Repeated Games with Present-Biased Preferences[*]

Hector Chade
Arizona State University

Pavlo Prokopovych
EERC, Ukraine

Lones Smith
University of Michigan

October, 2006

Abstract

We study infinitely repeated games with observable actions, where players have present-biased (so-called β-δ) preferences. We give a two-step procedure to characterize Strotz-Pollak equilibrium payoffs: compute the continuation payoff set using recursive techniques, and then use this set to characterize the equilibrium payoff set $U(\beta, \delta)$. While Strotz-Pollak equilibrium and subgame perfection differ here, the generated paths and payoffs nonetheless coincide.

We then explore the cost of the present-time bias. Surprisingly, unless the minimax outcome is a Nash equilibrium of the stage game, the equilibrium payoff set $U(\beta, \delta)$ is not separately monotonic in β or δ. While $U(\beta, \delta)$ is contained in payoff set of a standard repeated game with smaller discount factor, the present-time bias precludes any lower bound on $U(\beta, \delta)$ that would easily generalize the $\beta = 1$ folk-theorem.

[*]We have benefited from a conversation with Efe Ok at an early stage of the paper, the comments of Dan Silverman, and detailed helpful feedback from an associate editor and two referees. Hector Chade acknowledges the hospitality of the University of Michigan, where this work was started during his stay. This is based on chapter 1 of the PhD thesis of Pavlo Prokopovych. Lones Smith thanks the NSF for funding this work.

1 Introduction

In repeated decision-making, an agent with non-exponential time preferences must face the problem of dynamic inconsistency, namely: how should he currently behave if he knows that his future behavior might well undo his best laid current plans. If he can not precommit his future behavior, the best he can do, according to Strotz (1955–56), is to resign himself to the intertemporal conflict and choose "the best plan among those he will actually follow." That is, after any history of choices, the agent takes an action that maximizes his utility given his utility maximizing strategies of the future. Strotz-Pollak equilibrium, as formalized by Peleg and Yaari (1973), is then a strategy that is immune to one shot deviations, taking into account the intertemporal conflict. It has since also become conventional to interpret this as a subgame perfect equilibrium amongst temporally distinct selves: "a specification of a strategy for each player such that agent t's choice after any history of play is optimal given all other agents' strategies" (Kocherlakota, 1996).

The literature on such present-biased (or 'hyperbolic') discounting has so far been focused largely on decision makers acting either in isolation or in perfectly competitive environments. Examples include the savings consumption decision (Laibson, 1997, Rabin and ODonoghue 1999), the smoking decision (Gruber and Koszegi, 2001), the labor supply decision (Della Vigna and Paserman, 2003), and optimal taxation (Krusell, et al. 2002). This paper instead analyzes the strategic behavior of many hyperbolic discounters interacting in a non-cooperative game.

We study infinitely repeated games with observable actions, where players have β-δ preferences — i.e. with the quasi-geometric discount sequence $1, \beta\delta, \beta\delta^2, \ldots$. We first provide a two-step procedure to characterize the set of Strotz-Pollak equilibrium payoffs: First, a recursive methodology adapted from Abreu-Pearce-Stacchetti, 1986, 1990 (hereafter, APS) characterizes the set of continuation payoffs, using different weights on current versus future payoffs than are found in the incentive constraints. From this recursive set, the equilibrium payoffs are then computed. To ensure convex sets of payoffs, we then add a public randomization device.

Subgame perfection is the natural touchstone of credibility in repeated games.[1] In a strategic setting, the Strotz-Pollak equilibrium concept coincides under geometric discounting with the stronger notion of subgame perfection; however, the crucial one-shot deviation principle fails with time-inconsistency, so that these concepts

[1]We actually call it "sincere" subgame perfection, because a player at the beginning of the subgame analyzes the optimality of the strategy over the infinite horizon with his *current* preferences.

differ. Nevertheless, we show that the generated action paths and thereby payoffs nonetheless coincide. Unlike the single-agent case, it is therefore always possible to support a Strotz-Pollak equilibrium so that different "incarnations" of a player agree on the optimality of a strategy, given the other players' strategies: There is ultimately no intertemporal conflict among multiple selves.

A key result then separates this from geometric discounting: continuation payoffs are now a possibly proper subset of equilibrium payoffs. This wedge drives some interesting comparative statics, and allows us to calibrate the cost of the present-time bias in this strategic setting. We show that changing $\beta < 1$ matters more than shifts in the long-term discount factor $\delta < 1$. For suppose that an eternal one util flow discounted by the constant factor $\Delta < \delta$ has the same total present value as it does with β-δ preferences. Fixing Δ, the continuation and equilibrium payoff set shrinks if β falls and δ rises — despite the greater weight δ on continuation payoffs in the recursion. Thus, this comparative static cannot be proven by adapting the APS proof, and we provide an indirect argument by separation methods instead.

Surprisingly, the equilibrium payoff set need not be separately monotonic in β or δ. We show this by example, exploiting the non-coincidence of the worst equilibrium and continuation payoffs. We then rescue the "standard" APS payoff monotonicity for when the minimax point is a Nash equilibrium of the stage game, as is true, eg., with the Prisoner's Dilemma. In such games, the worst equilibrium and continuation payoffs clearly coincide, and the payoff sets expand as δ or β separately increase.

We finally try to understand the equilibrium payoff set by sandwiching it between those of standard repeated games with different discount factors. By a simple corollary of our comparative static, we bound it above using the discount factor Δ. But to really understand how much the present bias hurts payoffs, we seek a lower bound. For instance, might the payoff set implode with the present bias? The factor $\beta\delta$ clearly discounts all future payoffs less than the β-δ discounter does, and is a strong candidate for a lower bound. While this works in a symmetric class of prisoner's dilemma games, we show that it generally fails. An example proves that the present-time bias precludes any nontrivial lower bound set that could plausibly allow us to deduce a folk theorem from the standard one with geometric discounting.

The next section outlines the model and solution concepts. Section 3 develops the recursive characterization results, and relates Strotz-Pollak and subgame perfect equilibrium. We explore the effects of the present bias on the payoffs in Section 4.

2 The Model

We analyze infinitely repeated games with observable actions. The only departure from the standard such model is that we assume quasi-geometric (or β-δ) discounting. Preferences of this kind were first studied in decision theory by Phelps and Pollak (1968), and are sometimes called quasi-hyperbolic (e.g., see Laibson, 1996).

Denote the stage game by $G = (N, (A_i)_{i\in N}, (\pi_i)_{i\in N})$, where $N = \{1, \ldots, n\}$ is a finite set of players, $A_i = \{1, \ldots, I_i\}$ is player i's finite set of actions, and π_i is player i's payoff function from $A = \times_{i\in N} A_i$ to the real line $\mathbb{R}$. We shall assume that the stage game has a Nash equilibrium in pure strategies. Let V denote the set of feasible and weakly individually rational payoffs of G.

The repeated game $G^\infty(\beta, \delta)$ begins at stage 0, with the null history $h^0 = \varnothing$. At the beginning of stage k, each player observes the *history* h^k of actions chosen at all previous stages. If the action profile $a^k = (a_1^k, \ldots, a_n^k)$ is chosen at stage $k \in \{0, 1, \ldots\}$, then player i's stage-t ($t = 0, 1, \ldots$) discounted total payoff is:

$$\pi_i(a^t) + \beta \sum_{k=1}^{\infty} \delta^k \pi_i(a^{t+k}), \tag{1}$$

where β and δ (both between 0 and 1) are discount factors common to all players. The parameter β captures a player's bias for the present, for it implies that she assigns more relative weight to the stage-k payoff at stage k than she did at any stage prior to stage k. This present bias is the source of dynamic inconsistency in the model, since current and past preferences over 'today' versus 'tomorrow' differ.

Given discount factors β and δ, define the *effective* discount factor $\Delta(\beta, \delta)$ by:

$$1+\beta\delta+\beta\delta^2+\cdots = 1+\Delta(\beta,\delta)+\Delta(\beta,\delta)^2+\cdots = \frac{1}{1-\Delta(\beta,\delta)} \Rightarrow \Delta(\beta,\delta) = \frac{\beta\delta}{1-\delta+\beta\delta}$$

Critical throughout this paper is that $\Delta(\beta, \delta) < \delta$. Player i's stage-t discounted *average* payoff equals his stage-t discounted payoff scaled by $(1 - \Delta(\beta, \delta))$.

Let A^k denote the k-fold Cartesian product of A, and H^k the set of all stage-k action histories. Obviously, $H^k = A^k$. A *pure strategy* for player i is a map f_i from $H = \bigcup_{k=0}^{\infty} H^k$ to A_i. Since we restrict attention to pure strategies (see Abreu (1988)), we henceforth and throughout the paper drop the qualifier "pure".

Denote by $f_i^{|h}$ player i's strategy induced by f_i in the *subgame* after history

$h \in H$. By definition,

$$f_i^{|h}(h') = f_i(\{h, h'\})$$

for all $i \in N$, all $a_i \in A_i$, and all conjoined histories $\{h, h'\} \in H$, listing all actions in h followed by those in h'. As usual, we write $f^{|h} = (f_1^{|h}, \ldots, f_n^{|h})$.

Given a strategy profile $f = (f_1, \ldots, f_n)$ and a history $h \in H$, the discounted total payoff (1) of player i can be rewritten in discounted average form:

$$u_i(f^{|h} \mid \beta, \delta) = (1 - \Delta(\beta, \delta))\pi_i(f(h)) + \Delta(\beta, \delta)c_i(f^{|\{h, f(h)\}} \mid \delta),$$

where $c_i(f^{|\{h, f(h)\}} \mid \delta)$ is his naturally defined *continuation payoff.*

As mentioned in the Introduction, our solution concept is a straightforward extension of that found in the decision-theoretic literature on present-biased preferences to the infinitely repeated game setting under analysis.

Definition: *A strategy profile $f = (f_1, \ldots, f_n)$ is a* Strotz-Pollak equilibrium *if, for any history $h \in H$, player $i \in N$, and action $a_i \in A_i$,*

$$\begin{aligned} & (1 - \Delta(\beta, \delta))\pi_i(f(h)) + \Delta(\beta, \delta)c_i(f^{|\{h, f(h)\}} \mid \delta) \\ \geq \;& (1 - \Delta(\beta, \delta))\pi_i(a_i', f_{-i}(h)) + \Delta(\beta, \delta)c_i(f^{|\{h, (a_i', f_{-i}(h))\}} \mid \delta) \end{aligned}$$

Hence, a strategy profile is a Strotz-Pollak equilibrium if there are no profitable one-stage deviations. Intuitively, each player then finds it optimal, given his preferences at any stage, to carry out the actions his strategy specified for that stage.

Conditional on the players choosing their actions according to f, the stage-$(k-1)$ continuation payoff can be represented *along the equilibrium path* as follows:

$$c_i(f^{|h^k} \mid \delta) = (1 - \delta)\pi_i(f(h^k)) + \delta c_i(f^{|\{h^k, f(h^k)\}} \mid \delta).$$

As we later see, the recursive structure of the continuation payoffs allows us to adapt the APS technique to compute the Strotz-Pollak equilibrium continuation payoff set.

For the game $G^\infty(\beta, \delta)$, let $\Sigma^{S\&P}(\beta, \delta)$ denote the set of Strotz-Pollak equilibrium strategy profiles, $U(\beta, \delta)$ the set of Strotz-Pollak equilibrium payoffs, and $C(\beta, \delta)$ the set of Strotz-Pollak equilibrium continuation payoffs. Thus,

$$U(\beta, \delta) = \bigcup_{f \in \Sigma^{S\&P}(\beta, \delta)} u(f \mid \beta, \delta) \qquad \text{and} \qquad C(\beta, \delta) = \bigcup_{f \in \Sigma^{S\&P}(\beta, \delta)} c(f \mid \delta),$$

where $u(f|\beta, \delta) = (u_1(f|\beta, \delta), \ldots, u_n(f|\beta, \delta))$ and $c(f|\delta) = (c_1(f|\delta), \ldots, c_n(f|\delta))$.

3 Characterization Results

3.1 The Recursive Structure and Construction of $U(\beta, \delta)$

We now provide a constructive characterization of the set of Strotz-Pollak equilibrium payoffs, adapting the method of APS. Our two-step procedure can be succinctly explained as follows. First, we exploit the recursive structure of continuation payoffs and characterize the set $C(\beta, \delta)$ as the largest fixed point of an operator. Next, we construct the set $U(\beta, \delta)$ using continuation payoffs drawn from $C(\beta, \delta)$.

As in APS, fix a discount factor $\Delta \in (0, 1)$ and payoff set $W \subset \mathbb{R}^n$. Call a pair (a, κ), where $a \in A$ and κ is a function from A to continuation values W, Δ-*admissible* w.r.t. W if, for all players $i \in N$, and actions $(a_i, a_{-i}) \in A$ and $a'_i \in A_i$,

$$(1 - \Delta)\pi_i(a_i, a_{-i}) + \Delta\kappa_i(a_i, a_{-i}) \geq (1 - \Delta)\pi_i(a'_i, a_{-i}) + \Delta\kappa_i(a'_i, a_{-i}).$$

Let us define the following self-generating operator

$$B(W \mid \beta, \delta) = \{v \in V : v = (1-\delta)\pi(a) + \delta\kappa(a) \mid (a, \kappa) \text{ is } \Delta(\beta, \delta)\text{-admissible w.r.t. } W\}$$

A set $W \subset V$ is called *self-generating* if $W \subset B(W \mid \beta, \delta)$. For standard repeated games with $\beta = 1$, APS have shown that the self-generating sets are of interest precisely because they are candidate equilibrium payoff sets. Here, we focus on them for their recursive character. Just like the decision-theoretic setting (see Phelps and Pollak (1968)), but quite unlike in APS, the recursive structure applies only to the set of equilibrium *continuation* payoffs.

Theorem 1 (Characterization) *Fix* $0 < \beta, \delta < 1$.
(a) *If* $W \subset \mathbb{R}^n$ *is self-generating, then* $B(W \mid \beta, \delta) \subset C(\beta, \delta)$.
(b) Factorization*:* $C(\beta, \delta)$ *is the largest fixed point of the operator* $B(\cdot \mid \beta, \delta)$:

$$C(\beta, \delta) = \{v = (1-\delta)\pi(a) + \delta\kappa(a) \in V \mid (a, \kappa) \text{ is } \Delta(\beta, \delta)\text{-admissible w.r.t. } C(\beta, \delta)\}$$

Further, given $C(\beta, \delta)$*, the equilibrium payoff set* $U(\beta, \delta) \subset \mathbb{R}^n$ *satisfies:*

$$\begin{aligned} U(\beta, \delta) = \{v \in V : \quad & v = (1 - \Delta(\beta, \delta))\pi(a) + \Delta(\beta, \delta)\kappa(a) \\ & \text{and } (a, \kappa) \text{ is } \Delta(\beta, \delta)\text{-admissible w.r.t. } C(\beta, \delta)\}. \end{aligned}$$

(c) *The sets* $C(\beta, \delta)$ *and* $U(\beta, \delta)$ *are compact.*

The proof of (a) is similar to Theorem 1 of APS; the proof of (c) is standard. For (b), consider the collection of sets $\Gamma(\beta, \delta) = \{W \subset V : W \subset B(W \mid \beta, \delta)\}$. By definition, $C(\beta, \delta) \in \Gamma(\beta, \delta)$. As the stage game has a pure strategy Nash equilibrium by assumption, both $C(\beta, \delta)$ and $\Gamma(\beta, \delta)$ are non-empty. This collection of sets is partially ordered by the relation $\subset$. It is not difficult to see that any chain consisting of elements of $\Gamma(\beta, \delta)$ has an upper bound.By Zorn's Lemma, $\Gamma(\beta, \delta)$ has a maximal element; call it $W(\beta, \delta)$. Following APS, one can prove that $W(\beta, \delta) = C(\beta, \delta)$.

3.2 Strotz-Pollak Equilibrium and Subgame Perfection

In a Strotz-Pollak equilibrium, each player's strategy is a 'consistent plan' that takes into account — on and off the equilibrium path — the present bias that the player's preferences exhibit. In other words, at each stage of the game, each player's current 'self' finds it optimal to choose the actions specified by the strategy profile.

It is well known that in infinitely repeated games with observable actions and geometric discounting, the sets of Strotz-Pollak equilibria and subgame perfect equilibria coincide, by the one-stage deviation principle; hence, these equilibrium payoff sets coincide as well. We now analyze this 'equivalence' result when $\beta < 1$.

To this end, we call a strategy profile f a *sincere Nash equilibrium* of $G^\infty(\beta, \delta)$ if $u_i(f \mid \beta, \delta) \geq u_i(g_i, f_{-i} \mid \beta, \delta)$ for all $g_i : H \to A_i$ and all $i \in N$, and a *sincere subgame perfect equilibrium* (SPE) if $f^{|h}$ is a sincere Nash equilibrium of $G^\infty(\beta, \delta)$ for any history $h \in H$. The qualifier 'sincere' is added to emphasize that, in each subgame, a player evaluates the optimality of his strategy in all stages using his current preferences. This distinction is clearly immaterial if $\beta = 1$.

Notice that every sincere subgame perfect equilibrium is Strotz-Pollak, because the latter only checks for one-stage deviations at the beginning of each subgame. The following example shows that the converse is not true when $\beta < 1$. It also illustrates how to use Theorem 1 to construct Strotz-Pollak equilibria.

EXAMPLE 1: STROTZ-POLLAK EQUILIBRIUM NEEDN'T BE SUBGAME PERFECT. Consider an infinitely repeated prisoner's dilemma game with $\beta = \frac{1}{3}$ and $\delta = \frac{6}{7}$, so that $\Delta(\frac{1}{3}, \frac{6}{7}) = \frac{2}{3}$. The stage game is as follows:

	C	D
C	$1, 1$	$-1, 2$
D	$2, -1$	$0, 0$

It is easy to show that the set $W = \{(0,0), (1,1), (\frac{8}{7}, \frac{5}{7}), (\frac{41}{49}, \frac{44}{49})\}$ is self-generating and hence belongs to $C(\frac{1}{3}, \frac{6}{7})$. The vector $(0,0)$ can be supported by the action profile (D, D) and the continuation vector $(\kappa(C,C), \kappa(C,D), \kappa(D,C), \kappa(D,D)) = ((0,0), (0,0), (0,0), (0,0))$, which obviously constitutes an admissible pair. Likewise,

- $(1,1)$ can be supported by Nash reversion: $((C,C), ((1,1), (0,0), (0,0), (0,0)))$
- $(\frac{8}{7}, \frac{5}{7})$ can then be supported by $((D,C), ((0,0), (0,0), (1,1), (0,0)))$
- $(\frac{41}{49}, \frac{44}{49})$ can then be supported by $((C,D), ((1,1), (\frac{8}{7}, \frac{5}{7}), (0,0), (0,0)))$

In each case, it is straightforward to check that the pair suggested is admissible.

We shall use the set W of continuation payoffs to support the following Strotz-Pollak equilibrium path: (C,D), (C,D), (D,C), (C,C), (C,C), ... This path yields the average discounted payoff vector $u(f \mid \frac{1}{3}, \frac{6}{7}) = (\frac{11}{49}, \frac{62}{49})$ by the strategy profile f:

- Stage 0: the pair is $((C,D),\ ((0,0), (\frac{41}{49}, \frac{44}{49}), (0,0), (0,0)))$,
- Stage 1: after $h^1 = \{(C,D)\}$, the pair is $((C,D),\ ((1,1), (\frac{8}{7}, \frac{5}{7}), (0,0), (0,0)))$,
- Stage 2: after $\{(C,D), (C,D)\}$, $((D,C),\ ((0,0), (0,0), (1,1), (0,0)))$,
- Stage 3: after $\{(C,D), (C,D), (D,C)\}$, $((C,C),\ ((1,1), (0,0), (0,0), (0,0)))\ldots$

For instance, if a player deviates from the action profile (C,D) at stage 0, and, as a result, the stage-1 public history differs from $\{(C,D)\}$, the players would play (D,D) forever thereafter. A similar interpretation applies to subsequent stages.

It is not difficult to verify that (i) $u(f^{|\{(C,D)\}} \mid \frac{1}{3}, \frac{6}{7}) = (\frac{3}{7}, \frac{8}{7})$ and that (ii) this strategy profile is a Strotz-Pollak equilibrium. At stage 1, player 2 cannot benefit from deviating to C ($\frac{8}{7} > 1$) as he myopically attaches a relative weight $\frac{1}{3}$ to stage 1's payoff. But notice that the relative weight he attaches at stage 0 to stage 1's continuation payoff equals $\frac{6}{7}$, not $\frac{2}{3}$. As a result, if player 2, given his preferences at stage 0, plans to switch to C at stage 1, then his stage-0 expected payoff rises to $\frac{4}{3}$. Thus, this strategy profile is not a sincere subgame perfect equilibrium. $\square$

Inspired by this example, we now show that, unlike the single-agent case, we can always support a Strotz-Pollak equilibrium path by a strategy profile that meets the stronger requirement of sincere subgame perfection; in other words, there is no conflict between current and future 'selves' of any one player.

We shall say that a Strotz-Pollak equilibrium f obeys the *punishment property* if, for any player $i \in N$, history $h \in H$, and action $a_i \in A_i$, we have

$$c_i(f^{|h} \mid \delta) \geq c_i((a_i, f_{-i}(h)), f^{|\{h,(a_i,f_{-i}(h))\}} \mid \delta).$$

One normally thinks of deviations as securing immediate gain for later punishment. In fact, the opposite might sustain incentive compatibility — trading future gain against current losses. When this is not true, the punishment property obtains.

This property is violated in Example 1. Player 2's continuation payoff in stage 1 is $\frac{5}{7}$ if he plays R and 1 if he deviates and plays C. It is easy to check that if we replace $\kappa(C, C) = (1, 1)$ by $(0, 0)$ in stage 1, then the resulting Strotz-Pollak equilibrium obeys the punishment property and is sincere subgame perfect. Indeed,

Theorem 2 *If $f \in \Sigma^{S\&P}(\beta, \delta)$ has the punishment property, it is a sincere SPE.*

The proof is in the appendix, but the intuition is as follows. Everyone assigns more relative weight to current than continuation payoffs — i.e., $(1 - \Delta(\beta, \delta))$ vs. $(1 - \delta)$ — than he did at any earlier stage. Thus, player i's incentive constraints are met if δ replaces $\Delta(\beta, \delta)$, given the punishment property. This eliminates one-shot deviations planned at a 'future stage', as in Example 1. By induction, no one can benefit by deviating in any finite number of stages, or infinitely often, by standard reasoning. Thus, if there is never any temptation to defect this period, as in a Strotz-Pollak equilibrium, then no one ever wishes to *plan to* defect in the future.

Despite the couched message of Theorem 2, we now assert that any Strotz-Pollak equilibrium *outcome path* can be supported as a sincere subgame perfect one.

Corollary 1 *Every Strotz-Pollak equilibrium path is also a sincere SPE path, and therefore every Strotz-Pollak equilibrium payoff vector is a sincere SPE vector.*

The argument, in the appendix, is straightforward. By compactness of $U(\beta, \delta)$ and $C(\beta, \delta)$, there exists a worst Strotz-Pollak equilibrium payoff for each player. As in Abreu (1988), any Strotz-Pollak equilibrium path can, without loss of generality, be supported by threatening to switch to a player's worst such equilibrium if he deviates. Since an optimal penal code gives a player the lowest equilibrium continuation payoff if he deviates, it clearly satisfies the punishment property.

Quite unlike the single-agent literature, we can therefore always structure Strotz-Pollak equilibria so that, given the other players' strategies, no conflict arises between the different selves of a player about the optimality of his strategy.

An intuitive explanation of this important insight is as follows. In an infinitely repeated game, from the point of view of a player, the strategy profile of the other players defines an infinite-horizon decision problem. Different actions today may lead to different paths, and thus different payoffs, from tomorrow onward. Thus, given the other players' strategies, a player solves a *standard* single agent infinite-horizon decision problem but with β-δ preferences, which is subject to the usual time inconsistency issues: A player's optimal strategy *need not* coincide with the one that he would 'commit' to use in the future. This is *exactly* what Example 1 illustrates. The strategy of player 1 determines a Markov decision problem for player 2 with the following transition probabilities: an action different from D in stage 0 or different from C in stages 2 onwards induce a switch to a path with value 0, while an action different from D in stage 1 leads to a path with value 1. We saw that the optimal solution of player 2's decision problem exhibits time inconsistency.

Since we can always support a Strotz-Pollak equilibrium path by switching, in any period and for any deviation, to the worst possible equilibrium payoff for a deviant player, the aforementioned decision problem that each player solves can *always* be constructed as a stationary Markov decision problem that meets an additional 'consistency' property: In the SPE of the game that an agent plays with his future selves, there is *no conflict* between current and future selves. For instance, in Example 1 the time inconsistency problem of player 2 disappears if a deviation at any time from $D, D, C, C, ...$ induces a switch to the path with value 0.

It is important to emphasize that this is a *result* of our analysis, and it is due *precisely* to the repeated game structure we analyze. This cannot be assumed for a single agent problem, because in such a context, the 'states' and 'transition probabilities' of the infinite-horizon decision problem are *given* rather than *constructed* as in a repeated game setting. In other words, we link up the core simple strategy profile and optimal penal code characterization of Abreu (1988) with the reason why the adverse consequences of β-δ preferences are ameliorated in a strategic setting.

The next caveat is instructive. If we posited future-biased preferences, with $\beta > 1$ (for which our model is still well-defined), then Theorem 2 would fail. A reverse of the punishment property would be needed: Players who deviate should expect to be rewarded in the future, and punished immediately. But there is no general way of supporting subgame perfect equilibria failing the punishment property.

3.3 Public Randomization and Convex Payoff Sets

We now extend the results, adding a public randomization device to the stage game. This extension critically convexifies the sets of equilibrium and continuation payoffs.

A public randomization device P is a machine that, at the outset of each stage, randomly selects some $p \in [0,1]$ according to the uniform distribution and publicly informs the players of the realization. In such games, stage k's history $h^k \in H^k$ includes all past actions and public signals. Let $G_P^\infty(\beta,\delta)$ denote the game $G^\infty(\beta,\delta)$ extended by the public randomization device P. In $G_P^\infty(\beta,\delta)$, a *stage-k* strategy for player i is a Borel map f_i from $H^k \times [0,1]$ to A_i. Player i's expected average discounted payoff, conditional on any history $h \in H = \cup_{k=0}^\infty H^k$, then obeys:

$$u_i^P(f^{|h} \mid \beta,\delta) = \int_0^1 [(1-\Delta(\beta,\delta))\pi_i(f(h,p)) + \Delta(\beta,\delta)c_i^P(f^{|\{h,(f(h,p),p)\}} \mid \delta)]dp,$$

where

$$c_i^P(f^{|h} \mid \delta) = \int_0^1 [(1-\delta)\pi_i(f(h,p)) + \delta c_i^P(f^{|\{h,(f(h,p),p)\}} \mid \delta)]dp.$$

Strotz-Pollak equilibrium and sincere SPE easily extend to $G_P^\infty(\beta,\delta)$. Let $U_P(\beta,\delta)$ and $C_P(\beta,\delta)$ be the sets of Strotz-Pollak equilibrium and continuation payoffs.

Fix $\Delta \in (0,1)$ and $W \subset \mathbb{R}^n$. Let $a(p) : [0,1] \to A$ be Borel measurable and $\kappa(a,p) : A \times [0,1] \to W$ Borel measurable w.r.t. p for any $a \in A$. The pair (a,κ) is called (Δ,P)-*admissible* w.r.t. W if, for all $a_i' \in A_i$, $p \in [0,1]$, and $i \in N$,

$$(1-\Delta)\pi_i(a_i(p)) + \Delta\kappa(a_i(p),p) \geq (1-\Delta)\pi_i(a_i', a_{-i}(p)) + \Delta\kappa((a_i', a_{-i}(p)),p).$$

For any compact set $W \subset \mathbb{R}^n$, the set $B_P(W \mid \beta,\delta)$ is defined by[2]

$$B_P(W \mid \beta,\delta) = \{v \in \mathbb{R}^n : v = \int_0^1 [(1-\delta)\pi(a(p)) + \delta\kappa(a(p),p)]dp$$
$$\text{and } (a,\kappa) \text{ is } (\Delta(\beta,\delta),P)\text{-admissible w.r.t. } W\}.$$

A set $W \subset V$ is now *self-generating* if $W \subset B_P(W \mid \beta,\delta)$. We next extend the self-generation and factorization properties of Theorem 1:

Theorem 3 *Fix $\beta,\delta \in (0,1)$.*

[2]Since the set A is finite, $\kappa : A \times [0,1] \to W$ is a Caratheodory function; therefore, $\kappa(a(p),p) : [0,1] \to W$ is Borel measurable for any Borel measurable $a(p) : [0,1] \to A$ (see Aubin and Frankowska, 1990, for example). Thus, the integral is well-defined.

(*a*) *If* $W \subset \mathbb{R}^n$ *is self-generating, then* $B_P(W \mid \beta, \delta) \subset C_P(\beta, \delta)$.
(*b*) Factorization*:* $C_P(\beta, \delta)$ *is the largest fixed point of the operator* $B_P(\cdot \mid \beta, \delta)$ *:*

$$C_P(\beta, \delta) = \{v \in \mathbb{R}^n : v = \int_0^1 [(1-\delta)\pi(a(p)) + \delta\kappa(a(p), p)]dp \\ \text{and } (a, \kappa) \text{ is } (\Delta(\beta, \delta), P)\text{-admissible w.r.t. } C_P(\beta, \delta)\}.$$

Furthermore, $U_P(\beta, \delta) \subset \mathbb{R}^n$ *satisfies:*

$$U_P(\beta, \delta) = \{v \in \mathbb{R}^n : v = \int_0^1 [(1-\Delta(\beta, \delta))\pi(a(p)) + \Delta(\beta, \delta)\kappa(a(p), p)]dp \\ \text{and } (a, \kappa) \text{ is } (\Delta(\beta, \delta), P)\text{-admissible w.r.t. } C_P(\beta, \delta)\}.$$

(*c*) *The sets* $C_P(\beta, \delta)$ *and* $U_P(\beta, \delta)$ *are convex and compact.*

To see the convexity assertion, note that for any compact set $W \subset \mathbb{R}^n$, any $\beta \in (0, 1]$, and $\delta \in (0, 1)$,

$$B_P(W \mid \beta, \delta) = co\, B(W \mid \beta, \delta).$$

Since the convex hull of a compact set is compact, the proof of the compactness of $C_P(\beta, \delta)$ is similar to that of Theorem 4 in APS.

The logic underlying Corollary 1 extends here as well:

Corollary 2 *For any* $\beta, \delta \in (0, 1)$*, every Strotz-Pollak equilibrium payoff vector of* $G_P^\infty(\beta, \delta)$ *can be attained in a sincere SPE.*

When $\beta = 1$, we have $C_P(1, \delta) = U_P(1, \delta)$. A crucial difference, and the source of much novelty here, is that $C_P(\beta, \delta)$ is often a proper subset of $U_P(\beta, \delta)$ if $\beta < 1$.

Theorem 4 *For any* $\beta, \delta \in (0, 1)$, $C_P(\beta, \delta) \subset U_P(\beta, \delta)$.

Example 2 (found later) illustrates strict inclusion.

For a puzzling counter-intuition, observe that equilibrium payoffs in $U_P(\beta, \delta)$ place a *lower weight* on continuation payoffs than do payoffs in the set $C_P(\beta, \delta)$, because $\Delta(\beta, \delta) < \delta$. As a result, the APS monotonicity proof method fails, since it fixes the strategy played in the current stage game, and adjusts continuation payoffs. Any direct proof must be far more subtle, and we have found none. Our indirect proof owes to the following central lemma, with an argument by separation.

Lemma 1 *Let $\beta_1 < \beta_2$ and $\delta_1 > \delta_2$ satisfy $\Delta(\beta_1, \delta_1) = \Delta(\beta_2, \delta_2)$. Assume that $W \subset B_P(W \mid \beta_1, \delta_1) \subset \mathbb{R}^n$ is compact. Then $B_P(W \mid \beta_1, \delta_1) \subset B_P(W \mid \beta_2, \delta_2)$.*

Proof of Theorem 4: In Lemma 1, set $\beta_1 = \beta, \delta_1 = \delta, \delta_2 = \Delta(\beta, \delta) < \delta_1, \beta_2 = 1$, and choose $W = C_P(\beta, \delta)$. Observe that with this choice, $\Delta(\beta_1, \delta_1) = \Delta(\beta_2, \delta_2) = \Delta$, say. Since $C_P(\beta, \delta) = B_P(C_P(\beta, \delta)|\beta, \delta)$ and $U_P(\beta, \delta) = B_P(C_P(\beta, \delta)|1, \Delta)$, by Theorem 3 (b), Lemma 1 yields the desired inclusion $C_P(\beta, \delta) \subset U_P(\beta, \delta)$. $\square$

Proof of Lemma 1: By contradiction, assume there exists $x_1 \in B_P(W \mid \beta_1, \delta_1) \setminus B_P(W \mid \beta_2, \delta_2)$. Since both sets are convex, by the Separating Hyperplane Theorem, there exists $\gamma \in \mathbb{R}^n$ and $\varepsilon > 0$ with $\langle \gamma, x_1 \rangle \geq \langle \gamma, x_2 \rangle + \varepsilon$ for all $x_2 \in B_P(W \mid \beta_2, \delta_2)$.

Since $B_P(W \mid \beta_1, \delta_1)$ is compact, let $\bar{x}_1$ maximize $\langle \gamma, c \rangle$ over $c \in B_P(W|\beta_1, \delta_1)$. Put $\Delta = \Delta(\beta_1, \delta_1) = \Delta(\beta_2, \delta_2)$. Then for some pair (a, κ) (Δ, P)-admissible w.r.t. W,

$$\bar{x}_1 = \int_0^1 [(1 - \delta_1)\pi(a(p)) + \delta_1 \kappa(a(p), p)]dp \equiv (1 - \delta_1)\Pi + \delta_1 K,$$

thus defining Π and K. Setting $\bar{x}_2 = (1 - \delta_2)\Pi + \delta_2 K$, we have $\bar{x}_2 \in B_P(W \mid \beta_2, \delta_2)$. Therefore, using the definitions of $\bar{x}_1, x_1$, and $\bar{x}_2$, we have the chain of inequalities:

$$\langle \gamma, (1 - \delta_1)\Pi + \delta_1 K \rangle = \langle \gamma, \bar{x}_1 \rangle \geq \langle \gamma, x_1 \rangle \geq \langle \gamma, \bar{x}_2 \rangle + \varepsilon = \langle \gamma, (1 - \delta_2)\Pi + \delta_2 K \rangle + \varepsilon$$

Since $\delta_2 < \delta_1$, comparing the two weighted averages yields $\langle \gamma, K \rangle > \langle \gamma, \bar{x}_1 \rangle$. This contradicts the definition of $\bar{x}_1$, since $K = \int_0^1 \kappa(a(p), p)dp \in W \subset B_P(W|\beta_1, \delta_1)$. $\square$

4 The Cost of the Present-Time Bias

4.1 Monotonicity of the Equilibrium Payoff Set?

In the comparative static result in APS, the set of subgame perfect equilibrium payoffs expands in the discount factor δ. This result accords well with intuition, as it is natural to surmise that players should be able to support more equilibrium outcomes as they grow more patient and thereby weigh future payoffs more heavily. The recursive techniques of APS afford a simple yet elegant proof of this.

We now revisit this issue when $\beta < 1$. Unlike APS, more than one comparative static exercise is possible here. We explore how the set of Strotz-Pollak equilibrium payoffs changes when either or both of β or δ change. The analysis is more difficult than with geometric discounting, largely because of Theorem 4. As we show below, changes in either β or δ (or both) differentially affect each of the two sets $C_P(\beta, \delta)$ and $U_P(\beta, \delta)$, which wrecks havoc on the APS monotonicity intuition.

A. Compensated β-δ Changes. We first look at the relative importance of β and δ by studying games $G_P^\infty(\beta, \delta)$ with identical incentive constraints, holding $\Delta(\beta, \delta)$ fixed. The following result reveals an interesting trade-off between the two. In words, an increase in β accompanied by a Δ-compensating decrease in δ, yields more equilibrium payoffs. Just the same, an "inverse" monotonicity in δ obtains: An increase in δ holding $\Delta(\beta, \delta)$ fixed expands the range of equilibrium payoffs.

Proposition 1 *Let $\beta_1, \delta_1, \delta_2 \in (0, 1)$ and $\beta_2 \in (0, 1]$. Holding fixed the effective discount factor $\Delta(\beta_1, \delta_1) = \Delta(\beta_2, \delta_2)$, if $\beta_1 < \beta_2$, then $U_P(\beta_1, \delta_1) \subset U_P(\beta_2, \delta_2)$.*

Proof: Let $W \subset V$ be compact with $W \subset B_P(W \mid \beta_1, \delta_1)$. By Lemma 1, $W \subset B_P(W \mid \beta_1, \delta_1) \subset B_P(W \mid \beta_2, \delta_2)$. So if $W = C_P(\beta_1, \delta_1)$, then $C_P(\beta_1, \delta_1) \subset B_P(C_P(\beta_1, \delta_1) \mid \beta_2, \delta_2)$. Therefore, $C_P(\beta_1, \delta_1) \subset C_P(\beta_2, \delta_2)$ by Theorem 3 (b).

Finally, let $\Delta = \Delta(\beta_1, \delta_1) = \Delta(\beta_2, \delta_2)$. Anything (Δ, P) admissible w.r.t. $C_P(\beta_1, \delta_1)$ is (Δ, P) admissible w.r.t. $C_P(\beta_2, \delta_2)$. So $U_P(\beta_1, \delta_1) \subset U_P(\beta_2, \delta_2)$. $\square$

B. Uncompensated β-δ Changes. We now turn to comparative static analyses with respect to β or δ. The following result is a useful consequence of Proposition 1: Monotonicity in δ implies monotonicity in β.

Corollary 3 *If $U_P(\beta, \delta_1) \subset U_P(\beta, \delta_2)$ for any $\beta, \delta_1, \delta_2 \in (0, 1)$ with $\delta_1 < \delta_2$, then $U_P(\beta_1, \delta) \subset U_P(\beta_2, \delta)$ for any $\delta \in (0, 1)$ and $\beta_1, \beta_2 \in (0, 1]$, with $\beta_1 < \beta_2$.*

Proof: Let $\beta_1 < \beta_2 \leq 1$ and choose $\delta' < \delta$ with $\Delta(\beta_1, \delta) = \Delta(\beta_2, \delta')$. Then $U_P(\beta_1, \delta) \subset U_P(\beta_2, \delta')$ by Proposition 1, and $U_P(\beta_2, \delta') \subset U_P(\beta_2, \delta)$, since $\delta' < \delta$. $\square$

Following APS, one might expect that the set of equilibrium payoffs weakly expands in δ. Surprisingly, this well-known result from standard repeated games fails when $\beta < 1$. Indeed, the following example shows that $U_P(\beta, \delta)$ is not monotonic in β; by the contrapositive of Corollary 3, it is not monotonic in δ either.

Example 2: Non-Monotonicity in Either Discount Factor.

	L	M	R
U	$1, -1$	$-1, 0$	$-1, 0$
M	$1, -1$	$1, 1$	$1, -1$
D	$-1, 0$	$-1, 0$	$1, \frac{2}{11}$

It is clear that the minimax point of this game is $(1,0)$. Therefore, the line segment $[(1,0),(1,1)]$ in $\mathbb{R}^2$ contains both $U_P(\beta,\delta)$ and $C_P(\beta,\delta)$. Since (D,R) and (M,M) are Nash equilibria of the stage game, the payoff vectors $(1,\frac{2}{11})$ and $(1,1)$ belong to $C_P(\beta,\delta)\cap U_P(\beta,\delta)$, for all $\beta,\delta\in(0,1)$.

Let $\delta=\frac{7}{10}$ and $\beta_1=\frac{11}{21}$; hence, $\Delta(\beta_1,\delta)=\frac{11}{20}$.

We first show that, Player 2's worst Strotz-Pollak equilibrium continuation payoff vector is $(1,\frac{2}{11})$. Let $Y=[(1,0),(1,\frac{2}{11}))$. By contradiction, assume that $Y\cap C_P(\frac{11}{21},\frac{7}{10})\neq\varnothing$. For this to be possible, there must exist a continuation payoff vector $\kappa=(\kappa_1,\kappa_2)\in Y$ supported by an action profile with the stage-game payoff vector $(1,-1)$. But this can be dismissed on incentive compatibility grounds for player 2. For $\frac{3}{10}(-1)+\frac{7}{10}\kappa_2<\frac{2}{11}$ implies $\kappa_2<\frac{53}{77}$, while $\frac{9}{20}(-1)+\frac{11}{20}\kappa_2\geq 0$ implies $\kappa_2\geq\frac{63}{77}$. So, without loss of generality, we can assume that any deviation by player 2 is punished by the continuation payoff vector $(1,\frac{2}{11})$.

It is not difficult to see that the action profile (U,L) and the continuation value function κ with $\kappa(U,L)=(1,1)$ and $\kappa(a_1,a_2)=(1,\frac{2}{11})$ for all $(a_1,a_2)\neq(U,L)$ is $\Delta(\beta_1,\delta)$-admissible, with the corresponding equilibrium payoff vector $(1,\frac{1}{10})$. The admissibility of this pair is simple: if player 2 deviates to M or R, his expected payoff equals $\frac{1}{10}$. Therefore, $(1,\frac{1}{10})\in U_P(\frac{11}{21},\frac{7}{10})$. Also, it is easy to show that $\frac{1}{10}$ is the lowest equilibrium payoff for player 2, so that $U_P(\frac{11}{21},\frac{7}{10})=[(1,\frac{1}{10}),(1,1)]$.

For an illustration of Theorem 4, observe how

$$C_P\left(\frac{11}{21},\frac{7}{10}\right)=\left[\left(1,\frac{2}{11}\right),(1,1)\right]\subsetneq U_P\left(\frac{11}{21},\frac{7}{10}\right)=\left[\left(1,\frac{1}{10}\right),(1,1)\right]$$

Let $\beta_2=\frac{12}{21}$. The following argument shows that $(1,\frac{1}{10})\notin U_P(\beta_2,\delta_2)$.

- Observe that $\Delta(\beta_2,\delta)=\frac{4}{7}$ and that $C_P(\frac{11}{21},\frac{7}{10})=C_P(\frac{12}{21},\frac{7}{10})$. Thus, the worst continuation payoff for player 2 is the same as before.
- Any pair (a,κ) that delivers $\frac{1}{10}$ to player 2 uses the action profile (U,L) and thus is not $\Delta(\beta_2,\delta)$-admissible: player 2 earns $\frac{8}{77}$ by deviating from (U,L).
- The payoff vector $(1,\frac{8}{77})$ can be supported by the action profile (U,L) and continuation payoffs $\kappa(U,L)=(1,\frac{41}{44})$, $\kappa(a_1,a_2)=(1,\frac{2}{11})$ for all $a\neq(U,L)$.
- It is easy to see that $\frac{8}{77}$ is the lowest equilibrium payoff for player 2. Thus, $U_P(\frac{12}{21},\frac{7}{10})=[(1,\frac{8}{77}),(1,1)]$ and $(1,\frac{1}{10})\notin U_P(\beta_2,\delta_2)$.

Since $U_P(\frac{12}{21}, \frac{7}{10}) = [(1, \frac{8}{77}), (1, 1)]$ is a proper subset of $U_P(\frac{11}{21}, \frac{7}{10}) = [(1, \frac{1}{10}), (1, 1)]$, the equilibrium payoff set shrinks in β.[3] $\square$

The key to understanding the example is found in Theorem 4: Unlike with geometric discounting, the worst possible continuation payoff for a player need not coincide with the worst possible equilibrium payoff. Since (i) the increase in β does not affect the worst possible continuation payoff for player 2 in the example, and (ii) the worst possible equilibrium payoff for him must be supported in both cases by playing (U, L) in the current period (payoff -1), it easily follows that the worst possible equilibrium payoff must increase in this case.

Based on these considerations, a natural conjecture to explore is whether the monotonicity property holds when the worst possible equilibrium and continuation payoffs coincide. That $C_P(\beta, \delta)$ is actually monotone in β and in δ reinforces the plausibility of this conjecture — a result we show in the Appendix (see Claim 2).

An oft-studied class of games in which the worst possible equilibrium and continuation payoffs coincide are those where the minimax point is a pure strategy Nash equilibrium payoff of the stage game. The Prisoner's Dilemma is one such example. With this extra assumption, the equilibrium payoff set is indeed monotone.

Proposition 2 *Let the minimax point of G be a pure strategy Nash outcome of G.*
(a) *Let $\beta, \delta_1,\ \delta_2 \in (0, 1)$ and $\delta_1 < \delta_2$. Then $U_P(\beta, \delta_1) \subset U_P(\beta, \delta_2)$.*
(b) *Let $\beta_1, \beta_2,\ \delta \in (0, 1)$ and $\beta_1 < \beta_2$. Then $U_P(\beta_1, \delta) \subset U_P(\beta_2, \delta)$.*

The proof of part (a) is found in the Appendix. Corollary 3 then yields part (b).

Another natural conjecture is that, while separate monotonicity in β and δ fails, the set of equilibrium payoffs is monotonic in the effective discount factor $\Delta(\beta, \delta)$. As we shall see below in Example 3, no such result obtains.

4.2 Bounds on the Equilibrium Payoff Set

We now turn to the analysis of the following question: Given β and δ, can we bound $U_P(\beta, \delta)$ by sets of equilibrium payoffs of games with geometric discount factors?

This inquiry is important. First, it offers a different perspective on how severe is the failure of the monotonicity property, since geometric discounting is the natural benchmark. Does the payoff set collapse with $\beta < 1$? Second, it addresses how

[3] This example has been chosen for its simplicity. See the appendix for another example in which the same nonmonotone behavior occurs with efficient equilibrium payoffs.

persistent is the present bias cost at high discount factors: Can we deduce the folk-theorem for β-δ preferences from the standard one with geometric discounting?

A. Fixed $\boldsymbol{\beta, \delta < 1}$. Setting $\beta_2 = 1$ in Proposition 1 yields an easy upper bound.

Proposition 3 (Upper Bound) *For any* $\beta, \delta \in (0,1)$, $U_P(\beta, \delta) \subset U_P(1, \Delta(\beta, \delta))$.

How much does the set of equilibrium payoffs shrink with greater myopia (lower β)? Perhaps the monotonicity failure is not that strong. Observe that $\beta\delta$ is the largest discount factor that uniformly discounts payoffs by less than the β-δ decision maker:

$$(1, \beta\delta, (\beta\delta)^2, (\beta\delta)^3, \ldots) \leq (1, \beta\delta, \beta\delta^2, \beta\delta^3, \ldots)$$

Is $U_P(1, \beta\delta)$ then a lower bound on $U_P(\beta, \delta)$? We can show that the answer is yes in any symmetric prisoner's dilemma game (omitted example). But despite its intuitive appeal, the conjecture is generally wrong: Besides disproving the conjecture, the example also reveals that a slightly smaller β cannot always be offset by greater δ. This shows that the search for a lower bound set may prove a daunting task.

Example 3: Payoff Lower bound: $U_P(\beta, \delta)$ need not contain $U_P(1, \beta\delta)$. Let the stage game be:

	L	R
U	$1, -1$	$-1, 0$
D	$1, 1$	$0, 1$

The minimax point of this stage game is $(0, 0)$. Let us show that the set $W = \{(1, 0),\ (1, 1),\ (0, 0)\}$ is self-generating if $\beta = 1$ and $\delta = \frac{1}{2}$.

- The payoff vector $(1, 0)$ is supported by the action profile (U, L) and continuation values $(\kappa(U, L), \kappa(U, R), \kappa(D, L), \kappa(D, R)) = ((1, 1), (0, 0),\ (0, 0), (0, 0))$.
- The payoff vector $(0, 0)$ is supported by the action profile (U, R) and the continuation values $((0, 0), (1, 0), (0, 0), (0, 0))$.
- The payoff vector $(1, 1)$ is a Nash equilibrium outcome of the stage game.

Since W is self-generating, we have $W \subset U(1, \frac{1}{2})$, and, in particular, $(0, 0) \in U(1, \frac{1}{2})$.

We claim that $(0, 0) \notin U(\beta, \delta)$ for any $\beta, \delta \in (0, 1)$.[4] By contradiction, assume that $(0, 0) \in U(\beta, \delta)$ for some $\beta, \delta \in (0, 1)$. Because any continuation payoff vector

[4] Clearly, this claim is also valid if a public randomization device is added to the game.

lies in the set of feasible and individually rational payoffs, (U, R) is the only action profile that can be used to support the payoff vector $(0, 0)$, with the corresponding continuation payoff vector $\kappa(U, R)$ belonging to the half-open interval $((0, 0), (1, 0)]$. Denote $x^* = \max\{x \geq 0 : (x, 0) \in C(\beta, \delta)\}$, by compactness of $C(\beta, \delta)$. The continuation payoff vector $\kappa^* = (x^*, 0)$ can be supported only by a pair with the action profile (U, L). Let us show that the incentive constraints cannot hold for any such pair. Assume, by contradiction, that $(x^*, 0)$ can be supported by some pair $((U, L), \kappa')$. Then

$$0 = (1 - \delta)(-1) + \delta\kappa_2'(U, L),$$

and the corresponding incentive constraint for player 2 is

$$(1 - \Delta(\beta, \delta))(-1) + \Delta(\beta, \delta)\kappa_2'(U, L) \geq \Delta(\beta, \delta)\kappa_2'(U, R).$$

Since $\kappa_2'(U, R) \geq 0$, we have

$$(1 - \Delta(\beta, \delta))(-1) + \Delta(\beta, \delta)\kappa_2'(U, L) \geq 0,$$

which contradicts $\Delta(\beta, \delta) < \delta$. Thus, $(0, 0) \notin U(\beta, \delta)$ for any $\beta, \delta \in (0, 1)$. In particular, this holds for any β and δ with $\beta\delta = \frac{1}{2}$. Hence, $U_P(1, \frac{1}{2}) \not\subseteq U_P(\beta, \delta)$. $\square$

B. Large $\boldsymbol{\beta, \delta} \uparrow \mathbf{1}$. In the example, present-time bias so corrupted the equilibrium payoff set that no increase in the long-term discount factor δ can atone for the damage.[5] We now show more strongly that this precludes a lower bound that yields by corollary a folk-theorem for the β-δ case from the standard one. The argument is as follows. First, $U_P(1, \frac{1}{2}) \not\subseteq U_P(\beta, \delta)$ for all $\beta, \delta \in (0, 1)$. Since $U_P(1, \delta)$ is monotonic in δ and $(0, 0) \in U_P(1, \frac{1}{2})$, we have $U_P(1, \delta') \not\subseteq U_P(\beta, \delta)$ for all $\beta, \delta \in (0, 1)$ and all $\delta' \in (\frac{1}{2}, 1)$. Thus, there does not exist a function $\psi(\beta, \delta)$ such that $\lim_{\beta,\delta \to 1} \psi(\beta, \delta) = 1$ and $U_P(1, \psi(\beta, \delta)) \subset U_P(\beta, \delta)$ for all $\beta, \delta \in (0, 1)$. Hence, that $U_P(1, \psi) \uparrow V$ as $\psi \uparrow 1$ cannot be used to deduce $U_P(\beta, \delta) \uparrow V$ as $\beta, \delta \to 1$.

5 CONCLUSION

Time inconsistency is generally thought to be an important phenomenon, and worthy of study. The standard context for this has been a variety of decision theory

[5]The reader may wonder whether $U_P(1, \beta\delta) \subset U_P(\beta, \delta)$ when the minimax point of the stage game is a pure strategy Nash payoff vector. We omit an example that disproves this conjecture.

exercises with a state variable, such as arises in an addiction or consumption context. We have shifted this discussion to a wholly new framework, asking whether there are any economically important and novel implications for repeated games. We have *developed* a usable framework for exploring such games with β-δ preferences, that critically exploits the Abreu (1988) optimal penal code characterization. One may wonder whether the repeated game can offer any interesting and surprising implications for beta-delta preferences, even though the corresponding decision theory problem does not. With our tools, we feel that the answer is undeniably 'yes': We prove that *the most basic comparative static of repeated games literature — monotonicity in the discount factor — disappears.*[6] In its place, we deduce a compensated comparative static that only admits an indirect proof. We have also provided insights into the delicate nature of the β-δ folk theorem.

A Appendix

A.1 Non-Monotonic Efficient Equilibrium Payoffs

We now give an example showing that efficient equilibrium payoffs can be non-monotonic in either discount factor.

	L	M	R
U	$2\frac{2}{11}, -1$	$-1, 0$	$-1, 0$
M	$-1, 0$	$-1, 0$	$\frac{2}{11}, 1$
D	$-1, 0$	$1, \frac{2}{11}$	$-1, 0$

Notice that (M, R) and (D, M) are Nash equilibria of the stage game, and that $(\frac{2}{11}, 0)$ is the minimax point.

- Let $\delta = \frac{7}{10}$, $\beta_1 = \frac{11}{21}$, so that $\Delta(\beta_1, \delta) = \frac{11}{20}$.
- The payoff vector $(1\frac{2}{11} - \frac{1}{10}, \frac{1}{10})$ belongs to the Pareto frontier of the set of feasible and individually rational payoffs. Indeed, each payoff is larger than the minimax, and $(1\frac{2}{11} - \frac{1}{10}, \frac{1}{10}) = \frac{9}{20}(2\frac{2}{11}, -1) + \frac{11}{20}(1, \frac{2}{11})$.
- The action profile (U, L) is $\Delta(\beta_1, \delta)$-admissible, and yields equilibrium payoff $(1\frac{2}{11} - \frac{1}{10}, \frac{1}{10})$ and continuation payoffs $\kappa(U, L) = (\frac{2}{11}, 1)$, and $\kappa(a_1, a_2) = (1, \frac{2}{11})$ for any other action profile (a_1, a_2).

[6] This result only fails in contexts without public randomization.

- Let $\delta = \frac{7}{10}$, $\beta_2 = \frac{12}{21}$; then $\Delta(\beta_2, \delta) = \frac{4}{7}$.
- Any pair (a, κ) that yields $\frac{1}{10}$ to player 2 must use the action profile (U, L). But (U, L) is not $\Delta(\beta_2, \delta)$-admissible, for $\frac{1}{10} < \frac{3}{7}0 + \frac{4}{7}\frac{2}{11} = \frac{8}{77}$ (i.e., player 2 can gain by deviating from L).
- Thus, $\frac{1}{10}$ cannot be an equilibrium payoff for player 2 when $\beta_2 = \frac{12}{21}$, and hence $(1\frac{2}{11} - \frac{1}{10}, \frac{1}{10}) \in U_P(\beta_1, \delta)$ but $(1\frac{2}{11} - \frac{1}{10}, \frac{1}{10}) \notin U_P(\beta_2, \delta)$.

A.2 Punishment Property / SPE: Proof of Theorem 2

Assume, by contradiction, there exists a history $h^k \in H$ such that $f^{|h^k}$ is not a sincere Nash equilibrium of $G^\infty(\beta, \delta)$. Without loss of generality, $h^k = h^0$. Then there exist player i and strategy $g_i : H \to A_i$ such that $u_i(g_i, f_{-i} \mid \beta, \delta)) > u_i(f \mid \beta, \delta)$. For any $K \in \{0, 1, \ldots\}$, define the following function $g_i^K : H \to A_i$,

$$g_i^K(h) = \begin{cases} g_i(h) & \text{for all } h \in \bigcup\limits_{j=0}^{K} H^j, \\ f_i(h) & \text{for all } h \in \bigcup\limits_{j=K+1}^{\infty} H^j \end{cases}$$

Since the players' future payoffs are discounted, there exists $K \in \{1, \ldots\}$ such that $u_i(g_i^K, f_{-i} \mid \beta, \delta) > u_i(f \mid \beta, \delta)$. Let $h_g = \{a^0, a^1, \ldots, a^K, \ldots\}$ denote the path of action profiles generated by the strategy profile (g_i^K, f_{-i}), and $h_g^k = \{a^0, \ldots, a^{k-1}\}$, $k \in \{1, \ldots\}$. Player i's stage-K average discounted payoff $u_i((g_i^K, f_{-i})^{|h_g^K} \mid \beta, \delta)$ can be represented as follows:

$$(1 - \Delta(\beta, \delta))\pi_i(g_i^K(h_g^K), f_{-i}(h_g^K))) + \Delta(\beta, \delta)c_i(f^{|h_g^{K+1}} \mid \delta)$$

Since the strategy profile f is a Strotz-Pollak equilibrium of $G(\beta, \delta)$, the strategy profile $f(h_g^K)$ is a Nash equilibrium of the following reduced normal form game $(N, (A_i)_{i\in N}, ((1 - \Delta(\beta, \delta))\pi_i(\cdot) + \Delta(\beta, \delta)c_i(f^{|\{h_g^K, \cdot\}} \mid \delta))_{i\in N})$. Therefore,

$$\begin{aligned}(1 - \Delta(\beta, \delta))\pi_i(f(h_g^K)) &+ \Delta(\beta, \delta)c_i(f^{|\{h_g^K, f(h_g^K)\}} \mid \delta) \\ &\geq (1 - \Delta(\beta, \delta))\pi_i(g_i^K(h_g^K), f_{-i}(h_g^K)) + \Delta(\beta, \delta)c_i(f^{|h_g^{K+1}} \mid \delta).\end{aligned}$$

Assume that $g_i^K(h_g^K) \neq f_i(h_g^K)$. Because $\Delta(\beta,\delta) < \delta$, and since $c_i(f^{|\{h_g^K, f(h_g^K)\}} \mid \delta) \geq c_i(f^{|h_g^{K+1}} \mid \delta)$ by the punishment property, we conclude that

$$(1-\delta)\pi_i(f(h_g^K))+\delta c_i(f^{|\{h_g^K, f(h_g^K)\}} \mid \delta) \geq (1-\delta)\pi_i(g_i^K(h_g^K), f_{-i}(h_g^K))+\delta c_i(f^{|h_g^{K+1}} \mid \delta).$$

Therefore, player i improves his stage-0 expected payoff by employing the strategy $f_i(h_g^K)$ instead of $g_i^K(h_g^K)$ at stage K. That is,

$$u_i(g_i^{K-1}, f_{-i} \mid \beta, \delta) \geq u_i(g_i^K, f_{-i} \mid \beta, \delta).$$

Proceeding similarly, we get the following chain of inequalities $u_i(f \mid \beta,\delta) \geq u_i(g_i^0, f_{-i} \mid \beta,\delta) \geq \ldots \geq u_i(g_i^K, f_{-i} \mid \beta,\delta)$. This is a contradiction. $\square$

A.3 Strotz-Pollak and Sincere SPE: Proof of Corollary 1

Preamble. Let $Q(f)$ denote the infinite sequence of action profiles (the *path*) that results from conformity with the strategy profile $f = (f_1, \ldots, f_n)$ in the absence of deviations, and let $\Omega = A^\infty$ denote the set of paths. For each $Q \in \Omega$, let $Q = \{a^k(Q)\}_{k=0}^\infty$, where $\{a^k(Q)\}_{k=0}^\infty$ is the corresponding sequence of action profiles.

Let $Q_i \in \Omega$, $i = 0, 1, \ldots, n$. Following Abreu (1988), let $\sigma(Q_0, Q_1, \ldots, Q_n)$ denote the corresponding *simple strategy profile*. Here Q_0 is the initial path and Q_i is the punishment path for any deviation by player i after any history. Define continuation payoffs from path Q starting at stage k to be:

$$\varkappa_i(Q,k) = (1-\delta)\sum_{s=0}^{\infty}\delta^s\pi_i(a^{k+s}(Q)).$$

Claim 1 *A simple strategy profile $\sigma(Q_0, Q_1, \ldots, Q_n)$ is a Strotz-Pollak equilibrium if and only if*

$$\begin{aligned}(1-\Delta(\beta,\delta))\pi_i(a^k(Q_j)) + \Delta(\beta,\delta)\varkappa_i(Q_j, k+1)) \\ \geq (1-\Delta(\beta,\delta))\pi_i(a_i', a_{-i}^k(Q_j)) + \Delta(\beta,\delta)\varkappa_i(Q_i, k+1))\end{aligned}$$

for all $a_i' \in A_i$, all $i \in N$, all $j \in \{0\} \cup N$, and all $k = 0, 1, \ldots$.

This claim follows directly from the definition of a Strotz-Pollak equilibrium.

An *optimal penal code* for $G^\infty(\beta,\delta)$ is an n-vector of pure strategy profiles $(\underline{f_1}, \underline{f_2}, \ldots, \underline{f_n})$, where $\underline{f_i} \in \Sigma^{S\&P}(\beta,\delta)$ delivers the worst possible punishment for player i:

$$c_i(\underline{f_i} \mid \delta) = \underline{v_i}(\beta, \delta) = \min_{v \in C(\beta,\delta)} v_i.$$

Existence of an optimal penal code in $G^\infty(\beta, \delta)$ owes to compactness of $C(\beta, \delta)$.

Using Abreu's notation, Corollary 1 can be reformulated: If $f \in \Sigma^{S\&P}(\beta, \delta)$ and $(\underline{f_1}, \ldots, \underline{f_n})$ is an optimal penal code, then $\sigma = \sigma(Q(f), Q(\underline{f_1}), \ldots, Q(\underline{f_n}))$ is a sincere subgame perfect equilibrium.

THE PROOF. Obviously, the strategy profile $\sigma = \sigma(Q(f), Q(\underline{f_1}), \ldots, Q(\underline{f_n}))$ has the punishment property. It is a Strotz-Pollak equilibrium by Claim 1 and, therefore, is a sincere subgame perfect equilibrium by Theorem 2. $\square$

A.4 Proof of the Uncompensated Comparative Statics

Continuation Payoffs. The proof makes use of the following result:

Claim 2 *(a) If $\beta, \delta_1, \delta_2 \in (0,1)$ and $\delta_1 < \delta_2$, then $C_P(\beta, \delta_1) \subset C_P(\beta, \delta_2)$.*
(b) If $\beta_1, \delta \in (0,1)$, $\beta_2 \in (0,1]$, and $\beta_1 < \beta_2$, then $C_P(\beta_1, \delta) \subset C_P(\beta_2, \delta)$.

Proof of Claim: Let $W \subset \mathbb{R}^n$ be compact with $W \subset B_P(W \mid \beta, \delta_1)$. We first show that $W \subset B_P(B_P(W \mid \beta, \delta_1) \mid \beta, \delta_2)$. Let $x \in W$ be arbitrary. Since $x \in B_P(W \mid \beta, \delta_1)$, there exists a pair (a, κ^0) $(\Delta(\beta, \delta_1), P)$-admissible w.r.t. W such that

$$x = \int_0^1 [(1-\delta_1)\pi(a(p)) + \delta_1 \kappa^0(a(p), p)]dp.$$

For all $p \in [0,1]$, we have

$$w(p) = (1-\delta_1)\pi(a(p)) + \delta_1 \kappa^0(a(p), p) \in B(W \mid \beta, \delta_1).$$

Let us show that $w(p) \in B(B(W \mid \beta, \delta_1) \mid \beta, \delta_2)$ for all $p \in [0,1]$. Fix $\lambda = \frac{\delta_1(1-\delta_2)}{\delta_2(1-\delta_1)}$ and $p \in [0,1]$. Define the continuation value function $\kappa^1 : A \times [0,1] \to B(W \mid \beta, \delta_1)$

$$\kappa^1(a, p) = (1-\lambda)w(p) + \lambda \kappa^0(a, p).$$

It is clear that $w(p) = (1-\delta_2)\pi(a(p)) + \delta_2 \kappa^1(a(p), p)$. We need to verify that

$$(1-\delta_2)\pi(a(p)) + \beta\delta_2\kappa^1(a(p), p) \geq (1-\delta_2)\pi(a_i, a_{-i}(p)) + \beta\delta_2\kappa^1((a_i, a_{-i}(p)), p)$$

for all $a_i \in A_i$ and all $i \in N$. It is not difficult to see that the incentive constraints hold since, for all $a \in A$ and all $p \in [0,1]$,

$$(1-\delta_2)\pi(a)+\beta\delta_2\kappa^1(a,p) = \frac{(1-\delta_2)}{(1-\delta_1)}[(1-\delta_1)\pi(a)+\beta\delta_1\kappa^0(a,p)]+\frac{\beta(\delta_2-\delta_1)}{1-\delta_1}w(p).$$

Since $B(B(W \mid \beta,\delta_1) \mid \beta,\delta_2) \subset B_P(B_P(W \mid \beta,\delta_1) \mid \beta,\delta_2)$, we have $w(p) \in B_P(B_P(W \mid \beta,\delta_1) \mid \beta,\delta_2)$ and, thus, $x = \int_0^1 w(p)dp \in B_P(B_P(W \mid \beta,\delta_1) \mid \beta,\delta_2)$.

In the particular case when $W = C_P(\beta,\delta_1)$, any point of $C_P(\beta,\delta_1)$ belongs to $B_P(C_P(\beta,\delta_1) \mid \beta,\delta_2)$ since $B_P(C_P(\beta,\delta_1) \mid \beta,\delta_1) = C_P(\beta,\delta_1)$. Invoking the maximal property of $C_P(\beta,\delta_2)$, we conclude that $C_P(\beta,\delta_1) \subset C_P(\beta,\delta_2)$.

The second part of the statement follows from Corollary 3. □

Proof of Proposition 2. It is obvious that $\Delta(\beta,\delta_1) < \Delta(\beta,\delta_2)$. If $x \in U_P(\beta,\delta_1)$, then there exists a pair (a,κ^0) $(\Delta(\beta,\delta_1),P)$-admissible w.r.t. $C_P(\beta_1,\delta)$ such that

$$x = \int_0^1 [(1-\Delta(\beta,\delta_1))\pi(a(p)) + \Delta(\beta,\delta_1)\kappa^0(a(p),p)]dp.$$

WLOG, $\kappa_i^0((a_i,a_{-i}(p)),p) = \min_{w\in C_P(\beta,\delta_1)} \langle e_i, w\rangle = v_i^*$ for all $a_i \in A_i$, all $p \in [0,1]$, and all $i \in N$, where e_i is a vector with 1 in its i-th component and 0 otherwise. Note that even though, for each $p \in [0,1]$,

$$w(p) = (1-\Delta(\beta,\delta_1))\pi(a(p)) + \Delta(\beta,\delta_1)\kappa^0(a(p),p)$$

belongs to $U_P(\beta,\delta_1)$, $w(p)$ need not belong to $C_P(\beta,\delta_1)$, which is the case if $\beta = 1$.

We must show that $w(p) \in U_P(\beta,\delta_2)$ for all $p \in [0,1]$. Fix the constant $\lambda = \Delta(\beta,\delta_1)(1-\Delta(\beta,\delta_2)/(\Delta(\beta,\delta_2)(1-\Delta(\beta,\delta_1))) = \frac{\delta_1(1-\delta_2)}{\delta_2(1-\delta_1)}$ and $p \in [0,1]$. Define

$$\kappa^1(a,p) = \begin{cases} v^* \text{ if } a = (a_i,a_{-i}(p)) \text{ and } a_i \in A_i\backslash\{a_i(p)\} \text{ for some } i \in N, \\ (1-\lambda)w(p) + \lambda\kappa^0(a,p) \text{ otherwise.} \end{cases}$$

It is not difficult to see that

$$w(p) = (1-\Delta(\beta,\delta_2))\pi(a(p)) + \Delta(\beta,\delta_2)\kappa^1(a(p),p).$$

First, let us verify that

$$w_i(p) \geq (1-\Delta(\beta,\delta_2))\pi_i(a_i,a_{-i}(p)) + \Delta(\beta,\delta_2)v_i^*$$

for all $a_i \in A_i$and all $i \in N$. By contradiction, assume that there exist $i \in N$ and $a_i' \in A_i \backslash \{a_i(p)\}$ such that $w_i(p) < (1-\Delta(\beta,\delta_2))\pi_i(a_i',a_{-i}(p)) + \Delta(\beta,\delta_2)v_i^*$. Since $w_i(p) \geq (1-\Delta(\beta,\delta_1))\pi_i(a_i,a_{-i}(p)) + \Delta(\beta,\delta_1)v_i^*$ for all $a_i \in A_i$, we have

$$\kappa_i(a(p),p) > v_i^* > \pi_i(a_i',a_{-i}(p)) > \pi_i(a(p)).$$

Then $\pi_i(a_i,a_{-i}(p)) < v_i^*$ for all $a_i \in A_i$, versus the definition of a minimax point.

To finish the proof, we only need to show that $\kappa^1(a(p),p) \in C_P(\beta,\delta_2)$. Since

$$L = [(1-\delta_2)\pi(a(p)) + \delta_2\kappa^0(a(p),p), \kappa^0(a(p),p))] \subset C_P(\beta,\delta_1),$$

$L \subset C_P(\beta,\delta_2)$ by the claim.

Let $\kappa^1(a(p),p) \notin L$. It is then obvious that $\kappa^2(a(p),p)$ such that $\kappa^1(a(p),p) = (1-\delta_2)\pi(a(p)) + \delta_2\kappa^2(a(p),p))$ belongs to $[\kappa^1(a(p),p), \kappa^0(a(p),p))]$. Let us show that

$$(1-\Delta(\beta,\delta_2))\pi(a(p)) + \Delta(\beta,\delta_2)\kappa^2(a(p),p)) \geq (1-\Delta(\beta,\delta_2))\pi(a_i,a_{-i}(p)) + \Delta(\beta,\delta_2)v_i^*$$

for all $a_i \in A_i$and all $i \in N$. This is the case since

$$(1-\Delta(\beta,\delta_2))\pi(a(p)) + \Delta(\beta,\delta_2)\kappa^2(a(p),p)) = \frac{\delta_2-\Delta(\beta,\delta_2)}{\delta_2}\pi(a(p)) + \frac{\Delta(\beta,\delta_2)}{\delta_2}\kappa^1(a(p),p))$$

and $\Delta(\beta,\delta_2)/\delta_2 > \Delta(\beta,\delta_2)$. If $\kappa^2(a(p),p) \in L$, the result follows. Otherwise, consider $\kappa^3(a(p),p)$ with $\kappa^2(a(p),p) = (1-\delta_2)\pi(a(p)) + \delta_2\kappa^3(a(p),p))$. Proceeding similarly, it is not hard to see that there exists $l \in \{2,\ldots\}$ such that $\kappa^l(a(p),p)) \in L$, which completes the proof, since $\kappa^1(a(p),p)) \in C_P(\beta,\delta_2)$ by self-generation. □

References

Abreu, D. (1988) "Towards a Theory of Discounted Repeated Games," *Econometrica*, 56, 383-396.

Abreu, D., D. Pearce, and E. Stacchetti (1986) "Optimal Cartel Equilibria with Imperfect Monitoring," *Journal of Economic Theory* 39, 251-269.

Abreu, D., D. Pearce, and E. Stacchetti (1990) "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* 58, 1041-1063.

Aubin, J.-P. and H. Frankowska (1990) Set-Valued Analysis, Birkhauser Boston.

Della Vigna, S. and D Paserman (2003), "Job-Search and Impatience," *Jornal of Labor Economics*, forthcoming.

Goldman, S.M. (1980) "Consistent Plans," *Review of Economic Studies*, 47, 533-7.

Gruber, J. and B. Koszegi (2001), "Is Addiction Rational? Theory and Evidence," *Quarterly Journal of Economics* 116, 1261-1303.

Kocherlakota, N.R. (1996) "Reconsideration-Proofness: A Refinement for Infinite Horizon Time Inconsistency," *Games and Econ. Behavior*, 15, 33-54.

Krusell, P., B. Kuruscu, and A. Smith Jr. (2002) "Equilibrium Welfare and Government Policy with Quasi-Geometric Discounting," *Journal of Economic Theory* 105, 42-72.

Laibson, D.I. (1997) "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112, 443-477.

O'Donoghue, T. and M. Rabin (1999), "Doing it Now or Later," *American Economic Review* 89, 103-124.

Peleg, B. and Yaari, M.E. (1973) "On the Existence of a Consistent Course of Action when Tastes are Changing," *Review of Economic Studies*, 40, 391-401.

Phelps, E.S. and Pollak, R.A. (1968) "On Second-Best National Saving and Game-Equilibrium Growth," *Review of Economics Studies*, 35, 185-199.

Pollak, R.A. (1968), "Consistent Planning," *Review of Economic Studies*, 35, 201-208.

Strotz, R.H. (1955-56), "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 23, 165-180.