is a RubE.*

Proof. Let σ^* be a stationary SPE according to definition 1 in which player 1 always proposes x^*, player 2 always proposes y^* and each player $i \in I$ responds as indicated by acceptance rule $a_i : X \to \{0, 1\}$. Consider any $t \in \rho^{-1}(1)$ and $h \in H^{t-1}$ and suppose first that $a_2(x^*) = a_1(y^*) = 0$ so there is disagreement. Then, by SPE, a_2 must satisfy $a_2(x) = 1$ for any x with $x_2 > 0$ and the proposing player 1 self $(1, t)$ can increase her payoff from 0 under $\sigma^*(h) = x^*$ to $u_1\left(\frac{1}{2}\right) > 0$ by deviating to $\sigma_1(h) = \left(\frac{1}{2}, \frac{1}{2}\right)$.

Next suppose instead that $a_2(x^*) = 0$ and $a_1(y^*) = 1$ so the continuation outcome under σ^* is $(y^*, 1)$. Then, by SPE, a_2 must satisfy $a_2(x) = 1$ for any x with $x_2 > f_2(y_2^*, 1)$, whence $x_2^* \leq f_2(y_2^*, 1)$. Now argue that there exists an $\epsilon > 0$ such that the proposing player 1 self $(1, t)$ can increase her payoff above $d_1(1)\, u_1(y_1^*)$ under $\sigma^*(h) = x^*$ by proposing x' such that $x'_2 = f_2(y_2^*, 1) + \epsilon$ (recall that by impatience $f_2(y_2^*, 1) < 1$): if $f_2(y_2^*, 1) = 0$ then $f_2(y_2^*, 1) \leq y_2^*$, and otherwise

$f_2(y_2^*,1) < y_2^*$, so in any case $f_2(y_2^*,1) \leq y_2^*$, and because of $d_1(1) < 1$, continuity of u_i establishes the existence of $\epsilon > 0$ such that

$$u_1(1 - f_2(y_2^*,1) - \epsilon) > d_1(1)\, u_1(1 - y_2^*)$$

Apply a symmetric argument to conclude that σ^* must satisfy $a_2(x^*) = a_1(y^*) = 1$.

Finally, to prove that (x^*, y^*) must be a Rubinstein pair, note that $x_2^* < f_2(y_2^*,1)$ would contradict the optimality of $a_2(x^*) = 1$, and $x_2^* > f_2(y_2^*,1)$ would contradict the optimality of 1's proposing x^*, and in either case violate SPE, whence $x_2^* = f_2(y_2^*,1)$. A symmetric argument establishes $y_1^* = f_1(x_1^*,1)$ meaning (x^*, y^*) is a Rubinstein pair and σ^* is therefore a RubE. □

Observe the implication that any stationary SPE has immediate agreement in every round.

The next result justifies the terminology of RubE and, moreover, implies existence of stationary PCE, and, by consequence, also of stationary SPE, and of PCE and SPE generally.

Proposition 5. *Every RubE is a PCE.*

Proof. Take any Rubinstein equilibrium σ^R based on some Rubinstein pair (x^*, y^*), and consider any $t \in T$ odd and history $h \in H^{t-1}$, so it is player 1's round-t self's turn to propose. By adhering to the Rubinstein equilibrium, she obtains $U_1(x^*, 0) = u_1(x_1^*)$. Any other strategy's payoff is at most $\max\{d_1(1)\, u_1(y_1^*), d_1(2)\, u_1(x_1^*)\}$ because it results either in agreement in at least one more round where player 2 proposes y^* and player 1 accepts, or in agreement in at least two more rounds in a round where player 1 proposes some x with $x_2 \geq x_2^*$ since player 2 accepts, or in disagreement. The latter two outcomes are obviously no better than $(x^*, 0)$; and neither is the first, because by inequality 11,

$$d_1(1)\, u_1(y_1^*) \leq u_1(y_1^*) \leq u_1(x_1^*)$$

Next, consider any $x \in X$ and history (h, x) with h as before, so it is player 2's round-t self's turn to respond. Suppose first that $x_2 \geq x_2^*$ so any strategy σ_2 such that $\sigma_2(h, x) = 1$, and in particular σ_2^R, yields a payoff of $u_2(x_2) \geq u_2(x_2^*)$. Any other strategy σ_2 leads to a payoff of at most $\max\{d_2(1)\, u_2(y_2^*), d_2(2)\, u_2(x_2^*)\}$ because either there is agreement in a later round where 2 proposes on some y with $y_2 \leq y_2^*$ or there is agreement in a later round where 1 proposes on x^*, or disagreement. The latter two are obviously no better than $u_2(x_2^*)$; moreover, neither is the first because

$$u_2(x_2^*) = u_2(f_2(y^*, 1)) \geq d_2(1)\, u_2(y_2^*)$$

Second, suppose that $x_2 < x_2^*$ so σ_2^R yields a payoff of $d_2(1)\, u_2(y_2^*)$. Because $x_2^* = f_2(y^*, 1) > 0$, it follows that $d_2(1)\, u_2(y_2^*) = u_2(x_2^*)$. Therefore any alternative strategy σ_2 with $\sigma_2(h, x) = 1$ yields less. Any other strategy yields at most $\max\{d_2(1)\, u_2(y_2^*), d_2(2)\, u_2(x_2^*)\}$ which has been shown to not exceed $u_2(x^*)$ above.

A symmetric argument establishes that adhering to σ^R is optimal also for a proposing player 2 as well as a responding player 1. □

Corollary 1. *A PCE, and hence also a SPE, exist.*

Proof. A PCE exists because of lemma 2. The conclusion about SPE existence then follows from proposition 1. □

Moreover, since proposition 5 shows that every RubE forms a SPE, and a RubE is defined as a particular pair of stationary strategies, when combined with proposition 4, it also tells us that RubE is equivalent to stationary SPE.

Corollary 2. *A profile of strategies is a stationary SPE if and only if it is a RubE.*

Proof. Since a RubE is defined as a pair of stationary strategies based on a Rubinstein pair proposition 5 implies that every RubE is a stationary SPE (sufficiency). Proposition 4 is the converse (necessity). □

Because of lemma 2, proposition 5 tells us that a PCE, and by inclusion therefore also a SPE, exist. Uniqueness of stationary equilibrium thus is equivalent to uniqueness of a Rubinstein pair. The latter constitutes a combined restriction on the curvatures of the two players' utility functions. A sufficient condition in terms of individual preferences for this to hold is "increasing loss to delay" (Osborne and Rubinstein (1990, pp. 35-36)) which means that for each $i \in I$ and every $x_i \in [0,1]$, the "loss to delay" $x_i - f_i(x_i, 1)$ is increasing in x_i. The "loss to delay" for a given share is the additional compensation that makes a player willing to accept a one-period delay against the alternative of receiving this share immediately.[23]

Definition 9. A player i's preferences satisfy *increasing loss to delay* if, for any $x_i \in [0,1]$, $x_i - f_i(x_i, 1)$ is increasing in x_i.

The well-known and standard assumption of a differentiable concave u_i implies increasing loss to delay and thus a unique Rubinstein pair (Osborne and Rubinstein (1990, p. 35)); hence, this is true in particular of the case of linear utility functions u_i, which has by far received the most attention in the literature extending or applying Rubinstein's model, because then $f_i(x_i, 1) = d_i(1)\, x_i$ so the loss to one period's delay from the present is $x_i(1 - d_i(1))$.

Lemma 3. *If players' preferences satisfy increasing loss to delay then there exists a unique Rubinstein pair.*

Proof. Since only one-period delays are involved, this is simply reproducing Osborne and Rubinstein (1990, Lemma 3.2). □

[23] In his original paper, Rubinstein (1982, A-5 on p. 101) assumes only non-decreasingness of the loss to delay, which implies only that the set of Rubinstein pairs is characterised by a closed interval.

A.2 Present Bias and Uniqueness

Let f_i^U associate with any "rejection utility" U the minimal share a responding player i may accept.

$$f_i^U(U) = u_i^{-1}(\max\{u_i(0), U\})$$

Note that a player i's maximal possible rejection utility is $d_i(1)$: given a rejection the earliest best agreement delivers a share of one, hence utility $u_i(1) = 1$, in the round immediately succeding the rejection.

The first lemma lies at the core of the proof and says that for any continuation SPE outcome a respondent can expect upon rejection there exists an equilibrium with immediate agreement in which the proposer extracts all the benefits from agreeing earlier relative to the continuation outcome. In its statement z^0 denotes the outcome that is z after another round's delay; since this is payoff-relevant only for agreements (x,t), if $z = (x,t)$ then $z^0 = (x, 1+t)$.

Lemma 4. *For $\{i,j\} = I$ and any $z \in Z_i^*$, if $y_i = f_i^U(U_i(z^0))$ then $(y,0) \in Z_j^*$.*

Proof. Let σ^* be an SPE that induces z in G_i. Consider the following pair of strategies σ in G_j: j proposes y as in the statement, and i accepts a proposal y' if and only if $y_i' \geq y_i$. Upon rejection both continue play according to σ^* in G_i.

By construction of y via f_i^U, i's acceptance rule is optimal. By j's impatience, proposing y to have it accepted is then also optimal: among all proposals that i accepts j's share is maximal share in y, and rejection results in z^0 but $y_j = 1 - f_i^U(U_i(z^0)) \geq f_j^U(U_j(z^0))$. □

From this result immediately follows that disagreement is not an equilibrium outcome.

Corollary 3. *For any $i \in I$, $Z_i^* \subseteq X \times T$.*

Proof. Suppose to the contrary that for player i, there exists $z \in Z_i^* \cap D$ and note that this implies that also $z \in Z_j^*$. Let σ^* denote a SPE inducing z in G_i. Because $f_j^U(d_j(1)) < 1$ there exist proposals which player j accepts irrespective of what continuation outcome she expects; specifically, e. g. $x \in X$ with $x_j = \left(1 + f_j^U(1)\right)/2$ is such a proposal. Since $U_i(z) = 0 < u_i(x_i)$ such a proposal constitutes a profitable deviation for proposer i in G_i, a contradiction. □

The next lemma shows that no SPE with delay can yield a proposing player a payoff greater than all SPE without delay.

Lemma 5. *For any $i \in I$, $V_i = V_i^0$.*

Proof. Suppose $V_i > V_i^0$, implying that there exists an SPE agreement $(x,t) \in Z_i^*$ with $t > 0$ such that, for any $(x',0) \in Z_i^*$, $d_i(t)\, u_i(x_i) > u_i(x_i')$, and in particular $u_i(x_i) > u_i(x_i')$.

Accordingly, it must be that $(x,0) \in Z_j^*$. Applying lemma 4, for $y_j = f_j(x,1)$, $(y,0) \in Z_i^*$, whence $V_i^0 \geq u_i(y_i) \geq u_i(x_i)$, a contradiction.

Since $V_i \geq V_i^0$ by definition, the claim of the lemma is proven. □

The next lemma, in view of lemma 5, shows that a player's supremum SPE payoff when respondent is simply the once-discounted supremum SPE payoff when proposer, i. e. $W_i = d_i(1) V_i$. This relationship between a player's supremum SPE payoffs in her two different roles is the same as found under exponential discounting.

Lemma 6. *For any $i \in I$, $W_i^0 = d_i(1) V_i^0$ and $W_i = W_i^0$.*

Proof. It is straightforward to obtain the first equality:

$$\begin{aligned} W_i^0 &= \sup_{(x,0)\in Z_i^*} \{d_i(1) u_i(x_i)\} \\ &= d_i(1) \sup_{(x,0)\in Z_i^*} \{u_i(x_i)\} \\ &= d_i(1) V_i^0 \end{aligned}$$

For the second equality, suppose that $W_i < W_i^0$, saying that there exists $(x,t) \in Z_i^*$ with $t > 0$ such that, for any $(x',0) \in Z_i^*$, $d_i(1) u_i(x'_i) < d_i(1+t) u_i(x_i)$, and in particular $u_i(x'_i) < u_i(x_i)$. Now, $(x,0) \in Z_j^*$ must hold, and a construction similar to the one in the proof of lemma 5 can be employed to yield a contradiction. Since $W_i^0 \leq W_i$ by definition, also this part is thus proven. □

The next result relates the bounds on proposer and respondent SPE payoffs: the infimum SPE payoff of a proposer is simply the payoff resulting from immediate agreement when the respondent expects her supremum SPE payoff upon rejection; moreover, this statement holds true also when interchanging infimum and supremum.

Lemma 7. *For any $\{i,j\} = I$, $v_i = u_i\left(1 - f_j^U(W_j)\right)$ and $V_i = u_i\left(1 - f_j^U(w_j)\right)$.*

Proof. From the continuity and the increasingness of u_j it follows that

$$\begin{aligned} f_j^U(W_j) &= u_j^{-1}\left(\max\left\{u_j(0), \sup_{(x,t)\in Z_j^*}\{U_j(x,1+t)\}\right\}\right) \\ &= \sup_{(x,t)\in Z_j^*}\left\{u_j^{-1}\left(\max\{u_j(0), U_j(x,1+t)\}\right)\right\} \end{aligned}$$

Therefore, by continuity and increasingness of u_i,

$$u_i\left(1-f_j^U\left(W_j\right)\right) = u_i\left(\inf_{(x,t)\in Z_j^*}\left\{1-u_j^{-1}\left(\max\left\{u_j\left(0\right),U_j\left(x,1+t\right)\right\}\right)\right\}\right)$$
$$= \inf_{(x,t)\in Z_j^*}\left\{u_i\left(1-f_j\left(x_j,1+t\right)\right)\right\}$$

By a similar argument,

$$u_i\left(1-f_j^U\left(w_j\right)\right) = \sup_{(x,t)\in Z_j^*}\left\{u_i\left(1-f_j\left(x_j,1+t\right)\right)\right\}$$

Now, for the first equality, note that lemma 4 implies that

$$v_i \leq \inf_{(x,t)\in Z_j^*}\left\{u_i\left(1-f_j\left(x_j,1+t\right)\right)\right\} = u_i\left(1-f_j^U\left(W_j\right)\right)$$

It remains to show that this inequality cannot be strict. To do so, suppose to the contrary that there exists an outcome $z \in Z_i^*$ such that $U_i\left(z\right) < u_i\left(1-f_j^U\left(W_j\right)\right)$ and let σ^* be an SPE of G_i that induces it. Because $f_j^U\left(W_j\right) < 1$ must hold by impatience (there is at least one round's delay), the continuity of u_i guarantees existence of a proposal $x \in X$ such that $U_i\left(z\right) < u_i\left(x_i\right) < u_i\left(1-f_j^U\left(W_j\right)\right)$ which is accepted as $x_j > f_j^U\left(W_j\right)$ and thus constitutes a profitable deviation for i from σ^*, a contradiction.

For the second equality, note that lemma 4 implies that

$$V_i \geq \sup_{(x,t)\in Z_j^*}\left\{u_i\left(1-f_j\left(x_j,1+t\right)\right)\right\} = u_i\left(1-f_j^U\left(w_j\right)\right)$$

And since j rejects any proposal $x \in X$ with $x_j < f_j^U\left(w_j\right)$ it also follows that $V_i^0 \leq u_i\left(1-f_j^U\left(w_j\right)\right)$, which establishes the claim via lemma 5. □

Next, I will establish that there is no SPE with delay that is worse to the proposer than the worst SPE without delay, i. e. a result analogous to lemma 5 for a proposer's infimum payoffs.

Lemma 8. *For any $i \in I$, $v_i = v_i^0$.*

Proof. By lemma 4,

$$v_i^0 \leq \inf_{(x,t)\in Z_j^*}\left\{u_i\left(1-f_j\left(x_j,1+t\right)\right)\right\}$$

The proof of lemma 7 shows that

$$\inf_{(x,t)\in Z_j^*}\left\{u_i\left(1-f_j\left(x_j,1+t\right)\right)\right\} = v_i$$

Since, by definition, $v_i \leq v_i^0$ the claim follows. □

In general, however, only the first of the two properties of lemma 6 has an analogous version for infimum payoffs.

Lemma 9. *For any $i \in I$, $w_i^0 = d_i(1)\, v_i^0$.*

Proof. This is straightforward:

$$\begin{aligned} w_i^0 &= \inf_{(x,0)\in Z_i^*} \{d_i(1)\, u_i(x_i)\} \\ &= d_i(1) \cdot \inf_{(x,0)\in Z_i^*} \{d_i(1)\, u_i(x_i)\} \\ &= d_i(1)\, v_i^0 \end{aligned}$$

□

In view of lemma 8 this implies that $w_i^0 = d_i(1)\, v_i$. Yet, it may in general be the case that $w_i < d_i(1)\, v_i$ as a result of time-inconsistency. Present bias is sufficient to rule this out.

Theorem 4. *Suppose assumption 1 holds such that each player's preferences satisfy present bias according to definition 2. Then there exists a unique SPE if and only if there exists a unique Rubinstein pair, in which case this SPE is the RubE.*

Proof. From proposition 5 any Rubinstein pair has a corresponding SPE, proving the necessity of a unique Rubinstein pair as well as the claim that under uniqueness the SPE must be the RubE.

The main step towards establishing sufficiency is to show that present bias implies $w_i = w_i^0$. Suppose to the contrary that $w_i < w_i^0$, i. e. that there exists a continuation SPE agreement $(x,t) \in Z_i^*$ with $t > 0$ such that $d_i(1+t)\, u_i(x_i) < w_i^0$. By lemmata 9 and 8, this is equivalent to $d_i(1+t)\, u_i(x_i) < d_i(1)\, v_i$, which, moreover, implies that $v_i > 0$. The definition of v_i means that $v_i \le d_i(t)\, u_i(x_i)$ so $u_i(x_i) > 0$. Combining these yields a contradiction to present bias:

$$d_i(1+t)\, u_i(x_i) < d_i(1)\, d_i(t)\, u_i(x_i) \Rightarrow P_i(t) < P_i(0)$$

Recall lemmata 5 through 9, which imply that, for each $i \in I$, $\theta_i = \theta_i^0$ for any $\theta_i \in \{v_i, V_i, w_i, W_i\}$. In particular, therefore $v_i = v_i^0 = u_i(\underline{x}_i)$ and $V_i = V_i^0 = u_i(\overline{x}_i)$ for some $\{\underline{x}_i, \overline{x}_i\} \subseteq [0,1]$. Therefore, combining the respective equations obtained in those lemmata, any such $x_i^* \in \{\underline{x}_i, \overline{x}_i\}$ satisfies

$$\begin{aligned} u_i(x_i^*) &= u_i\left(1 - f_j^U\left(d_j(1)\, u_j\left(1 - f_i^U\left(d_i(1)\, u_i(x_i^*)\right)\right)\right)\right) \\ &\Leftrightarrow \\ x_i^* &= 1 - f_j^U\left(d_j(1)\, u_j\left(1 - f_i^U\left(d_i(1)\, u_i(x_i^*)\right)\right)\right) \end{aligned}$$

Now note that for any $k \in I$, $f_k^U(d_k(1)u_k(x_k)) = f_k(x_k, 1)$ so the above is equivalent to

$$x_i^* \;=\; 1 - f_j(1 - f_i(x_i^*, 1), 1)$$

Comparing this to the definition of a Rubinstein pair, it is easy to see that the two solutions to the last pair of equations—there is one per player—coincide with a Rubinstein pair. Assuming there is a unique one yields $v_i = V_i = u_i(x_i^*)$ from which uniqueness of SPE as the RubE is immediate by the players' impatience. □

Proposition 6. *Suppose assumption 1 holds. Then there exists a unique PCE if and only if there exists a unique Rubinstein pair.*

Proof. If there were more than one Rubinstein pair then there were more than one RubE, each of which is a PCE by proposition 5. This establishes necessity.

For sufficiency, interpret all payoff bounds on the basis of PCE, and note that all lemmata continue to apply. Recall that $v_i = v_i^0$ and $w_i^0 = d_i(1)v_i^0$. In a PCE a respondent i can ensure herself a payoff arbitrarily close w_i^0 by rejecting and controlling her subsequent proposal: for any $\epsilon > 0$, such a proposal x with $x_i = 1 - f_j^U(W_j) - \epsilon$ is accepted and yields a present value at the respondent stage of $d_i(1)u_i\left(1 - f_j^U(W_j) - \epsilon\right)$ which, by the continuity of u_i limits to w_i^0 as ϵ approaches zero. Hence $w_i = w_i^0$ and the fixed point argument of the proof of theorem 4 applies. □

In the more general case allowing for non-linear utility neither delay nor multiplicity imply intrapersonal conflict once there are multiple Rubinstein pairs and hence multiple RubE; this is true only for a unique Rubinstein pair. To see this point, take the multiplicity example of Rubinstein (1982, pp. 107-108), a version of which I reproduce here: players have preferences of the form $x_i - ct$ for each i, so the set of Rubinstein pairs is $\{(x, y) \in X \times X \,|x_1 - y_1 = c\}$, which is not a singleton, and where each element has a distinct RubE associated with it. Note that in the absence of any risk these preferences satisfy assumption 1 (see equation **??**), and, moreover, they are time-consistent. In particular, for $c \in (0, 1)$, both pairs (x, y) and (x', y') such that $(x_1, y_1) = (c, 0)$ and $(x_1', y_1') = (1, 1 - c)$ are Rubinstein pairs; let the associated RubE be denoted σ and σ', respectively, each of which is a PCE by proposition 5, and consider the following strategy profile:

- Round 1: player 1 demands the entire surplus, and player 2 accepts a proposal if and only if her share is at least $1 - c$, so there is a rejection and the game continues with
- Round 2:
  - if the previous offer to player 2 was positive then players continue with σ, resulting in immediate agreement on $y = (0, 1)$, and

– otherwise they continue with σ', yielding $y' = (1-c, c)$ without any (further) delay.

For $1 - 2c \geq c \Leftrightarrow \frac{1}{3} \geq c$ this is a PCE with delay.

A.3 General Characterisation Results

The crucial result to going beyond the sufficiency of present bias and finding necessary and sufficient conditions is the next lemma which permits an important simplification of the strategy space to consider that the theorem will exploit. To simplify its exposition I will first define a class of strategy profiles which attempt to implement a particular agreement by using extreme punishments. Described verbally, every such profile has the following structure: as long as both players have been complying with it and the agreement round has not been reached, a proposer claims the entire surplus and a respondent accepts only offers that yield her the maximal respondent payoff. Upon rejection of a different proposal the proposer is most severely punished and both players play the respondent's most preferred continuation SPE. In the agreement round the split to implement is proposed and the respondent accepts any split yielding her at least a share equal to that one; upon rejection of any such split the respondent's least preferred continuation SPE is played, and upon rejection of other splits the respondent's most preferred continuation SPE is played.

The following definition formalises this, where for $i \in I$, $e_i \in X$ will denote $x \in X$ such that $x_i = 1$, and for any $t \in T$, $h^t_{i,E} = (x^s, 0)^{t-1}_{s=0}$ is such that $x^s \in \{e_1, e_2\}$, $x^0 = e_i$ and $x^{s+1} \neq x^s$. Also, for any strategy profile σ, $t \in T$ and history $(h, x) \in H^t \times X$, let $\sigma^{0,(h,x)}$ denote the strategy profile that coincides with σ except, possibly, for the restriction that $\sigma^0_{\rho(t+1)}(h, x) = 0$.

Definition 10. For any $i \in I$ and G_i, and any $(x, t) \in X \times T$, a strategy profile $\sigma \in S$ is a *bang-bang strategy profile (BBSP) implementing* (x, t) *in* G_i if it satisfies the properties 1 through 5 below.

1. if $h = h^s_{i,E}$ for $s < t$ then
    - $\sigma_{\rho(s)}(h) = e_{\rho(s)}$ and,
    - for any $y \in X$, $\sigma_{\rho(s+1)}(h, y) = \mathbb{I}\left(y_{\rho(s+1)} \geq f^U_{\rho(s+1)}\left(W_{\rho(s+1)}\right)\right)$
2. if $h = h^t_{i,E}$ then
    - $\sigma_{\rho(t)}(h) = x$ and,
    - for any $y \in X$, $\sigma_{\rho(t+1)}(h, y) = \mathbb{I}\left(y_{\rho(t+1)} \geq x_{\rho(t+1)}\right)$
3. if $h = \left(h^s_{i,E}, y, 0\right) \neq h^{s+1}_{i,E}$ for $s < t$ and $y_{\rho(s+1)} < f^U_{\rho(s+1)}\left(W_{\rho(s+1)}\right)$ then

$$U_{\rho(s+1)}\left(z_{\left(h^s_{i,E}, y\right)}\left(\sigma^{0,\left(h^s_{i,E}, y\right)}\right)\right) > u_{\rho(s+1)}\left(y_{\rho(s+1)}\right)$$

4. if $h = \left(h^t_{i,E}, y, 0\right)$ then

   - if $y_{\rho(t+1)} \geq x_{\rho(t+1)}$ then

$$U_{\rho(t+1)}\left(z_{\left(h^t_{i,E},y\right)}\left(\sigma^{0,\left(h^t_{i,E},y\right)}\right)\right) \leq u_{\rho(t+1)}\left(x_{\rho(t+1)}\right)$$

   - if $y_{\rho(t+1)} < x_{\rho(t+1)}$ then

$$U_{\rho(t+1)}\left(z_{\left(h^t_{i,E},y\right)}\left(\sigma^{0,\left(h^t_{i,E},y\right)}\right)\right) > u_{\rho(t+1)}\left(x_{\rho(t+1)}\right)$$

5. $z(\sigma) = (x, t)$

The lemma below provides the main tool for characterising the temporal structure of Z^*_i and, consequently, the conditions which are necessary and sufficient for SPE to be unique under the more general preferences delineated by assumption 1.

Lemma 10. *For any $i \in I$ and $(x, t) \in X \times T$, $(x, t) \in Z^*_i$ if and only if there exists a BBSP implementing (x, t) in G_i which is an SPE of G_i.*

Proof. Sufficiency is obvious from property 5 of a BBSP implementing (x, t) in G_i.

For necessity, take any $(x, t) \in Z^*_i$ and begin with the first part of property 4. If respondent $\rho(t+1)$ preferred *every* continuation SPE outcome to immediate agreement on x then she would never accept it, a contradiction to (x, t) being an SPE outcome. Hence, there exists a continuation SPE with an outcome at least as good as immediate agreement on x, establishing existence of a strategy which satisfies the first part of property 4.

Moreover, $x_{\rho(t+1)} \leq f^U_{\rho(t+1)}\left(W_{\rho(t+1)}\right)$ must hold: in any SPE respondent $\rho(t+1)$ accepts any proposal y with $y_{\rho(t+1)} > f^U_{\rho(t+1)}\left(W_{\rho(t+1)}\right)$ for the reason that she prefers its immediate agreement over *any* continuation SPE outcome, meaning that proposer $\rho(t)$ would otherwise have a profitable deviation, e. g. proposing y with $y_{\rho(t+1)} = \frac{x_{\rho(t+1)} + f^U_{\rho(t+1)}\left(W_{\rho(t+1)}\right)}{2}$. Hence, for any y with $y_{\rho(t+1)} < x_{\rho(t+1)}$ it is true that $y_{\rho(t+1)} < f^U_{\rho(t+1)}\left(W_{\rho(t+1)}\right)$ and there exists a $(x', t') \in Z^*_{\rho(t+1)}$ such that $U_{\rho(t+1)}(x', 1 + t') > U_{\rho(t+1)}(y, 0)$. This yields that also the second part of property 4 can be satisfied. In fact, the very same argument applies to ensure existence of a strategy profile with property 3, and since properties 1 and 2 are feasible, all that remains to show in terms of existence is property 5.

Note that any strategy profile with properties 1 and 2 has property 5 if (and only if), for all $s < t$, it is true that $f^U_{\rho(s+1)}\left(W_{\rho(s+1)}\right) > 0$. Suppose then that $f^U_k(W_k) = 0$ for some $k \in I$ and let $l = 3 - k$; by definition of W_k this means that a respondent k accepts any proposal which specifies a positive share for her. Now, because to a proposer l any delayed agreement is worth at most

$d_l(1) < 1$ (recall the normalisation $u_l(1) = 1$), by continuity, $Z_l^* = \{(e_l, 0)\}$, which, by a similar argument, in turn implies that $Z_k^* = \{(x, 0)\}$ for $x \in X$ such that $x_l = f_l(d_l(1))$. Because for $t = 0$ the condition is vacuously true, this completes the proof that if (x, t) is an SPE outcome of G_i then there exists a BBSP implementing (x, t) in G_i.

To show that one can find an SPE among these strategy profiles, we only need to show that $(x, t) \in Z_i^*$ allows to rule out profitable deviations by respective proposer $\rho(s)$ after history $h = h_{i,E}^s$ for any $s \leq t$; by its construction, in any BBSP implementing (x, t) in G_i there are no other instances of profitable deviations. The following arguments demonstrate that if none of the BBSP implementing (x, t) in G_i were an SPE then $(x, t) \notin Z_i^*$.

First, consider such a case where $s < t$ and suppose $\rho(s)$ were to deviate to a split $y \neq e_{\rho(s)}$. If $y_{\rho(s+1)} \geq f_{\rho(s+1)}^U\left(W_{\rho(s+1)}\right)$ then this deviation would result in immediate agreement with a payoff to $\rho(s)$ of at most $u_{\rho(s)}\left(1 - f_{\rho(s+1)}^U\left(W_{\rho(s+1)}\right)\right)$. In regard of lemma 7, this *upper* bound on the deviation payoff equals $v_{\rho(s)}$ so such a deviation being profitable would require $U_{\rho(s)}(x, t - s) < v_{\rho(s)}$, implying $(x, t - s) \notin Z_{\rho(s)}^*$ and thus contradicting $(x, t) \in Z_i^*$.

Next, suppose there exists a deviation proposal y with $y_{\rho(s+1)} < f_{\rho(s+1)}^U\left(W_{\rho(s+1)}\right)$ such that, for any $(x', t') \in Z_{\rho(s+1)}^*$ with $U_{\rho(s+1)}(x', 1 + t') > u_{\rho(s+1)}\left(y_{\rho(s+1)}\right)$, it is true that $U_{\rho(s)}(x', 1 + t') > U_{\rho(s)}(x, t - s)$. Let y be such a proposal and recall lemma 6 which says that $W_{\rho(s+1)} = d_{\rho(s+1)}(1)\, V_{\rho(s+1)}^0$. By definition, for any $\epsilon > 0$ there exists $(x', 0) \in Z_{\rho(s+1)}^*$ such that $U_{\rho(s+1)}(x', 1) > W_{\rho(s+1)} - \epsilon$. Because existence of such a deviation y requires $W_{\rho(s+1)} > u_{\rho(s+1)}(0) \geq 0$, for $\epsilon \leq \left(1 - d_{\rho(s+1)}(1)\right) W_{\rho(s+1)}$ there exists $(x', 0) \in Z_{\rho(s+1)}^*$ such that $x'_{\rho(s+1)} > f_{\rho(s+1)}^U\left(\frac{W_{\rho(s+1)} - \epsilon}{d_{\rho(s+1)}(1)}\right) \geq f_{\rho(s+1)}^U\left(W_{\rho(s+1)}\right)$. For such an agreement, $u_{\rho(s)}\left(x'_{\rho(s)}\right) < u_{\rho(s)}\left(1 - f_{\rho(s+1)}^U\left(W_{\rho(s+1)}\right)\right)$. Recalling lemma 7, one observes that this implies, however, that $U_{\rho(s)}(x, t - s) < v_{\rho(s)}$ and thus $(x, t - s) \notin Z_{\rho(s)}^*$, a contradiction.

Finally, consider history $h_{i,E}^t$ and suppose proposer $\rho(t)$ were to deviate by proposing some split $y \neq x$: if $y_{\rho(t+1)} > x_{\rho(t+1)}$ then this deviation is also immediately accepted but yields the proposer a lower share, which cannot be profitable; and if $y_{\rho(t+1)} < x_{\rho(t+1)}$ then $y_{\rho(t+1)} < f_{\rho(t+1)}^U\left(W_{\rho(t+1)}\right)$ must hold so a similar argument to the one employed in the previous paragraph applies and ensures there exists a BBSP that deters this deviation. □

This lemma allows to characterise the set of SPE agreements and payoffs in terms of $(v_i)_{i \in I}$. Toward this end, define first, for each player i, a function $\hat{f}_i^U : [u_i(0), \infty) \to \mathbb{R}_+$ which extends f_i^U onto the domain of the entire non-negative real line such that

$$\hat{f}_i^U(U) = \begin{cases} f_i^U(U) & U \in \mathscr{U}_i \\ U & U > 1 \end{cases}$$

Corollary 4. *For any $i \in I$, if $f_i^U(W_i) = 0$ then $Z_j^* = (e_j, 0)$ and $Z_i^* = (x, 0)$ where $j = 3 - i$ and*

$x_i = 1 - f_j^U(d_j(1))$; *if* $f_1^U(W_1) \cdot f_2^U(W_2) > 0$ *then, for any* $i \in \{1,2\}$ *and* $t \in T \setminus \{0\}$,

$$(x,t) \in Z_i^* \Leftrightarrow x_i \in \left[\hat{f}_i^U\left(\frac{v_i}{d_i(t)}\right), 1 - \hat{f}_j^U\left(\frac{v_j}{d_j(t-1)}\right)\right]$$

Moreover, $w_i = \min\{P_i(t)|\, t \in T,\ t \le t_i\} \cdot v_i$.

Proof. The case of $f_i^U(W_i) = 0$ is straightforward (see the relevant argument in the proof of lemma 10).

Now suppose $f_1^U(W_1) \cdot f_2^U(W_2) > 0$ and take any $t \in T \setminus \{0\}$. If $(x,t) \in Z_i^*$ then $d_i(t)\,u_i(x_i) \ge v_i$ and $d_j(t-1)\,u_j(x_j) \ge v_j$ must clearly hold true, as otherwise there is a profitable deviation at one of the earliest two proposal stages. Noting that, for each $k \in I$, $v_k \ge u_k(0)$ and so $f_k^U(v_k) = u_k^{-1}(v_k)$, so the two implied inequalities for x_i yield the stated interval.

Next, take any $t \in T \setminus \{0\}$ and any x which satisfies the respective interval restriction. To have $(x,t) \in Z_i^*$ it is sufficient to prove existence of a BBSP implementing it in G_i which is a SPE. For existence only property 4 requires consideration because property 3 does not involve x and $f_1^U(W_1) \cdot f_2^U(W_2) > 0$ guarantees that property 5 can be met (properties 1 and 2 are feasible, hence unrestrictive). For property 4 it is sufficient that $x_{\rho(t+1)} \in \left[f_{\rho(t+1)}^U\left(w_{\rho(t+1)}\right), f_{\rho(t+1)}^U\left(W_{\rho(t+1)}\right)\right]$; if $\rho(t+1) = i$ then this is true because $\hat{f}_i^U\left(\frac{v_i}{d_i(t)}\right) \ge f_i^U(w_i)$ and $u_j\left(1 - f_i^U(W_i)\right) = v_j \le u_j\left(\hat{f}_j^U\left(\frac{v_j}{d_j(t-1)}\right)\right)$, and if $\rho(t) = i$ then it is true because $u_i\left(1 - f_j^U(W_j)\right) = v_i \le \hat{f}_i^U\left(\frac{v_i}{d_i(t)}\right)$ and $\hat{f}_j^U\left(\frac{v_j}{d_j(t-1)}\right) \ge f_j^U(w_j)$.

For SPE note simply that the interval restriction makes sure that no proposer has an incentive to deviate, which suffices since, by construction of a BBSP, existence alone implies optimal respondent behaviour.

Finally, to pin down w_i as a function of v_i and t_i note that the infimum proposer payoff is constant across all equilibrium delays (using corollary 3):

$$\begin{aligned} w_i &= \inf\{d_i(1+t)\,u_i(x_i)\,|(x,t) \in Z_i^*\} \\ &= \min\{d_i(1+t) \cdot \inf\{u_i(x_i)\,|(x,t) \in Z_i^*\}\,|t \in T_i\} \\ &= \min\left\{d_i(1+t) \cdot \frac{v_i}{d_i(t)}\,|t \in T_i\right\} \\ &= \min\{P_i(t)|\, t \in T,\ t \le t_i\} \cdot v_i \end{aligned} \tag{12}$$

□

This result allows one to prove a general theorem characterising uniqueness of SPE. It uses the following additional definitions, where $\psi_i : T \to (0,1)$ and $\pi_i : [0,1]^2 \times T \to \mathbb{R}_+$:

$$\begin{aligned} \psi_i(t) &= \min\{P_i(s) | \, s \in T, \; s \le t\} \\ \pi_i(v_i, v_j, t) &= \begin{cases} \hat{f}_i^U(v_i) + \hat{f}_j^U(d_j(1)v_j) & t = 0 \\ \hat{f}_i^U\left(\frac{v_i}{d_i(t)}\right) + \hat{f}_j^U\left(\frac{v_j}{d_j(t-1)}\right) & t > 0 \end{cases} \end{aligned}$$

Note that if i's preferences satisfy present bias then $\psi_i(t) = P_i(0) = d_i(1)$ for all $t \in T$. Also, if (x^*, y^*) is a Rubinstein pair as in definition 7 then $\pi_1(u_1(x_1^*), u_2(y_2^*), 0) = x_1^* + x_2^* = 1$ and $\pi_2(u_2(y_2^*), u_1(x_1^*), 0) = y_1^* + y_2^* = 1$. Importantly, the corollary also relates t_i to v_i and v_j.

Proposition 7. *Suppose assumption 1 holds. Then there exists a unique SPE if and only if there exists a unique solution $(\tilde{v}_i, \tilde{t}_i)_{i \in I} \in \times_{i \in I}(\mathscr{U}_i \times T)$ to the following system of equations where for each $i \in I$,*

$$\tilde{v}_i = u_i\left(1 - f_j^U\left(\psi_j(0)\, u_j\left(1 - f_i^U\left(\psi_i(\tilde{t}_i)\, \tilde{v}_i\right)\right)\right)\right) \tag{13}$$

$$\tilde{t}_i = \max\{t \in T \,|\pi_i(\tilde{v}_i, \tilde{v}_j, t) \le 1\} \tag{14}$$

In this case there is a unique RubE which then is the unique SPE.

Proof. First note that for $\tilde{t}_1 = \tilde{t}_2 = 0$ the system of equations 13-14 (for each i)–in what follows simply "the system"–reduces to one that is indeed equivalent to that defining a Rubinstein pair in terms of utilities. A Rubinstein pair exists by lemma 2 and, moreover, any such pair's associated proposer utilities yield $\tilde{t}_1 = \tilde{t}_2 = 0$ in the two equations 14, whence the system's set of solutions contains all those utilities obtained from Rubinstein pairs. From proposition 5, each of the latter is associated with a RubE that is an SPE. Hence, if SPE is unique then there can only be one solution with $\tilde{t}_1 = \tilde{t}_2 = 0$.

To prove the necessity part of the theorem it only remains to show that whenever there is a solution to the system with $\tilde{t}_i > 0$ for some $i \in I$ then there exists an SPE that is not a RubE. Suppose then there exists a solution $(\tilde{v}_1, \tilde{t}_1, \tilde{v}_2, \tilde{t}_2)$ with $\tilde{t}_i > 0$ for some $i \in I$. For each $i \in I$ let $\hat{t}_i \le \tilde{t}_i$ be such that $P_i(\hat{t}_i) = \psi_i(\tilde{t}_i)$. If $\hat{t}_1 = \hat{t}_2 = 0$ then $\tilde{t}_1 = \tilde{t}_2 = 0$ so let i be such that $\hat{t}_i > 0$. Consider agreements $(\hat{x}, \hat{t}_i)$ and $(\hat{y}, \hat{t}_j)$, where $\hat{x}_i = \hat{f}_i^U\left(\frac{\tilde{v}_i}{d_i(\hat{t}_i)}\right)$ and $\hat{y}_j = \hat{f}_j^U\left(\frac{\tilde{v}_j}{d_j(\hat{t}_j)}\right)$. It will be shown that both are SPE agreements, i. e. $(\hat{x}, \hat{t}_i) \in Z_i^*$ and $(\hat{y}, \hat{t}_j) \in Z_j^*$, by establishing that they are mutually self-enforcing: using them as continuation outcomes they are indeed equilibrium outcomes.

The key observation is that $(\hat{x}, \hat{t}_i)$ is weakly preferred by proposer i to satisfying j's demand when, subsequently, proposer j could push i down to her reservation share under continuation with

$(\hat{x}, \hat{t}_i)$ after another rejection (in fact, i is indifferent in the initial round):

$$d_i\left(\hat{t}_i\right) u_i\left(\hat{x}_i\right) \geq u_i\left(1 - f_j^U\left(d_j\left(1\right) u_j\left(1 - f_i^U\left(d_i\left(1 + \hat{t}_i\right) u_i\left(\hat{x}_i\right)\right)\right)\right)\right)$$

A similar point holds true about $(\hat{y}, \hat{t}_j)$ for proposer j.

Because $\hat{t}_i \leq \tilde{t}_i$ it is also true that $\hat{f}_i^U\left(\frac{\tilde{v}_i}{d_i(\hat{t}_i)}\right) + \hat{f}_j^U\left(\frac{\tilde{v}_j}{d_j(\hat{t}_i - 1)}\right) \leq 1$ and therefore $d_j\left(\hat{t}_i - 1\right) u_j\left(\hat{x}_j\right) \geq v_j = d_j\left(\hat{t}_j\right) u_j\left(\hat{y}_j\right)$. Hence, if both of $(\hat{x}, \hat{t}_i)$ and $(\hat{y}, \hat{t}_j)$ are SPE outcomes then they support $(\hat{x}, \hat{t}_i)$ as SPE outcome in G_i: for any $t < \hat{t}_i$, following a history $h_{i,E}^t$, $\rho(t)$ proposes $e_{\rho(t)}$ and respondent $\rho(t+1)$ accepts a proposal x if and only if

$$x_{\rho(t+1)} \geq \begin{cases} f_j^U\left(d_j\left(1\right) u_j\left(1 - f_i^U\left(d_i\left(1 + \hat{t}_i\right) u_i\left(\hat{x}_i\right)\right)\right)\right) & \rho(t) = i \\ f_i^U\left(d_i\left(1\right) u_i\left(1 - f_j^U\left(d_j\left(1 + \hat{t}_j\right) u_j\left(\hat{y}_j\right)\right)\right)\right) & \rho(t) = j \end{cases}$$

For $t = \hat{t}_i$, following a history $h_{i,E}^t$, proposer $\rho(t)$ proposes $\hat{x}$ and respondent $\rho(t+1)$ accepts a proposal x if and only if $x_{\rho(t+1)} \geq \hat{x}_{\rho(t+1)}$. A deviation by proposer i that is rejected is followed by j's proposing x such that $x_i = f_i^U\left(d_i\left(1 + \hat{t}_i\right) u_i\left(\hat{x}_i\right)\right)$ which is the smallest offer that i then accepts; if i rejects then an SPE implementing $(\hat{x}, \hat{t}_i)$ is played. A deviation by proposer j that is rejected is followed by i's proposing y such that $y_j = f_j^U\left(d_j\left(1 + \hat{t}_j\right) u_j\left(\hat{y}_j\right)\right)$ which is the smallest offer that j then accepts; if j rejects then an SPE implementing $(\hat{y}, \hat{t}_j)$ is played. It is clear that this construction supports $(\hat{x}, \hat{t}_i)$ as SPE outcome in G_i if $(\hat{x}, \hat{t}_i)$ and $(\hat{y}, \hat{t}_j)$ are indeed SPE outcomes. A similar construction can be devised to then also support $(\hat{y}, \hat{t}_j)$ as an SPE outcome. Thus the two are self-enforcing. The argument is complete and establishes a SPE with delay $\hat{t}_i$ in G_i, which is clearly not a RubE.

Sufficiency follows from the fact that $(v_i, t_i)_{i \in I}$ must be a solution to the system; hence, if there is a unique solution then it equals $(v_i, t_i)_{i \in I}$. □

While the uniqueness condition about the solutions to the system of (four) equations is not obviously interpretable in any useful way, by lemma 2, it requires a unique Rubinstein pair. For this case, the property that no other solutions exist is equivalent to the players' preferences not permitting constructions of self-supporting SPE outcomes with delay as the proof provides one. Technically, this is the novel phenomenon that may arise when preferences violate present bias.

A characterisation of SPE payoffs as well as outcomes is straightforward from this theorem on the basis of previous results. First note that existence of a solution $\tilde{v}_i$ to equation 13 for any $\tilde{t}_i \in T$ follows from the continuity of players' utility functions in a way similar to lemma 2. Now, for each $i \in I$, let B_i denote the set of pairs $(\hat{v}_i, \hat{t}_i) \in \mathscr{U}_i \times T$ such that, for some pair $(\hat{v}_j, \hat{t}_j) \in \mathscr{U}_j \times T$, $(\hat{v}_i, \hat{t}_i, \hat{v}_j, \hat{t}_j)$ solves the system of equations 13-14. Let $t_i^* = \max\{t \in T \,|\, \exists u \in \mathscr{U}_i,\ (u, t) \in B_i\}$, and let $v_i^* = \min\{u \in \mathscr{U}_i \,|\, \exists t \in T,\ (u,t) \in B_i\}$. Denote by $\mathscr{U}_k^{i,*}$ the set of SPE payoffs of player k in a subgame starting with player i's proposal.

Theorem 5. *Under assumption 1, $(v_i, t_i)_{i\in I} = (v_i^*, t_i^*)_{i\in I}$. Moreover, for $\{i,j\} = I$, the set of SPE payoffs in G_i is given by*

$$\begin{aligned}
\mathscr{U}_i^{i,*} &= \left[v_i^*, u_i\left(1 - f_j^U\left(\psi_j\left(t_j^*\right) v_j^*\right)\right)\right] \\
\mathscr{U}_j^{i,*} &= \left[u_j\left(f_j^U\left(\psi_j\left(t_j^*\right) v_j^*\right)\right), u_j\left(f_j^U\left(d_j\left(1\right) u_j\left(1 - f_i^U\left(\psi_i\left(t_i^*\right) v_i^*\right)\right)\right)\right)\right]
\end{aligned}$$

Proof. Take any $i \in I$ and note that by corollary 3 t_i^* is well-defined and, moreover, $t_i^* \le \bar{t}_i$. It is easily verified that $u_i\left(1 - f_j^U\left(\psi_j\left(0\right) u_j\left(1 - f_i^U\left(\psi_i\left(t\right) v_i\right)\right)\right)\right)$ is non-increasing in t so in fact $v_i^* = \min\{v_i \in \mathscr{U}_i \,|\, (v_i, t_i^*) \in B_i\}$. To establish that $(v_i^*, t_i^*)_{i\in I}$ is a solution to the system of equations 13-14 it needs to be shown that, for $\{i,j\} = I$, $t_i^* = \max\left\{t \in T \,\middle|\, \pi_i\left(v_i^*, v_j^*, t\right) \le 1\right\}$ but this follows from the fact that π_i is non-decreasing in each argument.

Each of $(v_i^*)_{i\in I}$ can be shown to be indeed an SPE payoff following the construction of SPE in the proof of theorem 7, whence $v_i \le v_i^*$. On the other hand, lemmata 6 and 7, when combined with equation 12 (from corollary 4), establish that a player i's infimum proposer payoff necessarily solve equation 13 which means that $v_i \ge v_i^*$. Hence we obtain $v_i = v_i^*$ and the payoff bounds follow from the relationships in lemmata 6 and 7.

Connectedness and $(t_i)_{i\in I} = (t_i^*)_{i\in I}$ are immediate consequences of corollary 4. Since the infimum proposer payoffs $(v_i)_{i\in I}$ are indeed SPE payoffs, closedness of the payoff intervals holds as well. □

For any two players i and i' with preferences representable as in assumption 1 such that $u_i = u_{i'}$ say that i' is *uniformly more patient* than i if, for all $t \in T\setminus\{0\}$, $d_{i'}(t) > d_i(t)$, or, equivalently, there exists a sequence $\epsilon(t)$ with $\epsilon(0) = 0$ and, for any $t \in T\setminus\{0\}$, $\epsilon(t) \in (0, \epsilon(t-1) + d_i(t-1) - d_i(t))$, such that, for any $t \in T$, $d_{i'}(t) = d_i(t) + \epsilon(t)$. Call any such sequence ϵ a *uniform patience increase of* d_i. In the bargaining game where i is replaced by i' against given opponent j, denote the resulting SPE payoff extrema and maximal SPE delays by applying to each the following pattern: $v'_{i'}$ is the minimal proposer payoff of i' and v'_j is the minimal proposer payoff of j.

Corollary 5. *Let $\{i,j\} = I$ and suppose $t_i > 0$. It is always possible to replace i with a player i' who is uniformly more patient than i such that $v'_j \le v_j$, $v'_{i'} \le v_i$ and $w'_{i'} < w_i$, which imply $[w_i, W_i] \subset [w'_{i'}, W'_{i'}]$, $[v_i, V_i] \subseteq [v'_{i'}, V'_{i'}]$, $[w_j, W_j] \subseteq \left[w'_j, W'_j\right]$, $[v_j, V_j] \subseteq \left[v'_j, V'_j\right]$ and $t_i \le t'_{i'}$ as well as $t_j \le t'_j$.*

Proof. Take $\hat{t} \le t_i$ such that $P_i\left(\hat{t}\right) = \psi_i(t_i)$ and let, for any $t \in T$ and any uniform patience increase ϵ of d_i, $P_i^\epsilon(t) = \frac{d_i(1+t)+\epsilon(1+t)}{d_i(t)+\epsilon(t)}$. Now choose ϵ as follows: $\epsilon(1) \in (0, 1 - d_i(1))$, for $t+1 \in T \setminus \left\{1, \hat{t}\right\}$, $\epsilon(t+1) = P(t)\,\epsilon(t)$ and for $t+1 = \hat{t}$, $\epsilon\left(\hat{t}\right) \in (0, P(t)\,\epsilon(t))$. Then, of course, $P_i^\epsilon(0) > P_i(0)$, but also for any $t+1 \in T \setminus \{1, \hat{t}\}$, $P_i^\epsilon(t) = P_i(t)$ and $P_i^\epsilon(\hat{t}) < P_i(\hat{t})$. Let i' be a player with $u_{i'} = u_i$ and $d_{i'}(t) = d_i(t) + \epsilon(t)$ for such a uniform patience increase. Therefore $\psi_{i'}(0) > \psi_i(0)$, implying $v_j' \leq v_j$ from inspection of h_j, and also $\psi_{i'}(t_i) < \psi_i(t_i)$, implying $v_{i'}' \leq v_i$ as well as $w_{i'}' < w_i$ from inspection of h_i. The remaining implications follow in a straightforward manner. □

Hence, there is always a way to make players more patient such that the set of supportable payoffs expands.

The second observation—that a player i may prefer not to be the initial proposer—means that $v_i^* < u_i\left(f_i^U\left(d_i(1)\,u_i\left(1 - f_j^U\left(\psi_j\left(t_j^*\right)v_j^*\right)\right)\right)\right)$. Since $v_i^* = u_i\left(1 - f_j^U\left(\psi_j(0)\,u_j\left(1 - f_i^U\left(\psi_i\left(t_i^*\right)v_i^*\right)\right)\right)\right)$ this is equivalent to

$$1 < f_j^U\left(d_j(1)\,u_j\left(1 - f_i^U\left(\psi_i\left(t_i^*\right)v_i^*\right)\right)\right) + f_i^U\left(d_i(1)\,u_i\left(1 - f_j^U\left(\psi_j\left(t_j^*\right)v_j^*\right)\right)\right) \tag{15}$$

Note that the right-hand side is the sum of the strongest threats, in share terms, that the two players have; the symmetry of this condition immediately reveals that $v_i^* < W_i^*$ and player i may not prefer to be the initial proposer if and only if this is true also about player j. To see that this is a possibility, suppose, without loss of generality, that $t_i^* > 0$ and let both players' one period discount factors approach one in a symmetric way, meaning that the RubE split converges to an equal split, and not that the right-hand side of the above inequality, by continuity, limits to

$$1 - f_i^U\left(\psi_i\left(t_i^*\right)v_i^*\right) + 1 - f_j^U\left(\psi_j\left(t_j^*\right)v_j^*\right)$$

Since, by the previous corollary, increasing players' one-period discount factors can only expand the sets of equilibrium outcomes and payoffs, there is still multiplicity and $f_i^U(\psi_i(t_i^*)v_i^*) < 1/2$, whence inequality 15 is satisfied in the limit.

Corollary 6. *Players' preferences are such that inequality 15 is satisfied if and only if for both $i \in I$, $v_i^* < W_i^*$.*

Proof. See the argument in the paragraph preceding the statement. □

B Theorem 3: an Example

While outside the class of preferences studied in this paper, it is particularly simple to illustrate the general result to example 1.1. The case of $k = 1$ is straightforward because of present bias so I focus on the novel phenomenon of multiplicity and delay due to Eve's dynamic inconsistency.

For Od it is certainly true that $w_1 = \alpha v_1$ irrespective of the maximal delay t_1; two rounds of backwards induction then yield that $v_1 = 1 - \beta(1 - \alpha v_1)$, i.e. $v_1 = x_1^*$. If the maximal delay when Eve proposes were zero then $w_2 = \beta v_2$ and there would be uniqueness of the stationary equilibrium as previously. If this maximal delay were positive, however, then Eve's worst threat would equal $w_2 = 0$ so there is an equilibrium in which Od achieves the maximal feasible payoff of 1 when proposing and the two steps of backwards induction yield Eve's worst payoff as a proposer of $v_2 = 1 - \alpha$. Indeed, the "residual" proposer advantage ensures she cannot obtain anything less than $1 - \alpha$ so the maximal delay Eve may experience as a proposer cannot exceed one period. It equals one if and only if, given one round's delay, the resulting most severe threats $v_1 = x_1^*$ and $v_2 = 1 - \alpha$ are sufficient to induce this delay, i.e. $(1 - \alpha)/\beta \leq 1 - x_1^*$ or, equivalently,

$$\frac{1}{1+\alpha} \leq \beta$$

Note that, as a function of α and β, v_2 in general has a discontinuity at the point where $(1+\alpha)^{-1} = \beta$ because for $(1+\alpha)^{-1} > \beta$, $v_2 = y_2^* \equiv (1-\alpha) / (1-\alpha\beta)$ but once β crosses the threshold of $(1+\alpha)^{-1}$ it becomes $v_2 = 1 - \alpha$. Hence, increasing β in fact can decrease Eve's worst payoff through the appearance of delay equilibria which exploit her then reduced minimal marginal patience.

For the sake of completeness, consider also the (sub-) game where Od makes the first move and proposes. The maximal equilibrium delay is at most two rounds and depends on parameters: it is positive if and only if $v_1/\alpha \leq 1 - v_2$, since $v_1 = x_1^*$, necessitates that $v_2 < y_2^*$ and hence delay in the (sub-) game when Eve is the initial proposer; in this case $v_2 = 1 - \alpha$ and the inequality becomes equivalent to

$$\frac{1+\alpha}{1+\alpha+\alpha^2} \leq \beta$$

This indeed implies existence of equilibrium delay when Eve moves first as the proposer; the maximal delay in "Od's game" then equals two if and only if the even stronger condition $v_1/\alpha^2 \leq 1 - (v_2/\beta)$ holds, and otherwise one. Note, however, that any delay that may occur in equilibrium when Od is the initial proposer is based on the multiplicity and the delay which arise jointly from Eve's time-inconsistency.

Finally, consider comparative statics; first, recall that Eve can be made "more patient" in the sense of a greater β while her worst payoff decreases because of the appearance of delay equilibrium; at the same time her best payoff, which arises in the stationary equilibrium, increases since it involves only the first-period discount factor. A limiting exercise where both α and β approach unity, but β approaches this limit sufficiently faster than α has $x_2^* = 1 - x_1^* \to 1$, and the sets of players' equilibrium payoffs then converge to the sets of feasible payoffs (which are all individually rational) since the non-stationary equilibrium remains intact and features the opposite extreme split.[24] It is also clear from this exercise that depending on which equilibrium is played, it may be that Od

[24] To be precise, the payoff pair which corresponds to Eve's obtaining the entire dollar is never an equilibrium payoff and thus not contained in the limit; it is, however, the only payoff pair with this property.

would prefer to be respondent initially rather than proposer.