The London School of Economics
and Political Science

Essays on Bargaining Theory and Welfare when Preferences are Time-Inconsistent

Sebastian Kodritsch

A thesis submitted to the Department of Economics of The London School of Economics and Political Science for the degree of Doctor of Philosophy, London, September 2013.

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others.

I declare that my thesis consists of 29945 words in total.

Abstract

This thesis studies time-inconsistent preferences. The first chapter introduces dynamically inconsistent time discounting into alternating-offers bargaining à la Rubinstein (1982), assuming players' utilities are linear in their share. For sophisticated players, it provides a characterisation of equilibrium uniqueness and a full characterisation of equilibrium outcomes. When players perceive a single period of delay from the present at least as costly as any such delay that takes place in the future, then equilibrium is unique and has immediate agreement. This property has a clear interpretation as present bias and is satisfied specifically by any form of hyperbolic and quasi-hyperbolic discounting. For violations of present bias within a short time horizon, which have recently been documented empirically, there exist multiple equilibria, where delayed agreement arises as an equilibrium outcome. Chapter two generalises the analysis to the entire class of separable time preferences for which, again, a full characterisation of equilibrium outcomes is obtained. Here, present bias can be combined with concavity of instantaneous utilities in order to be sufficient for uniqueness and immediate agreement. Chapter 3 is concerned with welfare properties of individual dynamic choice when preferences are time-inconsistent and there is perfect information. Applying the Pareto criterion to the sequence of temporal selves of the individual, it establishes two welfare properties of the standard solution concept for this case, which is Strotz-Pollak equilibrium: first, for finite-horizon problems without indifferences, Hammond's (1976) essential consistency is sufficient for choice to be Pareto-optimal. Second, if the problem satisfies a certain history-independence property, then the inefficiency of an equilibrium outcome implies it is Pareto-dominated by another equilibrium outcome, leading to an existence result for Pareto-efficient equilibrium.

Acknowledgements

I owe utmost gratitude to my supervisor Erik Eyster for being an unlimited and permanent source of inspiration, guidance and support. Many, in fact too many, years ago I presented him with my thoughts on the implications of expectations-based loss aversion for contract theory when, as the discussion had moved to the issue of ex-post bargaining, he posed the question "Do we even know what happens in Rubinstein-bargaining with (β, δ)-preferences?". Only a year or more later, I returned to this question, but this conversation marks the beginning of what culminated in this thesis.

Over the course of my work on bargaining, I discovered how it falls into a long tradition at LSE and STICERD: some of the foundational work of the 1980s on bargaining theory by Ken Binmore, Ariel Rubinstein, Avner Shaked, John Sutton and Asher Wolinsky appeared in the ICERD (now STICERD) discussion paper series. What a fateful thought that these giants on whose shoulders I am standing did (at least part of) their work on non-cooperative bargaining theory in the very same place on earth as I did!

I can hardly imagine a place more inspiring than LSE's department of economics, in particular the Theory Group of STICERD. It is characterised by a remarkable combination of curiosity and analytical rigour, of breadth in interest and depth in reasoning, within a truly international and diverse setting. Being part of this community has always felt like an incredible honour to me.

While I have certainly more or less benefitted from all members of LSE, the following have had a direct effect on the present work besides Erik Eyster: my advisor Balázs Szentes repeatedly challenged me in his uniquely candid way and prepared me for effective communication in the academic world. Michele Piccione was one of few who encouraged me to work on a characterisation result. Francesco Nava provided helpful comments on how to structure my thesis. Can Celiktemur pointed out the similarity between my approach to it and results in repeated games which helped me understand my work better.

I hope that all the colleagues who have made my time at LSE, STICERD and EOPP the most intellectually stimulating of my life know how much I appreciate them. Having procrastinated the completion of this work until the very deadline for submission, if I attempted to list all their names, I would be at too great a risk of forgetting some whom I feel I should include. I am convinced that I have spent time with some of the brightest, most energetic, wittiest, most knowledgeable and just most likeable people this world has to offer, and I am deeply grateful for it.

Last but most of all, I thank my family for their encouragement and support; Conny and Emil, the most wonderful gifts of my life, I owe you better times now.

Contents

Introduction

The study of dynamically inconsistent preferences has seen an enormous surge in economics within the past 15 years. Although in his pioneering analysis of such preferences, Robert Strotz had already presented anecdotal evidence of procrastination and commitment demand, suggesting the importance of self-control problems in economic decision making which are at odds with standard assumptions, it was the work of David Laibson in the mid- to late 1990s that provided the profession with a simple model of such dynamic inconsistency that since then has been used to explain a rich set of empirical phenomena: (β, δ)-discounting as a single-person reinterpretation of a functional form proposed by Edmund Phelps and Robert Pollak for imperfect intergenerational altruism.

An individual who recognises her dynamic inconsistency can act strategically to prevent certain future behaviours that she finds undesirable, e.g. invest in illiquid assets that prevent excessive consumption in the future. The presence of intertemporally conflicting objectives within the same decision maker raises the question of miscoordination of behaviour across time. Accordingly, the welfare properties of economic outcomes in the presence of dynamic inconsistency of preferences, viewed by some as a form of irrationality, deserve particular attention. This is a unifying theme of this work: while chapters 1 and 2 study the nature, in particular the efficiency, of outcomes when two parties with dynamically inconsistent time preferences bargain through time over how to share an economic surplus, chapter 3 investigates welfare properties of individual choice with time-inconsistent preferences. In the remainder of this introduction, I will briefly describe this work.

The aforementioned most popular model of self-control, (β, δ)-discounting, generates dynamic inconsistency through time preferences in a minimal way: the cost of delaying a future reward further are governed by parameter δ as in exponential discounting, whereas delaying immediate gratification involves additional impatience, captured through a second parameter β. While one focus of applied work has been on firm behaviour in view of such time-inconsistent consumers, analyses of strategic interaction by several time-inconsistent agents are rare. The core model of strategic bargaining theory developed by Ingolf Ståhl and Ariel Rubinstein provides, however, a setting in which time preferences matter: two parties jointly decide over how and when to divide a given economic surplus. Given the economic importance of bargaining in determining the terms of decentralised economic exchange and its prevalence in applied work ((re-) negotiation of contracts such as wage bargaining, household decision making, etc.), it appears especially important for further work with (β, δ)-discounting to have results about the outcome of bargaining when the parties have such time preferences.

Chapters 1 and 2 take up this question for rather general time preferences, including but not limited to (β, δ)-discounting; the more general chapter 2 investigates bargaining outcomes for the entire class of separable time preferences. The motivation for this is empirical: the success of the latter most popular model seems to rely on its capturing a very important feature of actual time preferences at the minimal cost of introducing one additional parameter; experimental work estimating time preferences in greater detail finds great qualitative heterogeneity of preferences which (β, δ)-discounting can, of course, not explain. From this point of view, my analysis can be seen as including various robustness checks.

Few analyses of bargaining with dynamically inconsistent time preferences are available, and all of them have been unduly restrictive in ruling out history-dependent behaviour. I argue (in chapter 1) that this assumption rules out whatever novel strategic implications dynamic inconsistency might have, and subsequently provide charaterisations of equilibrium outcomes for general strategies. The main findings can be summarised as follows.

(1) Present bias, which is a property of time preferences, where delaying a reward by one period from the immediate present is perceived (weakly) more costly than any one-period delay that takes place in the future, is sufficient for uniqueness of equilibrium (whenever instantaneous utility is concave). This is satisfied by all families of discounting typically used in economics: hyperbolic, exponential and (β, δ)-discounting. The unique equilibrium then displays the familiar properties of being "simple" (behaviour is history- and time-independent) and inducing immediate agreement (efficiency), with well-known comparative statics.

(2) For violations of present bias, where some player is more "sensitive" to a period of delay within the near future than a period of delay from the immediate present, novel equilibrium constructions, which do not require stationary equilibrium yield delayed agreement. Surprisingly, recent experimental designs that studied individual choices found clear evidence for such time preferences in domains which seem highly relevant for bargaining (money within short horizons).

Chapter 1 deals with the "textbook case" of separable time preferences with a representation where instantaneous utility is linear in the share obtained. It thus focuses on the novel phenomena that arise under dynamic inconsistency of time preferences in a set-up where the role of non-exponential discounting is most transparent, and it furnishes a simple example as well as applications and a discussion of empirical evidence on time preferences. In contrast, chapter 2 is more technically oriented and provides a characterisation of equilibrium outcomes for the entire class of separable time preferences. These have well-known axiomatisations and, moreover, (β, δ)-discounting has been found most convincing when interpreted in terms of consumption utility (rather than money directly). While the basic mechanisms discovered for time-inconsistent discounting in chapter 1 are similar in the more general case, the availability of formulae for the general case seems therefore important.

The final chapter, chapter 3, investigates welfare properties of sophisticated individual choice under general forms of time-inconsistency, in environments with perfect information. Using the standard solution concept to derive behaviour, which is Strotz-Pollak equilibrium, and applying an intrapersonal (equivalently, intertemporal) version of the classic Pareto criterion, it establishes two welfare properties: first, it shows that a property called essential consistency, which restricts the dynamic inconsistency of preferences, guarantees efficient choice (in this

Pareto sense). While essential consistency appears restrictive indeed, this result is obtained in a setting with arbitrary forms of history-dependence of welfare and serves to delineate those preferences which result in severe behavioural miscoordination across time (in the sense of Pareto-inefficiency). Second, if the decision problem satisfies a certain history-independence property, then the inefficiency of some equilibrium outcome implies that any outcome that Pareto-dominates it is also an equilibrium outcome. Thus, under standard technical assumptions, this result provides an existence result for efficient equilibrium outcomes. The two findings are applied to two influential analyses of choice under (β, δ)-discounting: procrastination/preproperation and overconsumption.

CHAPTER 1

Time-Inconsistent Discounting in Alternating-Offers Bargaining

1.1. Introduction

As a mechanism to share economic surplus, bargaining is pervasive in decentralised 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 Ståhl [1972] and firmly established in economics by the seminal work of Rubinstein [1982]. Building on Ståhl's disciplined formal description of the bargaining process as one where parties alternate in making and answering proposals, and extending it to situations without an exogenous deadline, Rubinstein [1982] reaches surprisingly sharp conclusions about how two completely informed and impatient parties share an economic surplus: under seemingly weak assumptions about the players' preferences, 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 by the proposer equals the smallest share the respondent accepts, given that upon a rejection, 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 a single-parameter extension of exponential discounting, the (β, δ)-model of "quasi-hyperbolic discounting": it emphasises a distinct time preference for intertemporal trade-offs involving the immediate present, a "present bias" governed by parameter β, leaving the long-run time preference over prospects which are in the future to satisfy exponential discounting with parameter δ. Thus this model captures one particular qualitative feature of hyperbolic discounting. In response to its success in applied work, also economists have turned to the experimental investigation of time preferences: while in the particular domain of single-dated monetary rewards, this literature has produced both (i) defence of exponential discounting, e.g. by Andreoni and Sprenger [2012], and (ii) further qualification of its violations

[1] Although Rubinstein [1982] works directly with preference relations, Fishburn and Rubinstein [1982] show that the axioms he imposes imply exponential discounting (see also Osborne and Rubinstein [1990, Section 3.3]).

for short delays, where increasing impatience has been observed, e.g. by Takeuchi [2011], when the domain of choice is actual consumption, the evidence for present bias remains strong, as most recently discussed and further confirmed by Augenblick et al. [2013].

In view of this evidence, one is naturally led to wonder whether the aforementioned results, which form the cornerstone of economists' thinking about bargaining, remain valid once time preferences take "richer" forms than exact 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 in view of the increased interest in economic applications using non-exponential discounting. 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?

To address all of these questions, I revisit the Rubinstein [1982] model for general separable and time-invariant preferences with a linear utility representation; more precisely, at any time during the bargaining, each player i evaluates a prospective division $x = (x_1, x_2)$ received after t periods of delay with utility $U_i(x,t) = d_i(t)\, x_i$, where d_i is a decreasing *discount function*. While the linearity of U_i in the share is restrictive, this case has received the greatest attention in the literature and, moreover, the resulting simplification allows to better focus on the role of non-exponential discounting. Observe that a player's such preferences 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 the well-studied case of exponential discounting. Hence this paper is most of all about time-inconsistent discounting.

I employ the standard equilibrium concept for games played by time-inconsistent players which is *Strotz-Pollak equilibrium (StPoE)*. 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. While some departures from exponential discounting have been analysed (see section 1.1.1), all of this work is restricted to stationary equilibrium. However, as I argue in section 1.3, the *assumption* of stationarity is particularly problematic under time-inconsistency: a stationary bargaining strategy is incapable of even *creating* preference reversals for the opponent. This paper is the first to characterise equilibrium 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 immediately present and future prospects are resolved, because these are governed separately by β. Nonetheless, for the present context, a player's discount function—her patience—for rewards delayed by $t+1$ periods, $d_i(t+1)$, can be usefully decomposed as the product of per-period discount factors, $d_i(t+1) = \prod_{\tau=0}^{t} P_i(\tau) = d_i(t)\, P_i(t)$, where $P_i(t) \equiv d_i(t+1)/d_i(t)$ measures what I (somewhat inappropriately) call a player i's *marginal patience*, her attitude to one additional period of delay for a given delay of t periods.

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,

$$P_i(0) \le P_i(t).$$

The unique equilibrium then indeed takes the simple stationary form described above, and attitudes to delay beyond one single period of bargaining turn out to be immaterial. The sufficient property of preferences can be interpreted as a weak form of *present bias*: it says that, for any given reward, an additional period's (or "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. Against the background of the increased interest in applications of the (β, δ)-model, this uniqueness result appears an important insight. Moreover, since present bias is a restriction on *individual* preferences, it also lends itself to empirical testing.

The second main result generalises the analysis to incorporate the possibility that present bias may fail to hold for some player. I obtain a *full characterisation of equilibrium outcomes* for the preferences assumed here. 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* (because players are impatient, any delayed agreement is Pareto-dominated by the corresponding division with immediate agreement). When a bargaining party's marginal patience falls below $P_i(0)$ within a short horizon of delays, 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-enforcing manner via non-stationary strategies. Rather surprisingly, the underlying property of preferences has recently been documented in several experimental studies of time preferences (e.g. Attema et al. [2010] and Takeuchi [2011]). A more general notion of bargaining power which emerges is the *minimal marginal patience* over a sufficiently long horizon from the present.[2]

1.1.1. Related Literature on Bargaining and Discounting. The literature studying variations of the Rubinstein-Ståhl model of bargaining is vast. The origin of this literature and my main inspiration is Rubinstein [1982]. His work extends the alternating-offers bargaining protocol proposed by Ståhl [1972] to an infinite horizon and a continuous agreement space. While his analysis dispenses with utility representations, the axioms imposed on the players' preferences necessarily imply an exponential-discounting representation where, moreover, the utility function is "not too convex".[3] For these preferences he characterises those surplus divisions that obtain in subgame-perfect Nash equilibrium (SPNE). Most textbook versions of the model as well as economists more generally have focused on his example of a linear utility function for which there exists a unique such equilibrium, which features stationary strategies and immediate agreement in every round (on as well as off the equilibrium path).[4] It is worthwile mentioning, however, that Rubinstein also covers preferences for which there is equilibrium multiplicity which in fact requires multiplicity of *stationary* SPNE; under certain conditions,

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

[3]These properties stem mainly from axioms 4 and 5, respectively. Osborne and Rubinstein [1990, Section 3.3] provide an illuminating discussion of (essentially) these axioms.

[4]Kreps [1990, Section 15.3], Fudenberg and Tirole [1991, Section 4.4] and Gibbons [1992, Section 2.1.D] are such examples of textbook treatments.

this even yields equilibrium delay. In contrast, the multiplicity and delay exhibited in this paper do not depend on multiplicity of stationary equilibria.

The theoretical literature since has gone on to study generalisations of this model regarding the protocol (e.g. Perry and Reny [1993]) or the surplus (e.g. Merlo and Wilson [1995]), but only recently have alternative time preferences been considered. Ok and Masatlioglu [2007] axiomatise preferences that imply forms of discounting that are more general than the ones I consider here in that they allow for particular forms of intertemporal non-transitivity; while there is separability, discounting emerges as relative to a particular intertemporal comparison instead of being absolute as a present-value calculation. The Rubinstein-Ståhl bargaining protocol serves as an application of this decision theory: proposition 2 claims that when the players' utility functions are strictly concave, then there is a unique "time-consistent" StPoE (p. 230), which is the familiar stationary equilibrium.[5] However, they do not define what they mean by "time-consistent" when it is used to qualify StPoE nor provide a proof, only indicating that the arguments of Shaked and Sutton [1984] apply.[6] Moreover, their arguments do not feature the strict concavity and would equally apply to the linear case dealt with here. My theorem 1.2 proves that for this case, without a refinement of StPoE, their proposition fails to hold when allowing for history-dependent strategies.

Noor [2011] generalises the exponential discounting model so the discount function depends on the size of the reward. This relaxation of separability also induces preference reversals of the type predicted by hyperbolic discounting and, additionally, can 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 to linear utility 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 only attitudes to delay of a single period, those preferences are indistinguishable from standard exponential preferences with non-linear utility.[7] While theoretically interesting, non-separability poses a conceptual challenge to the notion of time preference, because discounting then depends on the domain of choice and its units, and, consequently, a unitary measure of time preference becomes elusive.[8]

Akin [2007] studies bilateral bargaining with linear utility and (β, δ)-discounting. His focus is on naïveté about future preferences and learning from the opponent's rejection. Assuming stationary equilibrium conditional on beliefs,[9] 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 1.1 lends additional "credibility" to this analysis by showing that under sophistication there is indeed a unique stationary StPoE for such preferences.

[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 StPoE (see their footnote 15).

[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]).

[8] Frederick et al. [2002] discuss this issue in detail.

[9] The actual equilibrium concept is necessarily more involved than StPoE due to naïveté and learning.

In the context of an exogenous probabilistic risk of bargaining breakdown p after every round, abstracting from any "pure" time-discounting, one would set $\delta_i = 1-p$ to obtain results similar to those in Rubinstein [1982] (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 and concentrate their subsequent analysis on this equilibrium. Volij [2002] shows that when these preferences are restricted to being time-consistent, the model becomes equivalent to that of Rubinstein [1982].

In summary, besides the seminal work of Rubinstein [1982], all the papers discussed here assume stationary strategies in one way or another. As argued elsewhere, e.g. Osborne and Rubinstein [1990, p. 39], stationarity of strategies is problematic as an assumption in particular in bargaining. In the presence of time-inconsistency, there is an additional reason to be interested in non-stationary strategies, 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.

1.1.2. Outline. Section 1.2 defines the bargaining game, including the 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 subsequent proofs will use heavily. Section 1.3 studies stationary equilibrium and ends by arguing that stationarity, as an assumption on strategies, is particularly problematic in the analysis of bargaining with time-inconsistent preferences. This is followed by a section 1.4, pointing out the relationship between equilibrium delay and time-inconistency as well as laying the ground work for the subsequent analysis of uniqueness. Section 1.5 presents the first main result, which is the sufficiency of present bias for equilibrium uniqueness. This is generalised in section 1.6, 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 1.7 attempts to illustrate all the main results by means of a simple example, whereas section 1.8 sketches two "applications" of these results to bargaining environments where dynamic inconsistency is explicitly motivated from the respective environment. Section 1.9 concludes.

1.2. Model and Definitions

1.2.1. 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 $x \in \{(x_1, x_2) \in \mathbb{R}_+ | x_1 + x_2 = 1\} \equiv X$ to opponent player $3-\rho(t)$ (equivalently, offers the opponent a share $x_{3-\rho(t)}$), who responds by choosing $a \in \{0, 1\} \equiv A$, 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 on to the next round $t + 1$. 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 $H^{t-1} \equiv (X \times \{0\})^{t-1}$, and a terminal history ending in round t is an element of $H^{t-1} \times (X \times \{1\}) \equiv \mathcal{H}^t$; for completeness, let $H^0 \equiv \{h^0\}$. H^∞ denotes the set of non-terminal histories of infinite length.

A (pure) 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 $x \in X$,[10]

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

Let the space of all such functions of player i be denoted by Σ_i and define the space of strategy profiles $\Sigma \equiv \Sigma_1 \times \Sigma_2$.[11] 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 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 $h^s_\sigma \in H^s \cup (A \times H^{s-1})$ such that $(h, h^s_\sigma) \in \mathcal{H}^{t+s}$, 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 1.1. 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 $(h^t, x) \in H^t \times X$,

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

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. When "stationary" is used to qualify equilibrium, then this is supposed to mean that every player's strategy is stationary.

Note the normalisations regarding the size of the surplus and 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 of the protocol is that proposals are non-wasteful (players' shares add up to one). This is without loss of generality for the preferences assumed below where players only care about their own share which they want to maximise and obtain sooner rather than

[10]The restriction to pure strategies is standard in this model because it assumes away any risk in order to focus solely on the time dimension, as do the existing axiomatisations of time preferences.

[11]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.

later, and they can always choose to claim the entire cake. Due to their selfishness, a proposer who wants an offer accepted need not waste anything of what the opponent is willing to leave her, and, by claiming the entire cake, a proposer makes the "least acceptable" feasible offer anyways.

In what follows, I will use i to denote a typical element of I, "some player", and $j = 3 - i$ to denote "the other player", so $\{i, j\} = I$.

1.2.2. Preferences. In every round $t \in \mathbb{N}$ of bargaining, the domain of the players' preferences is assumed to be $(X \times \mathbb{N}_t) \cup \{D\}$, where $\mathbb{N}_t \equiv \{t' \in \mathbb{N} | t' \geq t\}$ and D is (perpetual) disagreement. Letting $T \equiv \mathbb{N}_0$ denote the set of possible delays of agreement, this domain can be expressed in terms of relative rather than absolute time as $(X \times T) \cup \{D\} \equiv Z$, which does not depend on t and will be referred to as the set of feasible outcomes (for any t). The preferences I consider in this paper are described by the following assumption.

ASSUMPTION 1. *In every round $t \in \mathbb{N}$, each player $i \in I$ has preferences over feasible outcomes Z represented by a function $U_i : Z \to [0, 1]$, which satisfies that*

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

where the function $d_i : T \to (0, 1]$ is decreasing with $d_i(0) = 1$ and $\lim_{\tau \to \infty} d_i(\tau) = 0$.

When using relative time, the domain of feasible outcomes as well as preferences are independent of absolute time. Impatience is captured by a decreasing "discount function" d_i: the more distant future receives less weight by the players. There are, however, no restrictions on the details of how much less and how this depends on the exact delay considered.

Although the various "selves" of a player "look alike" in terms of their preferences over the outcomes that are feasible when they are called upon to make a decision, corresponding to *the* individual's time preferences, dynamic inconsistency arises whenever different *periods* of delay carry different weights in the overall discounting. Define a function $P_i : T \to (0, \infty)$, which I will refer to as a measure of *marginal patience*, as follows:

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

This function is best interpreted in terms of an indifference condition: suppose $d_i(t) u = d_i(t+1) v$, so player i is indifferent between receiving u in t periods and waiting one additional period to receive v instead, where $v > u \geq 0$ and these are general (instantaneous) payoffs; then $u = P_i(t) v$ and $P_i(t)$ measures the minimal fraction of v that player i would find acceptable in order not to wait an additional period for obtaining v. A greater $P_i(t)$ means a larger such fraction and, accordingly, the $(t+1)$-th period of delay is less costly. 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.

Given the normalisation $d_i(0) = 1$, P_i encodes all information about d_i because for any $t \in T$,

$$d_i(t+1) = \prod_{\tau=0}^{t} P_i(\tau).$$

Now suppose for some $t \in T$ and some $s \in \mathbb{N}$, $P_i(t) \neq P_i(t+s)$, and take u and v to be such that i is indifferent between enjoying u in $t+s$ periods and $v > 0$ in $t+s+1$ periods, i.e. $u = P_i(t+s)v$. After s periods have elapsed, however, i will not be indifferent between the same consequences, since $u \neq P_i(t)v$. Constancy of P_i is therefore necessary for time-consistency. In fact, this is the defining property of exponential discounting which is well-known to be time-consistent. Therefore, the preferences studied here of a player i are time-consistent if and only if, for any $t \in T$, $P_i(t) = \delta$ with $\delta \in (0,1)$, i.e. $d_i(t) = \delta^t$.

I now define a property of preferences relating to this measure of marginal patience which turns out to be of great interest here.[12]

Definition 1.2. A player i's preferences satisfy *present bias* if, for any delay of $t \in T$ periods, $P_i(0) \leq P_i(t)$.

The significance of the property is clear: an individual with present bias considers a one-period delay most costly when it involves forgoing an immediate payoff. It is worthwhile pointing out that present bias is equivalent to any fixed delay of t periods being (weakly) more costly when it occurs from the immediate present than when it is added to an existing delay of one period: from cross-multiplication,[13]

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

As an example, exponential discounting, i. e. $d_i(t) = \delta^t$ for some $\delta \in (0,1)$, satisfies present bias in its weakest 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 present bias in a strict sense.

1.2.3. Equilibrium. In this section I introduce two equilibrium concepts for games with time-inconsistent players which are both adaptations of subgame-perfect Nash equilibrium (SPNE) and discuss them.[14] The first, Strotz-Pollak equilibrium, is the focus of the remaining analysis, and the purpose of the second—stronger—solution concept is only to clarify properties of the former. Underlying all definitions is the assumption that both protocol and preferences are common knowledge. This implies that the game has perfect information, and, in the terminology often used for the case of dynamically inconsistent preferences, players are perfectly

[12]The definition of "present bias" in Ok and Masatlioglu [2007, p. 225] is closely related, but somewhat stronger. Halevy [2008, Definition 1] introduces a concept identical to what I call here present bias but names 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 $\{s,t\} \subseteq T$,

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

[14]For an introduction to SPNE see e.g. the textbook on game theory by Osborne and Rubinstein [1994, Part II].

"sophisticated" about their own as well as their opponent's future preferences (Hammond [1976] is an early example of this usage in the context of individual choice).

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 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 $z_h(\sigma) = D$.

1.2.3.1. *Strotz-Pollak Equilibrium.* When a player's preferences over certain outcomes may change with the passage of time, 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 auxiliary game; for the origins of this concept, see Strotz [1955-1956] and in particular Pollak [1968].[15] Game-theoretically, the intrapersonal conflict is thus dealt with in exactly the same manner as interpersonal conflict. Specifically, this means that at any history of round t at which player i is to move, the self (i, t) of player i takes as given not only the behaviour of the opponent but also the behaviour of all other selves of player i; in other words, changing (i, t)'s strategy is then a "one-shot deviation". Adapting this idea to the present context results in the definition below.

Definition 1.3. A strategy profile σ^* is a *Strotz-Pollak equilibrium (StPoE, "equilibrium")* if, for any round $t \in \mathbb{N}$, history $h^{t-1} \in (X \times \{0\})^{t-1}$, proposal $x \in X$ and response $a \in \{0, 1\}$, the following holds:

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

This definition is really just an application of SPNE to the auxiliary game where the set of players is taken to be $I \times \mathbb{N}$. The well-known one-shot deviation principle guarantees that StPoE coincides with SPNE of the basic game played by I whenever players have time-consistent preferences (e.g. Fudenberg and Tirole [1991, Theorem 4.1], where continuity at infinity holds because of $\lim_{t\to\infty} d_i(t) = 0$). StPoE is the main concept I will use in this work: when referring to "equilibrium" I will mean StPoE.

1.2.3.2. *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 1.4. 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 $\sigma_j = \sigma_j^*$, 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 "commitment deviations", disregarding any commitment problems, whence PCE and SPNE also coincide under time-consistency of all players. Clearly, the test that a strategy profile has to pass in order to constitute a PCE is much stronger than that for StPoE.

[15]Further developments, in particular with regard to existence of StPoE, can be found in Peleg and Yaari [1973] and Goldman [1980].

PROPOSITION 1.1. *Any PCE is a StPoE.*

PROOF. Let σ^* be a PCE and restrict σ_i in definition 1.4 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. □

That, for time-inconsistent players, PCE is indeed stronger than StPoE will be demonstrated by means of the results further below (contrast proposition 1.4 and theorem 1.2). The observation that the two concepts lie at two opposite extremes in terms of commitment motivates the following terminology.[16]

DEFINITION 1.5. Any StPoE which is not a PCE is said to exhibit *intrapersonal conflict.*

In a PCE, *conditional on the opponent's strategy*, there is no intrapersonal conflict in the sense that, at any point in the game, no player would like to act differently in the future than under her PCE strategy. The condition is emphasised because the opponent's strategy determines the sets of feasible outcomes of a player, and there may be ways to limit a time-inconsistent player's choice sets such that there is no intrapersonal conflict in this sense (see the discussion in section 1.3).[17] It is this property that makes existence of a PCE remarkable; this turns out to be the case here for "stationary equilibria" of the type originally discovered in Rubinstein [1982] (see proposition 1.2).

1.2.4. Stationarity and 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, and time-inconsistency results from relating it across time. 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 equilibrium outcomes of G_i in relative terms as elements of Z will be referred to as Z_i^*. Proposition 2.2 below will ensure that both Z_1^* and Z_2^* are non-empty. Based on this set, I define the following payoff extrema, where the restriction to $Z_i^* \cap (X \times T)$, exluding disagreement, will also be justified below (see section 1.6): 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 equilibrium payoff of player i as the initial proposer in G_i, and W_i is the lowest upper bound on the equilibrium payoff of player i as the respondent conditional on rejection, i.e. the supremum continuation equilibrium 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

[16]Weaker refinements of StPoE 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 regarding future beliefs and behaviour, but there is yet to emerge a consensus on a viable alternative to StPoE. For such proposals, see Laibson [1994, Chapter 1], Ferreira et al. [1995], Kocherlakota [1996], Asheim [1997] and Plan [2010].

[17]For a trivial example consider an opponent who is able to dictate an outcome.

additional superscript of 0 indicate the restriction to immediate-agreement StPoE outcomes, e.g.

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

Note that both $W_i^0 = d_i(1) V_i^0$ and $w_i^0 = d_i(1) v_i^0$.

Moreover, I will introduce another lowest upper bound, the supremum equilibrium 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 may be used to relate and determine these variables, exploiting the stationarity property of the game which means that, as of the beginning of any two rounds t and $t+2$, the subgames have the same equilibrium outcomes. This has been effectively applied to the case of exponential discounters, where only $(v_i, V_i)_{i\in I}$ need to be considered because the equalities $W_i = d_i(1) V_i$ and $w_i = d_i(1) v_i$ that define a player i's best and worst threats, respectively, do not require further arguments. When allowing for time-inconsistent discounting preferences, this is not the case, however.

In any event, the game's stationarity property permits stationary strategies, and special attention will be given to equilibrium in such simple history- as well as time-independent strategies. The original construction of such equilibrium in this game goes back to Rubinstein [1982].

Definition 1.6. An equilibrium σ^* is a *Rubinstein equilibrium (RubE)* if for each $i \in I$, σ_i^* is a stationary strategy.

1.3. Stationary Equilibrium

The existing literature studying dynamically inconsistent time preferences in the bargaining protocol of Rubinstein and Ståhl, briefly reviewed in section 1.1.1, has constrained itself to an analysis of stationary strategies. The game's structure permits such simple strategies that ignore past play, and this section confirms previous results by establishing the existence and uniqueness of stationary equilibrium (RubE) also for the preferences considered here. At the same time, however, I show that the unique RubE actually forms a PCE. Based on this finding, I argue that in order to appreciate the implications of dynamically inconsistent time preferences in bargaining, one should include non-stationary strategies in the analysis.

Proposition 1.2. *There exists a unique RubE, which is given by the following strategy profile σ^R: for any $\{i,j\} = I$, $t \in \rho^{-1}(i)$, $h \in H^{t-1}$ and $x \in X$,*

$$\sigma_i^R(h) = x^{R,i},\ x_i^{R,i} = \frac{1-d_j(1)}{1-d_i(1)\,d_j(1)}$$

$$\sigma_j^R(h,x) = \begin{cases} 1 & x_j \geq 1 - x_i^{R,i} \\ 0 & x_j < 1 - x_i^{R,i} \end{cases}.$$

This equilibrium exhibits no intrapersonal conflict.

PROOF. Consider any stationary equilibrium σ. Because both strategies are stationary, it must induce either (i) agreement in every round or (ii) agreement only every other round or (iii) disagreement. Disagreement cannot be an equilibrium oucome, however. If it were, then any responding player, facing a payoff of zero upon rejection, would have to accept any positive share. Therefore all proposals would have to be such that the respective proposers demands the entire surplus. However, offering e.g. an equal split instead would constitute a profitable deviation. Hence σ induces agreement in at least every other round.

Consider therefore a round $\hat{t} \in \mathbb{N}$, with $\hat{t} > 2$, where agreement takes place, say on split $\hat{x} \in X$, and let $i \in I$ propose to $j \in I \setminus \{i\}$ in that round. Note that, by the stationarity of the strategies, $\sigma_i(h) = \hat{x}$ and $\sigma_j(h, \hat{x}) = 1$ for any history h with $h \in H^{t-1}$ for $t \in \rho^{-1}(i)$. In round $\hat{t} - 1$, i is the respondent and, facing any proposal $x \in X$ by proposer j, compares x_j to the payoff under rejection of $d_j(1)\hat{x}_j$. For σ to constitute a StPoE, it must be that she accepts x if and only if $x_i \geq d_i(1)\hat{x}_i$ and that j offers her exactly a share of $d_i(1)\hat{x}_i$.[18] Agreement must take place in every round.

Repeat this backwards-induction step once more to obtain that in round $\hat{t} - 2$, i must offer j a share of $d_j(1)(1 - d_i(1)\hat{x}_i)$. By stationarity,

$$1 - \hat{x}_i \equiv \hat{x}_j = d_j(1)(1 - d_i(1)\hat{x}_i) \Rightarrow \hat{x} = x^{R,i}.$$

This concludes the proof that σ^R is the unique RubE.

Next, I prove that σ^R is actually a PCE. First, observe that $0 < x_i^{R,j} = d_i(1)x_i^{R,i}$ for any $\{i, j\} = I$. Consider any round in which i is the proposer. Given σ_j^R, the set of feasible outcomes in terms of what i may achieve combines the following three cases: (i) any share in $\left[0, x_i^{R,i}\right]$ with any even number of periods of delay, including no delay, (ii) share $x_i^{R,j}$ with any odd number of periods of delay and (iii) disagreement D. Because of $x_i^{R,i} > x_i^{R,j} > 0 = U_i(D)$, mere impatience yields that immediate agreement on $x^{R,i}$ is i's optimal outcome.

Facing a proposal x, respondent j faces the following set of relevant outcomes: (i) immediate agreement with share x_j, (ii) any agreement with a share in $\left[0, x_j^{R,j}\right]$ and any odd number of periods of delay, (iii) agreement with share $x_j^{R,i}$ and any positive even number of periods of delay, (iv) disagreement D. Again, because $x_j^{R,j} > x_j^{R,i} > 0 = U_j(D)$, mere impatience yields that agreement on $x^{R,j}$ with one period of delay is most preferred among all of the feasible outcomes that involve delay; consequently, σ_j^R indeed yields the outcome that j finds optimal among all feasible outcomes. □

To understand this result, it is instructive to think about the textbook case of a finite horizon, which is not covered explicitly here, where backwards induction results in a unique equilibrium. In each round the proposer offers the opponent the present value of the unique continuation agreement who accepts it. Hence there is immediate agreement in any round and only the players' one-period discounting, $d_i(1)$, enters payoffs. The limits of the respective proposals and acceptance rules as the horizon becomes infinite exist and are independent of who moves last and of time. The resulting stationary strategies preserve the equilibrium property, which establishes existence of a stationary equilibrium. Once it is observed that there must be

[18] In case of indifference, equilibrium must involve acceptance here because otherwise an optimal proposal for i would not exist.

agreement in such an equilibrium, this also yields uniqueness. The proof provided is based on this idea but avoids explicit consideration of the finite-horizon case by exploiting stationarity directly.

The construction of σ^R is simple and familiar from the literature. And it is hardly surprising that this construction also provides an equilibrium here: the stationarity property of the game together with impatience means that at any point in the game, a player's problem effectively reduces to a two-period problem, where there is no scope for time-inconsistency to play out. Indeed, as proven, σ^R is in fact a PCE where each player's strategy creates a situation for the opponent in which the latter's time-inconsistency is entirely "neutralised".

Notwithstanding the appeal of this simple solution, *assuming* that both players adhere to a stationary strategy and thus, for any history of play, entirely disregard whatever behaviour they have observed, seems problematic and lacks a theoretical foundation. Osborne and Rubinstein [1990, p. 39] make this point very clearly in the context of bargaining by time-consistent players: assuming that player 1 adheres to a stationary strategy where she always proposes e.g. split $\left(\frac{3}{4}, \frac{1}{4}\right)$ implies that even after a long history in which player 1 proposed only equal splits instead (which were rejected), player 2 still expects her to offer her a share of $\frac{1}{4}$ next time.

In any case, the context of dynamically inconsistent time preferences provides further reason to move the analysis beyond stationary strategies: as I show below, constraining a player to a stationary strategy completely removes her of the opportunity to create—and potentially exploit—preference reversals of a (knowingly) time-inconsistent opponent (consider for instance the RubE of example section 1.7 for $k = 0$). To put it somewhat provocatively, this would be comparable to a study of imperfect competition under the assumption of marginal-cost pricing.

A preference reversal in this game must take the form that, subject to the feasible outcomes under the opponent's strategy, a player at some stage prefers some delayed outcome over the best immediate one (which she can implement "herself"), and yet, later, takes actions that induce another outcome that is worse than the originally envisaged one. Such reversals may take complex forms for general strategies of the opponent but 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: assume player 1 always proposes $\hat{x} \in X$ and follows acceptance rule a_1 such that $\hat{y}$ is the most preferred split for player 2 that she accepts.[19] 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, and
- either $(\hat{x}, 0)$ or $(\hat{y}, 1)$ when responding to player 1's proposal of $\hat{x}$.

Note that in order for player 2 to confront a 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:

$$d_2(1)\,\hat{x}_2 \overset{(>)}{\geq} \hat{y}_2 \quad \wedge \quad d_2(1)\,\hat{y}_2 \overset{(\geq)}{>} \hat{x}_2.$$

[19]While there may not exist a minimum of the set $\{y_1 \in [0,1] | a_1(y) = 1\}$, the continuity of preferences implicit in assumption 1 means that there exist values $\epsilon > 0$ such that the argument provided goes through with the sole modification of player 2's offering player 1 a share of $\hat{y}_1 + \epsilon$ instead.

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 actually intensifies when it becomes immediate and the latter is delayed instead.

To summarise the preceding discussion, while stationarity of equilibrium may be desirable as an eventual finding or appealing as a selection criterion, in the present context, its assumption renders time inconsistency uninteresting from a strategic point of view.

1.4. Time Inconsistency and Delay

The RubE has the property that it induces immediate agreement in any round. The dynamic inconsistency that manifests itself only when higher-order discounting enters equilibrium is strategically immaterial. The next result shows that whenever the payoff bounds are fully determined by immediate agreement equilibrium, then the RubE is in fact the only equilibrium. Thus for time-inconsistency to affect equilibrium there must be a delay equilibrium which yields a payoff extreme across all equilibria.

LEMMA 1.1. *The RubE is the unique equilibrium if and only if for both players $i \in I$, it is true that $V_i = V_i^0$, $v_i = v_i^0$, $W_i = W_i^0$ and $w_i = w_i^0$.*

PROOF. Necessity is clear by the properties of the RubE. For sufficiency, first observe that for both $i \in I$, $W_i^0 = d_i(1) V_i^0$ and $w_i^0 = d_i(1) v_i^0$ by definition, whence $W_i = d_i(1) V_i$ and $w_i = d_i(1) v_i$ because of the properties hypothesised. Lemmas 1.6 and 1.7, which are proven in the appendix, show that, for $\{i, j\} = I$, $V_i = 1 - w_j$ and $v_i = 1 - W_j$, respectively. Combining these yields

$$\begin{aligned} V_i &= 1 - d_j(1) v_j = 1 - d_j(1)(1 - d_i(1) V_i) \quad \Rightarrow \quad V_i = x_i^{R,i} \\ v_i &= 1 - d_j(1) V_j = 1 - d_j(1)(1 - d_i(1) v_i) \quad \Rightarrow \quad v_i = V_i. \end{aligned}$$

whence payoffs are unique and equal to the efficient RubE payoffs, implying that the RubE is the unique equilibrium. □

If it can be assumed that the payoff bounds are fully determined by immediate-agreement equilibrium, then the approach of Shaked and Sutton [1984] can be applied also to this case. The reason is that then each proposer's payoff extremes translate immediately into threat payoff extremes as $W_i = d_i(1) V_i$ and $w_i = d_i(1) v_i$. Suppose now that, say $v_i < v_i^0$ because there exists a delayed agreement $(x, t) \in Z_i^*$, $t > 0$, which yields i a payoff discretely less than any of the immediate equilibrium agreements, say $v_i^0 - \epsilon$ for some $\epsilon > 0$, so $x_i = \frac{v_i^0 - \epsilon}{d_i(t)}$. If i's time preferences are dynamically inconsistent, it is then not clear, however, whether this delayed equilibrium agreement also induces a worse threat for i as the respondent in G_j, $j = I \setminus \{i\}$, than the worst subsequent immediate equilibrium agreement, because the comparison is then $w_i^0 = d_i(1) v_i^0$ versus $d_i(t+1) x_i = \frac{d_i(t+1)}{d_i(t)} (v_i^0 - \epsilon)$.

The general problem in dealing with dynamically inconsistent time preferences, analytically, is that, despite the stationarity of the game, an additional period of delay to variously delayed agreements may change preferences over these. Thus, the relationship between an extreme payoff of a player as the proposer and the analogous one as a respondent is complicated.

Due to the players' impatience, which is understood to confer a natural strategic advantage upon the proposer, however, it is indeed true that $V_i = V_i^0$ and $W_i = W_i^0$, yielding $W_i = d_i(1)V_i$ (see appendix, where these are proven in lemmas 1.4 and 1.5, respectively). The greatest equilibrium payoff to a proposer is one where she extracts the maximal rent immediately, and the best threat a respondent has is based on that. Moreover, also $v_i = v_i^0$ (see appendix, lemma 1.7), because a proposer can always please the "most threatening" respondent immediately, whence the next result obtains.

PROPOSITION 1.3. *The RubE is the unique equilibrium if and only if $w_i = w_i^0$ for both players $i \in I$.*

PROOF. In this case, for any $\{i, j\} = I$, not only $v_i = 1 - d_j(1)V_j$ from combining lemmas 1.7 and 1.5, but also $V_j = 1 - d_i(1)v_i$, using lemma 1.6. Thus $v_i = V_i = x_i^{R,i}$ and $w_i = W_i = d_i(1)x_i^{R,i}$ for both $i \in I$, determining payoffs in any subgame for any player uniquely. Because of the players' impatience, since the two players' payoffs always add up to one, they cannot be obtained with delay; this yields unique offers in every round, which in turn pins down uniquely the acceptance rules, both as in the RubE. □

This proposition allows to focus the question of what kind of preferences yield uniqueness on a particular property of equilibrium, which the next section expoits.

1.5. Present Bias and Uniqueness

Applied work using strategic bargaining, e.g. wage-setting through negotiations by unions and firms or intra-household bargaining over how to share common resources, demands reliable predictions.[20] Under equilibrium multiplicity, the resulting ambiguity about the bargaining outcome feeds through all conclusions obtained from the model. 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. Ideally, such uniqueness can be guaranteed from properties of *individual* preferences which are more readily interpretable as well as testable.

At the same time, robustness is 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.

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 and empirically best-documented alternatives to exponential discounting.

THEOREM 1.1. *If each player's preferences satisfy present bias, then the RubE is the unique equilibrium.*

[20] For instance, Hall and Milgrom [2008] study the macroeconomic implications of strategic wage bargaining between workers and firms, and Chiappori et al. [2002] discuss the impact of various outside factors on household intra-household bargaining (they rely on a reduced-form structural model of household decision-making, however).

PROOF. This proof departs from proposition 1.3 and shows that $w_i = w_i^0$ indeed follows from present bias. Consider any player $i \in I$. First note that $w_i^0 = d_i(1)\, v_i^0$ by definition, so lemma 1.4 implies that $w_i^0 = d_i(1)\, v_i$.

Next, recall that present bias is equivalent to $d_i(1)\, d_i(t) \leq d_i(t+1)$ for all $t \in T$, which, starting from $w_i^0 = d_i(1)\, v_i$, implies the following inequality:

$$w_i^0 = d_i(1) \cdot \inf_{(x,t)\in Z_i^*} \{d_i(t)\, x_i\} \leq \inf_{(x,t)\in Z_i^*} \{d_i(t+1)\, x_i\} \equiv w_i.$$

Since $w_i \leq w_i^0$ holds by definition, this proves equality. □

Present bias ensures that a responding player, by rejecting, cannot obtain a worse payoff under a continuation equilibrium that has itself delay than under one with subsequent immediate agreement. Due to present bias, as next round's proposer, this player will be at least as impatient as the current round's respondent about a delayed outcome, so the threat of subsequent delay cannot confer an additional advantage to the proposing opponent in exploiting a present-biased player's time-inconsistency; in other words, the respondent is weakest—in terms of available "threats"—under subsequent immediate agreement, or $w_i = w_i^0$. In light of the discussion at the end of section 1.3, this theorem can be interpreted as follows: under present bias, it is never worthwhile creating preference reversals for the opponent through non-stationary strategies and delay, because full advantage of her time-inconsistency is taken already through immediate agreement in any round by means of stationary strategies.

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 the 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 effectively using the wrong model.

This section closes with a result which sheds further light on the RubE, and indirectly also on present bias in the bargaining context: the RubE is the only equilibrium which exhibits no intrapersonal conflict.

PROPOSITION 1.4. *The RubE is the unique PCE.*

PROOF. Recall that every PCE is a StPoE. Proposition 1.3 presents a sufficient condition for the RubE to be the unique StPoE, and proposition 1.2 says that the RubE is a PCE, whence that condition is in fact sufficient for the RubE to be the unique PCE.

Now simply note that a responding player i can always guarantee herself a payoff arbitrarily close to $w_i^0 = d_i(1)\, v_i^0$ (by definition) by committing to a subsequent proposal with j's share close enough from above to $W_j = 1 - v_i$ (by lemma 1.7 in the appendix), whence $w_i = w_i^0$. □

Note the following implication of this proposition for preferences where present bias is violated for at least one player and there is an equilibrium other than the RubE: such equilibrium necessarily exhibits intrapersonal conflict, meaning that some player at some stage would then prefer to change her own future actions. In fact, lemma 1.1 shows that every equilibrium other than the RubE must involve delay in some subgame (which may well be off the equilibrium path), so this applies in particular to the occurrence of delayed agreement, and this proposition therefore provides a sense in which any such equilibrium would arise purely from the dynamic inconsistency of some player's time preferences. The next section further investigates this possibility.

1.6. General Characterisation Results

Present bias restricts marginal patience at any delay. It is unnecessarily strong as a property sufficient for uniqueness, because the maximal delay in any subgame G_i, $i \in I$, can be bounded by the following simple rationality argument. Mere impatience guarantees existence of interior proposals by i which (any rational) respondent j immediately accepts: since the latter's rejection results in at least one period of delay and she cannot receive more than the entire cake, j's continuation payoff cannot exceed $d_j(1)$. Hence, even if proposer 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, a rational respondent j accepts any proposal x such that $x_j > d_j(1)$, whence the 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:

$$\bar{t}_i = \max\{t \in T \,|d_i(t) \geq 1 - d_j(1)\}. \tag{2}$$

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 this argument also establishes that $v_i > 0$, $W_i > 0$ and $V_i < 1$.

Even this simple argument involves both players' preferences, however. From proposition 1.3, the critical relationship between payoff bounds which potentially depends on the details of a player's preferences beyond the attitude to the first period of delay is that between the worst threat of a player i as a respondent, w_i, and i's lowest equilibrium payoff v_i as the initial proposer. While $w_i \leq d_i(1)\, v_i$ holds, because $v_i = v_i^0$, the main issue is whether and when a player's worst continuation payoff, w_i, might fall below the present value of the worst subsequent immediate agreement, w_i^0. The key to relating w_i to v_i is the introduction of the maximal delay t_i as an additional unknown: 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; 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.

Define, for each player i, the *minimal marginal patience* within the "equilibrium horizon" t_i of G_i as

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

This is based on the delay $t \leq t_i$ where an additional period of delay is perceived most costly because P_i reaches a minimum (by the argument at the beginning of this section, $t_i < \infty$, so a minimum exists). For a present-biased player i, this is the first period and $\delta_i(t_i) = P_i(0)$ irrespective of t_i. Also, define the *minimal cost of a delay* by t periods in G_i as

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

In each round, it is the proposer who has to have an incentive to not make any acceptable proposal, where the worst acceptable proposal yields the worst immediate payoff, i.e. v_i for both players $i \in I$. If a proposer is willing to wait t periods for a share, she is willing to wait any $t' < t$ periods, whence the first round before the agreement is the critical round. The idea of this "cost of delay" is that it corrsponds to the minimal total surplus that must be available in order to be able to promise each player sufficiently much after t periods of delay, when the players could obtain v_i and v_j, respectively, as the intermittent proposers. The promises, as of the initial round of G_i, must at least be $\frac{v_i}{d_i(t)}$ and $\frac{v_j}{d_j(t-1)}$, respectively, where the different denominators stem from the fact that i proposes first and j second (if at all).

The aforementioned key step in obtaining a general characterisation of uniqueness, payoffs and outcomes is the following lemma.

LEMMA 1.2. *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 $(x,t) \in Z$ where x satisfies the two inequalities and construct strategies as follows:[21] at any round $t' < t$, the respective proposer, say 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 $W_{j'} = 1 - v_{i'}$ by lemma 1.7. Upon rejection of a zero share, if $t' + 1 < t$ the same holds 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 of G_i, where the inequalities must imply that

[21] While the formulation assumes that the payoff extremes are indeed obtained in some equilibrium, this actually follows from the continuity of payoffs together with the compactness of action spaces.

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

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

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

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 because of alternating offers. Once $(v_i, t_i)_{i \in I}$ is known, the sets of equilibrium outcomes, $(Z_i^*)_{i \in I}$, can be characterised. To obtain the payoff bounds more readily, the familiar approach of Shaked and Sutton [1984], on which lemmas 1.3-1.7 in the appendix are based, can now be employed: these lemmas establish the relationships which allow the following backwards induction on v_i:

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

This expression for v_i is reminiscient to that in the RubE except that player i's worst equilibrium payoff may be lower than i's RubE payoff if $t_i > 0$ 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 $(\tilde{v}_i, \tilde{t}_i)_{i \in I}$:

$$\tilde{v}_1 = \frac{1 - d_2(1)}{1 - \delta_1(\tilde{t}_1)d_2(1)} \tag{3}$$

$$\tilde{t}_1 = \max\{t \in T | c_1(t|\tilde{v}_1, \tilde{v}_2) \leq 1\} \tag{4}$$

$$\tilde{v}_2 = \frac{1 - d_1(1)}{1 - \delta_2(\tilde{t}_2)d_1(1)} \tag{5}$$

$$\tilde{t}_2 = \max\{t \in T | c_2(t|\tilde{v}_2, \tilde{v}_1) \leq 1\}. \tag{6}$$

Existence of a solution to this system is guaranteed because the RubE payoffs together with zero delays, i.e. $(\tilde{v}_i, \tilde{t}_i)_{i \in I} = (x_i^{R,i}, 0)_{i \in I}$ constitute one indeed, as is easily verified.

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 a pair of self-enforcing payoff-delay outcomes, of which the RubE version $(x_i^{R,i}, 0)_{i\in I}$ is a special case. The necessity part of the theorem rules out any other solution.

THEOREM 1.2. *The RubE is the unique equilibrium if and only if the system of equations 3-6 has* $\left(\tilde{v}_i, \tilde{t}_i\right)_{i\in I} = \left(x_i^{R,i}, 0\right)_{i\in I}$ *as the unique solution.*

PROOF. Because both $(v_i, t_i)_{i\in I}$ and $\left(x_i^{R,i}, 0\right)_{i\in I}$ 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 $\left(x_i^{R,i}, 0\right)_{i\in I}$ 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 self-enforcing along the lines of the proof of the first part of lemma 1.2. The 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; the other cases follow from a similar argument.

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 a share equal to the present value of a continuation with $\left(x, \hat{t}_1\right)$ if $i = 1$, and $\left(y, \hat{t}_2\right)$ 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 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 respective cases of G_2), eventually 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. Thus, this pair of outcomes is self-enforcing. Because $\hat{t}_1 > 0$, this proves the necessity part. □

Recall lemma 1.2 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 self-enforcing 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 3-6 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 3-6, 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 3 and 5; 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 3-6, 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 3-6, and by the initial argument of this paragraph, $(v_i, t_i)_{i\in I} = (v_i^*, t_i^*)_{i\in I}$. Thus the following characterisation is obtained, where U_i^* and U_j^*, $\{i, j\} = I$, are the sets of equilibrium payoffs as of the initial round of subgame G_i.

THEOREM 1.3. *The set of equilibria satisfies the following properties:* $(v_i, t_i)_{i\in I} = (v_i^*, t_i^*)_{i\in I}$, *and, for* $\{i, j\} = I$,

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

PROOF. For $(v_i, t_i)_{i\in I} = (v_i^*, t_i^*)_{i\in I}$ see the argument in the paragraph preceding the theorem's statement. The sets of equilibrium outcomes and player payoffs follow from lemma 1.2 together with the relationships established in lemmas 1.3-1.7 in the appendix. □

Note that, in general, the multiplicity obtained here does not even rely on the existence of stationary equilibrium, 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 3-6 has, which is that of self-enforcing outcomes. This shows how certain forms of time-inconsistency invite non-stationary strategies, creating a role for them to exploit preference reversals and permitting equilibrium constructions reminiscient of those for repeated games, but despite the absence of the latter's punishment mechanisms from bargaining (the work of Busch and Wen [1995] and their discussion elucidate the relationship between repeated games and infinite-horizon alternating-offers bargaining). This stands in marked contrast to the available constructions for delayed agreement in extensions of the protocol of Rubinstein [1982], all of which involve stationary equilibrium.[22]

What kind of preferences permit delayed agreement then? Since equilibrium delay necessitates that at least one player violates present bias, these are rather unfamiliar, of course, to most economists. Qualitatively, the property of a player's time preferences that is conducive to delay is a sharp relative drop in the discount function and thus in marginal patience at a positive but small delay. Rather surprisingly, this feature is in line with the results of several recent experimental studies of time preferences in the domain of single monetary rewards for a majority of participants, e.g. by Attema et al. [2010] and Takeuchi [2011]. What appears to distinguish the designs of such studies is that they study time preferences for very short

[22]For an example, see Muthoo [1990]; Avery and Zemsky [1994] provide a synthesis.

horizons, down to days, which coincidentally is the more relevant case for bargaining as well. Nonetheless, at present, it seems premature to have much confidence in the reliability of these findings. However, in section 1.8, I sketch two specific bargaining environments where such preferences may arise "naturally" in reduced form.

I close this section with a remark on bargaining power. With exponential discounting, the notion of a player's bargaining power that the literature has adopted is her attitude to one period of delay; assuming linear utility, if a player's (exponential) discount factor is δ, then against any given opponent, her payoff is monotonically increasing in δ. With dynamic inconsistency of time preferences, different periods of delay are evaluated differently, i.e. marginal patience is not constant, however, so the question arises of what defines bargaining power in the context of the more general preferences analysed here. The characterisation suggests a generalisation, which is weaker, however, because the possibility of delay has to be accounted for: it is the minimal marginal patience within the equilibrium horizon. As can be seen in theorem 1.3, a greater $\delta_i(t_i^*)$ means that, no matter whether i is proposer or respondent, the minimal equilibrium payoff increases.

The next section serves to illustrate the main results by means of an example. At the same time, I wrote this section with the aim of providing the essence of the entire paper in the simplest form possible. And in its last subsection, it adds two observations to the general section within the context of the example: the worst equilibrium payoff of the initial proposer may decrease when that player's discount function is "increased". And, when considering all equilibria, it is possible that a player receives a greater payoff as initial respondent than as initial proposer against the same opponent.

1.7. An Example

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 that in any period, the players' preferences over delayed agreements $(x,t) \in X \times T$, $T = \mathbb{N}_0$, have the following representations, 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 are common knowledge—in particular, both players fully understand Eve's time preferences—and that players cannot commit to future actions, i.e. use Strotz-Pollak equilibrium ("equilibrium" in what follows). 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 only 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 anticipates her own best subsequent proposer payoff, and "vice versa". After two rounds of backwards induction from the maximal proposer payoff of a player, the resulting payoff must then again equal this maximal payoff, and similarly for the minimal payoff, because the subgames are formally identical. 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 subgame 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 period of time. Hence, the relationship between a player i's extreme proposer payoffs and analogously extreme respondent payoffs is more complicated: the latter need not simply equal the former multiplied by $d_i(1)$.

Adding the necessary 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 players' impatience and the resulting "natural" advantage of a proposer this is the best immediate agreement—the challenging part is the relationship between a player i's worst threat as a respondent (the lowest continuation equilibrium payoff) and i's lowest 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 occur in the future but is impatient about postponing agreement from the immediate present; intuitively, Eve's preferences display a form of *present bias.*[23] Her dynamic inconsistency implies the following: whereas Eve is indifferent about 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 if 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 subgame 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 1.1 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 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, where

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

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 $d_i(1)\, v_i$, and this allows to exploit the stationarity of the game via the backwards-induction argument of Shaked and Sutton [1984] for establishing uniqueness.

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

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

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 any proposal (immediate agreement on $(1, 0)$). Upon rejection, bargaining progresses to
- Round 2: Eve demands the entire dollar and Od accepts a proposal x if and only if $x_1 \geq \alpha = \frac{10}{11}$ (rejection). Upon rejection, bargaining continues through
- Round 3:
  - if, in the previous round, Eve did *not* demand the entire dollar, then play continues as from round 1 (immediate agreement on $(1, 0)$),
  - otherwise, play continues with the stationary equilibrium (immediate agreement on x^*).

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 this 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 in round 1.

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 payoff, two steps of backwards induction which involve only the single-period discount factors result in an increase 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.

Also 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 will not be as generous but instead prefer 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 1.2. Intuitively, whenever there is delayed agreement, say on split 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:

$$w_i \equiv \inf_{(x,t)\in Z_i^*} \{d_i(t+1)\, x_i\} = \inf_{t\le t_i} \left\{ d_i(t+1) \cdot \frac{v_i}{d_i(t)} \right\} = \min_{t\le t_i} \{P_i(t)\} \cdot v_i. \tag{7}$$

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.

The reasoning just given introduces the 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 surplus, i.e. $(v_i/d_i(t)) + (v_j/d_j(t-1)) \le 1$. Building on these results, a system of equations is obtained which theorem 1.2 studies to establish uniqueness of a solution to this system as both necessary and sufficient for uniqueness of equilibrium. Theorem 1.3 further generalises this result, producing a characterisation of equilibrium outcomes and payoffs when the system of equations may have multiple solutions.

Illustration of Theorem 1.3. 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 particular 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 the RubE would be the unique equilibrium. If this maximal delay were positive, however, then Eve's worst threat would equal $w_2 = 0$, whence there is an equilibrium in which proposer Od achieves the maximal feasible payoff of 1 and two steps of backwards induction yield Eve's minimal proposer payoff $v_2 = 1 - \alpha$. Indeed, the "residual" proposer advantage ensures she cannot obtain anything less than $1 - \alpha$, so the maximal delay that 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 subgame 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 subgame where 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 concurring multiplicity and delay which arise from Eve's time-inconsistency.

Observe that if Od had preferences identical to those of Eve, then the RubE is dispensable for the construction of equilibrium multiplicity and delay; let $\alpha \geq \frac{2}{3}$ and modify the strategies described above as follows for round 3, adding also a descriptions of rounds 4 and 5 for the novel cases:

- Round 3:
  - if, in the previous round, Eve did *not* demand the entire dollar, then play continues as from round 1, leading to immediate agreement on $(1,0)$,
  - otherwise, Od proposes $(1-\alpha,\alpha)$ and Eve accepts a proposal x if and only if $x_2 \geq \alpha$. Upon rejection here, bargaining goes on in round 4 where play continues as from round 1, but with roles reversed.

By symmetry, this describes an equilibrium, where the condition that $\alpha \geq \frac{2}{3}$ ensures the delay in round 2, because Eve then does not prefer giving away at least α in order to obtain an immediate share of $1-\alpha$ over forcing a rejection with an extreme demand that has present value of $\alpha\frac{1}{2}$.

Finally, consider the following comparative statics: if Eve is made "more patient" in the sense of a greater β, her worst payoff may decrease 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 her attitude to the first period of delay, which is β. 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 constructed here 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 would prefer to be respondent initially rather than proposer.

1.8. 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.

1.8.1. 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 [2011b]. 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 the terminology of this paper, this is decreasing marginal patience). 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. Building 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 splits $x \in X$ with delay $t \in T$ have the following representation, which—for the sake of simplicity—involves breakdown risk as the sole source of discounting:

$$U_i(x,t) = g_i\left(r^t\right) x_i. \tag{8}$$

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, $g_i\left(r^t\right) \equiv d_i(t)$, 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

[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.

utility (Quiggin [1982]) or cumulative prospect theory (Tversky and Kahneman [1992]) into the basic bargaining model and study their implications.[25]

A qualitative feature of probability weighting that appears widely accepted as empirically valid in the context of cumulative prospect theory is the 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 $\gamma \in (0,1]$:

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

If a player i' preferences have a representation as in equation 8, then they satisfy present bias (set $d_i(t) = g_i(r^t)$), as can easily verified, whence theorem 1.1 implies that the RubE, where $d_i(1) = g_i(r)$, is the unique equilibrium. Since increasing γ 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 1.2 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.

1.8.2. Imperfect Altruism in Intergenerational Bargaining. Suppose that several communities have access to a single productive resource. They decide over how to share it by means of bargaining. As long as these usage rights have not been settled, some surplus is forgone due to inefficient usage. 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.

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

[25]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".

equal to x_i,

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

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

The two parameters δ_i and γ_i are both assumed to 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 it still looks forward to a second period of lifetime.

Assume that in each round t, a new member of each community is nominated to join 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, or
- i always sends the old member to the bargaining table, 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 engaged 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 $\gamma_i < 1$, which 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 3-6, I propose a simple equilibrium construction similar to that in the first example of section 1.7. 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 for 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 an 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.

1.9. 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 separable linear time preferences, theorem 1.3, from which all other results could be derived. Reflecting both the genesis of this work and my anticipation of how the various implications would be received, I presented it as two main results. The first main result, theorem 1.1, establishes that if both players are most impatient about the first period of delay (in relative time), 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 and involve strategic bargaining 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 beyond the first period of delay from the immediate present are irrelevant to equilibrium; since the empirical evidence is 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 the first round's delay than a period's delay from the near future, then, in general, there may be multiplicity and inefficient delay, both based on such a player's preference reversals. The nature of these equilibria is novel, since their construction does not rely on the presence of stationary equilibrium. In fact, when there exist delay equilibria, the sets of equilibrium outcomes and payoffs may be generated entirely without stationary equilibrium.

The most recent experimental evidence on time preferences has indeed documented such violations of present bias which are most conducive to the existence of such delay equilibria.[26] What makes this evidence particularly relevant for the bargaining context is that these findings have been for short-horizon trade-offs on the domain of money rewards; sharing a monetary surplus is the classic bargaining example. It is still early to confidently judge the general validity of this qualitative property, but should it receive confirmation, this paper will constitute a first theoretical investigation of such preferences, and the equilibrium delay obtained may deserve wider interest.

In any event, this paper provides a first step towards the study of psychologically richer preferences in the basic strategic model of bargaining; the short section 1.8 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 may be investigated. More generally, the adaptation of the approach of Shaked and Sutton [1984] and its extension by ideas similar to those developed by Abreu [1988] for repeated games, which allowed me to obtain a full characterisation of equilibrium outcomes, may be useful in further theoretical research on bargaining with non-standard time preferences, or even for the study of other 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 1.5, it is the fact that such strict present bias causes a player's weakest delay attitude to fully enter the immediate agreement equilibrium that drives its uniqueness; section 1.3 might suggest, however, that it is the particular simplicity of such equilibrium under the alternating-offers protocol that may prevent preference reversals from playing a role for equilibrium.

Moreover, the assumption of full sophistication seems extreme from an empirical point of view. One may therefore ask how predictions change when players are assumed naïve, at least partially.[27] This introduces the potential of learning (and teaching) through delay, which may take different forms for present-biased and non-present-biased players.[28]

---

[26] For a list of studies which find such "increasing impatience", see the survey of Attema [2012].
[27] O'Donoghue and Rabin [2001] develop such a concept in the context of the quasi-hyperbolic (β, δ)-model.
[28] Akin [2007] studies this aspect for the quasi-hyperbolic (β, δ)-model.

Finally, it has been argued that Strotz-Pollak equilibrium goes too far in its assumption of fully non-cooperative selves and should be refined. In particular, in view of proposition 1.4, it is therefore an interesting question whether for plausible such refinements the multiplicity result disappears. In other words, what kind or how much of intrapersonal coordination is necessary to restore uniqueness more generally?

References

Dilip Abreu. On the theory of infinitely repeated games with discounting. *Econometrica*, 56 (2):383–396, 1988.

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

James Andreoni and Charles Sprenger. Estimating time preferences from convex budgets. *The American Economic Review*, 102(7):3333–3356, 2012.

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

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

Arthur E. Attema, Han Bleichrodt, Kirsten I. M. Rohde, and Peter P. Wakker. Time-tradeoff sequences for analyzing discounting and time inconsistency. *Management Science*, 56(11): 2015–2030, 2010.

Ned Augenblick, Muriel Niederle, and Charles Sprenger. Working over time: Dynamic inconsistency in real effort tasks. January 2013.

Christopher Avery and Peter B. Zemsky. Money burning and multiple equilibria in bargaining. *Games and Economic Behavior*, 7(2):154–168, 1994.

Ken Binmore, Ariel Rubinstein, and Asher Wolinsky. The nash bargaining solution in economic modelling. *The RAND Journal of Economics*, 17(2):176–188, 1986.

Albert Burgos, Simon Grant, and Atsushi Kajii. Bargaining and boldness. *Games and Economic Behavior*, 38(1):28–51, 2002.

Lutz-Alexander Busch and Quan Wen. Perfect equilibria in a negotiation model. *Econometrica*, 63(3):545–565, 1995.

Pierre-André Chiappori, Bernard Fortin, and Guy Lacroix. Marriage market, divorce legislation, and household labor supply. *Journal of Political Economy*, 110(1):37–72, 2002.

José Luis Ferreira, Ithzak Gilboa, and Michael Maschler. Credible equilibria in games with utilities changing during the play. *Games and Economic Behavior*, 10(2):284–317, 1995.

Peter C. Fishburn and Ariel Rubinstein. Time preference. *International Economic Review*, 23 (3):677–694, 1982.

Shane Frederick, George Loewenstein, and Ted O'Donoghue. Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2):351–401, 2002.

Drew Fudenberg and Jean Tirole. *Game Theory*. The MIT Press, 1991.

Robert Gibbons. *A Primer in Game Theory*. FT Prentice Hall, 1992.

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

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

Robert E. Hall and Paul R. Milgrom. The limited influence of unemployment on the wage bargain. *American Economic Review*, 98(4):1653–1674, 2008.

Peter J. Hammond. Changing tastes and coherent dynamic choice. *The Review of Economic Studies*, 43(1):159–173, 1976.

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

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

David M. Kreps. *A Course in Microeconomic Theory*. Princeton University Press, 1990.

David I. Laibson. *Essays in Hyperbolic Discounting*. PhD thesis, MIT, 1994.

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

Antonio Merlo and Charles Wilson. A stochastic model of sequential bargaining with complete information. *Econometrica*, 63(2):371–399, 1995.

Abhinay Muthoo. Bargaining without commitment. *Games and Economic Behavior*, 2(3): 291–297, 1990.

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

Ted O'Donoghue and Matthew Rabin. Choice and procrastination. *The Quarterly Journal of Economics*, 116(1):121–160, 2001.

Efe A. Ok and Yusufcan Masatlioglu. A theory of (relative) discounting. *Journal of Economic Theory*, 137(1):214–245, 2007.

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

Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.

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

Motty Perry and Philip J. Reny. A non-cooperative bargaining model with strategically timed offers. *Journal of Economic Theory*, 59(1):50–77, 1993.

Asaf Plan. Weakly forward-looking plans. May 2010.

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

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

Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.

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

Kota Saito. A relationship between risk and time preferences. February 2011b.

Avner Shaked and John Sutton. Involuntary unemployment as a perfect equilibrium in a bargaining model. *Econometrica*, 52(6):1351–1364, 1984.

Ingolf Ståhl. *Bargaining Theory*. EFI The Economics Research Institute, Stockholm, 1972.

Robert H. Strotz. Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, 23(3):165–180, 1955-1956.

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

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

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

Appendix

LEMMA 1.3. *For* $\{i, j\} = I$ *and any* $(x, t) \in Z_j^*$, *if* $y_j = d_j(1+t)\, x_j$ *then* $(y, 0) \in Z_i^*$.

PROOF. Consider G_i. Since next round's subgame G_i has equilibrium outcome $(x, t) \in Z_j^*$, this may be imposed as the continuation outcome following any rejection in the first round of G_i. In this case, it is optimal for j to accept i's initial proposal if and only if it leaves her a share of at least $d_j(1+t)\, x_j$. Given this, proposal y is optimal for i, since $1 - d_j(1+t)\, x_j > d_i(1+t)\, x_i$. □

This is a fundamental backwards-induction type lemma: for any equilibrium agreement of G_j, one can construct an immediate-agreement equilibrium of G_i which is based on i's offering j exactly the present value of the former agreement including the additional rejection period's delay which j accepts, being indifferent. In fact, this is the essence of the strategic advantage of being the initial proposer.

LEMMA 1.4. *For both* $i \in I$, $V_i = V_i^0$.

PROOF. From their definitions, clearly, $V_i \geq V_i^0$. Suppose $V_i > V_i^0$. This implies that there exists a $(x, t) \in Z_i^*$ with $t > 0$ such that $d_i(t)\, x_i > V_i^0$. Therefore $x_i > V_i^0$, and t must be odd so $(x, 0) \in Z_j^*$, $j = 3 - i$. But then, by lemma 1.3, $y \in Z_i^*$ for $y_i = 1 - d_j(1)\, x_j$, and therefore $V_i^0 \geq 1 - d_j(1)\, x_j > x_i$, a contradiction. □

No equilibrium with delay yields the proposer a greater payoff than her favourite immediate-agreement equilibrium. Stationarity rules this out for any even number of periods of delay, because this would be an immediate agreement as well. And backwards induction would lead to a contradiction if it were the case with an odd number of delay periods.

LEMMA 1.5. *For both* $i \in I$, $W_i = W_i^0$.

PROOF. First note that $W_i \geq W_i^0 = d_i(1)\, V_i^0$. Suppose now that $W_i > W_i^0$. Then there exists an outcome $(x, t) \in Z_i^*$ such that $d_i(1+t)\, x_i > d_i(1)\, V_i^0$, implying that $x_i > V_i^0$ and thus that t is odd, or $(x, 0) \in Z_j^*$. By lemma 1.3, $(y, 0) \in Z_i^*$ for $y_i = 1 - d_j(1)\, x_j > 1 - x_j = x_i > V_i^0$, a contradiction. □

Given the fact that a proposer's greatest equilibrium payoff is achieved in an immediate-agreement equilibrium, and that a proposer is able to extract some surplus from the respondent arising from the latter's impatience, any equilibrium with delay must in fact yield a payoff that is strictly lower. Due to impatience, a respondent's most preferred continuation equilibrium therefore has the best immediate equilibrium agreement. Combining lemmas 1.4 and 1.5, a relationship similar to that in the time-consistent case of exponential discounting between a player's best threat as a respondent and her greatest equilibrium payoff as a proposer is obtained: $W_i = d_i(1)V_i$.

LEMMA 1.6. *For* $\{i,j\} = I$, $V_i = 1 - w_j$.

PROOF. Consider G_i. In any equilibrium of this subgame, in any round where i proposes, in particular in the initial round, respondent j rejects any proposed split x with $x_j < w_j$ because she prefers any continuation equilibrium outcome to $(x,0)$. Therefore, in no equilibrium which has agreement in a round in which i proposes can she obtain a share greater than $1 - w_j$, and in particular $V_i^0 \leq 1 - w_j$, which, by lemma 1.4 is the same as $V_i \leq 1 - w_j$.

Now suppose $V_i < 1 - w_j \Leftrightarrow w_j < 1 - V_i$. Then there exists an outcome $(x,t) \in Z_j^*$ with $d_j(1+t)x_j < 1 - V_i$. By lemma 1.3, $(y,0) \in Z_i^*$ for $y_i = 1 - d_j(1+t)x_j > V_i$, a contradiction. □

The respondent's least preferred continuation equilibrium agreement yields immediately the proposer's best immediate equilibrium agreement, and by the previous lemma this is the proposer's best equilibrium agreement overall.

LEMMA 1.7. *For* $\{i,j\} = I$, $v_i = v_i^0 = 1 - W_j$.

PROOF. Consider G_i. In any equilibrium of this subgame, in any round where i proposes, in particular in the initial round, respondent j accepts any proposed split x with $x_j > W_j$ because she prefers $(x,0)$ to any continuation equilibrium outcome. Therefore, in no equilibrium of this subgame will i achieve a payoff of less than $1 - W_j$, so $v_i \geq 1 - W_j$.

Now suppose $v_i^0 > 1 - W_j \Leftrightarrow W_j > 1 - v_i^0$. Then there exists a continuation equilibrium outcome $(x,t) \in Z_j^*$ with $d_j(1+t)x_j > 1 - v_i^0$. By lemma 1.3, $(y,0) \in Z_i^*$ for $y_i = 1 - d_j(1+t)x_j < v_i^0$, a contradiction.

Hence, we have shown that $v_i \geq 1 - W_j \geq v_i^0$. Since $v_i^0 \geq v_i$ holds by definition, this establishes the claim. □

The respondent's most preferred continuation equilibrium determines the proposer's worst immediate equilibrium agreement. However, since the proposer may always offer a share that matches the respondent's favourite continuation equilibrium outcome's present value, her lowest equilibrium payoff cannot fall below the worst immediate equilibrium agreement.

CHAPTER 2

A Characterisation of Equilibrium Outcomes in Alternating-Offers Bargaining for Separable Time Preferences

2.1. Introduction

The Rubinstein-Ståhl model of bargaining (Ståhl [1972] and Rubinstein [1982]) forms the core of modern non-cooperative bargaining theory. This model posits an explicit dynamic bargaining protocol of alternating offers by two parties about how to split a fixed economic surplus and derives predictions about bargaining from the subgame-perfect Nash equilibrium of the extensive-form game that results after specifying the parties' preferences over possible agreements. The central aspect in determining these outcomes is the parties' relative attitudes to delay, their time preferences.

In view of recent empirical evidence, which casts doubt on the classic assumption of exponential discounting,[1] and the successful introduction of alternative decision models,[2] this paper provides a characterisation of equilibrium outcomes for general separable time preferences, as axiomatised by Fishburn and Rubinstein [1982] and Ok and Masatlioglu [2007]. Since bargaining is pervasive in economic modelling, it thus fills a gap in the literature that has become important with the surge of interest in applied work with non-exponential discounting.

In dealing with players' dynamic inconsistency, I assume common knowledge and thus full "sophistication" about preferences, and I employ the standard solution concept of Strotz-Pollak equilibrium (StPoE, and in what follows simply "equilibrium"), which is here equivalent to subgame-perfect Nash equilibrium for time-consistent players. There always exists a stationary such equilibrium (strategies are history- and time-independent) with immediate agreement in every round, which I call "Rubinstein equilibrium" (RubE). The two underlying agreements (one per player/round) form a pair which is self-enforcing when viewed as "threats", and equilibrium is unique if and only if this pair is the only fixed-point of an equation system which describes such pairs of self-enforcing threats. For the case of multiple fixed-points, there is one involving the "most severe threats", and it characterises the set of equilibrium outcomes and payoffs.

A main contribution of this paper is of a technical nature. Despite the stationarity of the game, the otherwise very effective backwards-induction type reasoning proposed by Shaked and Sutton [1984] cannot be employed as such because of the players' time-inconsistency. I develop an approach of proof to the characterisation results that is similar to that of Abreu [1988] for repeated games, where the strategy spaces considered are greatly simplified. For each

[1] See the surveys of Frederick et al. [2002] and Attema [2012]. A very recent contribution, which qualifies experimental results where exponential discounting could not be rejected, such as those of Andreoni and Sprenger [2012], is Augenblick et al. [2013].

[2] Again, Frederick et al. [2002] survey several such models, the most prominent being the (β, δ)-model with early applications in Laibson [1997] and O'Donoghue and Rabin [1999]. Recent such alternative models, which also nest or limit to exponential discounting are for instance Takeuchi [2011] and Pan et al. [2013].

pair of maximal threat payoffs as respondent, one for each player, this permits a description of all sustainable outcomes. The aforementioned fixed-point condition then requires that the maximal threats are self-enforcing. This approach seems promising for a characterisation of equilibrium outcomes in many other settings that satisfy a form of stationarity.

Related Literature. The work of Rubinstein [1982] was highly seminal and marks the beginning of a vast literature on non-cooperative bargaining theory. While the subsequent literature tended to focus on the case of linear instantaneous utility, which I analysed in chapter 1, the original model allowed for representations with non-linear (but not "too convex") utility functions and provided an example of multiplicity and delay, even for dynamically consistent time preferences; see Rubinstein [1982, pp. 107-8]. Osborne and Rubinstein [1990, p. 48] point out that uniqueness of a certain fixed point (see definition 2.2) characterises equilibrium uniqueness for any separable time preferences which satisfy time-consistent exponential discounting. My innovation over existing proofs allows to generalise this fixed-point condition in order to accomodate also any other separable time preferences.

According to my knowledge, there are only two papers which analyse bargaining for dynamically inconsistent time preferences: both Ok and Masatlioglu [2007] and Noor [2011] do restrict their respective analyses to stationary strategies, however. My results show that for some preferences this does not do full justice to their behavioural implications for bargaining and also characterise those preferences for which this simplification indeed yields the same prediction(s).

2.2. Game

Protocol, Histories and Strategies. The bargaining protocol is identical to that considered in Rubinstein [1982]. Two players $i \in \{1,2\} \equiv I$ bargain over which split $x \in X$ of a fixed surplus to implement, where $X \equiv \{(x_1, x_2)\,|0 \leq x_1 = 1 - x_2 \leq 1\}$. In each of potentially infinitely many rounds $t \in \mathbb{N}$, a player $\rho(t)$ proposes a split $x \in X$ to the other who then responds by either accepting, $a = 1$, or rejecting, $a = 0$. Upon acceptance, bargaining ends, with x being implemented; upon rejection, one period of time elapses until the next round, $t+1$, where roles are reversed, so player $\rho(t+1) = 3 - \rho(t)$ is the proposer and $\rho(t)$ responds. The initial proposer is player $\rho(1) = 1$.

As long as bargaining continues, action sets in each round are X for the proposer and $\{0,1\}$ for the respondent. These generate possible histories of play in the obvious way, with h_0 denoting the history as of round 1. Letting $H \equiv (X \times \{0\})$ with convention $H^0 = \{h_0\}$, a non-terminal history at the beginning of bargaining round $t \in \mathbb{N}$ is some $h \in H^{t-1}$. Call $(X \times \mathbb{N}) \cup \{D\}$ the set of "dated outcomes", which are equivalence classes of terminal histories: $(x,t) \in X \times \mathbb{N}$ denotes any history where there is agreement on split x in round t, and outcome D, which I call disagreement, captures any infinite history.

A player i's strategy σ_i maps the non-terminal histories at which i makes a choice to an available action: for instance, letting $h \in H^{t-1}$ be the history at the beginning of round t, if $i = \rho(t)$, then $\sigma_i(h) \in X$; otherwise, σ_i is defined for each such proposal from X and $\sigma_i(h,x) \in \{0,1\}$ (alternatively, $\sigma_i(h,\cdot) : X \to \{0,1\}$). Say σ_i is stationary, i.e. history- as well as time-independent, if, whenever i is called upon to propose, she makes the same proposal, and, whenever i is called upon to respond, she follows the same acceptance rule.[3]

Preferences. 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 it is reached, or how disagreement obtains; in other words, each player only cares about "dated outcomes" as defined above. Switching to the language of the decision theory of time preferences, think of the share of the surplus in an agreement as a "reward" and of its round as a calendar date. An individual having dynamically inconsistent time preferences means that her preference between two *dated* future rewards may change depending on the *date* at which she gets to make the choice.

As a consequence, I take as the primitive of a player i's preferences a sequence of *dated* preference orderings $\{\succeq_{(i,t)}\}_{t\in\mathbb{N}}$, where each element $\succeq_{(i,t)}$ is defined on the set of feasible outcomes at t, which is $Z_t \equiv (X \times \mathbb{N}_t) \cup \{D\}$ for $\mathbb{N}_t \equiv \{t' \in \mathbb{N} | t' \geq t\}$.[4] Now let $T \equiv \mathbb{N}_0$ denote the possible *delay* of an agreement. I assume that, *at any date t*, $\succeq_{(i,t)}$ has a separable utility representation

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

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 $u_i(1) = 1$.[5] The image of $[0,1]$ under u_i will be denoted $\mathscr{U}_i \equiv [u_i(0), 1]$.

Observe that this representation involves time only in relative terms, as delay from t, but not t itself. In terms of feasible *delayed* agreements, the domain at any round t is identically equal to $(X \times T) \cup \{D\} \equiv Z$, which I will refer to as the set of *outcomes*. This simplification is used in the following assumption that summarises the description of preferences.

Assumption 2. *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, at any round $t \in \mathbb{N}$, preferences $\succeq_{(i,t)}$ over feasible outcomes Z are represented by*

$$U_i(z) = \begin{cases} d_i(s)\, u_i(x_i) & z = (x, s) \in X \times T \\ 0 & z = D \end{cases}. \tag{9}$$

Equilibrium. In this paper, I assume there is perfect information and, in particular, that the players' preferences are common knowledge.[6] The equilibrium concept I use is that of Strotz-Pollak equilibrium which incorporates the assumption that players cannot commit to

[3] Formally, for any two rounds t and t' such that $\rho(t) = \rho(t') = \rho$, any two histories $h \in (X \times \{0\})^{t-1}$ and $h' \in (X \times \{0\})^{t'-1}$, if $i = \rho$, then $\sigma_i(h) = \sigma_i(h')$, and if $i \neq \rho$, then $\sigma_i(h, \cdot) = \sigma_i(h', \cdot)$.

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

[5] 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].

[6] In the context of a single decision maker who perfectly anticipates his future preferences, this has been termed "sophistication"; see e.g. Hammond [1976].

future behaviour; given common knowledge of strategies, every action of a player has to be optimal taking as given not only the opponent's strategy but also that player's own continuation strategy.[7] In case such an equilibrium satisfies the stronger property that at no stage, any player could gain from perfect commitment (there is no "intrapersonal conflict"), then I call it Perfect-Commitment equilibrium. To simplify the formal statement, define, for any history $h \in H^{t-1}$ at the beginning of a round $t \in \mathbb{N}$, $z_h(\sigma) \in Z$ as the outcome (in relative time) that obtains under strategy profile σ, and similarly, $z_{(h,x)}(\sigma)$ for any proposal $x \in X$.

Definition 2.1. A strategy profile σ^* is a *Strotz-Pollak equilibrium (StPoE, "equilibrium")* if, for any round $t \in \mathbb{N}$, history $h \in H^{t-1}$, proposal $x \in X$ and response $a \in \{0,1\}$, the following holds:

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

Such an equilibrium σ^* is a *Perfect-Commitment equilibrium (PCE)* if, moreover, for any σ with $\sigma_j = \sigma_j^*$, $j = 3 - i$,

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

Analytically, the defining property of StPoE is robustness against one-shot deviations. By the one-shot deviation principle (e.g. Fudenberg and Tirole [1991, Theorem 4.1], it is therefore equivalent to subgame-perfect Nash equilibrium (SPNE) when all players' preferences are time-consistent, because they satisfy continuity at infinity.[8] This is important for comparison with existing results about SPNE for time-consistent exponential discounting, which I may therefore interpret as results about StPoE.

2.3. Results

This section presents the formal results. First, my focus will be on stationary equilibrium, which is equilibrium in stationary strategies. The influential uniqueness result of Rubinstein [1982] for the case of time-consistent exponential discounting yielded such a stationary equilibrium, and existing work studying dynamically inconsistent time preferences in Rubinstein-Ståhl bargaining has restricted the strategy space to such simple strategies. The second section then extends the scope of the analysis to the full strategy space.

2.3.1. Stationary Equilibrium. The presentation of my results requires additional notation. 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 immediate share which is worth at least as much:

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

[7] Strotz-Pollak equilibrium was developed in the context of analyses of single-person decision problems with time-inconsistent preferences, as pioneered by Strotz [1955-1956] to which Pollak [1968] provided an illuminating response. For its use in strategic contexts, see for instance Rotemberg [1983] and, recently, Chade et al. [2008].

[8] The same holds true about PCE, of course, which checks for "full-strategy deviations".

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 on the set of shares $\left[0, u_i^{-1}(u_i(0)/d_i(t))\right]$ (possibly equal to the singleton $\{0\}$) and increasing on its complement in $[0,1]$. Moreover, $t > 0$ implies $f_i(x_i, t) \leq x_i$ where the inequality is strict for any $x_i > 0$ (impatience). Note that because $u_i(0)$ may be positive, the preferences considered here include the case where a player is impatient also about receiving a share of zero.

Definition 2.2. 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 x^* and rejecting it for agreement on y^* in the subsequent round, and, similarly, for player 1 for y^* and x^*. Note that

$$\begin{aligned} y_1^* &= f_1(x_1^*, 1) \\ &= u_1^{-1}(\max\{u_1(0), d_1(1)\, u_1(x_1^*)\}) \\ &\leq x_1^*. \end{aligned} \tag{10}$$

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

A Rubinstein pair only depends on the players' attitudes to a single round's delay. 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 $(f_i)_{i \in I}$:

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

Lemma 2.1. *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 stationary strategies from any Rubinstein pair.[9]

Definition 2.3. For any Rubinstein pair (x^*, y^*), a *Rubinstein profile* is a strategy profile σ^R that satisfies the following: for any round $t \in \rho^{-1}(1)$, division $x \in X$, histories $h \in H^{t-1}$ and $h' \in H^t$,

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

In fact, the following is true.

[9] $\mathbb{I}$ denotes the indicator function that evaluates to one if its argument is true and to zero otherwise.

PROPOSITION 2.1. *Every stationary StPoE is a Rubinstein profile.*

PROOF. Let σ^* be a stationary StPoE in which player 1 always proposes x^*, player 2 always proposes y^*, and each player $i \in I$ responds to every offer according to some 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 $a_2(x^*) = a_1(y^*) = 0$, implying disagreement. By StPoE, 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 StPoE, 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 $\epsilon > 0$ such that the proposing player 1 self $(1,t)$ can increase her payoff above that of $d_1(1)\,u_1(y_1^*)$ 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

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

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 StPoE, 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 Rubinstein profile. $\square$

Observe the implication that any stationary StPoE σ has immediate agreement in every round: since it is based on some Rubinstein pair (x^*,y^*), the outcome in any round $t \in \mathbb{N}$, as element of Z, is then $(x^*,0)$ if $\rho(t) = 1$, and it is $(y^*,0)$ if $\rho(t) = 2$.

The next result shows that in a Rubinstein profile, despite all possible forms that time-inconsistency may take under assumption 2, no player would want to change either their own action or their continuation behaviour at any point in the game, i.e. any Rubinstein profile is a PCE. By the mere impatience of the players, the particular structure of such a strategy profile, with agreement in every round and indifference of the respective respondent, effectively reduces the game to a sequence of two-period problems of either agreement now or agreement the next round. This, however, means that only the attitudes to a single period of delay matter and time-inconsistency cannot play a strategic role.

PROPOSITION 2.2. *Every Rubinstein profile is a PCE.*

PROOF. Take any Rubinstein profile 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 σ_1^R, 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 $x_2 \geq x_2^*$ which player 2 accepts, or in disagreement. The latter two outcomes are obviously no better than $(x^*,0)$; because by inequality 10, neither is the first, since

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

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 some y with $y_2 \leq y_2^*$, or there is agreement in a later round where 1 proposes x^*, or there is 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 above not to exceed $u_2(x^*)$.

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

Existence of Rubinstein profiles follows from lemma 2.1, whence this proposition establishes existence of StPoE; moreover, in combination with 2.1, it yields a characterisation of stationary StPoE as Rubinstein profiles.

Corollary 2.1. *StPoE exists.*

Proof. A PCE exists because of lemma 2.1, and every PCE is a StPoE. □

Corollary 2.2. *A profile of strategies is a stationary StPoE if and only if it is a Rubinstein profile.*

Proof. Since a Rubinstein profile is defined as a pair of stationary strategies based on a Rubinstein pair, proposition 2.2 implies that every Rubinstein profile is a stationary StPoE (sufficiency). Proposition 2.1 provides the converse (necessity). □

Because of this result, I will use the term "Rubinstein equilibrium" (RubE) for stationary StPoE in what follows.

Importantly, the set of stationary-equilibrium outcomes equals the set of Rubinstein pairs, so uniqueness of stationary equilibrium coincides with uniqueness of a Rubinstein pair. Without emphasising this in their uniqueness proof, Osborne and Rubinstein [1990, Chapter 3] point out that the same is true for the time-consistent exponential discounting preferences they assume. Since Rubinstein pairs are defined only via the players' attitudes to one period of delay, even if one knew the players' instantaneous-utility functions, stationary equilibrium could not reveal any dynamic inconsistency of discounting.

In anticipation of a more general analysis, uniqueness of a Rubinstein pair is, of course, necessary for a unique equilibrium. Regarding the players' preferences, this fixed-point uniqueness constitutes a combined restriction on the curvatures of the two players' utility functions, given

$(d_i(1))_{i\in I}$. A sufficient condition in terms of individual preferences is presented in Osborne and Rubinstein [1990, pp. 35-36] as "increasing loss to delay": for each player $i \in I$ and every share $x_i \in [0,1]$, the "loss to delay" $x_i - f_i(x_i, 1)$ is increasing in x_i. For a given share, the "loss to delay" is the additional compensation that makes a player willing to accept a one-period delay against the alternative of receiving this share immediately.[10]

DEFINITION 2.4. 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 standard assumption of a differentiable concave u_i implies increasing loss to delay (Osborne and Rubinstein [1990, p. 35]); hence, this is true in particular of a linear utility function u_i, where $f_i(x_i, 1) = d_i(1)\, x_i$, so the loss to one period's delay from the present is $x_i(1 - d_i(1))$, which is increasing in x_i due to impatience.

LEMMA 2.2. *If both 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]. □

2.3.2. General Analysis. The previous section has shown that there exist multiple stationary equilibria whenever there are multiple Rubinstein pairs. All of these equilibria exhibit immediate agreement (in every round) and are therefore efficient. However, such equilibrium multiplicity may entail equilibrium with delayed agreement when players use non-stationary strategies. This point was made by Rubinstein [1982, pp. 107-108] for time-consistent preferences with an example, a version of which I reproduce here: let players have preferences represented by $U_i(x,t) = x_i - ct$ for each player i and note that U_i corresponds to a positive monotonic transformation of an exponential-discounting representation, satisfying assumption 2.[11] The set of associated Rubinstein pairs is $\{(x,y) \in X \times X \,|x_1 - y_1 = c\}$. For $c < 1$, both of the 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 2.2, 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

[10] 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.

[11] Without uncertainty, positive monotonic transformations do not change preferences: assumption 2 then also covers preferences with "discrete costs of delay"; simply take the natural logarithm of the following representation

$$U_i(x,t) = \underbrace{\exp(-c_i(t))}_{=d_i(t)} \underbrace{\exp(\hat{u}_i(x_i))}_{=u_i(x_i)}, \tag{11}$$

where $c_i : T \to \mathbb{R}_+$ is a continuous function that increases from $c_i(0) = 0$ towards infinity as $t \to \infty$. The special case of time-consistent exponential discounting has $c_i(t) = ct$ for some $c > 0$.

– 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.

Many existing constructions of non-stationary equilibrium, in particular equilibrium with delayed agreement, in variations of the bargaining protocol of Rubinstein [1982] follow exactly this pattern (e.g. Van Damme et al. [1990]), and—in any case—necessitate the existence of a stationary equilibrium, as shown by the “canonical” treatment of Avery and Zemsky [1994]. Chapter 1 already proved the point that this need not be the case when time preferences are dynamically inconsistent. Here, I provide a characterisation of equilibrium outcomes for the more general time preferences of assumption 2, which, in particular, has to handle the above possibility of multiplicity of stationary equilibrium, and in this case can be used to check when multiplicity of such immediate-agreement equilibrium permits delay.

Preliminaries. The structure of the bargaining game satisfies a stationarity property: after any two histories to any two rounds which start with a proposal by the same player $i \in I$, the respective subgames are identical; denote this subgame by G_i (when referring to “subgame”, I will mean G_1 or G_2).[12] The respective sets of equilibrium outcomes, as subsets of Z, therefore coincide, and I will use Z_i^* to refer to this. An important role in the analysis will be assigned to the following bounds on equilibrium payoffs: V_i and v_i will denote the supremum and infimum, respectively, of U_i on Z_i^*. These are the tightest possible bounds on i's equilibrium payoffs as the initial proposer.

For any outcome $z \in Z$, let z^+ denote 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 $z^+ = (x, 1+t)$. In a similar manner, let $Z_i^{+,*} \equiv \{z^+ \in Z \,|\, z \in Z_i^*\}$, so W_i and w_i, as the supremum and infimum, respectively, of U_i on $Z_i^{+,*}$, are the tightest possible bounds on i's “rejection utility”, i.e. on the continuation equilibrium payoffs that i may obtain when rejecting a proposal as the respondent; I will refer to these also as best and worst threat points, respectively.

A superscript of zero on any of the above payoff bounds will mean it is derived from the restriction of Z_i^* or $Z_i^{+,*}$, respectively, to immediate-agreement equilibrium outcomes, so for instance $w_i^0 \equiv \sup\left\{U_i(x,1) \,\middle|\, (x,0) \in Z_j^*\right\}$.

While the idea of studying such payoff bounds is familiar from Shaked and Sutton [1984], another characteristic of Z_i^* plays a key role in my characterisation, namely the supremum delay of such an equilibrium agreement: $t_i \equiv \sup\{t \in T \,|\, \exists x \in X,\ (x,t) \in Z_i^*\}$.

I also introduce further notation to deal with the possibility of non-linear u_i and positive $u_i(0)$. Let f_i^U associate with any “rejection utility” U the minimal share a responding player i may accept, i.e. :

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

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.

[12] There are many other subgames, starting with the respondent's decision of whether to accept or reject a given proposal $x \in X$.

Finally, I make a few conventions in order not to keep notational track of histories and rounds up to the subgame I am analysing. Let G_i be a subgame starting in round t after history $\hat{h}$. I will abuse notation and identify $\rho(t)$ with $\rho(1)$, and $\hat{h}$ with h_0. For any history $h \in \cup_{t \in T} H^t$ in that subgame, $z_h(\sigma)$ will denote the outcome in Z that obtains in $G_{\rho(t+1)}$ when players follow the prescriptions of σ from h onwards, and $z_h^+(\sigma)$ will be the outcome that is $z_h(\sigma)$ delayed by one more round (e.g. if σ prescribes behaviour such that, after history h at the beginning of round $t+1$ in G_i, agreement is reached on split x after s more rounds, then $z_h(\sigma) = (x, s)$ and $z_h^+(\sigma) = (x, s+1)$).

The following proposition summarises intermediate results proven in the appendix, which do not hinge on properties of preferences beyond mere impatience. In this sense, under the general assumption 2, it captures the essence of the bargaining protocol which places the burden of choice over delay onto the respondent.

Proposition 2.3. *For* $\{i, j\} = I$, $V_i = 1 - f_j^U(w_j) = V_i^0$, $v_i = 1 - f_j^U(W_j) = v_i^0$, $W_i = d_i(1) V_i$, $w_i^0 = d_i(1) v_i$ *and* $t_i < \infty$.

Proof. See appendix. $\square$

Relative to time-consistent exponential discounting, the only potential difference is the possibility of $w_i < d_i(1) v_i = w_i^0$. Although the lowest equilibrium payoff that i may experience in G_i is achieved by the worst immediate equilibrium agreement ($v_i = v_i^0$), because i can always please the most demanding respondent immediately, the additional round's delay that i faces as respondent relative to Z_i^* (comparing preferences over Z_i^* and $Z_i^{+,*}$) may change her ranking of outcomes such that there is a worse subsequent equilibrium outcome which has delay in Z_i^*. This possibility arises from time-inconsistency and then drives a wedge in-between w_i and w_i^0 that feeds through all other payoff bounds. If, however, despite their time-inconsistency, preferences turned out such that this can be ruled out, so $w_i = w_i^0$, then there would be no difference to an analysis of the game for a corresponding time-consistent representation $d_i(1)^t u_i(x_i)$ instead.

Characterisation. The next lemma achieves a simplification of the space of equilibria to consider for a characterisation of $(Z_i^*)_{i \in I}$. Every such "simple equilibrium" implements an outcome (x, t) by always relying on the most extreme threat points (W_1, W_2) that the two players may entertain as respondents. Roughly speaking, it has the following structure, corresponding to properties 1-5 in the definition below: as long as both players have been complying with the strategies and the agreement round has not been reached, a proposer i claims the entire surplus and a respondent j accepts only those offers that yield her at least her threat point W_j (property 1). Because the maximal threat points need not correspond to an actual equilibrium payoff, property 3 is added. In the agreement round, the split to implement is proposed, and the respondent accepts any split that is at least as good as that (property 2); upon rejection of any such split, the respondent is punished with her least preferred continuation StPoE, and upon rejection of other splits, a proposer's least preferred continuation StPoE is played (proposition 2.3 shows that it is possible to do so with a continuation StPoE most preferred by the respondent); ideally, property 4 would specify respective payoffs $W_{\rho(s+1)}$ and $v_{\rho(s)}$, but the non-existence problem has to be dealt with which complicates it. Property 5 means this construction yields the desired outcome.

The following definition formalises this, using the following notation. For $i \in I$, $e_i \in X$ will denote $x \in X$ such that $x_i = 1$. Define $h^1_{i,E} \equiv h_0$ and, for any $t \in \mathbb{N} \backslash \{1\}$, $h^t_{i,E} = \left(h_0, (x^s, 0)_{s=1}^{t-1}\right)$ with $x^s \in \{e_1, e_2\}$, $x^1 = e_i$ and $x^{s+1} \neq x^s$.

Definition 2.5. For any $i \in I$ and G_i, and any $(x,t) \in X \times T$, a StPoE σ is a *simple implementation of* (x,t) *in* G_i if it satisfies 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 $0 < y_{\rho(s+1)} < f^U_{\rho(s+1)}\left(W_{\rho(s+1)}\right)$ then

$$U_{\rho(s+1)}\left(z^+_h(\sigma)\right) > u_{\rho(s+1)}\left(y_{\rho(s+1)}\right)$$

(4) if $h = \left(h^t_{i,E}, y, 0\right)$ then
  - if $y_{\rho(t+1)} \geq x_{\rho(t+1)}$ then

$$U_{\rho(t+1)}\left(z^+_h(\sigma)\right) \leq u_{\rho(t+1)}\left(x_{\rho(t+1)}\right)$$

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

$$U_{\rho(t+1)}\left(z^+_h(\sigma)\right) > u_{\rho(t+1)}\left(y_{\rho(t+1)}\right)$$

(5) $z_{h_0}(\sigma) = (x,t)$

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 StPoE to be unique. It uses lemma 2.7, which states that for any $\{i,j\} = I$, $v_i = u_i\left(1 - f^U_j(W_j)\right)$, and lemma 2.6 which states that $W_{\rho(s+1)} = d_{\rho(s+1)}(1)\, V^0_{\rho(s+1)}$, where, for any $i \in I$, V^0_i is defined as the supremum of U_i taken over all *immediate* agreements in Z^*_i; both lemmas are part of proposition 2.3 and proven in the appendix.

Lemma 2.3. *For any* $i \in I$, $(x,t) \in Z^*_i$ *if and only if there exists a simple implementation of* (x,t) *in* G_i.

Proof. Sufficiency holds by definition 2.5 (property 5 and equilibrium).

For necessity, take any $(x,t) \in Z^*_i$ and define a strategy profile so it satisfies properties 1 and 2. It is to be shown that there exist continuation equilibrium outcomes after deviations from the desired path $\left(h^t_{i,E}, x, 1\right)$ that ensure these two properties define optimal behaviour and do not conflict with property 5.

Begin with the first part of property 4. If respondent $\rho(t+1)$ preferred *every* continuation StPoE outcome to immediate agreement on x, then she would never accept it, a contradiction to (x,t) being an StPoE outcome in G_i. Hence, there exists $(x', t') \in Z^*_{\rho(t+1)}$ such that $U_{\rho(t+1)}(x', 1+t') \leq u_{\rho(t+1)}\left(x_{\rho(t+1)}\right)$.

Moreover, $x_{\rho(t+1)} \leq f^U_{\rho(t+1)}\left(W_{\rho(t+1)}\right)$ must hold: in any StPoE, 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 StPoE outcome, meaning that in no equilibrium proposer $\rho(t)$ could offer x, a contradiction. Hence, for any y with $y_{\rho(t+1)} < x_{\rho(t+1)}$, it is true that $y_{\rho(t+1)} < f^U_{\rho(t+1)}\left(W_{\rho(t+1)}\right)$, so 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 the second part of property 4.

In fact, the very same reasoning applies to ensure existence of a continuation equilibrium outcome so property 3 can be satisfied, so, thus far, we have shown that a strategy profile can be constructed which satisfies properties 1-4 with optimal respondent behaviour on path $\left(h^t_{i,E}, x, 1\right)$ where deviations are followed by some equilibrium play.

Now 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 $f^U_{\rho(s+1)}\left(W_{\rho(s+1)}\right) > 0$ which means if, (x, t) cannot be implemented "simply" (with extreme proposals and most demanding respondent behaviour), then it cannot be an equilibrium.[13]

To show that one can indeed find an StPoE 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 any history $h = h^s_{i,E}$, $s \leq t$, due to the choice of continuation equilibrium. The following arguments demonstrate this by contradiction.

First, consider such a case where $s < t$ and suppose $\rho(s)$ were to deviate to a split $y \neq e_{\rho(s)}$. If $y_{\rho(s+1)} \geq f^U_{\rho(s+1)}\left(W_{\rho(s+1)}\right)$ then this deviation would result in immediate agreement with a payoff to $\rho(s)$ of at most $u_{\rho(s)}\left(1 - f^U_{\rho(s+1)}\left(W_{\rho(s+1)}\right)\right)$. In regard of lemma 2.7, 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, consider deviation to some y with $y_{\rho(s+1)} < f^U_{\rho(s+1)}\left(W_{\rho(s+1)}\right)$. In order to be profitable for any choice of continuation equilibrium, y must satisy 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 consult lemma 2.6 of the appendix which says that $W_{\rho(s+1)} = d_{\rho(s+1)}(1)\, V^0_{\rho(s+1)}$. For any $\epsilon > 0$, there exists $(x', 0) \in Z^*_{\rho(s+1)}$ such that $U_{\rho(s+1)}(x', 1) > W_{\rho(s+1)} - \epsilon$ (by definition). Existence of such a deviation y requires $W_{\rho(s+1)} > u_{\rho(s+1)}(0) \geq 0$, for $\epsilon \leq \left(1 - d_{\rho(s+1)}(1)\right) W_{\rho(s+1)}$, so there exists $(x', 0) \in Z^*_{\rho(s+1)}$ such that:

$$\begin{aligned} x'_{\rho(s+1)} > f^U_{\rho(s+1)}\left(\tfrac{W_{\rho(s+1)} - \epsilon}{d_{\rho(s+1)}(1)}\right) \geq f^U_{\rho(s+1)}\left(W_{\rho(s+1)}\right) \\ \Rightarrow u_{\rho(s)}\left(x'_{\rho(s)}\right) < u_{\rho(s)}\left(1 - f^U_{\rho(s+1)}\left(W_{\rho(s+1)}\right)\right) \\ \Rightarrow U_{\rho(s)}(x, t-s) < v_{\rho(s)}, \end{aligned}$$

where the last implication follows from lemma 2.7 and itself implies that $(x, t-s) \notin Z^*_{\rho(s)}$, a contradiction.

Finally, consider history $h^t_{i,E}$ 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 $y_{\rho(t+1)} < f^U_{\rho(t+1)}\left(W_{\rho(t+1)}\right)$ must hold, so a similar argument to the one employed in the previous paragraph

[13] In fact, suppose 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 $x_l = f_l(d_l(1))$. Note also that for $t = 0$, property 2 yields property 5 already.

applies and ensures there exists a continuation equilibrium outcome that deters this deviation. $\square$

This lemma implies the following property of the temporal structure of equilibrium outcomes: comparing i's lowest equilibrium payoffs across delays in G_i (as of the initial round), they are constant. This yields a connection between v_i and w_i through the maximal delay t_i: while $t_i = 0$ implies $w_i = w_i^0 = d_i(1)\, v_i$ from proposition 2.3, for $t_i > 0$, $w_i = \delta_i(t_i)\, v_i$, where $\delta_i : T \to (0,1)$ with $\delta_i(t) \equiv \min\left\{ \frac{d_i(s+1)}{d_i(s)} \,|s \in T,\ s \le t \right\}$.

Due to the proposer's strategic advantage, w_i is irrelevant for the equilibrium outcomes in G_i, and, moreover, for delayed such agreements, also w_j is irrelevant, since there is an intermittent stage in which j proposes. This results in the following corollary.

Corollary 2.3. *For any $i \in I$ and any $t \in T$ with $t \le t_i$, $w_i = \delta_i(t_i)\, v_i$, and, moreover,*

$$\begin{aligned} (x,0) \in Z_i^* \quad &\Leftrightarrow \quad u_i^{-1}(v_i) \le x_i \le 1 - f_j^U(\delta_j(t_j)\, v_j) \\ (x,t) \in Z_i^*,\ t > 0 \quad &\Leftrightarrow \quad u_i^{-1}\left(\frac{v_i}{d_i(t)}\right) \le x_i \le 1 - u_j^{-1}\left(\frac{v_j}{d_j(t-1)}\right) \end{aligned}$$

Proof. The first step is to show that the minimal share for any equilibrium outcome with delay t equals

$$\min\{x_i \,|(x,t) \in Z_i^*\} = u_i^{-1}\left(\frac{v_i}{d_i(t)}\right). \tag{12}$$

The construction of a simple implementation implies that the minimum is actually reached and that the incentive problem only needs to be solved for the stages where i proposes, where i needs to be prevented from making acceptable offers before t is reached (note that there is then no further issue with opponent j's incentives as the proposer, because it is about the *minimal* share for i). This is solved for all rounds if and only if it is solved for the first round, where the relative delay to the given agreement is maximal. There, no profitable deviation proposal exists up to where i's share in the agreement means a present value of less than $1 - f_j(W_j) = v_i$, whence the minimal share x_i with delay t satisfies $d_i(t)\, u_i(x_i) = v_i$.

Recall then the definition of w_i and observe the following, using continuity of u_i:

$$\begin{aligned} w_i \quad &\equiv \quad \inf\{d_i(1+t)\, u_i(x_i) \,|(x,t) \in Z_i^*\} \\ &= \quad \min\{d_i(1+t) \cdot \inf\{u_i(x_i) \,|(x,t) \in Z_i^*\} \,|t \in T,\ t \le t_i\} \\ &= \quad \min\left\{d_i(1+t) \cdot \frac{v_i}{d_i(t)} \,|t \in T,\ t \le t_i\right\} \\ &= \quad \delta_i(t_i)\, v_i. \end{aligned} \tag{13}$$

This pins down V_i because $V_i = V_i^0 = u_i\left(1 - f_j^U(w_j)\right)$ from proposition 2.3, where the payoff is obtained in some equilibrium because v_j is. This gives the characterisation of immediate-agreement proposals in Z_i^* in terms of i's shares given v_i, v_j and t_j.

For any $(x,t) \in Z_i^*$ with $t > 0$, it must be that $(x, t-1) \in Z_j^*$, whence the characterisation follows from applying the same reasoning as at the outset of the proof to G_j. $\square$

While this result characterises $(Z_i^*)_{i \in I}$ in terms of $(v_i)_{i \in I}$ and $(t_i)_{i \in I}$, in combination with proposition 2.3, it also yields a system of equations that the payoff bounds must satisfy in terms of only two unknowns $(t_i)_{i\in I}$. However, these maximal delays are in turn pinned down by the payoff bounds $(v_i)_{i\in I}$: as long as the delay is such that one can find divisions that do not violate the minimum necessary for a simple implementation, as in equation 12, such a simple implementation exists. For instance, in G_1, a delay of two periods is an equilibrium if and only if there exists a division x such that $x_1 \geq u_1^{-1}\left(\frac{v_1}{d_1(2)}\right)$ and $x_2 \geq u_2^{-1}\left(\frac{v_2}{d_2(1)}\right)$.

To express this formally, first define a function $c_i(v_i, v_j, \cdot) : T \to \mathbb{R}_+$; for each player i, c_i measures the minimal total surplus that is required to be able to promise the proposers sufficiently large shares for a delay of t periods in terms of $(v_i)_{i\in I}$:

$$c_i(v_i, v_j, t) \equiv \begin{cases} \hat{f}_i^U(v_i) + \hat{f}_j^U(d_j(1)\, v_j) & t = 0 \\ \hat{f}_i^U\left(\frac{v_i}{d_i(t)}\right) + \hat{f}_j^U\left(\frac{v_j}{d_j(t-1)}\right) & t > 0 \end{cases},$$

where $\hat{f}_i^U$ continuously extends f_i^U onto the entire non-negative real line as

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

Then the following is true.

Corollary 2.4. *For any $i \in I$,*

$$t_i = \max\left\{t \in T \,|\, c_i(v_i, v_j, t) \leq 1\right\}.$$

Proof. Given equation 12, there exists a proposal $x \in X$ such that there is a simple implementation of (x, t) in G_i if and only if $c_i(v_i, v_j, t) \leq 1$. □

This closes the "system": given $(t_i)_{i\in I}$, all payoff bounds necessarily satisfy a system of equations which can be reduced to one in only $(v_i)_{i\in I}$, and the latter in turn determine $(t_i)_{i\in I}$ as above. Hence, these bounds $(v_i, t_i)_{i\in I}$ must be a fixed point of a system of four equations. This system contains that for Rubinstein pairs (expressed in payoff terms) as a special case, and uniqueness of the fixed point characterises uniqueness of equilibrium.[14]

Theorem 2.1. *There exists a unique StPoE if and only if there exists a unique solution $\left(\tilde{v}_i, \tilde{t}_i\right)_{i\in I} \in \times_{i\in I}(\mathscr{U}_i \times T)$ to the system of four equations which has, for each $i \in I$,*

$$\tilde{v}_i = u_i\left(1 - f_j^U\left(\delta_j(0)\, u_j\left(1 - f_i^U\left(\delta_i\left(\tilde{t}_i\right)\tilde{v}_i\right)\right)\right)\right) \tag{14}$$

$$\tilde{t}_i = \max\left\{t \in T \,|\, c_i(\tilde{v}_i, \tilde{v}_j, t) \leq 1\right\} \tag{15}$$

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

Proof. First note that for $\tilde{t}_1 = \tilde{t}_2 = 0$ the system of equations 14-15 (for each i)—in what follows simply "the system"—reduces to one that is indeed equivalent to that defining Rubinstein pairs in terms of utilities. A Rubinstein pair exists by lemma 2.1 and, moreover, any such pair's associated proposer utilities yield $\tilde{t}_1 = \tilde{t}_2 = 0$ in the two equations 15, whence the system's set of solutions contains all those utilities obtained from Rubinstein pairs.

[14]Note that if (x^*, y^*) is a Rubinstein pair as in definition 2.2, then $c_1(u_1(x_1^*), u_2(y_2^*), 0) = x_1^* + x_2^* = 1$ and $c_2(u_2(y_2^*), u_1(x_1^*), 0) = y_1^* + y_2^* = 1$.

First, consider sufficiency. If there is a unique solution, then it is indeed equal to $(v_i, t_i)_{i\in I}$. Because, by the above observation, it must correspond to a unique Rubinstein pair which has efficient payoffs, there must be immediate agreement in every round where these payoffs yield unique strategies, the associated RubE.

Moving toward necessity, note that if there is a unique StPoE, there cannot be two different solutions with $\tilde{t}_1 = \tilde{t}_2 = 0$, because this would mean there are two RubE.

To finish the necessity part of the theorem, I next show that whenever there is a solution to the system with $\tilde{t}_i > 0$ for some $i \in I$, there exists a StPoE that is not a RubE; because a RubE always exists, this would contradict uniqueness. Suppose then there exists a solution $\left(\tilde{v}_1, \tilde{t}_1, \tilde{v}_2, \tilde{t}_2\right)$ 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 $P_i\left(\hat{t}_i\right) = \delta_i\left(\tilde{t}_i\right)$. 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 $\left(\hat{x}, \hat{t}_i\right)$ and $\left(\hat{y}, \hat{t}_j\right)$, where $\hat{x}_i = \hat{f}_i^U\left(\tilde{v}_i / d_i\left(\hat{t}_i\right)\right)$ and $\hat{y}_j = \hat{f}_j^U\left(\tilde{v}_j / d_j\left(\hat{t}_j\right)\right)$. It will be shown that both are StPoE agreements, i.e. $\left(\hat{x}, \hat{t}_i\right) \in Z_i^*$ and $\left(\hat{y}, \hat{t}_j\right) \in Z_j^*$, by establishing that they are self-enforcing as a pair: using them as continuation outcomes, one can construct simple implementations of both.

The key observation is that $\left(\hat{x}, \hat{t}_i\right)$ 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 $\left(\hat{x}, \hat{t}_i\right)$ after another rejection (in fact, i is indifferent in the initial round):

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

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

Because $\hat{t}_i \le \tilde{t}_i$ it is also true that $\hat{f}_i^U\left(\tilde{v}_i / d_i\left(\hat{t}_i\right)\right) + \hat{f}_j^U\left(\tilde{v}_j / d_j\left(\hat{t}_j - 1\right)\right) \le 1$ and therefore $d_j\left(\hat{t}_i - 1\right) u_j\left(\hat{x}_j\right) \ge v_j = d_j\left(\hat{t}_j\right) u_j\left(\hat{y}_j\right)$. Hence, if both of $\left(\hat{x}, \hat{t}_i\right)$ and $\left(\hat{y}, \hat{t}_j\right)$ are StPoE outcomes, then they support $\left(\hat{x}, \hat{t}_i\right)$ as StPoE 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

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

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 $x_{\rho(t+1)} \ge \hat{x}_{\rho(t+1)}$. 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 a StPoE implementing $\left(\hat{x}, \hat{t}_i\right)$ 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 a StPoE implementing $\left(\hat{y}, \hat{t}_j\right)$ is played. It is clear that this construction supports $\left(\hat{x}, \hat{t}_i\right)$ as StPoE outcome in G_i if $\left(\hat{x}, \hat{t}_i\right)$ and $\left(\hat{y}, \hat{t}_j\right)$ are indeed StPoE outcomes. A similar construction can be devised to then also support $\left(\hat{y}, \hat{t}_j\right)$ as a StPoE outcome. Thus the two are self-enforcing. The argument is complete and establishes a StPoE with delay $\hat{t}_i$ in G_i, which is clearly not a RubE. □

While the uniqueness condition about the solutions to the system of (four) equations is not obviously interpretable in any useful way, by lemma 2.1, it requires a unique Rubinstein pair. For this case, the property that no other solutions exist is equivalent to the players' preferences not permitting the construction of a pair of self-enforcing StPoE outcomes with delay as the proof provides one. Technically, this possibility, illustrated in chapter 1, is the novel phenomenon that may arise when preferences are time-inconsistent, and that my approach accomodates through the distinction of respondent threats from analoguous proposer payoffs and the introduction of maximal delays.

A characterisation of StPoE 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 14 for any $\tilde{t}_i \in T$ follows from the continuity of players' utility functions in a way similar to lemma 2.1. Now, for each $i \in I$, let B_i denote the set of pairs $\left(\hat{v}_i, \hat{t}_i\right) \in \mathscr{U}_i \times T$ such that, for some pair $\left(\hat{v}_j, \hat{t}_j\right) \in \mathscr{U}_j \times T$, $\left(\hat{v}_i, \hat{t}_i, \hat{v}_j, \hat{t}_j\right)$ solves the system of equations 14-15. 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 StPoE payoffs of player k in G_i.

THEOREM 2.2. *Under assumption 2, $(v_i, t_i)_{i \in I} = (v_i^*, t_i^*)_{i \in I}$. Moreover, for $\{i,j\} = I$, the set of StPoE payoffs in G_i is given by*

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

PROOF. Take any $i \in I$ and note that by corollary 2.7 t_i^* is well-defined. It is easily verified that $u_i\left(1 - f_j^U\left(\delta_j(0)\, u_j\left(1 - f_i^U\left(\delta_i(t)\, v_i\right)\right)\right)\right)$ is non-increasing in t so in fact $v_i^* = \min\{v_i \in \mathscr{U}_i \,|(v_i, t_i^*) \in B_i\}$. In order to establish that $(v_i^*, t_i^*)_{i \in I}$ is a solution to the system of equations 14-15, it needs to be shown that, for $\{i,j\} = I$, $t_i^* = \max\left\{t \in T \,\middle|\, c_i\left(v_i^*, v_j^*, t\right) \leq 1\right\}$, but this follows from the fact that c_i is non-decreasing in each argument.

Each of $(v_i^*)_{i \in I}$ can be shown to be indeed an StPoE payoff following the construction of StPoE in the proof of theorem 2.1, whence $v_i \leq v_i^*$. On the other hand, the necessity of equation 14 means that $v_i \geq v_i^*$. Hence we obtain $v_i = v_i^*$, and the payoff bounds follow from the relationships in proposition 2.3.

Connectedness and closedness of the payoff intervalls as well as $(t_i)_{i \in I} = (t_i^*)_{i \in I}$ are immediate consequences of corollary 2.3. □

Discussion. Consider now the question of what an interesting sufficient condition for uniqueness at the level of individual preferences could be. First, it would have to ensure uniqueness of a Rubinstein pair. As discussed at the end of section 2.3, the standard assumption of concavity would be sufficient to do so. However, there may still be multiplicity if one can find other solutions to the system 14-15: this would be of the form that by increasing $\tilde{t}_i$ away from zero, $\delta_i\left(\tilde{t}_i\right)$—accordingly, i's worst respondent threat—drops sufficiently to feed through the other payoff extremes and permit a lower solution for i's worst proposer payoff $\tilde{v}_i$ that implies threats to i as the proposer which are consistent with delay $\tilde{t}_i$. Given a unique Rubinstein pair, if increasing t does not lower $\delta_i(t)$, then this cannot happen, and the RubE is indeed the unique equilibrium.

Hence, one important case of a sufficient condition for uniqueness is concavity of u_i together with $\delta_i(0) = \inf\{\delta_i(t) \,|t \in T\}$ for each $i \in I$. The discussion of section 1.2.2 of chapter 1 shows how this property can be interpreted as present bias (in a weak sense) and is satisfied by exponential discounting, (β, δ)-discounting and hyperbolic discounting, thus all of the most familiar time preferences.

Indeed, apart from the extra burden of further notation, chapter 1's insights qualitatively generalise to any case with a unique Rubinstein pair. In the case of multiple Rubinstein pairs that, as such, would not permit constructions of equilibrium delay, these interact with the possibility of delay through time-inconsistency (as described above), so there may still be delay equilibria.

On the other hand, time preferences which are conducive to delay are, roughly, such that the two players' δ_i's start high and then drop sharply at small but positive t, in combination with $f_i(\cdot, 1)$'s being very small for small shares but then increase very fast; in the terminology introduced at the end of section 2.3.1, the latter means there is relatively great loss to one period's delay for small shares turning into relatively small losses for large shares.

What about a meaningful notion of bargaining power in this class of preferences? In the context of general separable time preferences, the components which drive v_i are δ_i and f_i^U. Both a uniform (weak) increase of f_i^U and of δ_i improve i's bargaining outcome in the sense of increasing v_i weakly. Once more, this simply combines what has been known for the time-consistent exponential discounting case—see e.g. Osborne and Rubinstein [1990, Section 3.10.2]—with the insights of chapter 1.

Theorem 2.2, especially together with corollary 2.4, also implies the first characterisation of equilibrium outcomes for general separable time preferences with exponential discounting, without restrictions on the loss to delay and thus also covering preferences with "very convex" u_i in their representation. Applications with such non-standard u_i's may arise from the reduced form of certain "fairness" preferences and can then use this result.

Comparative Statics. Finally, I present two comparative statics results which generalise the discussion at the end of chapter 1's section 1.7. They serve to qualify two familiar results from the time-consistent exponential-discounting case within the context of general separable time preferences: first, the result that patience pays in bargaining and, second, the result that being the initial proposer is better than being the initial respondent (see, again, Osborne and Rubinstein [1990, Section 3.10.2-3]).

For any two players i and i' with preferences representable as in assumption 2 such that $u_i = u_{i'}$, say that i' is *uniformly more patient* than i if, for all $t \in T \setminus \{0\}$, $d_{i'}(t) > d_i(t)$. 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 a given opponent j, denote the resulting StPoE payoff extrema and maximal StPoE delays according to the following scheme: $v'_{i'}$ is the minimal proposer payoff of i' and v'_j is the minimal proposer payoff of j.

COROLLARY 2.5. *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 $v'_j \leq v_j$, $v'_{i'} \leq v_i$ 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 $t_i \leq t'_{i'}$ as well as $t_j \leq t'_j$.*

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

Hence, whenever delay equilibria exist, there is a sense in which players can be made more patient such that the sets of equilibrium payoffs and delays expand.

A second observation is that a player i does "not necessarily" prefer to be the initial proposer, or $v_i^* < u_i\left(f_i^U\left(d_i\left(1\right)u_i\left(1-f_j^U\left(\delta_j\left(t_j^*\right)v_j^*\right)\right)\right)\right)$; to be clear, the comparison is one of the worst proposer payoff and the best respondent payoff (this is how "not necessarily" is used). Since $v_i^* = u_i\left(1-f_j^U\left(\delta_j\left(0\right)u_j\left(1-f_i^U\left(\delta_i\left(t_i^*\right)v_i^*\right)\right)\right)\right)$, this is equivalent to

$$(16)\qquad 1 < f_j^U\left(d_j\left(1\right)u_j\left(1-f_i^U\left(\delta_i\left(t_i^*\right)v_i^*\right)\right)\right) + f_i^U\left(d_i\left(1\right)u_i\left(1-f_j^U\left(\delta_j\left(t_j^*\right)v_j^*\right)\right)\right).$$

Note that the right-hand side is the sum of the players' maximal threat points, in share terms; the symmetry of this condition immediately reveals that $v_i^* < W_i^*$, so 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 note that the right-hand side of the above inequality, by continuity, limits to

$$1 - f_i^U\left(\delta_i\left(t_i^*\right)v_i^*\right) + 1 - f_j^U\left(\delta_j\left(t_j^*\right)v_j^*\right).$$

Since, as is easily seen from the previous corollary, increasing players' one-period discount factors can only expand the sets of equilibrium outcomes and payoffs, there is still multiplicity and $f_i^U(\delta_i(t_i^*)v_i^*) < 1/2$, whence inequality 16 is satisfied in the limit.

COROLLARY 2.6. *Players' preferences are such that inequality 16 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. □

2.4. Conclusion

Based on a novel analytical approach to infinite-horizon alternating-offers bargaining, this paper characterised equilibrium outcomes for general separable time preferences without the restriction to stationary strategies that the small existing literature on this question had imposed. Qualitatively, the results combine familiar findings from the time-consistent exponential case where players' instantaneous utilities are non-linear in shares with those of chapter 1. Nonetheless, for applied research that involves bargaining in the context of various non-standard preferences covered here, the characterisation could be a useful result of reference.

Finally, the analytical approach, which has strong similarities to that of Abreu [1988] for repeated games promises to be fruitful for analyses also (i) of non-separable time preferences in the same environment (see the survey of Frederick et al. [2002] for general evidence and the work of Noor [2011] for the so-called "magnitude effect" in particular) or (ii) of dynamically inconsistent time preferences also in other stochastic games satisfying some version of stationarity.

References

Dilip Abreu. On the theory of infinitely repeated games with discounting. *Econometrica*, 56 (2):383–396, 1988.

James Andreoni and Charles Sprenger. Estimating time preferences from convex budgets. *The American Economic Review*, 102(7):3333–3356, 2012.

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

Ned Augenblick, Muriel Niederle, and Charles Sprenger. Working over time: Dynamic inconsistency in real effort tasks. January 2013.

Christopher Avery and Peter B. Zemsky. Money burning and multiple equilibria in bargaining. *Games and Economic Behavior*, 7(2):154–168, 1994.

Hector Chade, Pavlo Prokopovych, and Lones Smith. Repeated games with present-biased preferences. *Journal of Economic Theory*, 139(1):157–175, 2008.

Peter C. Fishburn and Ariel Rubinstein. Time preference. *International Economic Review*, 23 (3):677–694, 1982.

Shane Frederick, George Loewenstein, and Ted O'Donoghue. Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2):351–401, 2002.

Drew Fudenberg and Jean Tirole. *Game Theory*. The MIT Press, 1991.

Peter J. Hammond. Changing tastes and coherent dynamic choice. *The Review of Economic Studies*, 43(1):159–173, 1976.

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

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

Ted O'Donoghue and Matthew Rabin. Doing it now or later. *The American Economic Review*, 89(1):103–124, 1999.

Efe A. Ok and Yusufcan Masatlioglu. A theory of (relative) discounting. *Journal of Economic Theory*, 137(1):214–245, 2007.

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

Jinrui Pan, Craig S. Webb, and Horst Zank. Discounting the subjective present and future. April 2013.

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

Julio J. Rotemberg. Monetary policy and costs of price adjustment. *Journal of Economic Dynamics and Control*, 5(1):267–288, 1983.

Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.

Avner Shaked and John Sutton. Involuntary unemployment as a perfect equilibrium in a bargaining model. *Econometrica*, 52(6):1351–1364, 1984.

Ingolf Ståhl. *Bargaining Theory.* EFI The Economics Research Institute, Stockholm, 1972.

Robert H. Strotz. Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, 23(3):165–180, 1955-1956.

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

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

Appendix

LEMMA 2.4. *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 StPoE 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 $y_j = 1 - f_i^U(U_i(z^0)) \geq f_j^U(U_j(z^0))$. □

The first lemma provides a fundamental insight about the proposer advantage in bargaining: for any continuation StPoE outcome a respondent may obtain upon rejection, there exists an equilibrium with immediate agreement in which the proposer extracts all the benefits from agreeing earlier relative to that continuation outcome. From this result immediately follows that disagreement is not an equilibrium outcome.

COROLLARY 2.7. *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 StPoE 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 StPoE with delay can yield a proposing player a payoff greater than all StPoE without delay.

LEMMA 2.5. *For any* $i \in I$*,* $V_i = V_i^0$*.*

PROOF. Suppose $V_i > V_i^0$, implying that there exists an StPoE 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 2.4, 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. □

In view of lemma 2.5, it is shown below that a player's supremum StPoE payoff when respondent is simply the once-discounted supremum StPoE payoff when proposer, i. e. $W_i = d_i(1)\, V_i$. This relationship between a player's supremum StPoE payoffs in her two different roles is the same as found under exponential discounting.

LEMMA 2.6. *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 2.5 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 StPoE payoffs across players, based on the proposer advantage that is captured by lemma 2.4: the infimum StPoE payoff of a proposer is simply the payoff resulting from immediate agreement when the respondent expects her supremum StPoE payoff upon rejection; moreover, this statement holds true also when interchanging infimum and supremum.

LEMMA 2.7. *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,

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

By a similar argument,

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

Now, for the first equality, note that lemma 2.4 implies that

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

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 $U_i(z) < u_i\left(1 - f_j^U(W_j)\right)$ and let σ^* be an StPoE of G_i that induces it. Because $f_j^U(W_j) < 1$ must hold by impatience (there is at least one round's delay and $u_j(1) = 1$), the continuity of u_i guarantees existence of a proposal $x \in X$ such that

$U_i(z) < u_i(x_i) < u_i\left(1 - f_j^U(W_j)\right)$ which is accepted as $x_j > f_j^U(W_j)$ and thus constitutes a profitable deviation for i from σ^*, a contradiction.

For the second equality, note that lemma 2.4 implies that

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

And since j rejects any proposal $x \in X$ with $x_j < f_j^U(w_j)$ it follows that $V_i^0 \leq u_i\left(1 - f_j^U(w_j)\right)$ also, which establishes the claim via lemma 2.5. $\square$

Next, I will establish that there is no StPoE with delay that is worse to the proposer than the worst StPoE without delay, a result analogous to lemma 2.5 for a proposer's infimum payoffs.

Lemma 2.8. *For any $i \in I$, $v_i = v_i^0$.*

Proof. By lemma 2.4,

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

The proof of lemma 2.7 shows that

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

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

In general, however, only the first of the two properties of lemma 2.6 has an analogous version for infimum payoffs.

Lemma 2.9. *For any $i \in I$, $w_i^0 = d_i(1)\, v_i$.*

Proof. The following 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 2.8 this implies that $w_i^0 = d_i(1)\, v_i$. $\square$

CHAPTER 3

A Note on Choice and Welfare in Strotz-Pollak Equilibrium

3.1. Introduction

Strotz-Pollak equilibrium (StPoE) is the standard solution concept for intertemporal decision problems of individuals who have time-inconsistent preferences and perfectly know themselves. Dating back to the pioneering work of Strotz [1955-1956], this solution has been interpreted as the outcome of "consistent planning". Yet, a recurrent finding in applications is that outcomes thus obtained are inefficient according to the welfare criterion of Pareto-optimality when applied to the sequence of temporal selves of the decision maker; two well-known examples study the choices of (β, δ)-discounters in a timing problem (O'Donoghue and Rabin [1999, Proposition 5]) and in a consumption-savings problem (Phelps and Pollak [1968] or Laibson [1994, Chapter 1]), respectively. Such inefficient solutions represent instances of severe miscoordination of behaviour across time, which raises the question of what forms of dynamic inconsistency of preferences and environments permit or prevent this phenomenon.

This note presents welfare results about StPoE paths in general decision problems with perfect information. A main challenge in relating welfare rankings to equilibrium in general is the history-dependence of constraints as well as welfare. Nonetheless, allowing for arbitrary such history-dependence, the first result, proposition 3.1, provides a sufficient condition for "intrapersonal Pareto-optimality" of a StPoE path in finite-horizon problems without indifference: a limited form of intertemporal consistency of preferences, called "essential consistency" in reference to Hammond [1976] who originally advanced it, ensures this efficiency property. This result is illustrated and discussed with several examples of timing problems of a (β, δ)-discounter based on O'Donoghue and Rabin [1999].

Restricting the history-dependence inherent in the decision problem, corollary 3.1, relates welfare and multiplicity of StPoE paths by showing that under these restrictions, whenever a path supported by a StPoE is not intrapersonally Pareto-optimal, then any path that intrapersonally Pareto-dominates it, can also be supported by some StPoE. The welfare-rankability of multiple StPoE paths features prominently in various examples used to motivate refinements of StPoE—see Asheim [1997] and Kocherlakota [1996]—to which the result presented here adds a general observation. Together with proposition 3.2, which it is based upon, it also illuminates the occurrence of this phenomenon for the case of the consumption-savings problem of a (β, δ)-discounter introduced by Phelps and Pollak [1968] and rigorously analysed by Laibson [1994, Chapter 1].

3.2. Decision Problem

This section defines a general class of decision problems by a single decision-maker (DM) and the welfare criterion used throughout, and it presents two influential models from the literature which provide the running examples of this note.

3.2.1. Stages, Actions and Histories. There is a set of consecutive decision times $\mathcal{T} = \{t \in \mathbb{N}_0 \mid t < T\}$, where $T \in \mathbb{N}_0 \cup \{\infty\}$, at each of which a single DM takes an action a_t out of some non-empty, but possibly trivial (singleton), subset of a universal action space $\mathcal{A}$.[1] For any $t \in \mathcal{T}$, the set of all histories to time $t+1$ is denoted H^{t+1} and defined inductively from H^t via a mapping $A_t : H^t \to \mathcal{A}$, capturing constraints on actions at time t that evolve as a function of past choices: $H^0 \equiv \{\alpha\}$, where α is a parameter of the problem, and, for any $t \in \mathcal{T}$,

$$H^{t+1} = \left\{(h, a) \in H^t \times \mathcal{A} | a \in A_t(h)\right\}.$$

The set of terminal histories, called "paths", is then $H^T \equiv \Omega$, and the set of non-terminal histories, or—in what follows—simply "histories", is $\cup_{t\in\mathcal{T}} H^t \equiv \mathcal{H}$. It will be notationally convenient to also define a function $\tau : \mathcal{H} \cup \Omega \to \mathcal{T} \cup T$, such that, for any history $h \in \mathcal{H}$, $\tau(h) = t$ where $h \in H^t$, and, for any $\omega \in \Omega$, $\tau(\omega) = T$.

Generalising the above, for any $h \in \mathcal{H}$ and any time $t \geq \tau(h)$, define the set of histories to time t which are feasible after h, the "time-t continuations of h", denoted H^t_h, as follows: $H^{\tau(h)}_h \equiv \{h\}$ and, for any $t \geq \tau(h)$,

$$H^{t+1}_h = \left\{(h', a) \in H^t_h \times \mathcal{A} | a \in A_t(h')\right\}.$$

Accordingly, the set of paths feasible after h is $H^T_h \equiv \Omega_h$, and the set of histories feasible after h is $\cup_{t\geq\tau(h)} H^t_h \equiv \mathcal{H}_h$.

Finally, define the mapping $\eta : (\mathcal{H} \cup \Omega)^2 \to (\mathcal{H} \cup \Omega)$ to associate with any pairwise combination of histories or paths the longest history such that both are feasible: for any $(x, y) \in (\mathcal{H} \cup \Omega)^2$,

$$\eta(x, y) = \begin{cases} x & x = y \\ h & x \neq y, \{x, y\} \subseteq \mathcal{H}_h \cup \Omega_h, [\forall a \in A_{\tau(h)}(h), \{x, y\} \nsubseteq \mathcal{H}_{(h,a)} \cup \Omega_{(h,a)}] \end{cases}.$$

Note that this is well-defined, because whenever $x \neq y$, there is a unique history h with the required property; moreover, $\eta(x, y) = \eta(y, x)$. For any two histories h and h', whenever $\eta(h, h') = h$, then say h is a subhistory of h' and h' is a continuation history of h; and for any history h and path ω, if $\eta(h, \omega) = h$, then call h a history along ω.

3.2.2. (Pure) Strategies. A pure strategy of the DM is a function $s : \mathcal{H} \to \mathcal{A}$ with the property that, for any $h \in \mathcal{H}$, $s(h) \in A_{\tau(h)}(h)$; let $\mathcal{S}$ denote the set of such functions. For any $h \in \mathcal{H}$ and any time $t \geq \tau(h)$, define a mapping $\omega^t_h : \mathcal{S} \to H^t$ inductively as follows, where $\omega^{\tau(h)}_h(s) \equiv h$, and

$$\omega^{t+1}_h(s) = \left(\omega^t_h(s), s\left(\omega^t_h(s)\right)\right).$$

[1]The possibility of trivial action spaces at various dates allows to capture discrete-time problems where decision dates are not equidistant in time, or also problems where after some time no decisions are to be made any more, while there are still welfare effects.

Then, for any $h \in \mathcal{H}$, $s \in \mathcal{S}$ and date $t \geq \tau(h)$, $\omega_h^t(s)$ is the time-t continuation of h which results from following strategy s. Define $\omega_\alpha^t \equiv \omega^t$, so, in particular, $\omega^T(s)$ is simply the path under s.

For any $s \in \mathcal{S}$ and any $h \in \mathcal{H}$, denote the restriction of s to $\mathcal{H}_h$ by s_h and let $\mathcal{S}_h$ denote the set of functions thus obtained; elements of $\mathcal{S}_h$ will be called continuation strategies from h.

3.2.3. Preferences and Welfare Comparisons. At any time $t \in \mathcal{T}$, the DM has "preferences" over Ω which are represented by a function $U_t : \Omega \to \mathbb{R}$; note that, given domain Ω, U_t is allowed to vary with the particular history $h \in H^t$. Importantly, U_t goes beyond a representation of preferences in the usual sense: since it is defined for all paths at any time, two paths $\{\omega, \omega'\}$ may be compared even though there is no time-t history upon which both are actually feasible (formally, there does not exist any $h \in H^t$ such that $\{\omega, \omega'\} \subseteq \Omega_h$). Hence there is no choice experiment, not even under options with full commitment, that could elicit these "preferences". Thus U_t in fact measures the DM's welfare at time t for any path, and when feasible paths are compared, this implies preferences.

The welfare criterion I use throughout is a mere translation of the standard economic concept of Pareto efficiency into the language of dynamic paths and a single DM.

Definition 3.1. For any two paths $\{\omega, \omega'\} \subseteq \Omega$, ω *intrapersonally Pareto-dominates* (IP-dominates) ω' if, for any time $t \in \mathcal{T}$, $U_t(\omega) \geq U_t(\omega')$, and for some time $t' \in \mathcal{T}$, $U_{t'}(\omega) > U_{t'}(\omega')$. A path $\omega \in \Omega$ is *intrapersonally Pareto-optimal* (IP-optimal, "efficient") if there is no path $\omega' \in \Omega$ that IP-dominates ω.

3.2.4. Subproblems and Conventions. Denote any such decision problem by Γ; clearly, any history $h \in \mathcal{H}$ defines a decision problem of its own: by simply replacing h with α and times $t \geq \tau(h)$ with $t - \tau(h)$, it fits all the definitions above, and I will therefore denote this "subproblem" by $\Gamma(h)$. To simplify some of the notation here and in what follows, I make the convention that, for any history $h \in \mathcal{H}$ and any $t \in \mathcal{T}$, $\left(h, (a_s)_{s=t}^{t-1}\right) = h$. Moreover, when writing a history to some time t in explicit form as $(a_s)_{s=0}^{t-1}$, I usually omit α; however, $(a_s)_{s=0}^{-1} \equiv \alpha$.

3.2.5. Examples. This work focuses attention on two examples, which are among the most influential contributions to the analysis of decision making with time-inconsistent preferences. The first one is the model of O'Donoghue and Rabin [1999]: a (β, δ)-discounter chooses when to engage in a one-time activity before a deadline, where the activity yields immediate and delayed rewards as well as costs that vary with the timing of the activity. Real-life applications include the choice of when to prepare a report, visit a doctor for a medical check-up or go on a vacation.

Example 3.1. Let the "deadline" be $T < \infty$, set $\alpha = 0$ and $\mathcal{A} = \{0, 1\}$, where, for any $t \in \mathcal{T}$ and $h = \left(\alpha, (a_s)_{s=0}^{t-1}\right) \in H^t$, $z_t(h) \equiv \max\{a_s\}_{s=0}^{t-1}$ and $A_t(h) \equiv \{0, 1 - z_t(h)\}$. Action $a = 1$ at time t, when available, means that the DM performs the activity in period t; $A_t(h) = \{0\}$ if she has performed it in the past, though there still are welfare consequences to consider. The set of paths can be characterised by the timing of the activity: $\Omega = \mathcal{T} \cup T$, where $\omega = T$ is interpreted as performing the activity right at the deadline, when it must be done.[2] Let there be two non-negative functions $v : \Omega \to \mathbb{R}_+$ and $c : \Omega \to \mathbb{R}_+$, which define welfare together with a parameter β such that $0 < \beta \leq 1$, and distinguish two different types of problem.[3] First, a problem with immediate costs (and delayed rewards) is one where:

$$U_t(\omega) = \begin{cases} \beta\left(v(\omega) - c(\omega)\right) & t < \omega \\ \beta v(\omega) - c(\omega) & t = \omega \\ \beta v(\omega) & t > \omega \end{cases}.$$

The other type of this problem has immediate rewards (and delayed costs) instead:

$$U_t(\omega) = \begin{cases} \beta\left(v(\omega) - c(\omega)\right) & t < \omega \\ v(\omega) - \beta c(\omega) & t = \omega \\ -\beta c(\omega) & t > \omega \end{cases}.$$

Given how Ω is defined, the reward- and cost-schedules can be written as vectors of length $T+1$, so I will use the notation $v = (v_t)_{t=0}^{T}$ and $c = (c_t)_{t=0}^{T}$, where v_t and c_t are the reward- and cost-values, respectively, when the activity is performed in period t.

The second example is based on the formulation of Plan [2010, Example 4] of the following problem originally proposed by Phelps and Pollak [1968] and reinterpreted as well as further analysed by Laibson [1994, Chapter 1]: a (β, δ)-discounter with constant relative risk aversion facing a constant return on savings chooses a discrete consumption-savings path over an infinite time-horizon.[4]

Example 3.2. Let $T = \infty$, $\alpha = W_0 > 0$ and, for any $t \in \mathcal{T}$, $A_t = A = [0,1]$. W_0 is the DM's initial wealth and $a \in A$ is the fraction of wealth saved for the future in any period. With a constant gross interest rate of $R \geq 0$ and a given history $h = \left(W_0, (a_s)_{s=0}^{t-1}\right)$ to time $t \in \mathcal{T}$, wealth at time t equals $W_t = R^t \left(\prod_{s=0}^{t-1} a_s\right) W_0$. Preferences, and in fact welfare, are parameterised by (β, δ, ρ) with $0 < \beta \leq 1$, $0 < \delta < 1$ and $\rho < 1$, where the standard restriction

[2] For example, ignoring the initial history, if $T = 3$, then $\omega = 1$ is the path $(0,1,0)$ and $\omega = 3$ is the path $(0,0,0)$. See the discussion in O'Donoghue and Rabin [1999, p. 107, in particular footnote 12].

[3] The assumption about the (β, δ)-discounter that $\delta = 1$ is immaterial; see O'Donoghue and Rabin [1999, footnote 11] which shows that any "long-term discounting" can be incorporated in v and c.

[4] See also Barro [1999] for a variant of this problem in continuous time with more general time-varying time preferences and a neoclassical production technology, Krusell and Smith-Jr. [2003] who investigate stationary savings rules for more general (instantaneous) utility functions and savings technologies, or Bernheim et al. [2013] who extend this problem to the case of a credit constraint (a lower bound on assets at any time).

that $\delta R^{1-\rho} < 1$ is imposed:

$$\begin{aligned} U_t\left(W_0, (a_s)_{s=0}^{\infty}\right) &= \left((1-a_t) W_t\right)^{1-\rho} + \beta \sum_{s=t+1}^{\infty} \delta^{s-t} \left((1-a_s) R^{s-t} \left(\prod_{r=t}^{s-1} a_t\right) W_t\right)^{1-\rho} \\ &= W_t^{1-\rho} \underbrace{\left((1-a_t)^{1-\rho} + \beta \sum_{s=t+1}^{\infty} \delta^{s-t} \left((1-a_s) R^{s-t} \left(\prod_{r=t}^{s-1} a_t\right)\right)^{1-\rho}\right)}_{\equiv U\left((a_s)_{s=t}^{\infty}\right)} \\ &= \left(R^t \left(\prod_{s=0}^{t-1} a_s\right) W_0\right)^{1-\rho} U\left((a_s)_{s=t}^{\infty}\right). \end{aligned}$$

Note that this decision problem satisfies a history-independence property (see definition 3.5 below): action sets are constant and history enters welfare in a multiplicative manner, which means it does not affect the ranking of feasible continuation plans; the latter is always represented by the function $U : [0,1]^T \to \mathbb{R}$ as defined above.[5]

3.3. Choice and Welfare

3.3.1. Strotz-Pollak Equilibrium. Strotz [1955-1956] pioneered the analysis of a time-inconsistent DM's behaviour in the context of a deterministic continuous-time consumption problem. He suggested that a DM who correctly anticipates her future preferences, a "sophisticated" DM, would select "the best plan among those that he will actually follow" (Strotz [1955-1956, p. 173]), which Pollak [1968, Section 1] formalised for a discretised version of the original problem. Early generalisations of this definition can be found in Peleg and Yaari [1973, p. 395], Goldman [1979], pointing out the equivalence with (a particular application of) subgame-perfect Nash equilibrium (SPNE), and Goldman [1980], where the terminology of "Strotz-Pollak equilibrium" that the literature has adopted is introduced. Laibson [1994] describes the general solution as the SPNE of the "intrapersonal game" where each temporal self of the DM is defined to be a distinct non-cooperative player. The same approach has been applied to decision problems featuring imperfect recall (see Piccione and Rubinstein [1997] and other contributions to the same (special) journal issue).

Definition 3.2. A strategy $\hat{s} \in \mathcal{S}$ is a *Strotz-Pollak equilibrium* (StPoE) if, for any $h \in \mathcal{H}$ and $a \in A_{\tau(h)}(h)$,

$$U_{\tau(h)}\left(\omega_h^T(\hat{s})\right) \geq U_{\tau(h)}\left(\omega_{(h,a)}^T(\hat{s})\right).$$

A path $\hat{\omega} \in \Omega$ is a *Strotz-Pollak solution* (StPo-solution) if there exists a StPoE $\hat{s} \in \mathcal{S}$ such that $\omega^T(\hat{s}) = \hat{\omega}$.

StPoE requires that, at any history h, the DM best-responds to correct beliefs about future behaviour such that this behaviour, at any future history, is a best response to the same beliefs. The DM cannot commit to future actions but forms beliefs about them which, when shared at all histories, imply rational behaviour. As is clear from the definition as well as this description,

[5]Phelps and Pollak [1968] and Laibson [1994, Chapter 1] formulate this problem with absolute consumption as the action chosen in any period, subject to the wealth constraint, which is history-dependent.

if $\hat{s}$ is a StPoE of $\Gamma(\alpha)$, then, for any history $h \in \mathcal{H}$, $\hat{s}_h$ is a StPoE of $\Gamma(h)$ (the converse holds true as well, of course).

StPoE is an application of SPNE to the game with the same extensive form, but where a separate non-cooperative player acts at each decision time (equivalently, at each history, because only one history can be played to any given decision time). Thus, well-known existence theorems for SPNE apply, e.g. Harris [1985].[6] It shares the notion of "credibility" inherent in SPNE, where, fixing beliefs, the DM does not expect to take actions in the future that she would not find optimal once the contingency were to actually occur. Applied to a single DM with perfect self-knowledge, this could be termed loosely as ruling out that she "fool" herself.

3.3.2. Essential Consistency and Welfare. Recall example 3.1 with immediate rewards for $T = 2$, where $\beta = \frac{1}{2}$ and the reward- and cost-schedules are given by

$$\begin{aligned} v &= (0, 5, 1) \\ c &= (1, 8, 0) \,. \end{aligned}$$

This results in the following unique StPoE: since $U_1(1) = 5 - \frac{1}{2}8 > \frac{1}{2}1 = U_1(2)$, the DM in period 1 would engage in the activity. Therefore, it will actually be performed immediately: $U_0(0) = -\frac{1}{2}1 > -\frac{1}{2}(5-8) = U_0(1)$. Compare now the welfare consequences from this outcome to that if the DM waited until period 2 instead: $U_0(0) = U_1(0) = -\frac{1}{2}$, whereas $U_0(2) = U_1(2) = \frac{1}{2}$. The unique StPo-solution is therefore IP-dominated. The reason the DM does not wait initially, even though she would strictly prefer doing it in period 2 rather than now, is that she would otherwise do it next period; at that point, however, she would prefer the (then) immediate reward over waiting yet another period.

Clearly, these preferences are time-inconsistent, because the DM's preferences as of the initial period over doing it next period and doing it the period after that reverse once the next period arrives. In order to make terms precise, I provide a definition of the benchmark of time-consistency here.

Definition 3.3. (TC) Preferences are *time-consistent* if, for any history $h \in \mathcal{H}$ and any two paths $\{\omega, \omega'\} \subseteq \Omega_h$ (feasible after h),

$$U_{\tau(h)}(\omega) \geq U_{\tau(h)}(\omega') \Leftrightarrow U_0(\omega) \geq U_0(\omega') \,.$$

When preferences are time-consistent, there is a single utility function—without loss of generality, it is chosen to be U_0—that represents the DM's preferences over *feasible* paths at any history. Accordingly, if a path is optimal for the DM at the initial date 0, then it remains optimal for the DM at any history along that path; in particular, if, at the outset of the problem, the DM has a unique optimal path, then it remains uniquely optimal for the DM at any history along this path among all the paths feasible at that history.

In the example above, however, the nature of the violation of time-inconsistency is special. Notice the following intertemporal cycle: at $t = 1$, the DM prefers doing it in period 1 over doing it in period 2, whereas at $t = 0$, the DM prefers doing it in period 2 over doing it in

[6] See also Goldman [1980] who establishes existence of StPoE in a general class of finite-horizon problems, where Peleg and Yaari [1973] had initially raised concerns about non-existence despite "well-behaved" settings.

period 0 which is in turn preferred to doing it in period 1. This constitutes a violation of the following property.[7]

DEFINITION 3.4. (EC) Preferences are *essentially consistent* if, for any pair of histories $\{h, h'\} \subseteq \mathcal{H}$ and triple of paths $\{\omega, \omega', \omega''\} \subseteq \Omega$ such that $h = \eta(\omega, \omega') = \eta(\omega, \omega'')$ and $h' = \eta(\omega', \omega'') \in \mathcal{H}_h$,

$$U_{\tau(h')}(\omega') > U_{\tau(h')}(\omega'') \wedge U_{\tau(h)}(\omega) > U_{\tau(h)}(\omega') \quad \Rightarrow \quad U_{\tau(h)}(\omega) > U_{\tau(h)}(\omega''). \tag{17}$$

I formulate this consistency property in strict terms because I only use it in proposition 3.1 which rules out indifference. Moreover, it is thus identical to the property advanced by Hammond [1976], who showed—again for the case of no indifference—that essential consistency ensures the coincidence of sophisticated and naïve choices in finite decision trees ($T < \infty$ and $\mathcal{A}$ finite). It requires that sophisticated choice from $\{\omega, \omega', \omega''\}$ at history h not change when an alternative that is not chosen but still available at future history h' is removed.

Clearly, essential consistency is implied by time-consistency; however, it is indeed weaker: assuming no indifference and considering the same paths and histories as in the definition, when $U_{\tau(h)}(\omega) < \min\{U_{\tau(h)}(\omega'), U_{\tau(h)}(\omega'')\}$, it does not restrict preferences at history h over $\{\omega', \omega''\}$, nor when both $U_{\tau(h)}(\omega) > U_{\tau(h)}(\omega')$ and $U_{\tau(h')}(\omega') < U_{\tau(h')}(\omega'')$ are true, whereas time-consistency requires they coincide with those at history h'.

REMARK 3.1. If preferences are time-consistent, then they are essentially consistent, but the converse is not true.

PROOF. Suppose TC and consider histories and paths as in the definition of EC. Under TC, the antecedent in (17) is equivalent to

$$U_0(\omega') > U_0(\omega'') \wedge U_0(\omega) > U_0(\omega'),$$

which, by transitivity of $>$, yields $U_0(\omega) > U_0(\omega'')$, and applying TC once more gives $U_{\tau(h)}(\omega) > U_{\tau(h)}(\omega'')$.

For a counterexample to the converse, consider example 3.1 for $T = 2$ with immediate costs and $\beta = \frac{1}{2}$, where $v = (3, 3, 1)$ and $c = (2, 2, 1)$, so

$$\begin{array}{c} U_0(0) = -\frac{1}{2} < U_0(2) = 0 < U_0(1) = \frac{1}{2} \\ U_1(1) = -\frac{1}{2} < U_1(2) = 0 \end{array},$$

so these clearly violate TC. In contrast, EC is satisfied because—in the definition's notation—it must be that $\omega = 0$, whence the antecedent of (17) is vacuous here. □

In finite-horizon settings without any indifference essential consistency guarantees that the StPo-solution—there is a unique one by backwards induction—is IP-optimal. Alternatively put, if a StPo-solution is found to be inefficient by the Pareto-criterion, it must be that preferences violate essential consistency.[8] The proof of this result uses the following lemma, which exploits the structure of StPoE based on backwards induction when the horizon is finite.

[7] In the definition's notation, $\omega = 0$, $\omega' = 1$ and $\omega'' = 2$; these are compared as of $t = 0$ and $t = 1$.

[8] While IP-dominance, applied to example 3.1, compares present discounted utilities, O'Donoghue and Rabin [2001, Section V] demonstrate an even stronger "dominance" property: there exists another performance period which yields *instantaneous* utility at least as great in every period and greater in some period than the unique StPo-solution (with the above numbers, the respective sequences of instantaneous utilities for periods 0, 1 and

LEMMA 3.1. *Let $T < \infty$, suppose $\hat{\omega}$ is a StPo-solution and take any other path $\omega = (a_t)_{t=0}^{T-1} \in \Omega \setminus \{\hat{\omega}\}$. Then there exist an integer K with $0 < K \le T$ and a sequence of paths $(\omega_k)_{k=0}^{K}$ with $\omega_0 = \hat{\omega}$ and $\omega_K = \omega$ such that, for any $k \in \{0, \ldots, K-1\}$,*

$$\begin{aligned} \eta(\omega_{k+1}, \omega) &\in \mathcal{H}_{\eta(\omega_k,\omega)} \setminus \{\eta(\omega_k, \omega)\} \\ U_{\tau(\eta(\omega_k,\omega))}(\omega_{k+1}) &\le U_{\tau(\eta(\omega_k,\omega))}(\omega_k). \end{aligned}$$

PROOF. Let $\hat{s}$ be a StPoE such that $\hat{\omega} = \omega^T(\hat{s})$ and construct a sequence of paths $(\omega_0, \omega_1, \ldots)$ as follows: set $h_0 \equiv \alpha$, and iterate

$$\begin{aligned} \omega_k &\equiv \omega_{h_k}^T(\hat{s}) \\ h_{k+1} &\equiv \left(\eta(\omega_k, \omega), a_{\tau(\eta(\omega_k,\omega))}\right) \end{aligned}$$

until $\omega_k = \omega$, in which case set $K = k$. It is easily checked that this sequence satisfies $0 < K \le T$, $\omega_0 = \hat{\omega}$ and $\eta(\omega_{k+1}, \omega) \in \mathcal{H}_{\eta(\omega_k,\omega)} \setminus \{\eta(\omega_k, \omega)\}$.

Denote, for simplicity, $t_k \equiv \tau(\eta(\omega_k, \omega))$ and suppose now there is a $k \in \{0, \ldots, K-1\}$ such that $U_{t_k}(\omega_k) < U_{t_k}(\omega_{k+1})$. Letting $h = \eta(\omega_k, \omega)$, this would therefore imply that

$$U_{\tau(h)}\left(\omega_h^T(\hat{s})\right) < U_{\tau(h)}\left(\omega_{\left(h, a_{\tau(h)}\right)}^T(\hat{s})\right),$$

which contradicts that $\hat{s}$ is a StPoE. □

PROPOSITION 3.1. *Let $T < \infty$ and assume preferences exhibit no indifference in the sense that, for any time $t \in \mathcal{T}$ and any two paths $\{\omega, \omega'\} \subseteq \Omega_h$ (feasible after h) with $\omega \ne \omega'$, $U_t(\omega) \ne U_t(\omega')$ holds true. Then there is a unique StPo-solution, and if preferences are essentially consistent, it is IP-optimal.*

PROOF. Uniqueness of StPoE in this finite-horizon problem follows from backwards induction, since there is no indifference. Denote this unique StPoE by $\hat{s}$ and the associated unique StPo-solution by $\hat{\omega}$.

Take any path $\omega \ne \hat{\omega}$ and consider a sequence $(\omega_k)_{k=0}^{K}$ as in lemma 3.1; because there is no indifference, for any $k \in \{0, \ldots, K-1\}$, $U_{t_k}(\omega_k) > U_{t_k}(\omega_{k+1})$. In particular, $U_{t_{K-1}}(\omega_{K-1}) > U_{t_{K-1}}(\omega)$, and if $K = 1$, then $\omega_{K-1} = \hat{\omega}$, whence ω does not IP-dominate $\hat{\omega}$. If $K > 1$, since for any $k' > k$, $\omega_{k'} \in \Omega_{h_k}$, we can apply EC as follows:

$$\begin{aligned} U_{t_{K-1}}(\omega_{K-1}) > U_{t_{K-1}}(\omega) \quad &\wedge \quad U_{t_{K-2}}(\omega_{K-2}) > U_{t_{K-2}}(\omega_{K-1}) \\ &\Rightarrow \\ U_{t_{K-2}}(\omega_{K-2}) \quad &> \quad U_{t_{K-2}}(\omega). \end{aligned}$$

If $K = 2$ then $\omega_{K-2} = \hat{\omega}$, so this means ω does not IP-dominate $\hat{\omega}$. If $K > 2$ then apply EC once more:

$$\begin{aligned} U_{t_{K-2}}(\omega_{K-2}) > U_{t_{K-2}}(\omega) \quad &\wedge \quad U_{t_{K-3}}(\omega_{K-3}) > U_{t_{K-3}}(\omega_{K-2}) \\ &\Rightarrow \\ U_{t_{K-3}}(\omega_{K-3}) \quad &> \quad U_{t_{K-3}}(\omega). \end{aligned}$$

---

2 are $(0,0,0)$ for the StPo-solution and $(0,0,1)$ for performance in period 2 instead). Of course, their criterion is applicable only in discounted-utility models.

If $K = 3$ then $\omega_{K-3} = \hat{\omega}$, so this means ω does not IP-dominate $\hat{\omega}$. Since $K < \infty$, applying EC $K - 1$ times, this process will eventually yield $U_{t_0}(\hat{\omega}) > U_{t_0}(\omega)$, implying that ω does not IP-dominate $\hat{\omega}$. This is true for any $\omega \neq \hat{\omega}$, whence $\hat{\omega}$ is IP-optimal. $\square$

Discussion. Essential consistency rules out intertemporal cycles: when later the DM will be decisive about ω' versus ω'' in favour of ω' and now decides about $\{\omega\}$ versus $\{\omega', \omega''\}$, she does not prefer ω'' to ω and ω to ω'. The proof of proposition 3.1 shows that, for each alternative path that is not the unique StPo solution, there exists a time $t \in \mathcal{T}$ at which the DM prefers the solution to that path; in fact, this t is the first time the DM's action deviates from the alternative path. Considering the generality of the decision problem in terms of the history-dependence of welfare, this is a remarkable result, despite the strength of essential consistency.

For an illustration of this efficiency result when preferences are time-inconsistent, recall the special case of example 3.1 used in remark 3.1, where it was established that preferences are indeed essentially consistent. The unique StPo-solution is to wait until period 2 to perform the task, and this is IP-optimal: the time-0 DM prefers this path to doing it immediately, and the same is true at time 1 about doing it immediately then instead.

In a special case of example 3.1, essential consistency is also necessary.

Remark 3.2. In example 3.1 with $T = 2$, $\beta < 1$ and immediate rewards, the unique StPo-solution is IP-optimal if and only if preferences are essentially consistent.

Proof. Sufficiency follows from proposition 3.1, so suppose EC were violated. This means either (i) $U_1(1) > U_1(2)$ and $U_0(2) > U_0(0) > U_0(1)$ or (ii) $U_1(2) > U_1(1)$ and $U_0(1) > U_0(0) > U_0(2)$. However, (ii) cannot hold with immediate rewards because:

$$\begin{aligned} U_1(2) > U_1(1) \quad &\Leftrightarrow \quad \beta(v_2 - c_2) > v_1 - \beta c_1 \\ U_0(1) > U_0(2) \quad &\Leftrightarrow \quad \beta(v_1 - c_1) > \beta(v_2 - c_2), \end{aligned}$$

which implies $\beta v_1 > v_1$, a contradiction (since $\beta < 1$ and $v_1 \geq 0$).

Consider then case (i): the unique StPo-solution is that the activity is performed immediately. This path is IP-dominated by waiting to do it in period 2 whenever $U_1(2) > U_1(0)$, i.e. $\beta(v_2 - c_2) > -\beta c_0$; the latter is, however, an implication of $U_0(2) > U_0(0)$ because $v_0 \geq 0$:

$$U_0(2) > U_0(0) \Leftrightarrow \beta(v_2 - c_2) > v_0 - \beta c_0.$$

$\square$

For longer horizons, an essential inconsistency may be irrelevant to the StPo-solution. Informally, if one added a new initial period in which the DM prefers doing it immediately over any other outcome, this would result in an IP-optimal StPo-solution, irrespective of whether in the subproblem after waiting initially there is an essential inconsistency or not. Hence, essential consistency certainly needs to be weakened further for a characterisation of IP-optimality in example 3.1 with immediate rewards when $T > 2$, or even beyond to cope with both immediate rewards and immediate costs.

Significantly generalising proposition 3.1 to dealing with indifference would require first a modification of the notion of essential consistency and hardly appears promising.[9] Again using example 3.1 with $T = 2$ and $\beta < 1$, note that whenever at $t = 1$ the DM is indifferent between the two remaining feasible paths, she has a strict preference at $t = 0$ (with immediate costs for doing it in period 1 and with immediate rewards for doing it in period 2). Depending on v_0 and c_0, how this indifference at $t = 1$ translates into (expected) choice at $t = 1$ may determine behaviour at $t = 0$ and consequently result in two different StPo-solutions where one IP-dominates the other in a manner orthogonal to essential consistency.

Relatedly, when moving toward an infinite horizon, the assumption of no indifference becomes hardly defensible. Moreover, essential consistency loses its force as sequences constructed on the basis of the proof of lemma 3.1 become infinite.[10] Indeed, the work of Laibson [1994, Chapter 1, Section 3] shows that this extends to the case of even time-consistent preferences when payoffs are unbounded from below in example 3.2 (time-consistency there means $\beta = 1$): letting $\rho > 1$, any path can be supported as StPo-solution by the threat that, upon any past deviation, consumption would take place at a (constant) rate sufficiently close to one (the continuation payoff approaches negative infinity). Even with bounded payoffs, Plan [2010, Footnote 12] shows how, with infinite cascades of threats of ever lower savings rates, one can construct a StPoE such that at every *history*, adhering to it makes the DM better off than the stationary, constant-savings-rate StPoE proposed by Phelps and Pollak [1968] (the latter features undersaving and is used as the limiting savings rate of the punishment cascade).

3.3.3. History-Independence, Welfare and Multiplicity. The previous section presented a sufficiency result for the IP-optimality of a StPo-solution, and its discussion indicated how this welfare property may fail more generally. Relatedly, an argument used to discard particular StPo-solutions is that they are IP-dominated *by other StPo-solutions*: this phenomenon is shared by most examples that the literature introducing refinements of StPoE has produced, e.g. Kocherlakota [1996] or Asheim [1997]. While hardly made explicit, the argument seems to be that it reflects an implausible failure of coordination in that the beliefs arrived at are self-defeating: there is another "credible" path that IP-dominates the one resulting from those beliefs, so a "planning" DM will never coordinate future beliefs on such a strategy.

This section addresses the question of when this form of Pareto-rankable multiplicity obtains and thus also provides insights into existence of an IP-optimal StPo-solution. In order to be able to do so, I restrict the history-dependence inherent in the decision problem. Recall that welfare at any time is defined for all paths, whence also paths that are never altogether feasible are compared by the welfare criterion (see section 3.2.3). In contrast, for equilibrium choices, only comparisons of feasible paths matter. Without restrictions on the nature of history-dependence, welfare comparisons of feasible paths may not provide any information about welfare at other paths, and it is impossible to uncover implications for equilibrium properties from the welfare criterion in general. Since, to the best of my knowledge, this work is the first investigation

[9] One conclusion is immediate from lemma 3.1, however, when in the above definition of essential consistency (17) is instead formulated with weak preferences: no StPo-solution is strongly IP-dominated (IP-dominance with strict "preference" for every time t).

[10] See also the discussion of essential consistency in infinite trees by Hammond [1976, pp. 170-171].

of welfare of StPo-solutions beyond particular models, I consider the following rather strong properties.

DEFINITION 3.5. A decision problem satisfies *history-independence* if, for any time $t \in \mathcal{T}$ and any two histories $\{h, h'\} \subseteq H^t$, (i) $A_t(h) = A_t(h') \equiv A_t$, and, (ii), for any two sequences of continuation play $\left\{(a_s)_{s=t}^{T-1}, (a'_s)_{s=t}^{T-1}\right\} \subseteq \times_{s=t}^{T-1} A_s$,

$$U_t\left(h, (a_s)_{s=t}^{T-1}\right) \geq U_t\left(h, (a'_s)_{s=t}^{T-1}\right) \Leftrightarrow U_t\left(h', (a_s)_{s=t}^{T-1}\right) \geq U_t\left(h', (a'_s)_{s=t}^{T-1}\right).$$

It satisfies history-independence *even in a welfare sense* if (ii) is replaced by (ii*), for any continuation play $(a_s)_{s=t}^{T-1} \in \times_{s=t}^{T-1} A_s$, $U_t\left(h, (a_s)_{s=t}^{T-1}\right) = U_t\left(h', (a_s)_{s=t}^{T-1}\right)$.

History-independence of a decision problem means that, after any two histories to a particular date, the sets of feasible continuations are (i) identical (history-independent constraints) and (ii) ranked the same way (history-independent *preferences*).[11] It does not imply that welfare is unaffected by past choices, however, which is true only upon replacing (ii) with (ii*); clearly, the latter is stronger.[12] Example 3.2 illustrates this point, since initial wealth in any period (more precisely, a positive transformation of wealth), which is determined by past savings choices, enters the utility function multiplicatively, whereby it does not affect the rankings of continuation paths; thus (ii) holds whereas (ii*) is violated. This example also demonstrates that there are nonetheless important economic decision problems featuring dynamic constraints that (can be formulated so they) satisfy history-independence (see Plan [2010] for a closely related point).

The essence of history-independence is that, conditional on time, the DM's continuation behaviour can always ignore the past: any continuation play that is feasible at *some* history is feasible after *any* history, whence, if some continuation play constitutes a StPoE after that history, because of (ii) in definition 3.5, this is true after any other history to the same decision time; this is the content of the following lemma.

LEMMA 3.2. *Assume the decision problem satisfies history-independence. Take a strategy $s \in \mathcal{S}$ such that s_h is a StPoE of $\Gamma(h)$ for history $h \in \mathcal{H}$ and consider any history $h' \in \mathcal{H}$ with $\tau(h') = \tau(h) = \tau$. Then any strategy $s' \in \mathcal{S}$ such that, for any non-negative integer $k \leq T - \tau$ and $(a_t)_{t=\tau}^{\tau+k-1} \in \times_{t=\tau}^{\tau+k-1} A_t$,*

$$s'\left(h', (a_t)_{t=\tau}^{\tau+k-1}\right) = s\left(h, (a_t)_{t=\tau}^{\tau+k-1}\right)$$

satisfies that $s'_{h'}$ is a StPoE of $\Gamma(h')$.

PROOF. Suppose $s'_{h'}$ is not a StPoE of $\Gamma(h')$, so there exist a history $\hat{h}' \in \mathcal{H}_{h'}$ and an action $\bar{a} \in A_{\tau(\hat{h}')}$ such that

$$U_{\tau(\hat{h}')}\left(\omega^T_{(\hat{h}', \bar{a})}(s')\right) > U_{\tau(\hat{h}')}\left(\omega^T_{\hat{h}'}(s')\right).$$

[11]Note that (ii) relies on (i) to be well-defined; although one could define the history-independence of preferences independently to hold only when continuation plays are actually feasible under both histories, for the purposes here, this is unnecessary as (ii) is only considered in problems which satisfy (i) anyways.

[12](ii*) implies $U_t\left(h, (a_s)_{s=t}^{T-1}\right) = U_t\left(h', (a_s)_{s=t}^{T-1}\right)$ and $U_t\left(h, (a'_s)_{s=t}^{T-1}\right) = U_t\left(h', (a'_s)_{s=t}^{T-1}\right)$ from which (ii) follows.

Note that, by part (i) of history-independence of a decision problem as in definition 3.5, $\hat{h}' = \left(h', (a_t)_{t=\tau}^{\tau+k-1}\right)$ for some non-negative integer $k \leq T - \tau$, and consider $\hat{h} = \left(h, (a_t)_{t=\tau}^{\tau+k-1}\right)$. The definition of s' on $\mathcal{H}_{h'}$ via s on $\mathcal{H}_h$ implies that $\omega_{(\hat{h}',\bar{a})}^T (s')$ and $\omega_{(\hat{h},\bar{a})}^T (s)$ are identical from time τ onwards, and the same is true about the two paths $\omega_{\hat{h}'}^T (s')$ and $\omega_{\hat{h}}^T (s)$. Therefore, part (ii) of a decision problem's history-independence yields that

$$U_{\tau(\hat{h})} \left(\omega_{(\hat{h},\bar{a})}^T (s)\right) > U_{\tau(\hat{h})} \left(\omega_{\hat{h}}^T (s)\right).$$

This, however, contradicts the hypothesis that s_h is a StPoE of $\Gamma(h)$. □

Lemma 3.2 allows to establish a multiplicity result about history-independent decision problems, which is related to the welfare criterion in the discussion that follows, precisely in corollary 3.1 which uses definition 3.6.

PROPOSITION 3.2. *Assume the decision problem satisfies history-independence and let $\hat{\omega} = (\hat{a}_t)_{t=0}^{T-1}$ be a StPo-solution; then any other path $\tilde{\omega} = (\tilde{a}_t)_{t=0}^{T-1} \neq \hat{\omega}$ such that, for any time $t \in \mathcal{T}$,*

$$U_t(\tilde{\omega}) \geq U_t\left((\tilde{a}_s)_{s=0}^{t-1}, (\hat{a}_s)_{s=t}^{T-1}\right), \tag{18}$$

is also a StPo-solution.

PROOF. Take any StPoE $\hat{s}$ with $\omega^T(s) = \hat{\omega}$ and consider the strategy $\tilde{s}$ constructed as follows: whenever a history $h \in \mathcal{H}$ satisfies $\tilde{\omega} \in \Omega_h$, set $\tilde{s}(h) = \tilde{a}_{\tau(h)}$; then note that any other history can be written as $h = \left(h', \bar{a}, (a_t)_{t=\tau(h')+1}^{\tau(h')+k}\right)$ for some $k \in \mathbb{Z}$ with $0 \leq k \leq T - \tau(h') - 1$ and where $\eta(h, \tilde{\omega}) = h'$, in which case set

$$\tilde{s}(h) = \hat{s}\left((\hat{a}_t)_{t=0}^{\tau(h')-1}, \bar{a}, (a_t)_{t=\tau(h')+1}^{\tau(h')+k}\right).$$

This defines $\tilde{s}$ for every history h such that $\tilde{\omega} \notin \Omega_h$.

It will now be shown that $\tilde{s}$ is a StPoE and thus that $\tilde{\omega}$ is indeed a StPo-solution. Consider first any history h with $\tilde{\omega} \notin \Omega_h$ and note that there exist a history h' and an action $\bar{a} \in A_{\tau(h')}$ such that $\tilde{\omega} \in \Omega_{h'}$, $\tilde{\omega} \notin \Omega_{(h',\bar{a})}$ and $h \in \mathcal{H}_{(h',\bar{a})}$. Since, for $h'' = \left((\hat{a}_t)_{t=0}^{\tau(h')-1}, \bar{a}\right)$, $\hat{s}_{h''}$ is a StPoE of $\Gamma(h'')$, lemma 3.2 establishes that $\tilde{s}_{(h',\bar{a})}$ is a StPoE of $\Gamma(h', \bar{a})$; because $h \in H_{(h',\bar{a})}$, $\tilde{s}_h$ is therefore a StPoE of $\Gamma(h)$.

Now take a history h with $\tilde{\omega} \in \Omega_h$ and consider any $a \in A_{\tau(h)}$ with $a \neq \tilde{s}(h) = \tilde{a}_{\tau(h)}$. By definition of $\tilde{s}$, at all times $t \geq \tau(h)$, the actions on path $\omega_{(h,a)}^T(\tilde{s})$ are identical to those on path $\omega_{(h',a)}^T(\hat{s})$ when $h' = (\hat{a}_t)_{t=0}^{\tau(h)-1}$; using that $\tau(h') = \tau(h)$, since $\hat{s}$ is a StPoE, $U_{\tau(h)}\left(\omega_{(h',a)}^T(\hat{s})\right) \leq U_{\tau(h)}\left(\omega_{h'}^T(\hat{s})\right) = U_{\tau(h)}(\hat{\omega})$. The history-independence of preferences (property (ii) of definition 3.5) then yields that $U_{\tau(h)}\left(\omega_{(h,a)}^T(\tilde{s})\right) \leq U_{\tau(h)}\left(h, (\hat{a}_t)_{t=\tau(h)}^{T-1}\right)$. Combining this last inequality with $U_{\tau(h)}\left(h, (\hat{a}_t)_{t=\tau(h)}^{T-1}\right) \leq U_{\tau(h)}(\tilde{\omega})$ from the hypothesis of the proposition, one finally obtains $U_{\tau(h)}\left(\omega_{(h,a)}^T(\tilde{s})\right) \leq U_{\tau(h)}\left(\omega_h^T(\tilde{s})\right)$, completing the proof. □

Discussion. While it may appear that lemma 3.2 should immediately yield that if a path IP-dominates a StPo-solution, that path must be supportable by a StPoE as well—it could be based on the very same "threats"—this is not true in general. Consider the following simple example of a history-independent decision problem: $T = 2$, $A_0 = A_1 = \{0, 1\}$, $U_0(a_0, a_1) = -|a_0 - a_1|$

and $U_1(a_0, a_1) = 2a_0 - a_1$. At time $t = 1$, the DM prefers 0 over 1 irrespective of the previous action, whence she matches this action with $a_0 = 0$ at $t = 0$. Yet, this unique StPo-solution $(0, 0)$ is IP-dominated by $(1, 1)$. Note how this example relies on the history-dependence of *welfare* in the second period.

However, proposition 3.2 illuminates example 3.2: there is a unique StPoE with the property that consumption/saving takes place at the same rate irrespective of time and history. This "simple" equilibrium was first identified by Phelps and Pollak [1968, Part IV], who also showed that the resulting path is IP-dominated by other constant-rate paths of consumption/saving. Because we are comparing constant-rate paths, inequality 18 holds true: to see this, first note that when $a_s = \tilde{a} > 0$ for all $s \in \mathcal{T}$, then, for any $t \in \mathcal{T}$, $W_t = (R\tilde{a})^t W_0$ and (using the assumption that $\delta R^{1-\rho} < 1$)

$$
\begin{aligned}
U\left((a_s)_{s=t}^{\infty}\right) &= (1-\tilde{a})^{1-\rho}\left(1+\beta\sum_{s=t+1}^{\infty}\left(\delta\left(R\tilde{a}\right)^{1-\rho}\right)^{s-t}\right)\\
&= (1-\tilde{a})^{1-\rho}\left(1+\beta\left(-1+\sum_{s=0}^{\infty}\left(\delta\left(R\tilde{a}\right)^{1-\rho}\right)^{s}\right)\right)\\
&= (1-\tilde{a})^{1-\rho}\left(1+\frac{\beta\delta\left(R\tilde{a}\right)^{1-\rho}}{1-\delta\left(R\tilde{a}\right)^{1-\rho}}\right)\\
&\equiv V(\tilde{a})\,.
\end{aligned}
$$

Next, suppose a constant savings rate of $\tilde{a}$ IP-dominates a constant savings rate of $\hat{a}$. Because it is at least as good as of $t = 0$ when wealth is the same, this implies that $V(\tilde{a}) \geq V(\hat{a})$, which immediately yields inequality 18 where wealth is also identical in the comparison. Hence proposition 3.2 establishes that these other paths are StPo-solutions as well, although as such, they must be supported by more "complex" strategies involving history-dependence (see Laibson [1994, Chapter 1]).

Indeed, I conjecture that, more generally, example 3.2 satisfies the following "regularity" property.

Definition 3.6. A decision problem satisfying history-independence is *welfare-regular* if, whenever a path $\omega = (a_t)_{t=0}^{T-1}$ is not IP-optimal, there exists a path $\omega' = (a'_t)_{t=0}^{T-1}$ which IP-dominates ω and, moreover, is such that, for any time $t \in \mathcal{T}$,

$$U_t(\omega') \geq U_t\left((a'_s)_{s=0}^{t-1}, (a_s)_{s=t}^{T-1}\right). \tag{19}$$

Welfare regularity restricts the history-dependence of welfare: if a path ω is not IP-optimal, then there is some other path ω' that IP-dominates it, where as long as ω' has been followed, the DM would never prefer switching to continuation as under ω over staying on ω'. Observe the similarity of inequalities (18) and (19), and note that the example given at the outset of this discussion violates welfare-regularity. Of course, welfare-regularity is weaker than history-independence in a welfare sense.

Remark 3.3. If a decision problem satisfies history-independence even in a welfare sense, then it is welfare-regular.

PROOF. Simply note that when history-independence is satisfied in a welfare sense, in the above definition, $U_t\left(\left(a'_s\right)_{s=0}^{t-1}, \left(a_s\right)_{s=t}^{T-1}\right) = U_t(\omega)$, whence IP-dominance immediately yields the inequality. □

COROLLARY 3.1. *Assume the decision problem satisfies history independence and is welfare-regular. Then, if a StPo-solution is not IP-optimal, it is IP-dominated by another StPo-solution.*

PROOF. Let $\hat{\omega} = (\hat{a}_t)_{t=0}^{T-1}$ be a StPo-solution, where a path $\tilde{\omega} = (\tilde{a}_t)_{t=0}^{T-1}$ IP-dominates $\hat{\omega}$. Because the decision problem is welfare-regular, it is without loss of generality to choose $\tilde{\omega} = (\tilde{a}_t)_{t=0}^{T-1}$ such that inequality (18) holds true, whence it is a StPo-solution. □

Based on this corollary, I conjecture that every non-IP-optimal StPo-solution in example 3.2 is in fact IP-dominated by another StPo-solution (so that this is not only true about constant-rate paths).

In any case, this result immediately implies that if a decision problem satisfying history-independence which is welfare-regular has a unique StPo-solution, then this solution is IP-optimal. Moreover, under standard "well-behavedness" assumptions (e.g. compact action spaces and continuous utility functions), where IP-dominance of a path comes with the existence of an IP-optimal path that IP-dominates it, there then exists an IP-optimal StPo-solution.

3.4. Conclusion

This note addresses two important welfare phenomena in decision problems with time-inconsistent preferences: Pareto-inefficiency of StPo-solutions and IP-rankable multiplicity of such solutions. In a framework that allows for history-dependent welfare, my first result delineates the forms of intertemporal conflict inherent in preferences that yield inefficient outcomes in the Pareto-sense by showing that they must violate essential consistency whenever the horizon is finite and there is no indifference. Essential consistency is in fact necessary in a simple version of the "timing problem" analysed by O'Donoghue and Rabin [1999] where rewards are immediate and costs are delayed. While the discussion points out the likely obstacles to generalisations of these results even within the framework that this note assumes, because truncation is a popular approach to selection among multiple StPoE (see e.g. Laibson [1997]), finite-horizon results about welfare are also of interest for work on infinite-horizon problems.

The property of essential consistency was proposed by Hammond [1976] for a similar decision environment, where he discovered it to be sufficient for the coincidence of naïve and sophisticated choice. An interesting question is therefore the more general relationship between Pareto-efficiency and this invariance property of choice to various degrees of preference misprediction.

On the other hand, when some StPo-solution in a decision problem satisfying history-dependence fails to be IP-optimal, then this comes with IP-rankable multiplicity *of StPo-solutions* when the effects of past play on welfare satisfy a certain regularity property. The latter kind of multiplicity appears to have played a major role for the development of refinements of StPoE, but whereas the work in this area so far has relied mostly on rather abstract and specific examples (of the class described) to promote their own respective approaches, I thus organise them into a general insight. Moreover, beyond such abstract examples, my result applies also to the influential consumption-savings model of Phelps and Pollak [1968].

Maybe most importantly, however, the last result can be used to establish existence of IP-optimal StPo-solutions under standard technical assumptions. To the best of my knowledge, no such result has been available. Of course, its generalisation to a broader class of problems would be highly desirable for applications.

References

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

Robert J. Barro. Ramsey meets laibson in the neoclassical growth model. *The Quarterly Journal of Economics*, 114(4):1125–1152, 1999.

B. Douglas Bernheim, Debraj Ray, and Sevin Yeltekin. Poverty and self-control. January 2013.

Steven M. Goldman. Intertemporally inconsistent preferences and the rate of consumption. *Econometrica*, 47(3):621–626, 1979.

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

Peter J. Hammond. Changing tastes and coherent dynamic choice. *The Review of Economic Studies*, 43(1):159–173, 1976.

Christopher Harris. Existence and characterization of perfect equilibrium in games of perfect information. *Econometrica*, 53(613-628), 1985.

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

Per Krusell and Anthony A. Smith-Jr. Consumption-savings decisions with quasi-geometric discounting. *Econometrica*, 71(1):365–375, 2003.

David I. Laibson. *Essays in Hyperbolic Discounting*. PhD thesis, MIT, 1994.

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

Ted O'Donoghue and Matthew Rabin. Doing it now or later. *The American Economic Review*, 89(1):103–124, 1999.

Ted O'Donoghue and Matthew Rabin. Choice and procrastination. *The Quarterly Journal of Economics*, 116(1):121–160, 2001.

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

Edmund S. Phelps and Robert A. Pollak. On second-best national saving and game-equilibrium growth. *The Review of Economic Studies*, 35(2):185–199, 1968.

Michele Piccione and Ariel Rubinstein. On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior*, 20:3–24, 1997.

Asaf Plan. Weakly forward-looking plans. May 2010.

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

Robert H. Strotz. Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, 23(3):165–180, 1955-1956.