Renegotiating with One's Self

Karna Basu

Department of Economics, University of Chicago*

June 7, 2009

Abstract

In an infinite horizon setting, how does a sophisticated hyperbolic discounter decide when to purchase a durable good or a commitment product that is costly in the present but yields benefits in the future? In this paper, I describe a new type of agent, the semi-sophisticated hyperbolic discounter. This agent is aware of her time-inconsistency, but fails to anticipate that her future selves might renegotiate to a different equilibrium. When faced with a deadline, a semi-sophisticate will behave like a sophisticated hyperbolic discounter, but without a deadline (when she faces a multiplicity of equilibria), she will appear naïve. With a growing population, a government seeking to maximize social welfare will not offer the good everyday. A monopolistic firm might generate higher consumer surplus than a competitive market, since a monopolist can commit to a pattern of sale and non-sale prices, thus effectively creating deadlines.

*Email: karna@uchicago.edu. Website: http://home.uchicago.edu/~karna. I thank Abhijit Banerjee, Xavier Gabaix, Roger Myerson, John List, Sendhil Mullainathan, Patricia Cortes, Saurabh Bhargava, and seminar participants at the Federal Reserve Bank of Boston, Hunter College, Rutgers Newark, Royal Holloway, and Williams College for comments and suggestions.

1 Introduction

Hyperbolic discounting as a description of preferences has gained legitimacy in the development economics literature[1]. When hyperbolic discounters face high transactions costs and lack access to the complex contracts available to consumers in industrialized countries, what are the implications for market equilibrium and welfare? In this paper, I propose a distinction between two types of sophisticated hyperbolic discounters. *Full-sophisticates* always play a renegotiation-proof equilibrium with their future selves. *Semi-sophisticates* are sophisticated in the sense that they are aware of their future time-inconsistency. Yet, in situations with multiple equilibria, they fail to anticipate the possibility that their future selves will renegotiate to a different equilibrium. This generates several insights into markets for durable goods and commitment products.

Deadlines are salient: regardless of the distance of the deadline, semi-sophisticates perform well with deadlines, and badly without. In particular, I study environments where individuals must choose whether and when to perform tasks with immediate costs and delayed benefits. When faced with a deadline, a semi-sophisticate is indistinguishable from a sophisticated hyperbolic discounter, but when the deadline is lifted the same individual is indistinguishable from a naive hyperbolic discounter.

This model also provides a rationalization of sale pricing of durable goods that does not rely price discrimination. Monopolists who sell to semi-sophisticates will choose to vary prices, leading to sale periods and non-sale periods. When there are competitive firms, there are multiple equilibria. Competitive equilibrium can lead to lower social welfare than a monopolistic equilibrium, especially when the population is a mix of semi-sophisticated hyperbolic discounters and exponential discounters.

1.1 Motivation

The first question relates to equilibrium selection. Typically, hyperbolic discounters are classified as either *naive* or *sophisticated*.[2] A naive hyperbolic discounter believes that, from the next period onwards, she has time-consistent preferences. A sophisticated hyperbolic discounter is fully aware of her future time-inconsistency. Suppose this individual must decide when to complete a

[1] See, for example, Mullainathan (2004).

[2] O'Donoghue and Rabin (2001a) argue that people might also be *partially naive*; i.e. they underestimate the degree of future time-inconsistency. They show that even a little bit of naivete can significantly increase procrastination.

task with immediate costs and future benefits. Under finite horizons, as O'Donoghue and Rabin (1999) show, a sophisticated hyperbolic discounter will complete the task in accordance with her unique Subgame Perfect Nash Equilibrium. The outcome is less clear under an infinite horizon. When faced with multiple equilibria, which one will the agent choose? In any period, she might wish to "renegotiate" with herself if the previously selected equilibrium does not maximize her lifetime welfare. This renegotiation problem must be taken seriously because there is no one the agent needs to communicate with. Rather than renegotiate, she can simple re-choose an equilibrium for herself and her future selves. The fact that an agent is sophisticated about her preferences does not automatically provide a solution to the renegotiation problem. I separate sophisticated hyperbolic discounters into two groups: full-sophisticates restrict themselves to renegotiation-proof equilibria, while semi-sophisticates fail to recognize the possibility of renegotiation. I argue that semi-sophistication is not only psychologically realistic, but provides some new and potentially important insights into the second question.

The second question is this: How does a hyperbolic discounter decide when to purchase a durable good or a commitment product?[3] Even when she values the product, there is often no contract that commits her to adopting it by a given date. In this paper, I study how the take-up of these products is affected by the individual's preferences, beliefs about preferences, beliefs about equilibrium selection, and market structure.

A hyperbolic discounter's optimizing behavior at any time might conflict with what her past selves would want her to do. Commitment products (which essentially contract the individual's choice set in future periods) can therefore raise lifetime welfare. Once the product has been purchased, future selves behave more in accordance with early selves' objectives. However, the adoption of these products can be costly in itself. If adoption involves even a simple transaction cost, the agent might have an incentive to postpone it to the following period. So, it doesn't follow from the fact that a product increases the individual's lifetime utility, that she will indeed decide to purchase it when it is offered. In the case of commitment products, then, the question becomes: What can be done when I can't commit to commit? The problem is especially relevant in developing countries for two reasons: poor contracting and high transaction costs (in the form of geographic distance, time, and bureaucratic hurdles).

[3] I present examples of such products in Section 1.3.

1.2 Preview of Results

In my model, an infinitely-lived[4], sophisticated quasi-hyperbolic discounter has the option of purchasing a product, which provides delayed benefits and involves immediate costs that cannot be transferred to future periods. I assume these costs and benefits are such that the agent is willing to buy the good immediately, but would ideally like to postpone the purchase to the next period. Sophistication does not directly inform us about how an equilibrium is selected when there is more than one to choose from. I describe as "fully-sophisticated" those who play a renegotiation-proof equilibrium[5], and as "semi-sophisticated" those who choose an equilibrium without realizing that it is subject to renegotiation by their future selves. To put it another way, full-sophisticates understand that they must restrict themselves to equilibria that are not strictly dominated by other equilibria in the future. A semi-sophisticate is fully aware of her future time preferences but is naive about equilibrium reselection. In each period, she switches out of the equilibrium played by his past self if there is another equilibrium that gives her strictly higher continuation utility.

First, I show that this results in strong deadline effects. Semi-sophisticated agents appear sophisticated under finite horizons and naive under infinite horizons. Unlike either fully-sophisticated or naive agents, a semi-sophisticate does much better under a deadline, *however distant*, relative to when there is no deadline at all. This is because a deadline locks her into a unique equilibrium, while the absence of one allows her to perpetually renegotiate and never buy the good.

Second, I show that, with a growing population, a benevolent government that provides the good (priced at marginal cost) will fail to maximize consumers' welfare. To induce semi-sophisticates to buy the good, there must be stretches when it is not offered. As a result, agents born during those stretches will be forced to wait until the next date of availability. In this section, I also show that the government might offer the good in a similarly staggered manner to fully-sophisticated hyperbolic discounters as well. This can lead to more frequent purchases than with some renegotiation-proof equilibria under continuous access.

Finally, I study markets with profit-maximizing firms. Firms list their prices for every future period, based on which semi-sophisticated agents make their purchasing decisions. A monopolist will always vary prices over time. By keeping prices high in certain periods, the firm effectively

[4] The infinite horizon can also be interpreted as an inter-generational setting.

[5] I adapt renegotiation-proofness as defined in Farrell and Maskin (1989) and Kocherlakota (1996) to my setting.

generates deadlines. This allows semi-sophisticates to play a unique equilibrium and buy the good during sale periods. Under perfect competition, there are multiple equilibria. There is an equilibrium where firms coordinate on sale and non-sale periods, but also an equilibrium in which the good is always offered at marginal cost (and therefore never purchased). In the competitive equilibrium with sales, social welfare is higher than under monopoly, but in the equilibrium without sales it is strictly lower. Furthermore, if the population contains any exponential discounters, competitive firms can no longer maintain the equilibrium with sales.

1.3 Applications

The model makes predictions that are relevant to several applications.

A rosca (rotating savings and credit association) is an informal savings group that serves as a natural example. Gugerty (2007) provides evidence from Kenyan roscas that suggests they are an effective commitment device for hyperbolic discounters. In Basu (2008a), I find explicit conditions under which a rosca will survive as a commitment device even if social sanctions are ineffective. However, even when an agent knows she will never leave a rosca upon joining, this leaves open the question of when and if she will join at all. Given that there are initial time and administrative costs, a semi-sophisticated agent could indefinitely postpone entry. One way to deal with this is to coordinate ahead of time–either we all start the rosca tomorrow, or we don't start for another month. Alternatively, attempts to reduce the first-period cost (such as shortening the distance the person needs to travel to sign up) can prove more beneficial than other models would suggest.

This kind of reasoning can be extended to other commitment products as well. Ashraf, Karlan, and Yin (2006) offer women in the Philippines access to a bank account which makes savings illiquid until a certain date or target accumulation level. They find that hyperbolic discounting (elicited through hypothetical questions) is a significant predictor of take-up. The people who perceive the need for a commitment product must be sophisticated hyperbolic discounters. Of those, the ones who are semi-sophisticated and face high transaction costs will nevertheless not join. Here, a limited-time offer will create a unique equilibrium for semi-sophisticates, and generate higher takeup than an open offer would. The model should also help answer one of the questions they pose: "A natural question arises concerning why, if commitment products appear to be demanded by consumers, the market does not already provide them" (p. 638). While their focus is on the design of the product, I suggest here that patterns of access to the product also matter[6]. The analysis with competitive firms shows how markets could fail to deliver commitment even when individuals value it. It is possible to design experiments that vary the deadlines associated with a commitment product to test the strength of the semi-sophistication assumption.

This model also provides a framework for distinguishing between goods that have a natural deadline, and those that don't. Take, for example, a roof that needs to be fixed. Upon completion, this will yield a stream of future benefits. The constraint is that the roof must be fixed before winter sets in and makes outdoor work impossible. The model predicts that a naive hyperbolic discounter will fix the roof on the last possible day, an exponential discounter will do it on the first day itself, and a semi-sophisticated (and fully-sophisticated) hyperbolic discounter will do it sometime in between. Now consider the case where the agent resides in a tropical climate with perpetual summer. Under these circumstances, the semi-sophisticate, like the naif, will never fix the roof.

1.4 A Note on the Preferences

Properties of hyperbolic discounting have been investigated in several papers, including Phelps and Pollack (1968), Laibson (1997), Harris and Laibson (2001), and O'Donoghue and Rabin (2001a, 2001b). I suggest that semi-sophistication is interesting not only because of the unusual predictions it generates, but also because it is a new and reasonable description of preferences. This captures the idea that people are not under the foolish illusion that their future selves are exponential discounters, but are nevertheless frustrated about the choices they make and continue to make. The model allows people to be fully aware of their hyperbolic tendencies (so they are not naive), even as they make mistakes along another dimension (not realizing that future action sets also include choices over equilibria).

Choi, Laibson, Madrian, Metrick (2001) find that employees are much more likely to join and remain in 401(k) savings plans if these plans are offered as the default option. If we are to assume that, in either case, most people would prefer to be enrolled in 401(k) than not, but that there is a small cost to deviating from the default option, then the empirical results are

[6] In Basu (2008b), I suggest some other reasons for this in a finite horizon setting. I show that firms might not offer commitment if they can enforce lending contracts. Also, takeup of NGO-delivered commitment devices will be low if there is a moneylender nearby.

consistent with my model. When the 401(k) option is not a default, sophisticated hyperbolic discounters are failing to adopt it. In contexts such as this, semi-sophistication might provide an alternative interpretation for status-quo bias. Also, Cialdini (2001) writes about how the illusion of scarcity can make people more willing to purchase a good: "A great deal of evidence shows that items and opportunities become *more desirable* to us as they become less available" (emphasis mine). This could be taken as further indirect evidence of semi-sophistication.

2 Model Setup

The timing of the model is as follows. Time is discrete. There is one good. In period 0, a firm/government declares prices for all future periods. Starting in period 1, infinitely-lived individuals are born in every period. The individuals take the prices as given and plan their lifetime behavior (which depends of preferences and levels of sophistication as described below).

2.1 Preferences

The setup is a modified version of the model by O'Donoghue and Rabin (1999). An individual lives for infinite periods, with each period indexed t, $t \in (1, 2, 3, ...)$. Her per-period *instantaneous* utility is $u_t(y_t)$ and her lifetime utility in period t is $U^t(u_t, u_{t+1}, u_{t+2}, ...)$. This is the utility she wishes to maximize. The lifetime utility is a discounted sum of current and future instantaneous utilities. For a time-consistent agent (TC), it is simply:

$$\forall t, U^t = \sum_{\tau=t}^{\infty} \delta^{\tau-t} u_\tau$$

A quasi-hyperbolic agent overweights present utility relative to all future periods:

$$\forall t, U^t = u_t + \beta \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} u_\tau$$

In either case, $\beta \in (0, 1)$ and $\delta \in (0, 1)$.

Assume that, in the absence of any action, $u_t = 0$ in every period t (In other words, her consumption level is $y_t = 0$, and $u(0) = 0$). Also assume that $u'(y) > 0$. In this model, the agent's chance to improve her lifetime utility comes from the opportunity to purchase a product (or, equivalently, perform a task). At most one such purchase can be made over an entire lifetime, and only in a period when a 'seller' decides to offer it. Since the only variation in utility comes from the purchase of this product, we can rewrite continuation utility as a function of the period in which the purchase takes place, and the price in that period, p_τ. $\forall t, U^t(\tau, p_\tau)$ is the continuation utility from the perspective of time t, if the good is to be purchased in period $\tau \geq t$, with the special case of $\tau = 0$ if it will not be purchased from period t onwards.

I assume that the price of the good is not transferrable across periods. By restricting the impact of the monetary cost to the period in which is it incurred, I rule out saving and borrowing. The price could therefore be viewed as a transaction cost that the agent incurs. Suppose that the lowest possible transaction cost is given by $c > 0$. This means that a firm is not able to reduce the price below c, or that the marginal cost of producing the good is c. An interpretation of this is that the firm can choose the location of the sale: the further away from the agent the sale takes place, the higher the price p paid by the agent, and the higher the firm's profits.

The product yields a benefit of b in the period after which it is purchased. This benefit can be viewed as a one-period consumption benefit, or the discounted lifetime benefit from a durable good, or the future benefits of a commitment product. Then, for a hyperbolic discounter:

$$U^t(\tau, p_\tau) = \begin{cases} 0, \text{ if } \tau = 0 \\ u(-p_\tau) + \beta\delta u(b), \text{ if } \tau = t \\ \beta\delta^{\tau - t}\left[u(-p_\tau) + \delta u(b)\right], \text{ if } \tau > t \end{cases}$$

This notation can accommodate a probabilistic purchase of the good as well. For example, if there is μ probability of the good being purchased in period τ at price p, and v probability of it being purchased in period τ' at price p', the continuation utility in some period $t \leq \tau, \tau'$ is $\mu U^t(\tau, p) + vU^t(\tau', p')$.

I make the following assumptions about the product (restricting the parameters to the range of interest):

1. When the product is offered at marginal cost, it is better to purchase the it today than never at all (for all types of discounters): $u(-c) + \beta\delta u(b) > 0$

2. No indifference when prices are c: in any period t, for any $\tau > t$, either $U^t(t, c) > U^t(\tau, c)$ or $U^t(t, c) > U^t(\tau, c)$.

3. If the product is always offered at marginal cost, then in any period t, the hyperbolic agent would rather purchase it in $t + 1$ than in t: $u(-c) + \beta\delta u(b) < \beta\delta(u(-c) + \delta u(b))$

For simplicity, I will often refer to $U^t(\tau, c)$ as $U^t(\tau)$.

To make predictions about when an individual will buy the product, I classify individuals into types based on their time preferences and beliefs about preferences. As is standard in the literature, individuals with $\beta = 1$ are exponential discounters (TC) and those with $\beta < 1$ are hyperbolic discounters. Hyperbolic discounters can be naive (NH) or sophisticated (SH). A Naive hyperbolic discounter, in any period, believes that her future selves are exponential discounters. A sophisticated hyperbolic knows that her future selves discount the future exactly as she does; i.e. by placing relatively greater weight on instantaneous utility.

For sophisticated hyperbolic discounters, it is convenient to think of a person's lifetime behavior as a game with infinite agents, where each agent decides whether to purchase the product, given that it has not been purchased so far. The game ends in the first period that the product is purchased. Every player makes a decision at a different node. Figure 1 shows what the game would look like if the product was offered in every period (payoffs are in terms of continuation utilities from the point of view of each player, assuming prices are c).

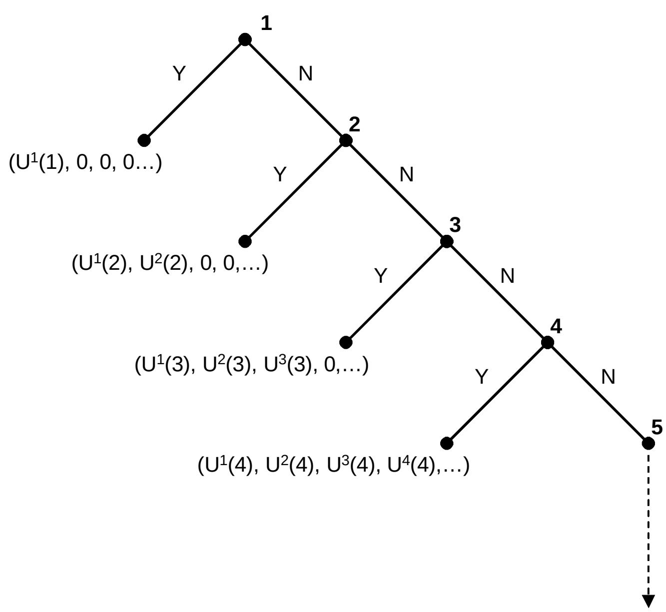

Figure 1: The dynamic game played by a sophisticated hyperbolic discounter's per-period selves when the good is offered in every period at price c.

In this paper, I draw a further distinction within the set of sophisticated hyperbolic discounters. These agents can be either semi-sophisticated or fully-sophisticated. Implications of this distinction are derived in the following sections. Semi-sophisticates are fully aware of their future preferences; hence they always play a Subgame Perfect Nash Equilibrium with their future selves. However, in any period, a semi-sophisticate chooses to play the equilibrium that maximizes her continuation utility, while failing to anticipate that future selves may switch to a different equilibrium. A full-sophisticate restricts herself to a strongly renegotiation-proof equilibrium, if available. We can see here that full-sophisticates and semi-sophisticates will behave similarly in games with a unique equilibrium, but their respective strategies can diverge if there are multiple equilibria.

2.2 "Optimal" Pure Strategies

In this and the following section, I formally define the equilibrium concept, or notion of optimization, for each type of agent. First, in this section, I restrict the agent to playing pure strategies. In any period, she either purchases the product with certainty, or she does not purchase it at all.

We can fully describe behavior by a *strategy* $s \equiv (s_1, s_2, s_3, ...)$, where s_t in any period t belongs to the action set $a_t = \{Y, N\}$. So a strategy is a description of what an individual in any time t would do, conditional on the good having not been purchased so far. Now, we need a way to determine which strategies constitute an "equilibrium" for each type of agent. O'Donoghue and Rabin use the concept of a *perception-perfect strategy* (PPS) in a finite-horizon model. This means that, in any period (assuming the good has not yet been bought), the optimizing individual would take future prices and the behavior of his future selves as given and make decisions according to the rules described below. A strategy, s, is a perception-perfect strategy if in every period t, the individual's optimal action is s_t given his beliefs about the actions played by her future selves.

For exponential discounters, the PPS is equivalent to a standard intertemporal maximization problem. Since there is no disagreement across selves, the agent will purchase the product in a given period as long as she cannot do better by purchasing it in the future instead.

Definition 1 *The optimal strategy for a TC is a strategy* $s^{TC} \equiv \left(s_1^{TC}, s_2^{TC}, ...\right)$ *such that* $\forall t$, $s_t^{TC} = Y$ *iff [*$U^t(t, p_t) \geq 0$ *and* $\forall \tau > t$, $U^t(t, p_t) \geq U^t(\tau, m_\tau)$*].*

A Naive hyperbolic makes a similar comparison. In any period, she will again purchase the product only if she cannot do better by purchasing it in the future. However, this is a flawed comparison, because she mistakenly believes that if it is now optimal for the product to be bought in some future period t, it will continue to be optimal for the period t agent. The definition below follows from the fact that a NH believes that, starting tomorrow, she is a TC.

Definition 2 *A PPS for a NH is a strategy* $s^{NH} \equiv \left(s_1^{NH}, s_2^{NH}, ...\right)$ *such that* $\forall t$, $s_t^{NH} = Y$ *iff* [$U^t(t, m_t) \geq 0$ *and* $\forall \tau > t$, $U^t(t, m_t) \geq U^t(\tau, m_\tau)$].

For a sophisticated hyperbolic discounter, a PPS is equivalent to a Subgame Perfect Nash Equilibrium with multiple selves. An agent will choose to purchase the good in any period as long as it is better than waiting for the next time it is purchased under a given equilibrium. In other words, in any period she understands that future periods will choose actions that are optimal from their own perspective, not hers.

Definition 3 *A PPS for a SH is a strategy* $s^{SH} \equiv \left(s_1^{SH}, s_2^{SH}, ...\right)$ *such that* $\forall t$, $s_t^{SH} = Y$ *iff* $U^t(t, m_t) \geq 0$ *and* $U^t(t, m_t) \geq U^\tau(\tau', m_{\tau'})$, *where* $\tau' \equiv \begin{cases} 0, \text{ if } \forall \tau > t,\ s_\tau^{SH} = N \\ \min_{\tau > t}\left\{\tau : s_\tau^{SH} = Y\right\}, \text{ otherwise} \end{cases}$.

Finally, it is useful to define $\tau_z(s^z)$ as the period in which the good actually gets purchased under a given strategy, s^z.

Definition 4 *For any* $z \in \{TC, SH, NH\}$,

$$\tau_z(s^z) \equiv \begin{cases} 0, \text{ if } \forall t,\ s_t^z = N \\ \min_t \{t : s_t^z = Y\}, \text{ otherwise} \end{cases}$$

$\tau_z(s^z)$ *is the period in which the product will be purchased, given a strategy* s^z.

2.3 "Optimal" Mixed Strategies

The definitions above can be extended to allow for mixed strategies. Behavior is now captured by a mixed-strategy profile $s \equiv (\rho_1, \rho_2, \rho_3, ...)$, where ρ_t in any period t is the probability that the player chooses Y: $\rho_t \in [0, 1]$. I will sometimes write $\rho_t = Y$ or $\rho_t = N$ when either of those strategies is played with certainty. (In the future, I drop superscripts when it is obvious which type of player we are referring to.

Definition 5 *A PPS for a TC is a strategy* $s^{TC} \equiv (\rho_1, \rho_2, \rho_3, ...)$ *such that* $\forall t$,

$$s_t^{TC} = \begin{cases} Y \text{ if } [U^t(t, m_t) > 0 \text{ and } \forall \tau > t, U^t(t, m_t) > U^t(\tau, m_\tau)] \\ N \text{ if } [U^t(t, m_t) > 0 \text{ or } \exists \tau > t \text{ s.t. } U^t(t, m_t) < U^t(\tau, m_\tau)] \\ x \in [0, 1], \text{ otherwise} \end{cases}$$

Definition 6 *A PPS for a NH is a strategy* $s^{NH} \equiv (\rho_1, \rho_2, \rho_3, ...)$ *such that* $\forall t$,

$$s_t^{NH} = \begin{cases} Y \text{ if } [U^t(t, m_t) > 0 \text{ and } \forall \tau > t, U^t(t, m_t) > U^t(\tau, m_\tau)] \\ N \text{ if } [U^t(t, m_t) > 0 \text{ or } \exists \tau > t \text{ s.t. } U^t(t, m_t) < U^t(\tau, m_\tau)] \\ x \in [0, 1], \text{ otherwise} \end{cases}$$

Definition 7 *A PPS for a SH is a strategy* $s^{SH} \equiv (\rho_1, \rho_2, \rho_3, ...)$ *such that* $\forall t$,

$$s_t^{SH} = \begin{cases} Y, \text{ if } U^t(t, m_t) > 0 \text{ and } U^t(t, m_t) > \sum_{\tau=t+1}^{\infty} \rho_\tau U^\tau(\tau, m_\tau) \prod_{r=t+1}^{\tau-1} (1 - \rho_r) \\ N, \text{ if } U^t(t, m_t) < 0 \text{ or } U^t(t, m_t) < \sum_{\tau=t+1}^{\infty} \rho_\tau U^\tau(\tau, m_\tau) \prod_{r=t+1}^{\tau-1} (1 - \rho_r) \\ x \in [0, 1], \text{ otherwise} \end{cases}$$

3 Basic Case: The Product is Offered at Marginal Cost

Now that we have a theory to limit the set of strategies that constitute a PPS for any type of agent, we can study outcomes in a simple setting. Assume that the good is offered at price c in every period. In this setting, it is natural to think of the good as a task instead. This task could be a job application, a trip to Ikea, or the first phone call to a romantic interest. By eliminating the role of markets, we are able to make some interesting predictions about outcomes for different types of agents.

I consider two cases here. First, look for PPSs when there is a deadline. Here, the task can be completed at cost c in any period up to a deadline T. The results here intuitively match those in O'Donoghue and Rabin (1999). An exponential discounter completes the task in period 1. A naive hyperbolic discounter completes the task in period T. Both fully-sophisticated and semi-sophisticated hyperbolic discounters play the unique PPS and complete the task by some date in between.

However, when we lift the deadline, outcomes diverge dramatically for full-sophisticates and semi-sophisticates. A full-sophisticate plays a renegotiation-proof PPS and performs the task with some probability in each period. However, a semi-sophisticate finds that, in each period, she can switch to a new PPS that allows her to postpone the task. As a result, she never completes the task. This highlights the fact that a semi-sophisticated agent appears sophisticated when there is a deadline, but naive when there isn't. The distance of the deadline is irrelevant; rather, it's existence plays a major role.

3.1 With Deadline

I assume that the product is offered in each period until some deadline. This restricts the action sets in the following way:

$$a_t = \begin{cases} \{Y, N\}, \text{ if } t \leq T \\ \{N\}, \text{ if } t > T \end{cases}$$

Alternatively, to be fully consistent with the strategies defined in the previous section, we could simply assume that the price from period T onwards is prohibitively high, so that it is never in any agent's interest to purchase the good after that date. To simplify notation, I write utility as $U^t(\tau)$, because the price is always c until the deadline T.

What will behavior look like in this case? For a TC, the unique PPS is:

$$s_t^{TC} = \begin{cases} Y, \text{ if } t \leq T \\ N, \text{ if } t > T \end{cases}$$

. This follows from A1. Then, $\tau_{TC} = 1$. She will always complete the task in the first period of her life. For a NH:

$$s_t^{NH} = \begin{cases} N, \text{ if } t \neq T \\ N, \text{ if } t = T \end{cases}$$

This follows from A1 and A3. Then, $\tau_{NH} = T$. In any period, she believes that her next period self will find it optimal to complete the task. Hence, she procrastinates until period T.

We can also uniquely determine the PPS for a SH. The intuition for this can be gained easily using backward induction. We know that the agent in period T will certainly complete the task if it has not yet been completed. Then, clearly it cannot be in the interest of the period $T-1$ agent to complete it, since she is able to attain her optimal continuation utility by leaving it for the next period. Now consider the agent in period $T-2$. She knows that, if she does not complete the task immediately, she will have to wait 2 periods. Suppose she finds it is best to complete it immediately. Then, period $T-3$ will certainly not complete it, and period $T-2$ certainly will, and so on. This gives a strategy where the agent plays Y every other period. If she were a bit more hyperbolic, we might have a PPS where she plays Y every 3 periods (using the same reasoning).

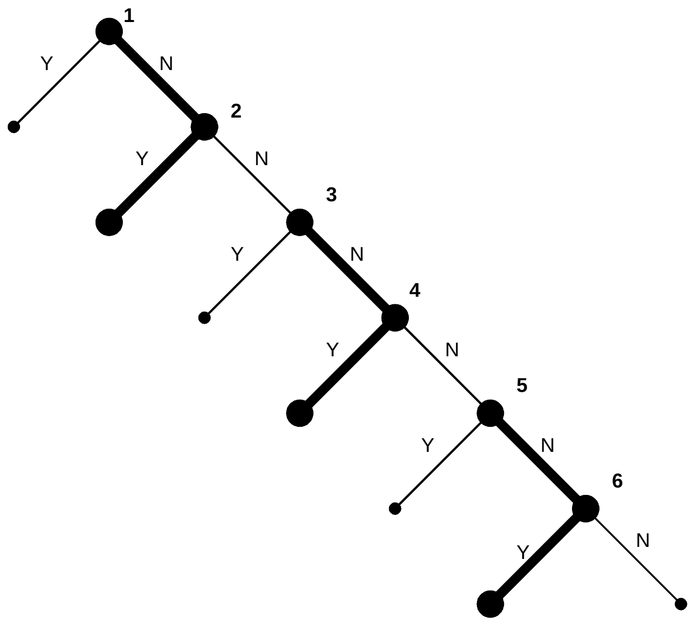

Figure 2: A possible PPS when the task cannot be completed after period 6.

Figure 2 provides a possible equilibrium (the thick lines are the equilibrium strategies). In this case, $T = 6$, and the unique PPS results in the task being completed by period 2.

Proposition 1 *When the task must be completed by a deadline, T, there is a unique PPS for a SH. This PPS satisfies, for some integer $\alpha > 1$,*

$$s_t^{SH} \equiv \begin{cases} Y, \text{ if } t \in \{t \geq 1 : t = T - k\alpha\}, \text{ where } k \text{ is any integer} \geq 0 \\ N, \text{ otherwise} \end{cases}$$

Proof. Given the deadline restriction, $s_t = N$ for $t > T$.

By A1, $U^T(T) > U^T(0)$. So, $s_T = Y$.

By A3, $U^{T-1}(T) > U^{T-1}(T-1)$

For any $i \geq 1$, $U^{T-i}(T) = \beta\delta^i(u(-c) + \delta u(b))$. Clearly, U^{T-i} is decreasing in i and $\lim_{i\to\infty} U^{T-i}(T) = 0$. Therefore (given A2), $\exists i^*$ s.t. $[U^{T-i^*}(T) < U^{T-i^*}(T-i^*)$ and $\forall i < i^*$,

$U^{T-i}(T) > U^{T-i}(T-i)$].

If $i^* > T-1$, then $\forall t < T$, $s_t = N$, and the unique equilibrium has been found by construction.

Alternatively, if $i^* \leq T-1$, then $\forall T - i^* < t < T$, $s_t = N$ and $s_{T-i^*} = Y$.

Repeating the above steps by backward induction, we can characterize the unique PPS of this problem, with $\alpha \equiv i^*$. ■

In Section 4, I specify a utility function and provide an explicit solution for α (Equation 2).

3.2 Without Deadline

Now we assume that $\forall t$, $a_t = \{Y, N\}$. The task can be completed in any period at price c. Since the TC is time consistent and given A1, there is a unique PPS for TC: $\forall t$, $s_t^{TC} = Y$. So, $\tau_{TC} = 1$. She will again complete it in the first period of her life. Since the NH consistently mis-predicts her future preferences, and given A3, she ends up never purchasing the product: $\forall t$, $s_t^{NH} = N$. So, $\tau_{TC} = 0$.

The interesting case is that of the SH. Since there are multiple PPSs in this case, we need a theory of equilibrium selection to predict outcomes. Sophistication is an assumption about beliefs about preferences: in any period, the individual is fully aware of her future intertemporal preferences. Yet that does not pin down her beliefs about what her future selves will actually do. I now list all possible equilibria.

3.2.1 Equilibria for Sophisticates

Recall that, under a deadline, a pure-strategy PPS involved completing the task every α periods. Even without a deadline, any pure-strategy PPS must have the same property. However, there are now α such PPSs: one in which the agent plays Y in period 1 and then waits another α periods, one in which the agent plays N in period 1 and Y in period 2 and then waits another α periods, and so on. It is important to note here that, in any period, there exists a PPS that involves a certain purchase in the next period.

Proofs of the following propositions are in the Appendix.

Proposition 2 *When there is no deadline, there are exactly α pure-strategy PPSs (α is as defined in Proposition 1). They are:*

$$1)\ s_t = \begin{cases} Y,\ \text{if } t = 1 + k\alpha, \text{ where } k \text{ is any nonnegative integer} \\ N,\ \text{otherwise} \end{cases}$$

$$2)\ s_t = \begin{cases} Y,\ \text{if } t = 2 + k\alpha, \text{ where } k \text{ is any nonnegative integer} \\ N,\ \text{otherwise} \end{cases}$$

and so on, up to:

$$\alpha)\ s_t = \begin{cases} Y,\ \text{if } t = k\alpha, \text{ where } k \text{ is any nonnegative integer} \\ N,\ \text{otherwise} \end{cases}$$

We can also solve for mixed-strategy PPSs. For an agent to play a mixed strategy in any period, she must be indifferent between completing the task today and leaving it for the future. To have the agent be similarly indifferent in each period, there must be some fixed probability with which she completes the task in every period starting with $t = 2$. This probability, ρ, is uniquely defined for the following reason: if the agent plays Y with probability $\rho' > \rho$ in any period, her previous self loses the incentive to play a mixed strategy and prefers to play N. If the agent plays Y with probability $\rho'' < \rho$ in any period, her previous self now has a strict incentive to play Y.

Proposition 3 *There are infinite fully mixed-strategy PPSs. However, there is some $\rho \in (0, 1)$ such that every fully mixed strategy PPS involves playing Y with probability ρ for all $t > 1$.*

So, we end up with a set of fully mixed-strategy PPSs that all look like $(\rho_1, \rho, \rho, \rho, ...)$. Period 1 is free to mix in any way she likes since this does not affect the choices of her future selves. Note that ρ must satisfy the following condition:

$$u(-c) + \beta\delta u(b) = \frac{\rho\beta\delta}{1 - (1 - \rho)\delta}\left[u(-c) + \delta u(b)\right]$$

The LHS is the individual's continuation utility from completing the task immediately, while the RHS is the continuation utility from not completing the task. Solving explicitly for ρ, we get[7]:

$$\rho = \frac{(1 - \delta)(u(-c) + \beta\delta u(b))}{\delta u(-c)(\beta - 1)} \tag{1}$$

There can also be partly-mixed-strategy PPSs. Among these are equilibria that follow pure strategies until some time $\hat{t}$ and then switch to fully mixed-strategies forever. These can take

[7] The comparative statics are as we would expect: ρ rises in β, δ, $u(b)$, and $u(-c)$.

the following form: follow one of the pure-strategy equilibria until some time $\hat{t}$ and switch to a strategy of ρ from time $\hat{t}$ +2. In particular, if $\hat{t} > 1$, it must have the property that $s_{\hat{t}} = Y$ or $s_{\hat{t}-1} = Y$. If neither of these conditions is satisfied, it will be the case that $s_{\hat{t}} = s_{\hat{t}-1} = N$. But since player $\hat{t}$ is indifferent, it must be that player $\hat{t} - 1$ prefers Y. Also, it is clear that one of the pure-strategy equilibria must be played in the first $\hat{t}$ periods, because if not, at least one of the first $\hat{t} - 1$ players will have an incentive to deviate from the given strategy.

3.2.2 Equilibrium Selection

We have seen above that a sophisticated hyperbolic discounter can choose from several equilibria. I now formally define semi-sophisticated and fully-sophisticated hyperbolic discounters, and show that this allows us to predict how each type of agent will behave when faced with the option to complete a task under an infinite horizon.

Definition 8 *A sophisticated hyperbolic discounter is* ***semi-sophisticated (SS)*** *if she is sophisticated about her preferences but naive about her equilibrium-reselection tendencies. In any period t, she selects her optimal equilibrium regardless of the equilibrium her previous selves played.*

A SS makes a different kind of mistake than a NH does. A SS wrongly believes that her future selves' strategy sets are merely $\{Y, N\}$ when in fact the effective strategy sets also include a choice over available equilibria. So, in each period, she is indeed playing an equilibrium, but she fails to recognize that the might be mis-predicting future behavior (since the next period self could play according to a different equilibrium).

A fully-sophisticated hyperbolic discounter, on the other hand, chooses from a smaller set of equilibria. I apply the definitions of renegotiation-proofness developed for repeated games by Farrell and Maskin (1989), and for hyperbolic discounters by Kocherlakota (1996).

Definition 9 *A SPNE is* ***weakly renegotiation-proof*** *if it contains no two subgames such that equilibrium continuation strategies in one of them provides a strictly higher continuation utility than the other.*

Definition 10 *A weakly renegotiation-proof SPNE is* ***strongly renegotiation-proof*** *if none of its subgame equilibria provides a strictly lower continuation utility than a subgame equilibrium of any other weakly renegotiation-proof SPNE.*

Definition 11 *A sophisticated hyperbolic discounter is* ***fully-sophisticated (FS)*** *if she plays a strongly renegotiation-proof equilibrium.*

This means that a FS restricts herself to a set of equilibria in which each has the following property: there is no other equilibrium in this set that provides a higher continuation utility in any period. This is a reasonable description of full-sophistication. Since the agent is always aware that certain equilibria would tempt future selves to renegotiate, she eliminates those from the set that she chooses from.

Now we are in a position to make predictions about behavior. A hyperbolic discounter maximizes her utility in any equilibrium where the task is performed in the following period. A SS, being sophisticated, will be knowledgeable about the set of equilibria available to her. Among these, she will choose an optimal equilibrium. We have shown above that there are several possible equilibria that start with $s_1 = N$ and $s_2 = Y$ (and thus maximize the intertemporal utility of the period 1 agent). In period 2, however, the equilibrium played by her in period 1 is no longer optimal. She can do strictly better by playing any of the equilibria that contain $s_2 = N$ and $s_3 = Y$. Continuing this equilibrium reselection over time, this means that, in the absence of a deadline, the semi-sophisticated hyperbolic discounter will never perform the task: $\tau_{SS} = 0$.

On the other hand, a FS will only choose an equilibrium that is strongly renegotiation-proof.

Proposition 4 *The set of strictly renegotiation-proof equilibria includes all SPNE that give the agent a continuation utility of $U^t(t)$ in every period, and no other SPNE.*

Proof. First, any equilibrium strategy that gives any player t a utility higher than $U^t(t)$ is not weakly renegotiation-proof. To see this, consider a continuation strategy (that must be a SPNE itself): $(s_t, s_{t+1}, s_{t+2}, ...)$. Suppose this strategy gives player t some continuation utility $\hat{U} > U^t(t)$. Given our set of available equilibria, this can only be the case if there is some $\tilde{t}$ greater than t but less than $t+\alpha+1$ with $s_{\tilde{t}} = Y$. For player $\tilde{t}$, then, the equilibrium being played is strictly dominated by the following continuation equilibrium: $(s_{\tilde{t}} = s_t, s_{v+1} = s_{t+1}, ...)$. Since this is also a subgame equilibrium of the game, this equilibrium cannot be weakly renegotiation-proof.

This leaves us with equilibria that give each player no more than the utility she could get from completing the task in that period. Clearly, there is no equilibrium that gives any less utility either (since the player could always do better by simply completing the task immediately). So, the only equilibria that are candidates for weak renegotiation-proofness are those that yield an identical utility of $U^t(t)$ in every subgame. And in fact, all such equilibria are both weakly and strongly renegotiation-proof (since there is never a subgame of any of these equilibria that yields a higher continuation utility). The following equilibria are therefore strongly renegotiation-proof:

$(\rho_1, \rho, \rho, \rho, ...)$ where ρ is as defined previously and $\rho_1 \in [0, 1]$

$(Y, \rho_2, \rho, \rho, \rho, ...)$ where ρ is as defined previously and $\rho_2 \in [0, \rho]$ ■

We have seen here that, even though semi-sophisticates and full-sophisticates would display identical choices under a deadline, they behave very differently when the deadline is lifted. In the next subsection, I think about the welfare implications of these outcomes.

3.3 Welfare

The assessment of welfare for people with time-inconsistent preferences is fundamentally problematic since there is no natural measure of the "true" intertemporal utility function. I define welfare from the perspective of both the period 1 player and the hypothetical period "0" player. The period 1 welfare is simply the period 1 agent's continuation utility.

Welfare from the perspective of the period 0 agent, simply denoted *welfare*, is period 1's continuation utility with β set to equal 1. This notion of welfare does not privilege the present-biased preferences of the period 1 agent. It can be interpreted as the benevolent objectives of an agent's parents or the government, or even the objectives of the agent herself just before she starts making decisions.

In this model, then, welfare is maximized when the task is completed in period 1, while the hyperbolic discounter's continuation utility is maximized when the task is completed in period 2. The exponential discounter attains her optimal level of welfare regardless of deadline.

Definition 12 *For any player in period 1, let her maximized feasible continuation utility be $\bar{U}$. If the actual attained utility is U^*, her* ***utility loss*** *is defined as $L \equiv \bar{U} - U^*$.*

Definition 13 *For any player, let the maximized feasible non-hyperbolic welfare be $\bar{W}$. If the actual attained non-hyperbolic welfare is W^*, her* ***welfare loss*** *is defined as $L_W = \bar{W} - W^*$.*

For a NH, both the existence and length of deadline are crucial in determining her welfare. Her maximized utility is $U^1(2)$. However, since she will perform the task only on the day of the deadline, T, her utility loss is $L = U^1(2) - U^1(T)$. In the limit as T approaches infinity, this ratio approaches $U^1(2)$. When there is no deadline, her entire potential utility gain is lost. The same intuition holds for her welfare loss.

For a sophisticated hyperbolic discounter, the utility loss is relatively low as long as there is a deadline. The interesting point here is that, regardless of how distant the deadline is, her utility loss is bounded above by $U^1(2) - U^1(1)$. When the deadline is lifted, her beliefs about equilibrium reselection come into play. If she is fully-sophisticated, her utility loss is now exactly $U^1(2) - U^1(1)$. This is the same as the upper bound on utility loss when a deadline exists. For the same reasons, we can see that, both with and without a deadline, the fully-sophisticated welfare loss is limited.

However, when the agent is only semi-sophisticated and there is no deadline, her utility loss and welfare loss are the same as for a naive hyperbolic discounter. For this type of agent, the welfare loss in the limit as the deadline goes to infinity can be much smaller than the welfare loss when there is actually no deadline. A semi-sophisticated hyperbolic discounter behaves like a sophisticate as long as there is a deadline, however far, but behaves like a naif as soon as the deadline is lifted.

4 Government Provision

In the next 3 sections, I study the implications of semi-sophistication in different market settings. Before proceeding, it is useful to make some further assumptions. Suppose there is a growing population that consists only of semi-sophisticated hyperbolic discounters. In each period, g infinitely-lived agents are born. Let us also simplify the utility function to $u(y) = y$. This does not affect the results of the following sections. Now, if the good is bought in any period τ at price p, $u_t = -p + \beta\delta b$ and $u_{\tau+1} = b$. For a good bought in period τ at price p, the welfare is $W(\tau, p) = \delta^{\tau-1}[-p + \delta b]$ and period 1's continuation utility is:

$$U^1(\tau, p) = \begin{cases} [-p + \beta\delta b], \text{ if } \tau = 1 \\ \beta\delta^{\tau-1}[-p + \delta b], \text{ if } \tau > 1 \end{cases}$$

Assume there is a well-meaning government that wishes to maximize total individual welfare by providing these products at marginal cost. The government discounts future generations according to some discount rate, $\delta_g \in (0, 1)$. *Total welfare* is the discounted sum of individual welfare (welfare for all individuals born on day 1, plus δ_g discounted welfare for all individuals born on day 2, and so on). *Total utility* is the discounted sum of individual continuation utility. Assume the government can choose when to make the product available, but cannot discriminate across individuals. In period 0, the government declares the dates in which it will offer the product.

Clearly, total utility is maximized when each individual purchases the product in the second period of her life. Even before attempting to characterize the government's optimal strategy, we can see that this will never be achieved. Similarly, total welfare is maximized when each individual buys the product in the first period of her life. This will not be achieved either.

To maximize total utility or total welfare, the good would have to be offered everyday (since the population is increasing). As we have seen in the previous analysis, the semi-sophisticated hyperbolic discounter will never buy it since she can always play a pure strategy PPS under which she postpones purchase. In any period, such a strategy involves playing Y tomorrow and every α periods from then. Since we have simplified the utility function, α can be derived explicitly as the smallest integer that satisfies:

$$-c + \beta\delta b \geq \beta\delta^{\alpha} \left[-c + \delta b\right] \tag{2}$$
$$\Rightarrow \alpha \geq \frac{\ln\left[\frac{-c+\beta\delta b}{-\beta c+\beta\delta b}\right]}{\ln \delta}$$

Since the maximized welfare and maximized utility are not attainable, we can ask what the government's best option is. What pattern of availability should the government offer to maximize social welfare? We see below that the government can do no better than with a pattern in which the good is offered every α periods. This locks semi-sophisticated hyperbolic discounters into a unique equilibrium. As a result, some individuals have to wait longer than others to buy the good.

Proposition 5 *The maximum total utility achievable by the government is given by*

$$\frac{g\left[U^1(2) + \delta_g U^1(1) + \delta_g^2 U^1(\alpha) + \delta_g^3 U^1(\alpha - 1) + ... + \delta_g^{\alpha-1}(3)\right]}{1 - \delta_g}$$

This can be achieved through any of the following patterns of offers: the good is offered in periods $2, 2+\alpha, 2+2\alpha, 2+3\alpha$ *and so on, and at any time* t*, there is some* $t' > t$ *starting at which the* *good is not offered for $\alpha - 1$ consecutive periods.*

Proof. Consider any period t in which the surviving individuals have a strict incentive to purchase it. Then, for individuals born in $t-2, t-3$, and so on until $t-(\alpha-1)$, the equilibrium that involves playing N in the period of birth (N, N, N, and so on, followed by the best possible equilibrium starting in period t) is superior to any equilibrium that involves playing Y in the period of birth. So, in any period in which an individual has a strict incentive to purchase the product, there will be no take-up for the previous $\alpha - 1$ periods. Suppose instead that, in the best possible equilibrium, the incentive to purchase in t is weak, not strong. Now consider an agent in $t-1$. This person's best equilibrium involves N in $t-1$ and Y in period t. Similarly, for people born in $t-2, t-3$, and so on until $t-(\alpha-1)$, the equilibrium of playing N until period $t-1$ and then playing Y in t is better than purchasing it in their current period. Once again, there will be no purchase for $\alpha - 1$ periods before t. Finally: how frequently can we have a period in which the agent purchases the good? Clearly, this can be achieved by offering the good at least every α periods, as long as there are occasionally $(\alpha - 1)$ consecutive periods when the good is not offered. ■

The proposition above can be easily modified to deal with total welfare as well (offer starting in period 1 rather than 2). It is interesting to note that the same pattern will also allow naive hyperbolic discounters to buy the product when otherwise they would not.

This also suggests some variations on this setting that would allow the government to achieve higher levels of welfare. Clearly, if the population was not growing, the government could simply offer the good in period 1. With a growing population, this leaves out all newborns starting in period 2. However, in cases where the offer can be conditioned on birthdate, the government can immediately maximize total welfare. Secondly, in cases where it is at all feasible for the government to reduce prices below c, it would make sense to lower them enough (so that hyperbolic discounters optimally purchase the good immediately), while making up the revenue difference by taxing the entire population.

4.1 Government Provision to Full-Sophisticates: An Aside

Recall that the following equilibria are strongly renegotiation-proof: (1) $(\rho_1, \rho, \rho, \rho, ...)$ with $\rho_1 \in [0, 1]$ and (2) $(Y, \rho_2, \rho, \rho, \rho, ...)$ with $\rho_2 \in [0, \rho]$. First, we can explicitly solve for ρ:

$$-c+\beta\delta b = \frac{\rho\beta\delta}{[1-(1-\rho)\,\delta]}\,[-c+\delta b] \tag{3}$$
$$\Rightarrow \quad \rho = \frac{(1-\delta)\,(-c+\beta\delta b)}{\delta c\,(1-\beta)}$$

If fully-sophisticated agents play an equilibrium that involves Y in period 1, the government could offer the product everyday and maximize total welfare. However, it is also possible that individuals play a fully mixed strategy. Consider the following strongly renegotiation-proof PPS: $(\rho, \rho, \rho, \rho, ...)$. An interesting result that emerges from the construction of ρ is that the government might be able to raise welfare further by offering the product every α periods rather than everyday.

First suppose the good is offered everyday. The probability that an individual who has not purchased it so far will do so on any given day is ρ. Now suppose the good is offered every α periods. At any time, all remaining people will purchase the good in the next period that it is offered. Therefore, the probability that an individual who has not purchased it so far will do so on any given day is $\frac{1}{\alpha}$. The lemma below shows that if α satisfies Equation 2 with equality, then $\rho < \frac{1}{\alpha}$.

The intuition for this is the following. First, suppose that α does, in fact, satisfy Equation 2 with equality. Then:

$$-c+\beta\delta b = \beta\delta^{\alpha}\,[-c+\delta b]$$

Now suppose, under the mixed strategy, the expected number of days an individual had to wait for the good to be purchased was exactly α. Then, his continuation utility from waiting would be greater than $\delta^{\alpha}\,[-c+\delta b]$, because the probability of buying it early is overweighed (due to discounting) relative to the probability of buying it late. But then this means that the RHS of Equation 3 is greater than the LHS – waiting is strictly better than purchasing now. So this mixing strategy cannot be a PPS. Therefore, the only acceptable value of ρ must be one that leads to an expected wait that is longer than α.

Lemma 14 *If α satisfies Equation 2 with equality and ρ satisfies Equation 3, then $\rho < \frac{1}{\alpha}$.*

Proof. Since α satisfies Equation 2 with equality,

$$\begin{aligned} -c+\beta\delta b &= \frac{\rho\beta\delta\left[-c+\delta b\right]}{\left[1-(1-\rho)\,\delta\right]}=\beta\delta^{\alpha}\left[-c+\delta b\right] \\ &\Rightarrow \quad \rho=\frac{\delta^{\alpha-1}-\delta^{\alpha}}{1-\delta^{\alpha}} \end{aligned}$$

Since $\delta<1$, ρ is increasing in δ. Applying L'Hopital's rule, we get $\lim_{\delta=1}\rho=\frac{1}{\alpha}$. Therefore, $\rho<\frac{1}{\alpha}$ for $\delta<1$. ■

The reason this intuition does not always hold is that time is discrete, and therefore α might be larger than the non-integer value that satisfies Equation 2 with equality. If α is sufficiently large, the mixed strategy equilibrium could result in more frequent purchases than a staggered provision.

5 Monopolist Seller

Consider the case of an infinitely-lived monopolist or some infinitely-lived oligopolists. One can think of a setting like a shopping mall. Assuming the stores in the mall sell similar items (winter coats), and the mall is the only place one can buy coats, it might be possible for the firms to behave strategically and sell to all agents. Suppose, first, that there is just one monopolist that has a discount factor of δ_f. This monopolist must post a menu of prices in period 0: $(p_1,p_2,p_3,...)$. As we have seen before, if the firm offers the good at a fixed price above marginal cost in each period, then semi-sophisticates will never buy the good. To encourage agents to buy the product, the firm will have to occasionally raise the prices to a prohibitively high level.

This leads us to a new rationalization of sales. Among existing theories, Sobel (1984) proposes a model where each firm has monopoly power over a group of impatient buyers who value the product highly, and all firms compete for access to a set of patient buyers with lower valuations of the product. In this model, firms keep prices high as long as they can make higher profits from their core clients than they can by selling to the patient agents. Once enough patient clients develop, prices are lowered.

In our model, price variation can be justified not as a form of price discrimination, but as a way to generate deadlines. These deadlines come in the form of high prices, and help to lock semi-sophisticated hyperbolic discounters into a unique equilibrium. While the actual pattern of prices offered by a firm will depend on the tradeoff between frequency of sales and price at which sales are made (higher price at time of purchase means that these sale periods must be more staggered), we can see that there is always a credible pricing rule that produces strictly higher profits than flat pricing.

Definition 15 *A stream of prices is* ***credible*** *if, at any time t, given the future stream of prices, the firm has no incentive to change the current price.*

Proposition 6 *With a growing population of semi-sophisticated buyers, there is always a credible pricing strategy with sales that is more profitable for the firm than any strategy with constant prices.*

Proof. First, it is easy to see that a firm will have zero profits with constant prices. If the firm sets a price p in each period, it must be the case that $c \leq p \leq \beta\delta b$. By A1,

$$\beta\delta[-c+\delta b] > -c+\beta\delta b$$
$$\Rightarrow [c][1-\beta\delta] > b[\beta\delta-\beta\delta^2]$$

For any $p > c$, the condition above always holds at any price the firm sets. So, for any fixed price strategy, the optimal SPNE for the buyer will involve purchasing the good in the next period. Therefore, the good will never be bought. Regardless of the price the firm sets, it will always have zero profits.

Second, we can show that regardless of δ_f, there is at least one pricing strategy with sales that yields positive profits and is credible. Consider the following pricing rule: the good is offered at price p every α periods (where α is defined in Equation 2) and at $\bar{p}$ in all other periods. Let $\bar{p} > \beta\delta b$, so an agent would never purchase at that price. Let $p = c+\varepsilon$, where ε satisfies:

$$-(c+\varepsilon)+\beta\delta b = \beta\delta^{\alpha}\left[-(c+\varepsilon)+\delta b\right] \tag{4}$$

By A2, it must be the case that $\varepsilon > 0$. Since Equation 4 is satisfied, the agent will always buy the good when it is offered at the low price.

These prices are always credible. The firm has no incentive to lower prices in high price periods since agents will not buy anyway. It has no incentives to lower prices in sale periods since all agents are anyway buying at $c+\varepsilon$. Finally, it has no incentive to raise prices in sale periods since this will break Equation 4 and all consumers will choose to wait until the next sale period. ■

6 Competitive Markets

Consider m infinitely-lived competitive firms, each able to produce the product at a constant marginal cost, c. In period 0, each firm i posts a menu of prices $\{p_t^i\}$. The firms have a discount rate of δ_f. At any time t, the firm's discounted value of profits is given by:

$$\Pi_t^i = \pi_t^i + \delta_f \pi_{t+1}^i + \delta_f^2 \pi_{t+2}^i + ...$$

The timing is as described before. After the firms publicize their respective menus of prices, each individual plays according to her optimal equilibrium in each period. Observe that the relevant set of prices for the consumer is simply the lowest posted price for each period t. Let us denote this menu of prices $\{p_t\}$. A competitive equilibrium is defined as a set of prices with the following property: there is no period t in which any firm is able to strictly increase Π_t^i by deviating from its proposed price. Note that any firm that posts $p_t^i > p_t$ will sell 0 units in period t. Suppose in some period t, that q firms each set some price p_t. Then, each firm's expected profit from that period is given by the total expected profits divided by p.

First, it is clear that firms cannot make profits in competitive equilibrium. If there is any period t in which a firm makes a profit by selling with positive probability at $p_t > c$, every other firm has an incentive to sell at some $p_t - \varepsilon$.

Proposition 7 *(1) There is always a competitive equilibrium in which semi-sophisticates never buy the good. (1) There is also always a competitive equilibrium in which semi-sophisticates buy the good every α periods (as defined in Equation 2) at price c.*

Proof. (1) Consider a pricing strategy $p_t^i = c$ for every firm for every period t. This is a competitive equilibrium, since no firm can raise profits by changing its price in any period. As proved before, semi-sophisticates will never purchase the good under such a pricing strategy.

(2) Consider a pricing strategy $p_t^i = c$ for every firm in $t \in (1, 1+\alpha, 1+2\alpha, ...)$ and $p_t^i > \beta\delta b$ in all other periods. This is a competitive equilibrium. No firm can make a sale by lowering prices in a high price period. No firm can make positive profits by changing prices in a low price period. ■

This proposition shows us that competition can lead to both the worst and the best possible outcomes. In one equilibrium, no individual ever buys the good. In another equilibrium, total welfare is equal to the highest total welfare under government provision. While this leads to an ambiguous prediction about the relative merits of monopoly and competition, it is nevertheless interesting to note that competitive equilibrium can lead to strictly lower total welfare than a monopolistic market.

This comparison becomes more favorable to monopoly if the population also contains a small fraction of exponential discounters. As long as this fraction is sufficiently small, the monopolist will still find it in its interest to sell to semi-sophisticates by using deadlines. However, regardless of the fraction of exponential discounters, there will no longer be a competitive equilibrium in which semi-sophisticates buy. To see this, suppose μ fraction of the population are exponential discounters. In competitive equilibrium, if there is some period in which a semi-sophisticate purchases the good with a positive probability, then there must be a period t in which the lowest market price is $p_t > c$. Then, in period t, a firm can strictly raise its profits by lowering its price by some ε that satisfies: $\mu g\left(p_t - \varepsilon - c\right) > \frac{\mu g(p_t - c)}{m}$. So prices must collapse to c in each period.

7 Conclusion

In this paper, I have proposed a new way to classify sophisticated hyperbolic discounters based on their beliefs about equilibrium selection in the future. I argue that semi-sophistication is a reasonable way to think about individuals who are aware of their preferences but still make mistakes. By investigating the behavior of semi-sophisticates when they must purchase products that have immediate costs and delayed benefits, I have shown that deadline effects can be very strong. This also provides an alternate theory of sale pricing of durable goods. Finally, we have seen that it is possible for competitive equilibria to lead to lower welfare than markets with a monopolist.

There are some natural directions for continued research. In my model, I have assumed homogenous preferences. It would be useful to carefully study market equilibria with multiple types of individuals – varying the discount rates as well as the levels of time-inconsistency. The model contains some stark predictions about the takeup of commitment devices under different settings, and it would be instructive to test these predictions in the field. Also, there might be interesting insights to be gained from a model that allows people to learn from their mistakes and move from semi-sophistication to full-sophistication.

8 Appendix

Proposition 2: When there is no deadline, there are exactly α pure-strategy PPSs (α is as defined in Proposition 1). They are:

1) $s_t = \{$ Y, if $t = 1 + k\alpha$, where k is any nonnegative integer
N, otherwise

2) $s_t = \{$ Y, if $t = 2 + k\alpha$, where k is any nonnegative integer
N, otherwise

and so on, up to:

α) $s_t = \{$ Y, if $t = k\alpha$, where k is any nonnegative integer
N, otherwise

Proof. Because of the way α is defined, it is clear that all of the above are SPNE.

Now, suppose there is another pure-strategy PPS. For it to be distinct from the PPSs above, it must have the property that either:

a) $\exists t$ s.t. $s_t = s_{t+1} = s_{t+2} = \ . \ . \ . = s_{t+\alpha} = N$

or b) $\exists t, t'$ s.t. $0 < t' - t < \alpha$ and $s_t = s_{t'} = Y$

If (a) is the case, it cannot be a PPS because in period t she can do better by choosing $s_t = Y$. If (b) is the case, it cannot be a PPS because in period t she can do better with $s_t = N$. Therefore there are exactly α pure-strategy PPSs. ■

Proposition 3: There are infinite fully mixed-strategy PPSs. However, there is some $\rho \in (0, 1)$ such that every fully mixed strategy PPS involves playing Y with probability ρ for all $t > 1$.

Proof. In any mixed-strategy equilibrium, it must be the case that at some time t, the player is indifferent between Y and N. Given A2, this is possible only if the next time Y is played with a positive probability, it is not played with certainty. And this is only possible if the next time Y is played with a positive probability, it is again not played with certainty. And so on. So, any mixed strategy equilibrium must have some period t after which Y is never played with certainty. Specifically, the condition for indifference is:

$$\begin{aligned} & u(-c) + \beta\delta u(b) \\ = \; & \rho_2\beta\delta[u(-c) + \delta u(b)] + (1-\rho_2)\rho_3\beta\delta^2[u(-c) + \delta u(b)] \\ & + (1-\rho_2)(1-\rho_3)\rho_4\beta\delta^3[u(-c) + \delta u(b)] + ... \end{aligned}$$

If $\rho_1 = \rho_2 = \rho_3 = ... = \rho$, then the condition is:

$$u(-c) + \beta\delta u(b) = \frac{\rho\beta\delta\left[u(-c) + \delta u(b)\right]}{1-(1-\rho)\delta}$$

By A1 and A3, there must exist a value of ρ such that the above equality is satisfied. Can there be other kinds of fully mixed-strategy equilibria? The second player's indifference condition is:

$$\begin{aligned} & u(-c) + \beta\delta u(b) \\ = \; & \rho_3\beta\delta[u(-c) + \delta u(b)] + (1-\rho_3)\rho_4\beta\delta^2[u(-c) + \delta u(b)] \\ & + (1-\rho_3)(1-\rho_4)\rho_5\beta\delta^3[u(-c) + \delta u(b)] + ... \end{aligned}$$

Equating the right-hand sides of the first and second players' indifference conditions yields:

$$u(-c) + \beta\delta u(b) = \frac{\rho_2\delta\left[u(-c) + \delta u(b)\right]}{1-(1-\rho_2)\delta}$$

Applying the same conditions to future players, we get:

$$\begin{aligned} & u(-c) + \beta\delta u(b) \\ = \; & \frac{\rho_2\delta\left[u(-c) + \delta u(b)\right]}{1-(1-\rho_2)\delta} \\ = \; & \frac{\rho_3\delta\left[u(-c) + \delta u(b)\right]}{1-(1-\rho_3)\delta} \\ = \; & \frac{\rho_4\delta\left[u(-c) + \delta u(b)\right]}{1-(1-\rho_4)\delta} = ... \end{aligned}$$

This determines a unique $\rho_2 = \rho_3 = \rho_4 = ... = \rho$. Therefore, if a player at time t is indifferent, then it must be that all future players choose Y with probability ρ. The fully mixed-strategy equilibria of this game must be of the kind $(\rho_1, \rho, \rho, \rho, ...)$ where $\rho_1 \in (0,1)$ and ρ is determined uniquely as shown above. ■

References

[1] AMBEC, S., and TREICH, N. (2007), "Roscas as Financial Agreements to Cope with Self-Control Problems", *Journal of Development Economics*, 82 (1), 120-137.

[2] ANDERSON, S., and BALAND, J. (2002), "The Economics of Roscas and Intrahousehold Resource Allocation", *Quarterly Journal of Economics*, 117 (3), 963-995.

[3] ASHRAF, N., GONS, N., KARLAN, D., and YIN, W. (2003), "A Review of Commitment Savings Products in Developing Countries", *Asian Development Bank Economics and Research Department Working Paper #45.*

[4] ASHRAF, N., KARLAN, D., and YIN, W. (2006), "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines", *Quarterly Journal of Economics* 121(2), 635-672.

[5] BASU, K. (2008a), "Hyperbolic Discounting and the Sustainability of Rotating Savings Arrangements", (mimeo, University of Chicago)

[6] BASU, K. (2008b), "The Provision of Commitment in Informa Banking Markets: Implications for Takeup and Welfare", (mimeo, University of Chicago)

[7] BAUER, M., CHYTILOVA., and MORDUCH, J (2008), "Behavioral Foundations of Microcredit: Experimental and Survey Evidence from Rural India", (mimeo, NYU).

[8] BESLEY, T., COATE, S., and LOURY, G. (1993), "The Economics of Rotating Savings and Credit Associations", *American Economic Review*, 83, 792-810.

[9] BERTRAND, M., KARLAN, D., MULLAINATHAN, S., SHAFIR, E., and ZINMAN, J. (2005), "What's Psychology Worth? A Field Experiment in the Consumer Credit Market", (mimeo).

[10] BERTRAND, M., MULLAINATHAN, S., and SHAFIR, E. (2004), "A Behavioral Economics View of Poverty", *American Economic Review Papers and Proceedings*, 419-423.

[11] CHOI, J., LAIBSON, D., MADRIAN, B. and METRICK, A. (2004), "For Better or For Worse: Default Effects and 401(k) Savings Behavior", in David Wise (ed), **Perspectives in the Economics of Aging**, The University of Chicago Press.

[12] CIALDINI, R., (2001), "The Science of Persuasion", *Scientific American*, 284 (2).

[13] CONNING, J. and UDRY, C. (2007), "Rural Financial Markets", in *Handbook of Agricultural Economics, Vol. 3, Agricultural Development* by Evenson, R., Pingali P., and Schultz T. (eds.), North-Holland.

[14] DAGNELIE, O. and LeMAY, P. (2007), "Rosca Participation in Benin: a Commitment Issue", (mimeo).

[15] DELLAVIGNA, S. and MALMENDIER, U. (2004), "Contract Design and Self-Control: Theory and Evidence", *Quarterly Journal of Economics*, 119(2), 353-402.

[16] FARRELL, J. and MASKIN, E. (1989), "Renegotiation in repeated games", *Games and Economic Behavior*, 1(4), 327-360.

[17] GUGERTY, M. (2007), "You Can't Save Alone: Commitment in Rotating Savings and Credit Associations in Kenya", *Economic Development and Cultural Change*, 55, 251-282.

[18] HARRIS, C. and LAIBSON, D. (2001), "Dynamic Choices of Hyperbolic Consumers", *Econometrica*, 69(4), 935-957.

[19] KOCHERLAKOTA, N. (1996) "Reconsideration-Proofness: A Refinement for Infinite Horizon Time Inconsistency", *Games and Economic Behavior*, 15, 33-54.

[20] KRUSELL, P., and SMITH, A. (2003), "Consumption-Savings Decisions with Quasi-Geometric Discounting", *Econometrica*, 71, 365-375.

[21] LAIBSON, D. (1997), "Golden Eggs and Hyperbolic Discounting", *Quarterly Journal of Economics*, 62, 443-479.

[22] LOEWENSTEIN, G. and PRELEC, D. (1992), "Anomalies in Intertemporal Choice: Evidence and an Interpretation", *Quarterly Journal of Economics*, 107 (2), 573-597.

[23] MORDUCH, J. (1999), "The Microfinance Promise", *Journal of Economic Literature*, 37, 1569-1614.

[24] MULLAINATHAN, S. (2004), "Development Economics Through the Lens of Psychology", (mimeo).

[25] O'DONOGHUE, T., and RABIN, M. (1999), "Doing it Now or Later", *American Economic Review*, 89, 103-124

[26] O'DONOGHUE, T., and RABIN, M. (2001a), "Choice and Procrastination", *Quarterly Journal of Economics*, 116(1), 121-160.

[27] O'DONOGHUE, T., and RABIN, M. (2001b), "Self Awareness and Self Control", in Roy Baumeister, George Loewenstein, and Daniel Read (eds), **Now or Later: Economic and Psychological Perspectives on Intertemporal Choice**, Russell Sage Foundation Press (forthcoming).

[28] PHELPS, E., and POLLACK, R. (1968), "On Second-Best National Saving and Game-Equilibrium Growth", *Review of Economic Studies*, 35, 185-199.

[29] SOBEL, J. (1984), "The Timing of Sales", *Review of Economic Studies*, 51(3), 353-368.

[30] THALER, R. and BERNARTZI, S. (2004), "Save More Tomorrow: Using Behavorial Economics in Increase Employee Savings." *Journal of Political Economy*, 112, 164-187.