Models of Limited Self-Control: Comparison and Implications for Bargaining*

Shih En Lu[†]

May 2016

Abstract

This paper compares two models of limited intertemporal self-control: the linear-cost version of Fudenberg and Levine's dual-self model (2006) and the quasi-hyperbolic discounting model. The main distinction between the two frameworks can be formulated as whether agents care about future self-control costs: dual selves do, while quasi-hyperbolic discounters do not. The dual-self model is applied to a bargaining game with alternating proposals where players negotiate over an infinite stream of payoffs, and it is shown that, in subgame-perfect equilibrium, the first proposer's payoff is unique and agreement is immediate. By contrast, Lu (2016) shows that with quasi-hyperbolic discounters, a multiplicity of payoffs and delay can arise in equilibrium.

Keywords: Self-Control, Bargaining, Time Inconsistency, Dual Self, Quasi-Hyperbolic Discounting

JEL Codes: C78, D90

---

*DOI: dx.doi.org/10.1016/j.econlet.2016.06.012. Portions of this paper were previously part of "Self-Control and Bargaining."

[†]Department of Economics, Simon Fraser University, Burnaby, BC V5A 1S6, Canada. Email: shihenl@sfu.ca.

1 Introduction

To explain preference reversals that are likely caused by inconsistent preferences over time,[1] economists and psychologists have put forth the idea that immediate rewards are disproportionately more appealing than rewards in the near, but not immediate future. *Quasi-hyperbolic* discounting, where the sequence of discount factors is $1, \beta\delta, \beta\delta^2, \beta\delta^3, ...$ with $\beta, \delta \in (0, 1)$, is often used to capture this extra weight put on immediate payoffs.[2]

An alternative framework for studying limited self-control is Fudenberg and Levine's (2006) dual-self model.[3] It postulates that each agent is comprised of a sequence of short-run selves interacting with the world and a long-run self that may, at a cost, influence the short-run self.[4] Each short-run self lasts only one period and cares only about the immediate payoff, while the long-run self discounts the future with a standard exponential function.

This paper studies the relation between the dual-self model and quasi-hyperbolic discounting. Proposition 1 in Section 2 shows that sophisticated quasi-hyperbolic agents[5] can be understood as dual selves with self-control costs linear in the amount of immediate utility forgone, but modified such that the long-run self no longer cares about the costs of influencing future short-run selves, even though she is aware of them.[6] Therefore, the main distinction between the two frameworks is whether agents care about their future self-control costs.

Section 3 shows that this difference can have a large impact on equilibrium predictions in games. The example used is the alternating-offer bargaining game proposed by Lu (2016), where an *infinite stream* of unit-surpluses is divided, unlike in Ståhl (1972) and Rubinstein (1982). Each offer allocates the entire stream. The game ends when an offer is accepted;

---

[1] Frederick, Loewenstein and O'Donoghue (2002) provide an overview of some experimental findings.

[2] Phelps and Pollak (1968) first proposed this discount function to study intergenerational saving, and Laibson (1997) applied it to individual intertemporal decision-making. See, for example, Angeletos *et al.* (2001) and Laibson, Repetto and Tobacman (2007) for empirical support, and Gul and Pesendorfer (2005) and Montiel Olea and Strzalecki (2014) for axiomatic foundations.

[3] Many other dual-self models have been proposed, *e.g.* Thaler and Shefrin (1981), Bénabou and Pycia (2002), Loewenstein and O'Donoghue (2004), Bernheim and Rangel (2004), Benhabib and Bisin (2005) and Brocas and Carillo (2008). In this paper, Fudenberg and Levine's model is used due to its generality (notably, it applies to situations with an infinite horizon, unlike some of the models listed above) and its tractability. Also, Fudenberg and Levine show that, with linear self-control costs (as assumed in this paper), their model satisfies the axioms from Gul and Pesendorfer (2001).

McClure, Laibson, Loewenstein and Cohen (2004) show, through functional magnetic resonance imaging (fMRI), that there are two distinct brain systems governing discounting, which provides a motivation for dual-self models.

[4] More precisely, whenever it is an agent's turn to move, the long-run self acts first by deciding whether and how much to change the short-run self's preferences; the latter self then moves in the main game.

[5] Sophistication means that agents know (and are not mistaken about) their future selves' preferences.

[6] This "long-run self" would therefore be more accurately described as a sequence of forward-looking agents. For brevity, however, the term "long-run self" is used.

Xue (2008) shows that quasi-hyperbolic discounting can be obtained as the result of cooperative bargaining between a myopic self and a more patient time-consistent self.

when an offer is rejected, that period's surplus is lost. The fact that both current and future surpluses are divided corresponds to many economic situations (*e.g.* employment), and is important for teasing out the effects of limited self-control: when proposing, agents are tempted to demand more of the current surplus (*e.g.* in the form of a signing bonus) in exchange for future surplus.[7] Proposition 2 describes subgame-perfect equilibrium (SPNE) play between dual selves with equal discount factor δ. Here, agreement is always immediate, and the first proposer's payoff is unique and continuous in the self-control parameters; Lu (2016) shows that neither is true with quasi-hyperbolic discounters.

2 Relation between the Dual-Self Model and Quasi-Hyperbolic Discounting

Fudenberg and Levine (2006) propose a *dual-self* model where: (i) each person acts through a sequence short-run selves that each cares only about utility in the current period, and (ii) a forward-looking long-run self, before the short-run self plays in each period, can take actions affecting how the short-run self's choice determines current utility. They show that under mild assumptions,[8] their dual-self model has an equivalent reduced form where the long-run self directly takes actions to maximize aggregate utility at time t given by

$$U_t = \sum_{\tau=t}^{\infty} \delta^{\tau-t}(u_\tau - C_\tau),$$

where u_τ is the utility of the short-run self at time τ, and C_τ is the self-control cost incurred by the long-run self at time τ. This paper adopts its most tractable form, where the cost to the long-run self of making the short-run self take action a when the state variable is y, denoted $C_t(y, a)$, is linear in the difference in short-run utility, $u_t(y, .)$, caused by the change:

$$C_t(y, a) = \gamma[\sup_{a'} u_t(y, a') - u_t(y, a)],$$

where $\gamma > 0$.

[7] It can be shown that with dual-self agents whose discount factors are δ_i and whose self-control costs are linear with coefficients γ_i, the SPNE is unique and the same as with exponential agents whose discount factors are $\delta_i/(1 + \gamma_i)$. Kodritsch (2014) shows that, with quasi-hyperbolic agents, the same holds with "effective" exponential discount factors $\beta_i\delta_i$. Therefore, in SPNE, self-control problems, as defined in either the quasi-hyperbolic or the dual-self framework, cannot be separated from time-consistent impatience in complete-information Rubinstein-Ståhl bargaining.

[8] Namely, self-control is costly, the long-run self is able to make the short-run self take any action, utility is continuous in both selves' actions, and the long-run self can break ties faced by the short-run self at arbitrarily small cost.

To relate quasi-hyperbolic discounting and the dual-self model, define, as a technical device, the following modified version of the dual self:

Definition: A *selfish dual self* is a dual self whose long-run self's utility at time t is $U_t = -C_t + \sum_{\tau=t}^{\infty} \delta^{\tau-t} u_\tau$, where u_τ is the utility of the short-run self at time τ, and C_t is the self-control cost incurred by the long-run self at time t.

The difference between the regular dual self and the selfish dual self is that the latter does not care about self-control costs $C_{t+1}, C_{t+2}, ...$ incurred by future versions of himself, and therefore only cares about the presence of future temptation if it affects future choices. The preferences of the long-run self are therefore time-inconsistent themselves. Proposition 1 shows that the utility of the selfish dual self directly relates to that of the quasi-hyperbolic discounter.

Proposition 1: Suppose an agent chooses from a set of streams $\mathbf{u}^k$ of expected utility, where u_t^k denotes the expected utility from stream k at time t. Let the valuation of $\mathbf{u}^k$ by a quasi-hyperbolic agent with discount function $1, \beta\delta, \beta\delta^2, ...$ be u_{QH}^k, and let the valuation of $\mathbf{u}^k$ by a selfish dual self with linear self-control cost coefficient γ be u_{DS}^k. Then, if $\beta = \frac{1}{1+\gamma}$, $u_{DS}^k = -\gamma \sup_{k'}\{u_0^{k'}\} + (1+\gamma)u_{QH}^k$ for all k.

Proof: Denote the current period as period 0. We have $u_{QH}^k = u_0^k + \beta \sum_{t=1}^{\infty} \delta^t u_t^k = u_0^k + \frac{1}{1+\gamma} \sum_{t=1}^{\infty} \delta^t u_t^k$.

The dual self's self-control cost of choosing stream k is $C_0^k = \gamma(\sup_{k'}\{u_0^{k'}\} - u_0^k)$. It follows that

$$\begin{aligned} u_{DS}^k &= -\gamma(\sup_{k'}\{u_0^{k'}\} - u_0^k) + \sum_{t=0}^{\infty} \delta^t u_t^k \\ &= -\gamma \sup_{k'}\{u_0^{k'}\} + (1+\gamma)u_0^k + \sum_{t=1}^{\infty} \delta^t u_t^k \\ &= -\gamma \sup_{k'}\{u_0^{k'}\} + (1+\gamma)u_{QH}^k. \blacksquare \end{aligned}$$

Since u_{DS}^k is an affine transformation of u_{QH}^k, Proposition 1 states that, under the parametrization $\beta = \frac{1}{1+\gamma}$, the quasi-hyperbolic discounter and the selfish dual self have the same preferences. By ignoring future self-control costs, the selfish long-run self, just like the quasi-hyperbolic agent, treats two future periods in a time-consistent way, but treats today and tomorrow differently than two future periods.[9] Example 1 in the Appendix illustrates the result.

[9] I thank an anonymous referee for suggesting this remark.

3 Bargaining between Dual Selves

3.1 The Game

Two players with transferable utility bargain in discrete time over an infinite stream of unit surpluses. The game starts in period 0, and in each even (odd) period t, player 1 (2) proposes an allocation $(\mathbf{x}, \mathbf{1}-\mathbf{x})$, where $\mathbf{x} = (x^t, x^{t+1}, ...) \in [0,1]^\infty \equiv X$ is the stream of payoffs demanded by the proposer. The opponent then chooses between acceptance or rejection. In the former case, the proposal is enacted, and the game ends. In the latter case, the surplus from period t vanishes, and the game continues in period $t+1$. Thus, letting H^t denote the set of all possible histories at the start of period t, a pure strategy for player 1 is a pair of functions (f, g) where $f : \cup_{k=0}^\infty H^{2k} \to X$ and $g : \cup_{k=0}^\infty (H^{2k+1} \times X) \to \{accept, reject\}$, and a pure strategy for player 2 is a pair of functions (f, g) where $f : \cup_{k=0}^\infty (H^{2k} \times X) \to \{accept, reject\}$ and $g : \cup_{k=0}^\infty H^{2k+1} \to X$.

SPNE is defined in the usual way, with each period's self considered independently. That is, each self optimizes taking the strategies of other selves' (whether of the same agent or of the opponent) as given.

3.2 Play between Dual Selves

Suppose players have the same δ, but potentially different coefficients γ_1 and γ_2 in their linear self-control cost functions.

Proposition 2: A SPNE of the bargaining game exists, and in any SPNE, player 1's offer in period 0 is accepted, and player 1's aggregate payoff v_1 is as follows:

Case I: $1 + \max\{\gamma_1, \gamma_2\} \leq \delta(1 + \delta(1 + \min\{\gamma_1, \gamma_2\}))$

a) If $\gamma_1 < \gamma_2$, $v_1 = \frac{1+(\gamma_2-\gamma_1)}{1-\delta^2}$. Player 2 obtains all of the period-0 surplus.

b) If $\gamma_1 > \gamma_2$, $v_1 = \frac{1-\delta(\gamma_1-\gamma_2)}{1-\delta^2}$. Player 1 obtains all of the period-0 surplus.

Case II: $1 + \max\{\gamma_1, \gamma_2\} > \delta(1 + \delta(1 + \min\{\gamma_1, \gamma_2\}))$

a) If $\gamma_1 < \gamma_2$, $v_1 = \frac{1}{1-\delta} - \frac{\delta(1+\gamma_1)}{1+\gamma_2-\delta^2(1+\gamma_1)}$. Player 2 only obtains period-0 surplus $\frac{\delta}{1+\gamma_2-\delta^2(1+\gamma_1)} = \frac{\delta}{1+\gamma_2} \frac{1}{1-\delta^2 \frac{1+\gamma_1}{1+\gamma_2}}$.

b) If $\gamma_1 > \gamma_2$, $v_1 = 1 + \frac{\delta^2}{\frac{1+\gamma_1}{1+\gamma_2} - \delta^2}$. Player 1 obtains all of the period-0 surplus.

In both cases, if $\gamma_1 = \gamma_2 = \gamma$, then $v_1 = \frac{1}{1-\delta^2}$.

Proof: See Appendix.

The proof of Proposition 2 follows Shaked and Sutton's (1984) proof of SPNE uniqueness in the Rubinstein (1982) game. The argument is modified to account for the different implications of offering current or later surplus in terms of self-control costs.

When the player with better self-control (lower γ) proposes, she offers future surplus only if the current surplus is insufficient to meet the opponent's reservation value. Suppose $\gamma_1 < \gamma_2$, and consider a thought experiment where players are endowed with the surpluses from the periods where they propose. For every unit of current surplus player 1 offers when proposing, player 2 would incur a self-control cost of γ_2 by turning down the offer. Therefore, if player 1 offers the entire current surplus, she can ask for future payoffs with present value $1+\gamma_2$ in return. However, player 1 then incurs self-control cost γ_1: she could instead obtain the entire current surplus by offering all future surplus. Therefore, player 1's gain from trade is $\gamma_2 - \gamma_1$ whenever she proposes in Case Ia, where the difference in γ between the agents is small enough (and δ is large enough) that player 1 offers the entire current surplus to player 2.[10]

3.3 Comparison with Quasi-Hyperbolic Discounting

SPNE payoff multiplicity does not arise with dual selves[11] because, for any aggregate payoff v in period t, the associated reservation value in period $t-1$ is always δv. Therefore, whether v is achieved using time-t or later surplus does not matter for a player's bargaining power in earlier periods. By contrast, with quasi-hyperbolic agents, the value of surplus from period t is discounted by $\beta\delta$ from the perspective of self $t-1$, while for surplus from later periods, the extra discounting that self $t-1$ applies relative to self t is δ. As a result, an agent with aggregate payoff v in period t can have a reservation value in period $t-1$ ranging from $\beta\delta v$ to δv, depending on the source of v. As shown by Lu (2016), this potential multiplicity in reservation values sustains SPNE payoff multiplicity with quasi-hyperbolic agents for a large range of parameter values. The root of this difference can be traced to Proposition 1: dual selves care about future self-control costs, unlike quasi-hyperbolic agents, and since future self-control costs are part of the future payoff v, not caring about them leads to multiple possible reservation values for a single v in the quasi-hyperbolic case.

To quantitatively compare the predictions from the quasi-hyperbolic discounting and dual-self frameworks, Table 1 restates the payoffs with quasi-hyperbolic discounters from Lu (2016) using the selfish dual selves model from Section 2, with the parametrization

[10] The same thought experiment can be extended to interpret the payoffs in Case II as well.

[11] The equilibrium payoff for player 2, however, is not unique at $\gamma_1 = \gamma_2$, which distinguishes this case from the equal patience case in the traditional exponential framework. The reason is as follows: since both players are indifferent about trade, any amount $x \in [0,1]$ of the period-0 surplus can be traded, which results in player 1 incurring a self-control cost of γx. Therefore, the total surplus can be anywhere in $[\frac{1}{1-\delta} - \gamma, \frac{1}{1-\delta}]$, depending on how much period-0 surplus player 1 trades away. In other words, if trading away period-0 surplus, player 1 comes out even by charging her self-control cost to player 2, who suffers. By contrast, if player 1 keeps the period-0 surplus, neither party exerts self-control, which results in a Pareto improvement.

$\beta = \frac{1}{1+\gamma}$ that makes it equivalent to quasi-hyperbolic discounting. The focus is on the case $\min_i \beta_i = \frac{1}{1+\max_i \gamma_i} \geq \frac{1}{\delta(1+\delta)}$, which includes the relevant range of parameter values for most applications (for example, this condition is satisfied whenever $\beta, \delta > 0.76$), and guarantees that neither player accepts obtaining only part of the current surplus and no future surplus. Figure 1 plots player 1's normalized payoff share (payoffs from Table 1 multiplied by $(1-\delta)$) as a function of her self-control parameter $\beta_1 = \frac{1}{1+\gamma_1}$, for $\gamma_2 = \frac{1}{2}$, $\delta = 0.95$, and $\beta_1 \geq \frac{1}{\delta(1+\delta)} \approx 0.54$.

Table 1: Player 1's SPNE payoffs when $1 + \max\{\gamma_1, \gamma_2\} \leq \delta(1+\delta)$	Selfish Dual Self (Lu (2016)'s Proposition 2 restated)	Dual Self (Proposition 2)
$\gamma_1 \in (0, (1-\delta)\gamma_2)$	$\frac{1+\gamma_2}{1-\delta^2} - \gamma_1$	$\frac{1+\gamma_2-\gamma_1}{1-\delta^2}$
$\gamma_1 \in [(1-\delta)\gamma_2, \gamma_2]$	$\left[\frac{1+\gamma_2}{1-\delta^2} - \frac{\gamma_1}{1-\delta}, \frac{1+\gamma_2}{1-\delta^2} - \gamma_1\right]$	$\frac{1+\gamma_2-\gamma_1}{1-\delta^2}$
$\gamma_1 \in [\gamma_2, \frac{1}{1-\delta}\gamma_2]$	$\left[\frac{1-\delta\gamma_1}{1-\delta^2}, \frac{1-\delta\gamma_1}{1-\delta^2} + \frac{\delta}{1-\delta}\gamma_2\right]$	$\frac{1-\delta(\gamma_1-\gamma_2)}{1-\delta^2}$
$\gamma_1 \in (\frac{1}{1-\delta}\gamma_2, \infty)$	$\frac{1-\delta\gamma_1}{1-\delta^2}$	$\frac{1-\delta(\gamma_1-\gamma_2)}{1-\delta^2}$

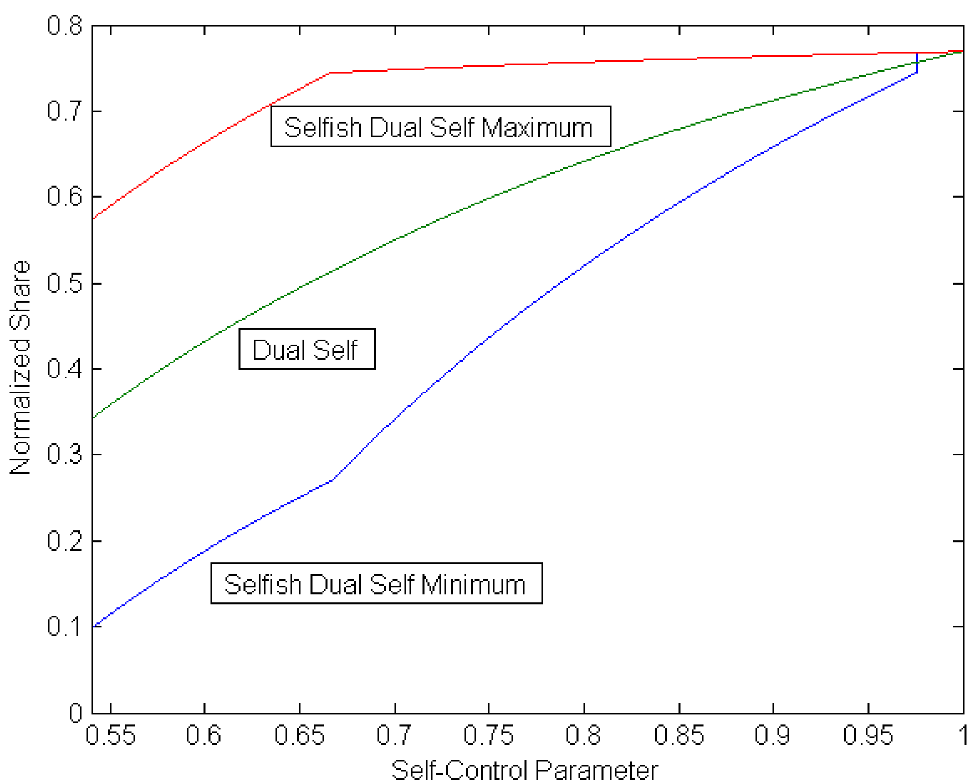

Figure 1: Player 1's SPNE normalized share vs. $\frac{1}{1+\gamma_1}$, for $\gamma_2 = \frac{1}{2}$ and $\delta = 0.95$

When $\overline{v_1} > \underline{v_1}$ in the selfish-dual-self case, $\overline{v_1} - \underline{v_1} = \frac{\delta}{1-\delta}\min\{\gamma_1, \gamma_2\}$. Off path, every period, the player with lower self-control cost obtains only future surplus in her best SPNE, and obtains the entire then-current surplus in her worst SPNE.[12] In the former case, this player incurs aggregate future self-control costs $\frac{\delta}{1-\delta}\min\{\gamma_1, \gamma_2\}$; in the latter case, she forgoes the same amount of actual discounted future surplus. As the selfish dual self ignores the former, her payoff is higher in that case.

When multiple SPNE payoffs exist for selfish dual selves, the dual-self payoff is within the range of selfish-dual-self payoffs. This follows from the previous paragraph: because either player could obtain only future surplus in the continuation and thereby incur self-control costs ignored by earlier selves, either player could be better off than in the dual-self model.

The SPNE from Proposition 2 are Markov perfect equilibria (MPE): players' strategies depend only on t. With quasi-hyperbolic agents, agreement is immediate in MPE and (except when $\beta_1 = \beta_2$) MPE payoffs are unique, just like with dual selves. However, as Table 2 shows, large differences in MPE predictions remain.

Table 2: Player 1's MPE payoffs	Selfish Dual Self (Lu (2016)'s Proposition 1 restated)	Dual Self (Proposition 2)
Case I: $1 + \max\{\gamma_1, \gamma_2\} \leq$	$\delta(1+\delta)$	$\delta(1 + \delta(1 + \min\{\gamma_1, \gamma_2\}))$
$\gamma_1 < \gamma_2$	$\frac{1+\gamma_2}{1-\delta^2} - \gamma_1$	$\frac{1+\gamma_2-\gamma_1}{1-\delta^2}$
$\gamma_1 > \gamma_2$	$\frac{1-\delta\gamma_1}{1-\delta^2}$	$\frac{1-\delta(\gamma_1-\gamma_2)}{1-\delta^2}$
$\gamma_1 = \gamma_2 = \gamma$	$\left[\frac{1-\delta\gamma}{1-\delta^2}, \frac{1+\gamma}{1-\delta^2} - \gamma\right]$	$\frac{1}{1-\delta^2}$
Case II: $1 + \max\{\gamma_1, \gamma_2\} >$	$\delta(1+\delta)$	$\delta(1 + \delta(1 + \min\{\gamma_1, \gamma_2\}))$
$\gamma_1 < \gamma_2$	$\frac{1}{1-\delta} - \frac{\delta(1+\gamma_1)}{1+\gamma_2-\delta^2}$	$\frac{1}{1-\delta} - \frac{\delta(1+\gamma_1)}{1+\gamma_2-\delta^2(1+\gamma_1)}$
$\gamma_1 > \gamma_2$	$1 + \frac{\delta^2}{1+\gamma_1-\delta^2}$	$1 + \frac{\delta^2}{\frac{1+\gamma_1}{1+\gamma_2} - \delta^2}$
$\gamma_1 = \gamma_2 = \gamma$	$\left[1 + \frac{\delta^2}{1+\gamma-\delta^2}, \frac{1}{1-\delta} - \frac{\delta(1+\gamma)}{1+\gamma-\delta^2}\right]$	$\frac{1}{1-\delta^2}$

As Lu (2016) notes, the player i with higher β acts as if $\beta_i = 1$ in MPE. The selfish dual self formulation provides a simple explanation. Because the long-run self ignores future self-control costs, the player with lower γ does not worry that rejecting an offer will result in self-control costs next period, when she would relinquish at least part of the then-current

[12] More precisely, Lu (2016) shows that the worst SPNE payoff for player i is achieved by an SPNE where, off path, player i always obtains the then-current surplus. This is possible in equilibrium even when i has better self-control (meaning that it is inefficient for i to obtain current surplus) because:

- when i proposes, j receives her best continuation value minus the cost of inefficiency $\gamma_j - \gamma_i$, and if i attempts to profitably deviate to an efficient offer, j rejects and obtains her best continuation value;
- when j proposes, any profitable deviation by j is punished by rejection and jumping to j's worst SPNE.

Other equilibria, for example with delay, may also achieve i's worst SPNE payoff.

surplus. Therefore, when receiving offers containing no current surplus, she plays as if $\gamma = 0$. The same applies when proposing: she offers as little future surplus as possible.

When a dual self plays against an exponential discounter with the same discount factor, it is immaterial whether the dual self is selfish: in the MPE of every subgame starting with a proposal, the dual self will either demand the entire surplus (if proposing) or accept the offer (if not), thus incurring no self-control cost. Thus, payoffs from the first column (except when $\gamma_1 = \gamma_2$) can be obtained from the second column by setting the lower γ to 0 and subtracting player 1's on-path self-control cost when $\gamma_1 < \gamma_2$.[13]

By contrast, with standard dual selves, future self-control costs triggered by a rejection matter for the player with lower γ. Therefore, the equilibrium outcome depends on both γ's, and, when $\gamma_1 \neq \gamma_2$, the player with better self-control is better off in the selfish-dual-self/quasi-hyperbolic model than in the dual-self model. Figure 2, the analog of Figure 1 for MPE payoffs, illustrates these results.

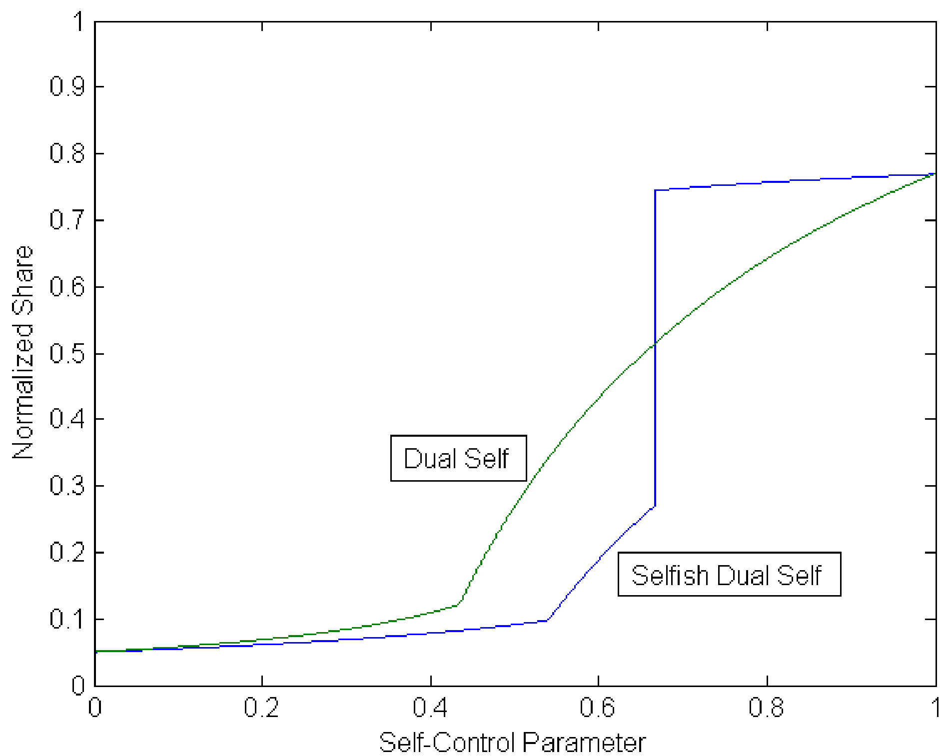

Figure 2: Player 1's MPE normalized share vs. $\frac{1}{1+\gamma_1}$, for $\gamma_2 = \frac{1}{2}$ and $\delta = 0.95$

[13] This cost is γ_1 in Case I since player 1 offers the entire period-0 surplus, and $\frac{\delta\gamma_1}{1+\gamma_2-\delta^2}$ in Case II since she offers $\frac{\delta}{1+\gamma_2-\delta^2}$ of the period-0 surplus.

Similarly, the boundary between Cases I and II in the first column corresponds to that in the second column with $\min\{\gamma_1, \gamma_2\} = 0$.

While player 1's payoff is continuous in the dual-self model, it makes a large jump at $\frac{1}{1+\gamma_1} = \frac{2}{3}$ in the selfish-dual-self model: with dual selves, both players' self-control problems hurt their continuation payoffs, but with quasi-hyperbolic discounters, this holds only for the agent with worse self-control. Therefore, while small differences in self-control matter little for dual selves' welfare, they can have a dramatic impact on quasi-hyperbolic discounters' welfare.[14]

4 Conclusion

This paper shows that quasi-hyperbolic discounting has a straightforward interpretation as a reduced-form dual-self model where each period's agent does not care about future self-control costs. Therefore, unlike for Fudenberg and Levine (2006) dual selves, quasi-hyperbolic discounters' utility is not recursive. In a bargaining setting, this difference has a large impact on equilibrium predictions.[15] Generally speaking, the intuition for results from the dual-self model is closer to the standard intuition from exponential discounting, and thus more familiar, than that from quasi-hyperbolic discounting. It may therefore be important to study which of these frameworks better models limited intertemporal self-control.

5 Acknowledgements

I am grateful to Attila Ambrus, David Freeman, John Friedman, Lisa Kahn, Sebastian Kodritsch, David Laibson, Qingmin Liu, Marciano Siniscalchi, Andrea Wilson, Leeat Yariv, three anonymous referees, and especially Drew Fudenberg for helpful conversations and suggestions. I would also like to thank seminar participants at Harvard University, and participants at the "Logic, Game Theory and Social Choice 8" and "8th Pan-Pacific Conference on Game Theory" joint conference and the 90th Annual Conference of the Western Economic Association International for their questions and comments. I acknowledge the support of the Simon Fraser University President's Research Start-Up Grant.

[14] When $\frac{1}{1+\gamma_1} > \frac{2}{3}$, the selfish-dual-self payoff climbs very slowly. Changes in γ_1 in that range have no effect on the implemented proposal, unlike in the dual-self model. The small positive slope is due to the reduction in player 1's on-path self-control cost as $\gamma_1 \to 0$.

[15] An earlier version of this paper showed that similar qualitative differences can arise when the size of surplus varies across periods or when the players have different δ's (though, as noted in Lu (2016), with quasi-hyperbolic discounting, MPE payoff discontinuity in the latter case would typically occur at points other than $\beta_1 = \beta_2$).

If utility in each period were non-transferable, Lu (2016) shows that MPE payoffs with quasi-hyperbolic agents become continuous under certain conditions. However, it remains true that, in MPE, the agent with better self-control is worse off in the dual-self model than the quasi-hyperbolic model: the off-path self-control costs incurred by this agent only affect her reservation value in the dual-self case.

6 Appendix: Example 1 and Proof of Proposition 2

Example 1 is a one-player "procrastination game" due to O'Donoghue and Rabin (1999), who analyze it with the quasi-hyperbolic model. Fudenberg and Levine (2006) use it to illustrate the dual-self model. The game is solved below using both models to illustrate Proposition 1.

Example 1: There are 4 periods (labeled 1-4), and the agent must complete a task by the end of the last one. The costs of performing the task increase as time progresses: they are 3, 5, 8 and 13 respectively. Suppose that the quasi-hyperbolic agent (A) has $\beta = 1/2$, and the dual self (B) has $\gamma = 1$, so that $\beta = \frac{1}{1+\gamma}$; both agents have $\delta = 1$.

In period 4, both agents complete the task and incur cost 13.

In period 3, player A values the cost of acting at 8, and the cost of waiting at $1/2 \times 13 = 6.5$, so she waits. Player B incurs cost 16 if she acts (8 from the action and 8 in self-control costs) and 13 if she waits, so she also waits. Notice that B's costs are exactly $1+\gamma$ times A's, as the maximum utility in period 3 is 0 (this is not the case in period 4, where the players' utilities are the same).

In period 2, player A compares cost 5 if she acts to cost $1/2 \times 13 = 6.5$ if she waits; therefore she acts. Agent B has cost 10 if she acts versus cost 13 if she does not, so she also acts. Again the utilities differ by a factor of 2.

In period 1, player A suffers a disutility of 3 by acting and $1/2 \times 5 = 2.5$ by waiting, so she waits. Player B incurs cost 6 by acting and 10 by waiting, so she acts. Notice that if player B does not care about her self-control cost in period 2, she would value the cost of waiting at 5. In that case, her costs would again be exactly twice agent A's, and she also would wait.

Proof of Proposition 2: It is always possible for the proposer i to get the entire current pie by offering everything after the current period. The opponent j clearly would do worse by rejecting such an offer, so it will be accepted. Thus offering share y_0 of the current pie carries self-control cost $\gamma_i y_0$ if i expects j to accept the offer, and γ_i if i expects j to reject the offer.

If the receiver accepts the offer, he is maximizing his then-current share and therefore exerts no self-control; if he rejects the offer, the self-control cost is $\gamma_j y_0$.

It follows that i can reduce j's reservation value by $\gamma_j y_0$ at a cost of $\gamma_i y_0$.

Following Shaked and Sutton (1984), let $\underline{v_k}$ and $\overline{v_k}$ be the infimum and supremum of player k's aggregate SPNE payoff when proposing first, and suppose that $\gamma_i \leq \gamma_j$.

Because player i can ensure that her offer will be accepted by player j by offering $\delta\overline{v_j} - \gamma_j y_0$, where $y_0 = \min\{1, \delta\overline{v_j} - \gamma_j y_0\}$ is the share of the current surplus offered (so $y_0 =$

$\min\{1, \frac{\delta\overline{v_j}}{1+\gamma_j}\})$, we have

$$\underline{v_i} \geq \frac{1}{1-\delta} - \delta\overline{v_j} + (\gamma_j - \gamma_i)\min\{1, \frac{\delta\overline{v_j}}{1+\gamma_j}\}.$$

In order for player j's offer to be accepted, he must at least offer $\delta\underline{v_i} - \gamma_i y_0$. Since offering y_0 triggers self-control cost $\gamma_j y_0$, the best that can be done is to set $y_0 = 0$. Alternatively, player j may make a rejected offer. Then player j's payoff next period will be no more than $\delta\overline{v_j}$, so this course of action cannot lead to payoff $\overline{v_j}$. Thus

$$\overline{v_j} \leq \frac{1}{1-\delta} - \delta\underline{v_i}.$$

Similar reasoning leads to the following equations for $\overline{v_i}$ and $\underline{v_j}$:

$$\begin{aligned} \overline{v_i} &\leq \frac{1}{1-\delta} - \delta\underline{v_j} + (\gamma_j - \gamma_i)\min\{1, \frac{\delta\underline{v_j}}{1+\gamma_j}\} \\ \underline{v_j} &\geq \frac{1}{1-\delta} - \delta\overline{v_i} \end{aligned}$$

Solving these equations yields $\underline{v_i} \geq \overline{v_i}$ and $\underline{v_j} \geq \overline{v_j}$, which implies payoff uniqueness. Simple algebraic manipulations lead to the expressions in the statement of Proposition 2. Checking existence is straightforward.

It remains to be shown that there cannot be delay in SPNE. Note from the equations above that player 1 can guarantee herself payoff arbitrarily close to $\frac{1}{1-\delta} - \delta v_2$ by offering slightly more than required in period 0. If instead there were delay, player 2's payoff in period 1 must be v_2, which implies that player 1's payoff is no more than $\delta(\frac{1}{1-\delta} - v_2)$ from the perspective of period 0. Therefore, player 1 has a profitable deviation in any strategy profile featuring delay. ■

7 References

Angeletos, G.-M., D. Laibson, A. Repetto, J. Tobacman, and S. Weinberg. 2001. "The Hyperbolic Buffer Stock Model: Calibration, Simulation, and Empirical Evaluation." *Journal of Economic Perspectives*, 15(3): 47-68.

Bénabou, Roland, and Marek Pycia. 2002. "Dynamic Inconsistency and Self-Control." *Economics Letters*, 77: 419-424.

Benhabib, Jess, and Alberto Bisin. 2005. "Modeling Internal Commitment Mechanisms and Self-Control: A Neuroeconomics Approach to Consumption-Saving Decisions." *Games and Economic Behavior*, 52: 460-492.

Bernheim, B. Douglas, and Antonio Rangel. 2004. "Addiction and Cue-Triggered Decision Processes." *American Economic Review*, 94(5): 1558-1590.

Brocas, Isabelle, and Juan D. Carrillo. 2008. "The Brain as a Hierarchical Organization." *American Economic Review*, 98(4): 1312-1346.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue. 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature*, 40(2): 351-401.

Fudenberg, Drew, and David Levine. 2006. "A Dual Self Model of Impulse Control." *American Economic Review*, 96(5): 1449-1476.

Gul, Faruk R., and Wolfgang Pesendorfer. 2001. "Temptation and Self Control." *Econometrica*, 69: 1403-1436.

Gul, Faruk R., and Wolfgang Pesendorfer. 2005. "The Revealed Preference Theory of Changing Tastes." *Review of Economic Studies*, 72, 429-448.

Kodritsch, Sebastian. 2014. "On Time Preferences and Bargaining." Discussion paper, WZB Berlin Social Science Center.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*, 112(2): 443-477.

Laibson, David, Andrea Repetto, and Jeremy Tobacman. 2007. "Estimating Discount Function from Lifecycle Consumption Choices." *Mimeo*.

Loewenstein, George, and Ted O'Donoghue. 2004. "Animal Spirits: Affective and Deliberative Processes in Economic Behavior." *Mimeo*.

Lu, Shih En. 2016. "Self-Control and Bargaining." *Journal of Economic Theory*, forthcoming.

McClure, Samuel M., David I. Laibson, George Loewenstein, and Jonathan D. Cohen. 2004. "Separate Neural Systems Value Immediate and Delayed Monetary Rewards." *Science*, 306: 503-507.

Montiel Olea, José Luis, and Tomasz Strzalecki. 2014. "Axiomatization and Measurement of Quasi-Hyperbolic Discounting." *Quarterly Journal of Economics*, 129: 1449-1499.

O'Donoghue, Ted, and Matthew Rabin. 1999. "Doing It Now or Later." *American Economic Review*, 89(1): 103-124.

Phelps, E. S., and Robert Pollak. 1968. "On Second-Best National Saving and Game-Equilibrium Growth." *Review of Economic Studies*, 35(2): 185-199.

Rubinstein, Ariel. 1982. "Perfect Equilibrium in a Bargaining Model." *Econometrica*, 50: 97-110.

Shaked, Avner, and John Sutton. 1984. "Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Game." *Econometrica*, 52: 1351-1364.

Ståhl, Ingolf. 1972. *Bargaining Theory*. Stockholm: Stockholm School of Economics.

Thaler, Richard H., and Hersh M. Shefrin. 1981. "An Economic Theory of Self-Control." *Journal of Political Economy*, 89(2): 392-410.

Xue, Licun. 2008."The Bargaining within." *Economics Letters*, 101: 145-147.