Past and Future: Backward and Forward Discounting

Debraj Ray, Nikhil Vellodi and Ruqu Wang[2]

November 2022

Abstract. We study a model of time preference in which both current consumption and the memory of past consumption enter "experienced utility" — or the felicity — of an individual. An individual derives *overall* utility from her own felicity and the anticipated felicities of future selves. These postulates permit an agent to anticipate future regret in current decisions, and generate a set of novel testable implications in line with empirical evidence. The model can be applied to disparate phenomena, including present bias, equilibrium savings behavior, anticipation of regret, and career concerns.

1. Introduction

We study time preferences in which agents discount their payoffs backwards as well as forwards, deriving their overall utility from past, present and future. Visually, they live under a "payoff umbrella" that moves with their own temporal location. This idea, that agents might consider past consumption in evaluating current utility, goes back to the seminal work of Strotz (1956).[1] Yet, whether the past is incorporated into utility or not, it is sunk. It could generate regret or happiness

[2]Ray: NYU and University of Warwick; Vellodi: PSE; Wang: Queen's University. Ray developed an early version of this idea under a Guggenheim Fellowship, 1997-98, and acknowledges support under NSF Grants SES-1629370 and SES-1851758. Vellodi acknowledges support from the EUR grant ANR-17-EURE-0001 and the hospitality of the Simons Institute for the Theory of Computing at UC Berkeley. Wang acknowledges support from SSHRC of Canada. We thank Simone Galperti, Matt Jackson, George Loewenstein, Gregor Smith, Ennio Stacchetti, Leeat Yariv and various seminar participants for useful comments.

[1]In an earlier incarnation of the current paper, Ray and Wang (2000) based a theory of "parental influence" on the notion of memory utility, arguing that an earlier self might take such "anticipated memories" into account in her own intertemporal decision-making. Caplin and Leahy (2004) take a planner's perspective to the presence of memory, arguing that it must affect the determination of the social discount rate. Gilboa, Postlewaite and Samuelson (2016) develop a theory of "memorable consumption," wherein individuals consume some goods in a lumpy way that can explained by the anticipation of "the future utility flows generated by memorable consumption goods."

without affecting any actions. While we might enjoy (or discount) our memories, we could still just maximize "forward discounted utility." But matters are more delicate when it comes to *anticipated* future feelings of happiness or regret. An agent at date t could appreciate that her future self at date $T > t$ might view events differently. Indeed, by "backward discounting," a payoff at T will be given more weight (by the T-self) than the same payoff at t. If the current self values her future self, she will take such feelings into account, even if they enter quite dimly into her felicity today. For instance, an agent might work hard for her retirement, knowing well that her "retirement self" will look back on her life with satisfaction; indeed, discounting the current phase of hard effort. The promise of a future self enjoying happy memories might provide added incentive for planning a holiday. The fear of regret upon becoming addicted might lead to the foregoing of addictive pleasures today, and so on.

Similar ideas extend to interactions across *distinct* individuals. Parents might urge their children to work harder, knowing that once grown, children will value their adult outcomes more (and their juvenile escapades less) than they do now. A money manager might engage in more future-oriented behavior than that warranted by the discount factor of the individual who hires her, knowing that future feelings of satisfaction might elicit a positive recommendation. Or a government might engage in populist expenditure in the run-up to an election, anticipating that the public will recall recent experiences with greater clarity.

We study such interactive situations within and across individuals by allowing for individuals to derive utility both from their own felicity as well as the felicities of others. In some of these situations we study the same individual as she moves through time, internalizing her "anticipated memories" via a concern for her future self. We are particularly interested in weight placed on a distinguished future self — a *retirement* or *terminal self* who looks back on her life and derives happiness or regret from it. In other situations the individuals concerned are separate entities, but possess intertwined utilities, either because of direct empathy (or envy), or because the experienced utility of another might generate payoff-relevant actions for the individual.[2] We specifically seek to explore how the *discounting* of the past might affect the actions of the present.

To this end, Section 2 introduces a model of interaction across different selves. Our emphasis is firmly on the twin discounting of past and future, running as they do in opposite directions, and on the particular features generated by this formulation. In this sense we are both less general than, but go beyond, the usual theory of individuals connected by interdependent utilities, as developed by Strotz (1956), Pearce (1983, 2008), Hori and Kanaya (1989), and several others. For instance, if Asha, who discounts forward, cares about a future Asha who discounts backward, then her "effective" discount factor is *never* geometric, even though she might possess an intrinsic discount factor that is geometric both forward and backward.

Section 3 takes up these and other implications. While our departure from the standard theory of geometric discounting is deliberately minimal, it generates several testable predictions regarding rates of time preference; see Proposition 1. Some of these predictions are shared with hyperbolic discounting models (see, e.g., Ainslie 1991, Loewenstein and Prelec 1992, Loewenstein and Thaler 1989, Laibson 1997, 1998, and O'Donoghue and Rabin 1999). It is also in line with models of collective decision-making with time-inconsistency, such as Gollier and Weitzman (2010), Zuber (2011), and Jackson and Yariv (2015). But Proposition 1 generates additional predictions that do not arise with hyperbolic discounting. For instance, it predicts a profile of smaller local discount rates in middle age, relative to both the early adult and final phases, an observation that finds empirical support; see, e.g., Harrison et al (2002) and Read and Read (2004). We detail these findings, and

[2]The placement of weight on the utility of others is often viewed as pure or non-paternalistic altruism. Barro (1978) and Loury (1981) interpreted value functions in growth settings as the lifetime utility of future generations, with investments interpreted as bequests. Ray (1987) took these ideas to multiple descendants, an exercise which necessitated a game-theoretic rather than an optimization-based approach. Pearce (1983, 2008), Hori and Kanaya (1989) and Bergstrom (1999) study altruistic interactions across individuals not necessarily located on a timeline, providing conditions under which those interactions could be reduced to a representation on consumption streams. Ray ⓡ Vohra (2020) and Vásquez and Weretka (2020) study the equilibrium interactions of individuals connected through interdependent utility. Galperti and Strulovici (2017) provide axiomatic foundations for such "pure" utility interdependence, in the sense of being defined "on the total utility (rather than the mere consumption utility) of future generations."

refer to empirical and experimental studies that provide evidence to support our claims.

The remainder of the paper develops a host of applications, ranging from intra- to inter-personal interactions, and across a variety of seemingly disparate economic settings. Taken together, they showcase the flexibility and portability of the backward discounting framework.

Section 4 nests these preferences into an otherwise standard life-cycle consumption problem. We study outcomes both when agents can commit to consumption plans ahead of time and when such commitment is impossible. With commitment, plans reflect current preferences sharply; with sufficient weight placed on future selves, agents burden their middle-aged selves with the greatest saving rate, relative to the classical forward-looking benchmark. Without commitment, sophisticated agents strategically anticipate the decisions of their future selves and engage in higher rates of saving earlier in life. To solve the equilibrium problem, we employ a limiting argument similar to that in Björk et al. (2017) in which we allow our agent to commit to consumption plans over vanishingly small time intervals. We develop these arguments in Propositions 2–5.

Section 5 discusses *inter*personal interactions. Section 5.1 discusses a famous anti-smoking campaign that was deliberately aimed at making individuals more aware of their future selves, asking them to incorporate the possibility of future regret into their current decision-making. Section 5.2 notes that under backward discounting, a focus on end states could be incorporated into welfare functions, a point that has been made by Caplin and Leahy (2004). Section 5.3 discusses situations in which end states are not intrinsic welfare objectives but are instrumental goals. These typically arise when one agent has a say in decision-making on behalf of another, and conflict arises even though both agents fundamentally have the same preferences. We consider the example of a financial advisor who is investing and paying out dividends for current consumption on behalf of a client, and who will receive a recommendation from that client at the end of the interaction. Should the advisor maximize the present discounted value of payoffs for her client, or the end state value? The two exercises are very different. This section also considers a related example: parenting. Our theory provides a parsimonious explanation of generational conflict: a parent may want her child to be more forward-looking because she, in effect, is maximizing the utility of her child's *adult* self.

Finally, Section 5.4 develops an alternative interpretation of backward discounting as fading memory. Here, the salience of recent outcomes derives from the possibility that recently acquired information is more likely to be retained and utilized when making instrumental choices. (Stochastic forgetting is no different from backward discounting.) We embed this intuition into a simple sender-receiver setting, wherein a sender can exert hidden effort to distort the realization of signals that stochastically depend on some underlying state of nature θ and the sender's effort. At a final date, a receiver takes an action that induces a payoff for the sender that is increasing in the receiver's estimate of θ. The novel twist is that our receiver forgets signals acquired further in the past. This leads the sender to *backload* their efforts, as early attempts might be forgotten. Applications to electoral settings are immediate— a politician seeking re-election might optimally concentrate their efforts toward the voting deadline. (It is even enough for a politician to *think* that way, whether or not such a higher-order belief is indeed true.) We close with a telling exchange between ex-US president Donald Trump and journalist Bob Woodward regarding a supposed "Covid plan":[3].

TRUMP: You will see the (Covid) plan over the next 4 weeks...

WOODWARD: This is what —

TRUMP: You will see the plan, Bob. I've got 106 days. That's a long time. You know, if I put out a plan now, people won't even remember it in 100. I won the last election in the last week.

In summary, what makes our model compelling is that two simple postulates simultaneously generate a number of different findings and predictions in different economic and cultural contexts. We hope that the reader will view the current work as a first step to a deeper exploration of these matters.

[3]See https://tinyurl.com/3e9msnub

2. Intrapersonal Interaction

2.1. **Preliminaries.** A person lives in continuous time from 0 to N, and $\vec{c} = \{c_s\}_{s=0}^N$ is a consumption path over her lifetime. Let π be her instantaneous payoff function defined on consumption at every date.

Our first postulate is that each self at each date t discounts both the past and the future. Specifically, a person located at date t discounts future payoffs at rate ρ_f, and past payoffs at rate ρ_b. This forward and backward discounting system gives us the lifetime *felicity* ϕ^t of the person's date-t self:

$$\phi^t(\vec{c}) = \int_0^t e^{-\rho_b(t-s)}\pi(c_s)ds + \int_t^N e^{-\rho_f(s-t)}\pi(c_s)ds \tag{1}$$

The discounting system in (1) is "piece-wise geometric" in s, around the time location t of the current self. For example, a "current self" at date $t+a$ puts a weight of $e^{-\rho_b a}$ on $\pi(c_t)$ and 1 on $\pi(c_{t+a})$. In contrast, a "current self" at date t puts a weight of 1 on $\pi(c_t)$ and $e^{-\rho_f a}$ on $\pi(c_{t+a})$. Thus, the self at date $t+a$ may regret the consumption choices made by the self at date t. Such regret is of no functional consequence if the self at date t places no weight on the self at date $t+a$, but to the extent that such weight *is* placed and that regret anticipated, it will affect decision-making. That brings us to our second posuulate.

2.2. **Different Selves.** We will assume that the person values the preferences of "future selves." If $w(t, A)$ is the weight placed by self t on sets of agents $A \subseteq [t, N]$, then the overall *utility* U^t experienced by self t is

$$U^t(\vec{c}) = \int_{[t,N]} \phi^s(\vec{c})dw(t, ds). \tag{2}$$

We make two remarks on our formulation. First, we could have placed weight on past selves in addition to future selves. In terms of our results, it makes little difference if we do or not (see the Online Appendix).[4] Second, we could have defined utilities over the *utilities* of others rather than over their felicities, as in Pearce (1983, 2008) and Ray (1987) but following natural dampening conditions as in Hori and Kanaya (1989), Bergstrom (1999) and Pearce (2008), these reduce to the felicity-based form in (2).

[4] Additionally, the inclusion of past selves raises philosophical difficulties — for instance, if we deviate from a plan, will those selves "feel" the change? We avoid these conceptual tangles.

2.3. **The Retirement Self.** Even at this highly abstract level, it should be clear that the discounting across periods implied by (2) will not be geometric, and — to the extent such discounting isn't entirely fixed in calendar time — that such non–geometric discounting will lead to dynamic inconsistency in behavior. But we have more structure in mind. Specifically, we are going to presume that there is a special self $T \leqslant N$ — to be referred to as a "retirement" self — such that a self at any date $t \leqslant T$ places an atom of weight α on her current self, an atom β on that retirement self, and uniform, atomless "weight density" $\omega \geqslant 0$ on all other future selves. Then (2) reduces to:

$$U^t(\vec{c}) = \begin{cases} \alpha\phi^t(\vec{c}) + \beta\phi^T(\vec{c}) + \omega \int_t^N \phi^s(\vec{c})ds & \text{if } t \leqslant T \\ \alpha\phi^t(\vec{c}) + \omega \int_t^N \phi^s(\vec{c})ds & \text{if } t > T. \end{cases} \tag{3}$$

We maintain the following assumption throughout the main text:

[A] $\rho_f\alpha$ **and** $\rho_b\beta$ **both exceed** ω**.**

That is, the current and retirement selves have privileged importance — the "density equivalents" of their weights exceed those on any of the other selves.

Variations are explored in the Online Appendix. The agent might place weight only on selves between t and T, or also on past selves as mentioned above. The outcomes are qualitatively similar. We've also presumed that (α, β, ω) are unchanged as the individual ages. These could be affected, as selves "originally" in the future move progressively to the past with age. The Online Appendix discusses this case as well. But setting aside these relatively minor issues, the presumption that there is a distinguished future self for whom (or around whom) the individual's weight spikes upward, *is* important. Without this restriction the results in Proposition 1 are generally invalid; for details, see the Online Appendix.

2.4. **The Implied Discount Rate.** With the above formulation in place, an individual located at date t will experience an overall payoff of $d(t,s)\pi(c_s)$ for consumption at date s, where $d(t,s)$ is an "effective discount factor" obtained by aggregating over all valued selves for whom date s matters. Lifetime utility for the date-t self is therefore given by an obvious rewriting of (2):

$$\int_0^N d(t,s)\pi(c_s)ds. \tag{4}$$

We are particularly interested in discounting at ages $t < T$, and at all dates $s \geqslant t$ contemplated by those t-selves. Some integration shows that:

$$d(t,s) = \alpha e^{-\rho_f(s-t)} + \beta e^{-\rho_b(T-s)} + \frac{\omega}{\rho_b}\left[1 - e^{-\rho_b(N-s)}\right] + \frac{\omega}{\rho_f}\left[1 - e^{-\rho_f(s-t)}\right]$$

when $t \leqslant s \leqslant T$, and by

$$d(t,s) = \alpha e^{-\rho_f(s-t)} + \beta e^{-\rho_f(s-T)} + \frac{\omega}{\rho_b}\left[1 - e^{-\rho_b(N-s)}\right] + \frac{\omega}{\rho_f}\left[1 - e^{-\rho_f(s-t)}\right]$$

when $t \leqslant T \leqslant s \leqslant N$. (For $T < t \leq s \leq N$, simply drop the retirement self from the equation above.) These discount factors generate local rates of impatience — equivalently, local discount rates — at any instant of time from the vantage point of any other instant. Specifically,

$$i(t,s) = \lim_{\epsilon \to 0}\left[\frac{1}{\epsilon}\ln\frac{d(t,s)}{d(t,s+\epsilon)}\right] = -\frac{d_s(t,s)}{d(t,s)}$$

measures the degree of *local impatience at* s that person t feels regarding choices at date s. In the standard model, $i(t,s) = \rho_f$. We refer to $i(t,t)$ as *current impatience.*

3. Time Preference and Time Consistency

Unless the individual puts all weight *only* on her current self or only on her retirement self T, our model implies time-inconsistency in her intertemporal decisions. This is almost immediate from examining the collection of effective discount factors $d(t,s)$, which are not geometric, and additionally not fixed in calendar time. Indeed, our model generates a special case of the general class of non-geometric preferences studied by Strotz (1956).

As an example, suppose that $\rho_b = \rho_f$, $t = 30$, $T = 70$, $\alpha \in (0,1)$, and $\beta = \alpha$, with no weight on any other selves. Notice that preferences over the very near future (say from age 30 to $30 + \epsilon$) are governed largely by the preferences of the *current* self — the comparison $1 : e^{-\rho_f \epsilon}$ matters far more than the comparison $e^{-40\rho_b} : e^{-(40-\epsilon)\rho_b}$, the latter being the implied weight of the 70 year old retirement self. However, with respect to future delays — say a consumption comparison across ages 50 and $50 + \epsilon$ — the present bias exerted by the current self is increasingly compensated by the future bias exerted by the future self, and the two effects cancel, *from the vantage point of the thirty-year old.* So with $\alpha = \beta$ and $\rho_f = \rho_b$, the projected instantaneous discounting of a thirty-year old, looking twenty years into the future, is around zero.

3.1. **Some Testable Implications.** The implied bias towards current consumption relative to near-future consumption naturally invites comparison with hyperbolic discounting models (see Laibson 1997, 1998, and Harris and Laibson 2001). The well-known hyperbolic class of preferences accounts for the following empirical regularity: discounting seems to be more active for time delays that are situated in the immediate future, whereas delays situated at a more distant date are viewed more neutrally — see Thaler (1981), Ainslie (1991), Loewenstein and Prelec (1992), Laibson (1997) and O'Donoghue and Rabin (1999). Our model exhibits this phenomenon by combining forward and backward discounting, generating a system of non-geometric "effective" discount factors. But the model makes additional predictions that go beyond hyperbolic discounting:

Proposition **1.** [1] *At each non-retirement age $t < T$, local impatience $i(t, s)$ declines in future dates s over all $s \in [t, T)$;*

[2] *If $T - t$ is large enough, local impatience $i(t, s)$ at future dates s close to but smaller than T fall below 0 — that is, discounting turns negative.*

[3] *For each non-retirement age t, local impatience $i(t, s)$ jumps up as s crosses T;*

[4] *Before retirement, an individual displays less current impatience with age: $i(t, t)$ is decreasing in t. But that very same current impatience jumps up as she crosses retirement: $i(t, t)$ discontinuously increases at $t = T$.*

[5] *After retirement, current impatience grows with age, approaching ρ_f as the end of life approaches. It is approximately ρ_f after retirement, if ω is small enough.*

Part 1 of the Proposition states that from the vantage point of an unretired person at date t, relative impatience across adjacent dates in the future declines as the future is made more distant, being highest for choices between "today" and "tomorrow". This is where our model looks like hyperbolic discounting, with time-inconsistent present bias.

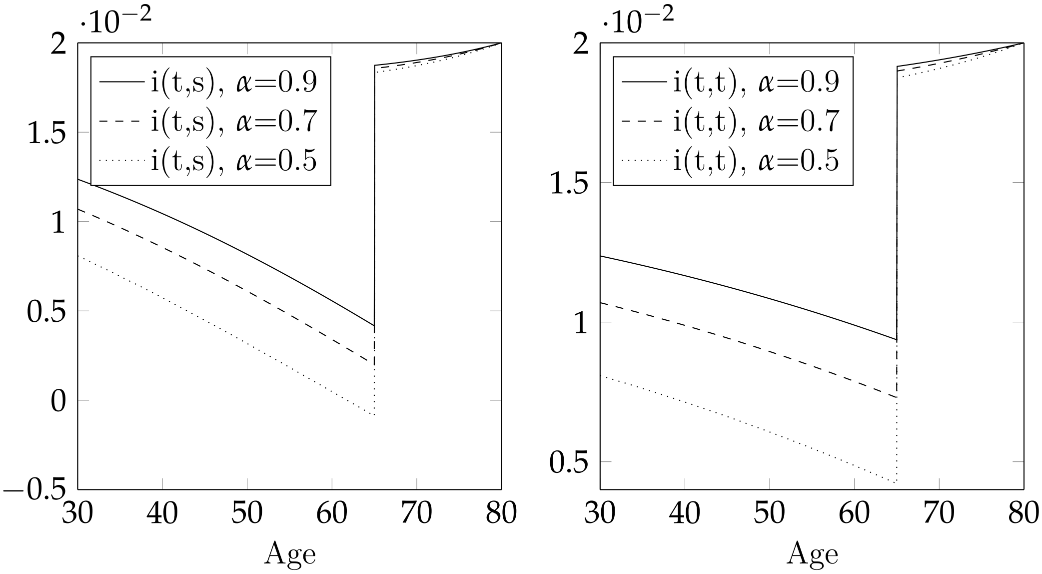

(a) $i(t, s)$, various s **(b)** $i(t, t)$, various t

Figure 1. Local and Instantaneous Rates of Impatience for $t = 30$, $\rho_f = \rho_b = 0.02$, $\beta = 0.3$, $\omega = 0.001$ and Various Values of α.

But Parts 2–5 of the Proposition go beyond hyperbolic discounting. While a person at date t is generally impatient (in a perfectly standard way) over adjacent periods in the vicinity of the present, she might actually exhibit *negative* discounting for periods well into the future. This is the assertion in Part 2 of the Proposition. From the vantage point of the present, a young person plans to make real sacrifices in her middle age to provide for her post-retirement self, but is unprepared to make those sacrifices just yet.

However, Part 3 of the Proposition states that as this future contemplated date s crosses the retirement age T, then discounting again reverts to its old ways. After all, both t and T will discount adjacent dates at s (which exceeds both T and t) in exactly the same way, which induces a discontinuity as s crosses T. A young person not only plans great sacrifices in middle age, she plans to enjoy them just after crossing retirement.

Parts 4 and 5 of the Proposition state that these youthful anticipations are only imperfectly realized. These parts concern *current* levels of impatience $i(t, t)$ for both unretired and retired individuals. The proclivity of a person to be impatient over current and immediately adjacent future choices is attenuated for ages approaching the retirement age T.

Yet T is truly a watershed age: as the individual slips over into retirement, her current impatience jumps upwards. Part 5 of the Proposition states that her instantaneous rate of impatience climbs as she "runs out" of future selves to care about. (If ω is close to zero, there are no future selves to care about at all, and her impatience will be roughly geometric and equal to the forward rate of discount ρ_f.) Figure 1 summarizes these findings, highlighting the obvious role that the weight α plays — a greater relative weight on the current self leads to greater impatience throughout life, as the countervailing force of future selves becomes less powerful.

In summary, our results in this section depend on two postulates. One is that individuals place weight on a special retirement self. The second is that that self is expected to look back in life, experiencing the afterglow (or aftershock) of her memories. The two postulates together precipitate the rich, non-geometric movements in discount rates in Proposition 1, and in particular they privilege the current self over the future self, just as in hyperbolic discounting. Specifically, by utilizing the expression obtained for $i(t,s)$ in proving the Proposition, and setting $\alpha = 1, \beta = \omega = 0$, we see that $i(t,s) = \rho_f$. The same occurs when $\rho_b \to \infty$, so that the retirement self is now only forward-looking. These observations reveal the *necessity* of both of our central behavioral postulates in delivering the patterns of local rates of impatience observed in Proposition 1.

Time-inconsistency is defined, of course, not by local discount rates but more generally by a tendency to revise a pre-planned course of action. In this more basic sense, time-inconsistency still exists even as $\rho_b \to \infty$ and backward discounting is shut down, but *only* over choices located across dates that straddle the retirement threshold. To retain the more nuanced patterns in Proposition 1, we will need ρ_b to stay away from infinity, so that the retirement self has memories. See the end of Section 4 for further discussion on the special case $\rho_b = \infty$.[5]

[5]The fact that some time-inconsistency remains even as $\rho_b \to \infty$ is in line with the findings of papers such as Ray (1987), Pearce (2008), Gollier and Weitzman (2010) and Jackson and Yariv (2015), where weights placed on standard forward-looking agents can yield a pattern of collective choices that is not time-consistent.

3.2. **Empirical Support for Time-Varying Discount Rates.** Proposition 1 uncovers a rich set of predictions delivered by our backward discounting framework. We now document how various of these find support from existing empirical studies. Our model shares with models of generalized hyperbolic discounting the prediction of diminishing local impatience, $i(t,s)$ in s, as long as $t \leq s < T$. But in addition, the backward discounting model exhibits a reversal in this trajectory when s crosses T; see part 3 of the Proposition. Also in contrast, current impatience $i(t,t)$ is constant under generalized hyperbolic discounting, while it is declining in the backward discounting model up to retirement (part 4 of Proposition 1), after which it jumps up. These are nuanced predictions: they parallel hyperbolic discounting up to a point, but suggest further tests. Our framework applies to time horizons that are longer than those typically considered under hyperbolic discounting, despite remaining analytically valid over shorter horizons.

Read and Read (2004) study discount rates in a sample drawn from a UK population. They write (p. 30–31):

> "We observed systematic but relatively complex relationships between discounting and age. The major trends were for the elderly to discount the most, and for the middle-aged to discount less than either the elderly or the young. The effects were particularly strong for discounting over long delays …The overall pattern of discounting matches that predicted by Sozou and Seymour (2003) with patience increasing until middle age, at which point it falls."[6]

In addition to finding that the young and the old discount more heavily than the middle-aged, Read and Read also document an additional result of direct relevance here. They find that the young appear to fit a hyperbolic discounting model, at least over the short-to-medium run, whereas the old fit a standard geometric discounting function. Were our agent to place no weight on future selves beyond T, or little weight on selves other than the current and retirement selves, our model would predict precisely this behavior: beyond retirement age, the conflict

[6]Sozou and Seymour (2003) present a model of discounting under uncertain mortality rates and decreasing fertility that generates such a profile. Our model, while consistent with the profile and the predictions of Sozou and Seymour, also generate the other predictions in Proposition 1, which Sozou and Seymour do not obtain — though in fairness, they did not intend to.

			Demographic			
Young < 29	Middle 30-40	Middle 41-50	Old >50	Retired	All	
Discount Rate	28.7	28.4	25.1	30.0	38.7	28.2
Annual %	(0.96)	(0.87)	(1.07)	(1.26)	(1.03)	(0.54)

Table 1. Selected estimates of annual discount rates in percentages from Harrison et al. (2002), Table 3. Standard errors in parentheses.

of interests between the different selves disappears, reducing the model to the standard geometric setting, as described in part 5 of our Proposition.

Harrison et al. (2002) investigate current impatience $i(t,t)$ in a sample of Danish individuals across a rich set of demographic indicators. They observe that retirement "is associated with a discount rate over 12 percentage points higher than otherwise," but at the same time, there is a general tendency for discount rates to come down with pre-retirement age. These results are consistent with the predictions of Part 4 of Proposition 1. Table 1 shows selected results from this study.[7] There is a decline of discount rates into middle age, rising again in old age, and with retired individuals discounting at a higher rate than the overall average.

Finally, it is worth stressing that the novelty of the predictions generated by our framework is that they directly pertain to rates of impatience, rather than patterns of behavior. To elaborate on this point, consider an alternative and entirely plausible argument based on the mapping from current actions to future consequences. When young, individuals might find it hard — possibly due to cognitive limitations or limited information — to internalize and understand the (possibly stochastic) mapping between current choices and future outcomes. As such, their choices might exhibit relative short-termism. As they age, these limitations are mitigated, and they make choices that appear more patient. Finally, towards the end of life, short-termism reappears due to a simple deadline effect (this last effect is present in our setting as well; see Section 4). Such a model would presumably generate dynamic choice paths similar to ours. But this alternative story can be separated from ours: the agents, if interviewed, would still report time-consistent

[7]Their stratified results, quoted here, control for possibly correlated covariates.

preferences with a constant rate of impatience. If a direct interrogation of rates of impatience is not methodologically acceptable, the two pathways can still be separated by offering the agent choices with sharply defined intertemporal consequences. The agent in the alternative story would respond to such well-defined choice situations by applying a geometric discount factor, which can then be estimated. An agent who internalizes backward discounting would continue to exhibit the same patterns of behavior as reported in Proposition 1.

4. The Retirement Self and Equilibrium Savings Plans

Time-inconsistency, as described in the previous section, only describes a *proclivity* towards behavior revision and not necessarily an *actual* revision of behavior. The latter will depend on whether that inconsistency is anticipated and dealt with by the agent. A sophisticated individual who understands these tensions fully might either make binding commitments for the future or anticipate such revisions in strategic behavior.

In Section 3, we saw how the combination of backward discounting and future weights give rise to a rich profile of local rates of impatience that enjoy empirical support, in particular a non-monotone preference for the future (Proposition 1 parts 2 and 3). Given these observations, how might such non-monotonicities translate into *actual* consumption decisions? We now examine this question.

We proceed by examining outcomes under two distinct notions of optimality. Under the first, agents have access to savings commitment technologies that help them implement optimal consumption plans.[8] Under the second, agents have no such commitment ability but do take the choices of future selves into account. The idea is to treat each self as a separate player in a dynamic game, and solve for the subgame perfect equilibria of this game (Strotz 1956, Phelps and Pollak 1968).

For this section, we work with a simpler formulation in which a person places weight on just two selves: α on her current self at date t, and $\beta = 1 - \alpha$ on a retirement self located at some fixed $T \leqslant N$.[9] We presume that the local payoff

[8]Many such technologies exist. See Ashraf et al (2003) for an overview.

[9]For economy in writing, we presume the agent at $t > T$ continues to value her retirement self. This is irrelevant, as backward discounting ceases to have any instrumental effect by then.

indicator is logarithmic, and that there is a common discount rate of ρ for both forward and backward discounting.[10]

We also keep the asset accumulation process as simple as possible. There is a constant interest rate of r on both borrowing and lending. In addition, the agent receives an an exogenous income stream $\{y_s\}$ over the dates $[0, N]$.[11] Let M_t be the present value of current and future income earnings at date t, i.e. $M_t = \int_t^N e^{-r(s-t)} y_s ds$. Let F_t denote financial assets at t. Then, if c_t is consumed at date t, we have $\dot{F}_t = rF_t + y_t - c_t$ and $\dot{M}_t = rM_t - y_t$. It is immediate that total asset wealth at date t — call it $A_t \equiv F_t + M_t$ — must evolve according to

$$\dot{A}_t = \dot{F}_t + \dot{M}_t = r(F_t + M_t) - c_t = rA_t - c_t. \tag{5}$$

An agent located at date t seeks to choose $\{c_s\}$ over dates in the interval $[t, N]$ to maximize the expression in (4). Noting that all decisions up to date t have already been made, this is tantamount to maximizing:

$$\alpha \int_t^N e^{-\rho(\tau - t)} \ln c_\tau d\tau + (1-\alpha) \int_t^N e^{-\rho|\tau - T|} \ln c_\tau d\tau. \tag{6}$$

4.1. **Commitment Consumption and Savings.** First, we maximize (6) evaluated at $t = 0$ by choosing a committed consumption plan over $[0, N]$. As such, we will refer to it as the *commitment* plan. The solution is summarized in the following result; see Appendix B for a detailed proof:

Proposition **2.** *The commitment consumption profile viewed from date 0 is given by*

$$c_t(A) = \left[\frac{\alpha e^{-\rho t} + (1-\alpha) e^{-\rho|T-t|}}{\alpha e^{-\rho t} a_t + (1-\alpha) p_t} \right] A \equiv \lambda_t A \tag{7}$$

for $A \geqslant 0$, where a_t and p_t in this expression are given by:

$$a_t = \rho^{-1} \left[(\rho - 1) e^{-\rho(N-t)} + 1 \right] \tag{8}$$

[10]The analysis that follows can be conducted for all constant-elasticity utility functions, for weights placed on other selves. We discuss rates that vary across forward and backward discounting later in the analysis.

[11]See Appendix A for an extension of our analysis to endogenous labor supply.

and

$$(9)\qquad p_t = \begin{cases} \rho^{-1}e^{-\rho(t-T)}[(\rho-1)e^{-\rho(N-t)}+1] & \textit{for } t > T \\ \rho^{-1}\left\{[(\rho-1)e^{-\rho(N-T)}+1]+[1-e^{-\rho(T-t)}]\right\} & \textit{for } t < T, \end{cases}$$

With the help of (8) and (9), we can explore the commitment consumption rate λ_t. There is an intrinsic tendency for this rate to drift upward simply by virtue of the finite-horizon nature of the problem. We can "benchmark" this drift by setting $\alpha = 1$ in the problem above, whereupon the situation reduces to a perfectly standard life cycle problem. Using (7) and (8) with $\alpha = 1$, we see that the consumption rate sequence $\{\bar{\lambda}_t\}$ for this benchmark problem solves

$$(10)\qquad \bar{\lambda}_t = \frac{1}{a_t}.$$

This benchmark helps to quantify the extent of our departure from the standard model. Using (7) and (10), form the commitment agent's *discrepancy ratio*

$$(11)\qquad \zeta_t \equiv \frac{\lambda_t}{\bar{\lambda}_t} = \frac{\alpha + (1-\alpha)e^{-\rho(|t-T|-t)}}{\alpha + (1-\alpha)(p_t/a_t e^{-\rho t})},$$

and observe that a value of $\zeta_t = 1$ implies that at date t there is no difference between the consumption rates predicted by the two models. That would be true of all post-retirement selves: formally, substitute the value of a_t (from (8)) and p_t (from (9)) for $t \geqslant T$ in (11) to see that $\zeta_t = 1$ for all $t \geqslant T$.

On the other hand, if the commitment discrepancy ratio is below 1, then backward discounting predicts a higher savings rate, and this effect is directly reflected in the extent to which that discrepancy ratio falls below unity. That is what happens when $t < T$. Using the observation[12]

$$(12)\qquad e^{-\rho(T-t)} < \frac{p_t}{a_t} \text{ for all } 0 \leqslant t < T,$$

and combining it with (11), we conclude that consumption ratios — including the ratio at the current date — are lowered for all periods up to the age of the retirement self. This in itself is not surprising, as the agent places some weight on the retirement self at all dates, so that her effective rate of impatience is always lower than ρ.[13] What is of greater interest is how the *extent* of that divergence

[12] This is easy enough to establish by direct computation, using (8) and (9).

[13] If the agent increases the weight on his future self, the savings rate increases at each date.

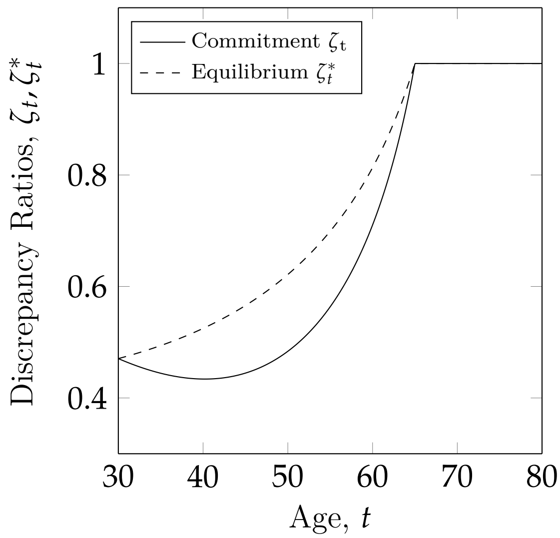

Figure 2. Commitment and equilibrium discrepancy ratios relative to the standard model, $\rho = 0.05$, $t = 30$, $N = 80$, $T = 65$, $r = 0.03$ and $\alpha = 0.5$. This graph shows the extent to which consumption ratios are depressed relative to the life-cycle model, in both the commitment and equilibrium versions of our framework.

changes with time. At what point over the agent's lifetime do we observe maximal (committed) divergence from the standard model?

Not surprisingly, the answer to this question depends on the weight that the agent attaches to the future self. If this weight is high enough, the agent plans most of her savings in the here and now, with consumption ratios rising over time relative to the standard model. Otherwise, the agent postpones the bulk of her planned savings to middle age, consuming more now and at retirement.

Proposition 3. *There exists $\hat{\alpha} \in (0, 1]$ such that if $\alpha \leqslant \hat{\alpha}$, the commitment discrepancy ratio ζ_t always increases in t; while if $1 > \alpha > \hat{\alpha}$, ζ_t first decreases and then increases in t.*

Figure 2 numerically illustrates Proposition 3. The solid line plots ζ_t for $\rho = 0.5$, $r = 0.03$ and $\alpha = 0.5$. Life starts at 30, terminates at 80, with retirement planned at 65. Then, $\hat{\alpha} \approx 0.19$, so that any positive weight below 81% on the retirement self leads to the U-shaped pattern in the discrepancy of consumption ratios. Most savings — relative to the standard model — are planned for middle age.

4.2. **Equilibrium Consumption and Savings.** We turn now to the situation in which our agent cannot commit to optimal consumption plans, and furthermore is well aware that their future selves will not in general honor plans made today. Following Strotz (1956) and Phelps and Pollak (1968), we treat each self as a separate player in a noncooperative game, and analyze its equilibria. The continuous time formulation, with a different player at every instant of time, is a convenient limit of the following natural discretization. Break the time horizon into sub-intervals of length Δ. Allow each of these intervals to be controlled by only one agent. For that short interval, each agent is in effect solving a problem similar in nature to the commitment problem described in Proposition 2. Consider an equilibrium profile thus generated and pass to the limit along a sequence of such equilibria as $\Delta \to 0$. We view the limiting profile as an equilibrium of the continuous time formulation.[14] It turns out that there is a unique limiting profile:

PROPOSITION **4.** *The equilibrium consumption profile, i.e., the unique limit of equilibria in the discretization described above, is given by*

$$c_t^*(A) = \left[\frac{\alpha + (1-\alpha)e^{-\rho|T-t|}}{\alpha a_t + (1-\alpha)p_t}\right] A \equiv \lambda_t^* A \tag{13}$$

for $A \geqslant 0$, where a_t, p_t satisfy (8) and (9).

The solutions to the a_t's and the p_t's are precisely what they were under commitment. So any difference between the commitment and equilibrium problems can be traced solely to a simple comparison of the formula in (13) with its predecessor (7). Recalling that consumption ratios in the standard model are given by $\bar{\lambda}_t = 1/a_t$, we can form an equilibrium version of the commitment discrepancy ratio by

$$\zeta_t^* \equiv \frac{\lambda_t^*}{\bar{\lambda}_t} = \frac{\alpha + (1-\alpha)e^{-\rho|t-T|}}{\alpha + (1-\alpha)(p_t/a_t)}, \tag{14}$$

For dates that are less than T, (12) continues to apply. Combining (12) with (14), it is easy to see that consumption ratios — relative to the standard model — are again lowered for all dates up to the age of the retirement self. But we can say a bit more now than we could for the commitment exercise: as long as $t < T$, the equilibrium discrepancy ratio *increases monotonically over time.* That is, the greatest impetus to the rate of savings comes at the earliest dates, no matter what

[14] This approach is similar in spirit to that taken in Björk et al. (2017).

the weight on the retirement self. To check this, it is sufficient to see that $e^{-\rho(T-t)}$ increases over time, while p_t/a_t declines, and to then apply these findings to (14).[15]

It is worth reiterating that this equilibrium push towards early savings is not necessarily what an agent might want, if she could commit to a future plan. Indeed, Proposition 3 tells us that in general, this is *not* something that the agent wants. Rather, she would like future incarnations of herself to do the bulk of the savings. But the "unravelling" effect of subgame perfection implies that the young agent internalizes the heightened propensity of their middle-aged self to save, generating the monotone trend just described. Indeed, this striking contrast between the commitment and equilibrium outcomes highlights how the two different notions of optimality can either amplify or subdue the underlying time preferences. We summarize this discussion as:

Proposition **5.** *Suppose that an agent places weights α on herself and $1-\alpha$ on a future self of fixed age T. Then*

[1] *In both the commitment and equilibrium versions of the problem, consumption rates (out of permanent income) are lower relative to those for the standard problem, at every date $t < T$. (For $t \geqslant T$, there is no such discrepancy.)*

[2] *A larger weight on the future self depresses consumption rates even further (at each date), in both the commitment and equilibrium versions.*

[3] *For $t < T$, the commitment consumption rates are smaller than their equilibrium counterparts. But the equilibrium discrepancy ratio ζ_t^* monotonically increases over time, so that the bulk of excess equilibrium savings is carried out at early dates.*

To facilitate our exposition, we've taken $\rho_f = \rho_b = \rho$. A possible price paid for this simplification is that we do not examine the independent effects of forward and backward discounting in this exercise. While a full analysis of this issue is beyond the scope of the analysis as conducted, it is easy to see the echoes of some observations made previously in Section 3. Time inconsistency is, of course, endemic when

[15]It is immediate that $e^{-\rho(T-t)}$ increases with t. To check the claim for p_t/a_t, note that this expression equals $1 + \frac{1-x(t)}{(\rho-1)Dx(t)+1}$, where $x(t) = e^{-\rho(T-t)}$ and $D = e^{-\rho(N-T)}$. It is easy to check that this last expression is a monotonically decreasing function of x, while x is itself an increasing function of t.

$\rho_f < \infty$, but still remains across binary choices located on either side of T even as $\rho_b \to \infty$. Though such comparisons played no role when discussing *local* rates of impatience (Proposition 1), in the present life-cycle setting such comparisons can lead to persistent distortions. To see this, it is readily verified (proof omitted) that in the limit as $\rho_b \to \infty$, and from the viewpoint of some anterior date s, the optimal consumption plan at some later date $t > s$ is

$$c_t(A,s) = \left[\frac{\alpha e^{-\rho(t-s)}}{\alpha e^{-\rho(t-s)} a_t + (1-\alpha)p}\right] A, \tag{15}$$

for some constant $p > 0$. Because s enters the right hand side of (15), time-inconsistency is an immediate consequence. However, almost as immediate from this expression (combined with the solution in the benchmark problem) is the observation that the commitment consumption discrepancy ratio is now *monotone* over time. Thus it is truly the combination of *finite* backward discounting (or an agent with memories) and weights placed on such agents that gives rise to the non-monotone profiles uncovered both in Propositions 1 and 3.

5. Interpersonal Interaction

Sections 3 and 4 explored the implications of our basic postulates on intra-personal interactions, showing among other things how they generate a rich profile of time preferences. We next turn to *inter*-personal interactions, and uncover how backward discounting can link decision makers located at distinct moments in time and generate patterns of behavior that speak to a range of applications. Our aim here is to showcase the ability of our backward discounting framework to organize observed phenomena across disparate settings. Common to all these applications — and indeed a central theme throughout our analysis — is the notion that agents located in the present anticipate the backward discounting of future agents and distort their behavior accordingly.

Not all the applications we list below rely on the presence of every component of our model. As already noted, a core idea is that weight is placed on two or more selves at once, but it is not always necessary for the future self to look backward at all, that is, we could set $\rho_b = \infty$. This is an important special case in which every agent has standard forward-looking payoffs, but some agents place weight on the utilities of others; see, for instance, Pearce (1983, 2008), Ray (1987, 2014), Hori and Kanaya (1989), Bergstrom (1999) and Galperti and Strulovici (2017). For instance, such a setup is enough to appreciate the points in Sections 5.1 and 5.2. But as we note below, the case $\rho_b < \infty$ further enriches our understanding of the applications in Sections 5.3 and 5.4.

5.1. **The Generation of Intrapersonal Empathy.** In 2013, Gerry Collins — a businessman, soccer coach and father of three — became the face of the Irish Health and Safety Executive's anti-smoking campaign, QUIT. He had been diagnosed with terminal lung cancer. Collins appeared in numerous advertisements, lamenting his past decisions to continue smoking. Collins died in 2014. The Health and Safety Executive estimates that over 100,000 individuals had been affected by Collins' campaigns, and had joined a QUIT program in response.[16]

We have seen how agents who place some weight on future selves might engage in more prudent behavior than standard forward discounters, foregoing current consumption to a greater degree. Turning from this positive angle to a policy-oriented one, a natural task might then be to identify contexts within which individuals could be *encouraged* to empathize with their future selves, in order to improve lifetime outcomes. The Collins campaign is a good example.

The classic reference for a theory of addiction is Becker and Murphy (2004).[17] In their model, agents rationally trade off current consumption gains against future addiction losses. Standard policy interventions in this setting involve taxes (Becker, Grossman and Murphy 1994, Gruber and Köszegi 2001). In such a context, the internalization of backward discounting via a concern for future selves would provide powerful incentives to attenuate addictive behavior, as future selves not only value current consumption less, but are manifestly harmed by the future stock of addictive consumption. As such, an intervention that could serve to complement taxes might be to launch campaigns that serve to explicitly demonstrate the regret experienced by one's future self, thus combining both the future weight and backward discounting aspects of our model.

[16] See https://tinyurl.com/5r3edsnc.

[17] In their model of memory utility, Gilboa, Postlewaite and Samuelson (2016) discuss habit formation and addictive behavior within a rational, time-consistent framework.

This is not quite the same thing as reducing individual discounting of the future. If that were the case, such a policy would also affect choices between the present and the immediate future, including choices that are presumably inconsequential for the well-being of the future self. Rather, the policy highlights the plight of the future self *per se*, along with that self's attendant memories and regrets. It is akin to increasing the weight on that self, and reversing the tradeoffs in time — the pleasures of addiction are now in the distant past, while the price is to be paid "in the present."[18] The argument here also suggests that the two channels of potential influence — an increase in patience versus a heightened regard for a privileged terminal self — can be empirically distinguished.

5.2. **End-State Satisfaction as Welfare Objective.** As Caplin and Leahy (2004) have observed, the phenomenon of backward discounting can give rise quite naturally to an argument for greater patience on the part of a social planner, compared to the discount factor of the agents in the economy. Suppose that generation τ derives lifetime felicity V_τ from a consumption stream $\vec{c}$. A social planner might want to choose $\vec{c}$ to maximize some weighted sum of these generational felicities,

$$\lambda_\tau V_\tau,$$

where the λ's are positive weights that integrate to unity. Assume that each generation exhibits a mixture of backward and forward discounting, so that:

$$V_\tau(\vec{c}) = \int_\tau^\infty e^{-\rho_f(s-\tau)}\pi(c_s)ds + \int_0^\tau e^{-\rho_b(\tau-s)}\pi(c_s)ds,$$

where ρ_f is the forward rate and ρ_b is the backward rate. Then it is easy to see that the planner will exhibit a degree of patience that exceeds the patience implied in ρ_f. The planner will also be time-inconsistent, in just the same way as an agent imbued with empathy for future selves might be; see Strotz (1956).

5.3. **End-State Satisfaction as Instrumental Goal.** The creation of intrapersonal empathy, or the valuation of end states, is not just a task for institutional policy makers. Apart from welfare evaluations, and to the extent that end-state

[18] The situation is more complex when the activity is socially harmful but nevertheless brings future rewards. Examples might include miscreant behavior, such as crime or corruption. In this case, a weight on future selves might *exacerbate* the problem — agents now foresee the future rewards promised by such activities and thus engage in the activity at increased current levels.

satisfaction can directly influence an individual's actions, other individuals who might benefit from those actions might seek to maximize memory utility rather than forward utility. That is an instrumental goal rather than an intrinsic one.

Suppose that a financial advisor manages the funds of an individual over the time span $[t, T]$, after which that individual writes a letter of recommendation. We presume that the advisor's task is to generate a consumption path $\vec{c}$ for the individual over this time period. Let $\pi(c)$ be a smooth, concave payoff function for the individual and $\Psi(A)$ some utility of terminal assets A_T. Then

$$V_t = \int_0^T e^{-\rho_f t} \pi(c_s) ds + \Psi(A_T)$$

is conventionally discounted total felicity at the start of date t.

For instance, if ρ_f equals the rate of interest r, the solution to maximizing V is simple: provide constant consumption to the individual for any feasible A_T, and then choose A_T to solve an auxiliary maximization problem. But that isn't what our financial advisor will do if the final recommendation is based on the end-state satisfaction of the individual. Rather, she would prefer to maximize

$$V_T = \int_0^T e^{-\rho_b (T-s)} \pi(c_s) ds + \Psi(A_T)$$

instead of V_t, where ρ_b is the discount rate used by the individual to evaluate her end-state payoff. This gives rise to a tension between the objective function of the individual located at date t and that of the advisor, who has a different goal. One way to resolve the tension is to think of a weight $\alpha \in (0, 1)$ which resolves this tension to some degree, so that the advisor effectively seeks to maximize

$$\alpha V_t + (1 - \alpha) V_T$$

subject to the feasibility constraints $\dot{A}_s = r A_s - c_s$ at every instant, which generates the end-point constraint $A_T = A e^{rT} - \int_0^T e^{r(t-\tau)} c_\tau d\tau$. This is an immediate application of our baseline model.

If the advisor can engage in her task unhindered, then $\alpha = 0$, and her behavior is quite different from the constant-consumption path when $\rho_f = r$. Under the same circumstances, consumption must rise over time, and both terminal consumption c_T and the value of terminal assets A_T are larger than in the forward-optimal problem. Our individual's date 0 self isn't going to like this plan of action, but his date-T self will be quite delighted at the restraint shown by his financial advisor (which will presumably translate into a splendid endorsement).

More generally, note how this effect could manifest itself in active disagreement and "advice-giving," even when there is no differential information across the advisor and her client. (In practice there will be differential information, of course, but that is not the point here.) That disagreement will be manifested not just in different objectives but in possible time-inconsistency of behavior, as the advisor initially seeks to please her client's current self, hoping to cut back on consumption at an intermediate stage in the relationship; see Propositions 2 and 3. If such commitment cannot be maintained and the advisor is sophisticated, Propositions 4 and 5 will apply instead, but in all cases with $\alpha > 0$, there will be a divergence from the client optimum at any date $t < T$.

In a similar vein, consider intergenerational conflicts between a parent and child, where a crucial tradeoff might arise between effort placed into education today and future payoffs from that education. Parent and child may have exactly the same preference *structure*, disagreeing only on account of the time period where they are located. In effect, a parent takes actions to maximize the utility of their child viewed from the vantage point of a later age, thereby discounting more heavily, say, the efforts that go into learning, while discounting more lightly the later rewards that might be reaped on the labor market. Such disagreements could simply emanate from distinct parent-child preferences, but this less parsimonious explanation begs the question of why those preferences are distinct to begin with.[19] Backward discounting provides a natural mechanism for such conflicts, without requiring assumed differences in preferences from the outset.

5.4. **Imperfect Memory and Career Concerns.** We close this section with an application that shows how our backward discounting framework can be used explicitly to model memories that fade. A sender (male) and receiver (female) interact over a time interval $[0, T]$, during which the sender lays the groundwork for signaling the value of a hidden state θ. These signals will arrive at the end of the interaction, but the sender can affect signal realizations by exerting effort throughout $[0, T]$. Such effort is hidden from the receiver, but the receiver is

[19] See, e.g., Doepke and Zilibotti (2017) on models of parenting where parents and children are assumed to have different preferences.

well aware that the sender could have been engaged in signal manipulation, and rationally accounts for such manipulation. In short, just as in a standard career concerns setting (Holmström 1999), the sender faces a "hamster wheel" problem. He attempts to fool the receiver into believing θ is higher by exerting greater effort, but the receiver correctly anticipates those efforts and subtracts them away.

At T, the receiver takes an action. We don't model this, except to presume that the sender's resulting payoff is a smooth, increasing and concave function U of the receiver's posterior estimate of θ. The sender chooses an effort schedule to maximize the expected value of U net of any effort costs.

As an example, suppose that θ measures the true quality of an incumbent politician or party. This quality is currently unknown to both the politician and to voters, but at election time (T) voters will form an estimate of it, using all the signals available at that time. The politician would like to influence these signals, and expends costly resources to do so. Voters know this, of course, and in equilibrium they will account for — and filter out — the strategic deployment of political resources. Backward discounting will make its appearance in the form of decaying voter memory, as earlier interactions fade from public consciousness, and consequently so do any signals that could have arisen from them.

Suppose that signals arrive according to a time-homogeneous Poisson process with parameter $\lambda > 0$. If a signal z_t arrives at t, it will only become visible to either agent at T,[20] with value

$$z_t = \theta + a_t + \epsilon_t, \tag{16}$$

where $\epsilon_t \sim \mathcal{N}(0, \tau_\epsilon^{-1})$ is exogenous noise injected into the observation, and a_t is the effort exerted by the sender at t. If receiver sees this signal, she will know when it arrived — it could be the memory of some incident at an earlier date t that is recalled at T. She will also know the value z_t, but what she makes of it will depend on her belief regarding sender effort at that date.

The connection to backward discounting arises from the possibility that such incidents can fade from the receiver's memory. Formally, suppose that signals are discarded according to a Poisson process with parameter ρ_b, so that if signal z was

[20]The assumption of visibility at the terminal date eliminates the complication — unnecessary in our view — that the sender's effort schedule is conditioned on historical information.

born at t, the receiver updates her prior $\mu_0 \mapsto \langle \mu_0, z \rangle$ at time T with probability $e^{-\rho_b(T-t)}$, but retains her prior with complementary probability $1 - e^{-\rho_b(T-t)}$. The older is the signal, the greater the chance it is forgotten. Notice that this specification corresponds *exactly* to the backward geometric discounting expressions used throughout the analysis.

There is a flow opportunity cost of applying effort, which can be modeled in more than one way. One possibility is that the sender has a fixed effort budget $A > 0$ and can draw down on this reserve at a maximal flow rate $\bar{a} > 0$. To avoid trivialities we presume that $\bar{a}T > A$. Then feasible effort paths satisfy:

$$\int_0^T a_t \, dt \leqslant A, \qquad a_t \in [0, \bar{a}] \text{ for each } t. \tag{17}$$

The following proposition tells us that there is a unique equilibrium, in which *all effort is fully backloaded.*

Proposition **6.** *With $\rho_b > 0$ and with a fixed effort budget, there is a unique equilibrium. Under it, effort is fully-backloaded:*

$$a_t = \begin{cases} \bar{a} & \text{if } t \geqslant \bar{t} \\ 0 & \text{if } t < \bar{t}, \end{cases} \tag{18}$$

where $\bar{t} \in (0, T)$ solves $\bar{a}(T - \bar{t}) = A$. On the other hand, if $\rho_b = 0$, any effort schedule that satisfies (17) with equality constitutes an equilibrium.

Thus all efforts by the sender *must* be concentrated at the very end of the interaction, to the extent permissible by the upper bound $\bar{a}$ on flow effort. Returning to our example, this is a particularly stark expression of the idea that a politician's re-election efforts are often bunched up against the proximate deadline of an upcoming election. It isn't that this strategy fools the voting public. Everyone in this setting understands that the politician will channel his strategic efforts in this way. The public rationally adjusts for the heightened pre-election expenditure of resources. And yet — despite the fact that all effort is fully screened out — no other configuration of effort can constitute an equilibrium.

If the backward discounting induced by fading memory is dropped from the story, so that $\rho_b = 0$, then the model as stated loses all its predictive power. *Any* allocation of effort across time, adding up to the budget A, can be an equilibrium. This is stated in the last part of Proposition 6.

A more interesting observation arises when considering the opposite extreme $\rho_b = \infty$. Now Proposition 6 comes as no surprise at all. What better can a politician do when dealing with a memoryless voter, than to cynically inject all his resources into an election boom? But the Proposition actually holds without having to rely on this stark extreme. Extreme backloading occurs even when voters have memories.

To understand this further, note that the politician can always entertain a surprise deviation from any strategy, and so the marginal gains from such a deviation must be equalized across all dates, evaluated at the putative equilibrium. But this necessary condition for a politician's best response cannot be satisfied when there is positive effort before $\bar{t}$. A reallocation of that effort into the future must yield a net gain, because more recent signals are more likely to be remembered by the voting public. This argument would of necessity remain obscured were we to assume that $\rho_b = \infty$ to begin with.

We note that the sharpness of our results could be tempered under other specifications of the model. For instance, suppose there is a cost function $c(a)$ for flow effort a at every instant of time, with $c'(a) > 0$, $c''(a) > 0$ and $c'(0) = 0$. We assume that the sender does not discount forwards over $[0, T]$.[21]

PROPOSITION **7.** *An equilibrium exists. If $\rho_b > 0$, then the sender's equilibrium effort schedule is strictly increasing over $[0, T]$. On the other hand, if $\rho_b = 0$, then the sender's equilibrium effort schedule is constant over $[0, T]$.*

6. Another Interpretation of the Retirement Self

We close the paper with some reflections on sociological foundations for our model. In particular, we speculate why some individuals might internalize their retirement self in a special way. One hypothesis is that retirement is the imprint of parental influence: that the empathy for one's retirement self is akin to providing for a parent, or the representation of one.

[21] We know already that forward discounting by the sender, together with additively time-separable costs at every date, will lead to increasing effort over time.

It is well known that a parent's influence on their children is highly significant (see, e.g., Hess and Torney (1967), Bandera (1977) and Moschis (1987), among many others). Among other things, parents have a significant impact on such outcomes as their children's choice of career (Dryler 1998), their focus on academic excellence (Salili 1994), their perception of leadership (Harris and Hartman 1992), their political attitudes (Hess and Torney 1967), and on characteristics such as home ownership (Henretta 1984). A smaller economics literature on parental influence does exists; see, e.g., Becker and Mulligan (1997), Bisin and Verdier (1998, 2000) and Doepke and Zilibotti (2017). As an example, Weinberg (2001) connects progeny behavior to the child-rearing practices of parents. Parents with higher incomes are more able to mold their children's behavior through pecuniary incentives, while parents with lower incomes have to rely on (less effective) non-pecuniary mechanisms, such as corporal punishment.

Parents certainly appear to influence individual attitudes to consumption and savings. Moschis (1987, p. 77) summarizes thus: "[T]here appears to be reasonably good supportive evidence that the family is instrumental in teaching young people basic rational aspects of consumption. It influences the development of rational consumption orientations related to a hierarchy of economic decisions delineated by previous writers ...: spending and saving, expenditure allocation, and product decisions, including some evaluative criteria."

In our opinion, this form of influence extends beyond *deliberate* attempts by parents to inculcate rudimentary notions of financial budgeting and other values in their children. The appropriate channel of influence may be more akin to what Hess and Torney (1967) have termed *anticipatory socialization*: the acquisition of attitudes and values about adult roles that have only limited relevance for the child but serves as a basis for subsequent adult behavior.[22] At the cost of some simplification, we might interpret this as stating that a particular personality is nothing more than the weighted combination of other personalities in the "cognitive neighborhood" of the individual in question. In a similar vein, the proxy parent or retirement self in this model can also be replaced by "role models." People look up to their role models, who presumably are older in most cases. Furthermore, such role models often serve as implicit adjudicators of one's life choices, and as such, their effect on

[22] As Ward (1974) describes it, such influence might be embodied in "implicit often unconscious learning for roles which will be assumed sometime in the future".

the current self might be well described by a mode of backward discounting. The weight of $1-\alpha$ on future selves then measures the influence of these role models.

7. Summary

In this paper, we introduce a framework in which lifetime individual felicities are derived from both present and past consumption streams. Each of these streams is discounted, the former forward in the usual fashion, the latter backward. We have argued that these two notions of discounting can interact in novel and interesting ways. In particular, we have shown that an individual who places weight on her current self as well as some future "retirement self" displays present bias in consumption. But there are additional implications generated by the model: an individual gets more patient as she grows older during her working lifetime, but impatience jumps up again following the age of retirement. These results survive even if the individual cares for *all* future selves, as long as there is some sufficient weight on this particular stock-taking retirement self. We apply this model to consumption-savings decisions in a standard life-cycle model, as well as discussing its application in other wide-ranging settings.

Our paper lies quite far from a general and comprehensive study of backward discounting. But we hope that it takes a small step in introducing the reader to the potentially important aspects of such a study.

Appendix A. Variable Labor Supply

In this section, we extend the analysis of Section 4 to allow for endogenous income flows through labor choice. Suppose that we interpret the income stream $\{y_t\}$ as a stream of *potential* incomes at every date. However, the actual income of the individual is determined by her choice of labor supply $\ell_t \in [0,1]$, whch generates $y_t\ell_t$ of income. Suppose that the one-period payoff function is given by

$$\ln(c_t) + \eta \ln(1-\ell_t),$$

where $\eta > 0$ is the weight put on leisure. It is easy to see that the evolution of financial assets is given by $\dot{F}_t = rF_t + y_t\ell_t - c_t$, while we now interpret $\dot{M}_t = rM_t - y_t$ as the evolution of *potential* labor income in present value terms. So total asset wealth $A_t = M_t + F_t$, now to be interpreted as the sum of financial assets and the present value of potential income, evolves according to

$$\text{(A.1)} \quad \dot{A}_t = \dot{F}_t + \dot{M}_t = r(F_t + M_t) + y_t\ell_t - c_t - y_t = rA_t - [c_t + y_t(1-\ell_t)] = rA_t - d_t,$$

where $d_t \equiv c_t + y_t(1-\ell_t)$ is a composite variable that will replace c_t. To achieve this replacement, consider the "mini-problem" of solving, at any date t,

$$\text{Maximize } \ln(c_t) + \eta \ln(1-\ell_t), \text{ subject to } d_t \equiv c_t + y_t(1-\ell_t) \text{ and } \ell \geq 0,$$

where we temporarily regard d_t as a given parameter. For ease of exposition suppose that $\ell = 0$ is never binding.[23] Standard methods show that $\eta c_t = y_t(1-\ell_t) = \eta d_t/(1+\eta)$. Therefore indirect utility at date t (call it μ) as a function of d_t is easily seen to be

$$\mu(d_t) \equiv \ln\left(\frac{d_t}{1+\eta}\right) + \eta \ln\left(\frac{\eta d_t}{(1+\eta)y_t}\right) = (1+\eta)\ln(d_t) - \ln\left(\frac{(1+\eta)^{1+\eta} y_t^\eta}{\eta^\eta}\right).$$

Apart from the extra additive term in the right hand side and the multiplication of the log function by $1+\eta$, this problem has now been reduced to the baseline problem studied by us, with only the notional replacement of c by d. The extra terms might look complicated but they make *no* difference to any form of behavior, with or without commitment. So the policy function must be of the form

$$d_t = \zeta_t A_t,$$

where in the baseline problem d_t is replaced by c_t, and ζ_t by either λ_t or λ_t^*, depending on whether we are studying the commitment solution or the equilibrium solution. *Exactly* the same coefficients as before will be employed in this extended problem, except that they will generate the composite variable $d_t = c_t + y_t(1-\ell_t)$ and not consumption c_t. This observation, along with the solution to the mini-problem, allows us to back out the labor supply dynamics as

$$\ell_t = 1 - \eta\frac{c_t}{y_t} = \frac{\eta d_t}{(1+\eta)y_t} = \frac{\eta \zeta_t A_t}{(1+\eta)y_t}.$$

The first equality tells us that in this extension, labor supply moves in the opposite direction to the consumption rate out of current income. The last inequality relates labor supply fully to the parameters of the model, by replacing consumption by its equilibrium fraction of assets, which we already know how to calculate.

[23]This *could* happen for some parameters and at some dates (not all), because even if no labor is supplied at some date, consumption could still be positive via past saving or borrowing.

Appendix B. Proofs

B.1. **Proof of Proposition 1. Part 1**. Recall that the local rate of impatience at time s, from the vantage point of date t, is given by $i(t,s) = -d_s(t,s)/d(t,s)$. Simple differentiation of d shows that for all $s \in [t, T)$,

$$i(t,s) = \left[\frac{(\rho_f \alpha - \omega)e^{-\rho_f(s-t)} - \rho_b \beta e^{-\rho_b(T-s)} + \omega e^{-\rho_b(N-s)}}{(\alpha - \frac{\omega}{\rho_f})e^{-\rho_f(s-t)} + \beta e^{-\rho_b(T-s)} - \frac{\omega}{\rho_b}e^{-\rho_b(N-s)} + \frac{\omega}{\rho_f} + \frac{\omega}{\rho_b}} \right] \tag{B.1}$$

By [A], we know that $\rho_b \beta - \omega > 0$, and hence $\rho_b \beta e^{-\rho_b(T)} - \omega e^{-\rho_b(N)} > 0$, and that $\rho_f \alpha - \omega > 0$. Hence $i(t,s)$ is a ratio of a decreasing function of s over an increasing one, and is therefore decreasing in s.

Part 2 If the conditions of the statement hold, then $e^{-\rho_f(s-t)}$ can be made arbitrarily small, while $e^{-\rho_b(T-s)}$ approaches unity. Thus, the numerator of $i(t,s)$ becomes negative, while the denominator is always positive.

Part 3 For $s \geqslant T$, note that

$$i(t,s) = \left[\frac{(\rho_f \alpha - \omega)e^{-\rho_f(s-t)} + \rho_f \beta e^{-\rho_f(s-T)} + \omega e^{-\rho_b(N-s)}}{(\alpha - \frac{\omega}{\rho_f})e^{-\rho_f(s-t)} + \beta e^{-\rho_f(s-T)} - \frac{\omega}{\rho_b}e^{-\rho_b(N-s)} + \frac{\omega}{\rho_f} + \frac{\omega}{\rho_b}} \right]$$

To see that $i(t,s)$ jumps up at $s = T$, simply compare this expression to (B.1), and note that the only difference between these expressions at $s = T$ is the second term of the numerator, which turns from being strictly negative to strictly positive.[24]

Part 4 For $t < T$, we can write $i(t,t)$ as

$$i(t,t) = \left[\frac{(\rho_f \alpha - \omega) - \rho_b \beta e^{-\rho_b(T-t)} + \omega e^{-\rho_b(N-t)}}{(\alpha - \frac{\omega}{\rho_f}) + \beta e^{-\rho_b(T-t)} - \frac{\omega}{\rho_b}e^{-\rho_b(N-t)} + \frac{\omega}{\rho_f} + \frac{\omega}{\rho_b}} \right], \tag{B.2}$$

and once again, Condition A guarantees that the numerator is decreasing in t and the denominator is increasing. For $t > T$,

$$i(t,t) = \left[\frac{(\rho_f \alpha - \omega) + \rho_f \beta e^{-\rho_f(t-T)} + \omega e^{-\rho_b(N-t)}}{(\alpha - \frac{\omega}{\rho_f}) + \beta e^{-\rho_f(t-T)} - \frac{\omega}{\rho_b}e^{-\rho_b(N-t)} + \frac{\omega}{\rho_f} + \frac{\omega}{\rho_b}} \right] \tag{B.3}$$

[24]The two denominators are generally different as well, but at $s = T$ they are both the same.

and again, we can easily compare (B.2) and (B.3) to confirm the positive jump in $i(t,t)$ as t crosses T.

Part 5 Differentiating (B.3) with respect to t, we see that $i(t,t)$ is increasing if $\rho_f(\rho_b+\omega) > \omega(\rho_b+\rho_f)e^{-\rho_b(N-t)}$, which holds by virtue of [A]. Finally, setting $\omega = 0$ yields the expression

$$i(t,t) = \left[\frac{\rho_f\alpha + \rho_f\beta e^{-\rho_f(t-T)}}{\alpha + \beta e^{-\rho_f(t-T)}}\right] = \rho_f \text{ for all } t \geq T.$$

■

B.2. **Proof of Proposition 2.** An agent located at date 0 seeks to maximize, for any date t and any asset A_t at t,

$$e^{-\rho t}\alpha\int_t^N e^{-\rho s}\ln c_s ds + (1-\alpha)\int_t^N e^{-\rho|s-T|}\ln c_s ds \tag{B.4}$$

by choosing $\{c_s\}$, given some initial asset A_t and the law of motion (5) for assets. Let

$$V(A_t,t) = \int_t^N e^{-\rho(s-t)}\ln(c_s)ds \tag{B.5}$$

and

$$W(A_t,t) = \int_t^N e^{-\rho|T-s|}\ln(c_s)ds \tag{B.6}$$

be the values generated for each of the terms in the maximand by the supremum choice of consumption plan. That is, $V(A_t,t)$ records the forward discounted utility experienced from time t onwards, while $W(A_t,t)$ records the utility experienced at time t discounted backward and forward from time T, corresponding to the supremum choice of consumption plan.[25]

We guess and verify that these functions inherit the logarithmic utility structure, i.e. there exist coefficients a_t, b_t, p_t, q_t such that

$$V(A_t,t) = a_t\ln(A_t) + b_t \tag{B.7}$$

and

$$W(A_t,t) = p_t\ln(A_t) + q_t, \tag{B.8}$$

[25] As the supremum value of the maximand is approached, it is easy to check that forward and backward utility integrals converge to well-defined limits, so that these values are well-defined.

where the coefficients a_t and p_t satisfy the differential equations:

$$\rho a_t = 1 + \dot{a}_t \tag{B.9}$$
$$\dot{p}_t = -e^{-\rho|T-t|}$$

with boundary values $a_N = 1$, $p_N = e^{-\rho(N-T)}$, and the coefficients b_t and q_t are continuously differentiable with respect to t, with $b_N = q_N = 0$.

With these guesses in hand, the time t problem (viewed from time 0) can be solved for all t. To do so, observe that the agent's supremum value from date t onward at initial asset A_t can be written as $e^{-\rho t}\alpha V(A_t, t) + (1-\alpha)W(A_t, t)$. Moreover,

$$e^{-\rho t}\alpha V(A_t,t) + (1-\alpha)W(A_t,t) = \sup_{\{c_s\}} e^{-\rho t}\alpha\Big[\int_t^{t+h} e^{-\rho(s-t)}\ln(c_s)ds + e^{-\rho h}V(A_{t+h}, t+h)\Big]$$
$$+(1-\alpha)\Big[\int_t^{t+h} e^{-\rho|T-s|}\ln(c_s)ds + W(A_{t+h}, t+h)\Big],$$

for all h such that $t + h \leqslant N$, where the law of motion for A is given by $\dot{A}_s = rA_s - c_s$ for $s \geqslant t$. It follows that

$$\sup_{\{c_s\}} e^{-\rho t}\alpha\Big[\int_t^{t+h} e^{-\rho(s-t)}\ln(c_s)ds + e^{-\rho h}V(A_{t+h}, t+h) - V(A_t, t)\Big] +$$
$$(1-\alpha)\Big[\int_t^{t+h} e^{-\rho|T-s|}\ln(c_s)ds + W(A_{t+h}, t+h) - W(A_t, t)\Big] = 0$$

for all such $h > 0$. Dividing by h and taking the limit $h \to 0$, the above expression becomes

$$0 = \sup_{c_t} \alpha e^{-\rho t}\big[\ln c_t + \dot{A}_t V_A(A_t, t) + V_t(A_t, t) - \rho V(A_t, t)\big]$$
$$+ (1-\alpha)\big[e^{-\rho|T-t|}\ln c_t + \dot{A}_t W_A(A_t, t) + W_t(A_t, t)\big]$$

where $\dot{A}_t = rA_t - c_t$ and the subscripts on V denote suitable partial derivatives. (In taking this limit, we note that $V(A_t, t)$ and $W(A_t, t)$ are both differentiable, given their conjectured forms.[26]) Combining this expression with the formulae for $V(A_t, t)$ and $W(A_t, t)$ in (B.7) and (B.8), and taking first order conditions, we see

[26]Observe that a_t and b_t have continuous derivatives. While $\dot{p}_t$ and $\dot{q}_t$ are possibly non-differentiable at T, (B.9) tells us that $\dot{p}_t$ is continuous at T, and hence that $\dot{q}_t$ is continuous at T.

that that the supremum is attained at

$$\alpha e^{-\rho t}\left[\frac{1}{c_t}-\frac{a_t}{A_t}\right]+(1-\alpha)\left[e^{-\rho|T-t|}\frac{1}{c_t}-\frac{p_t}{A_t}\right]=0,$$

so that

$$c_t=\left[\frac{\alpha e^{-\rho t}+(1-\alpha)e^{-\rho|T-t|}}{\alpha e^{-\rho t}a_t+(1-\alpha)p_t}\right]A_t=\lambda_t A_t \tag{B.10}$$

which verifies our asserted optimal policy at the instant t. But this is true only conditional on the conjectured forms of V and W, and the proof must now be completed by showing that V and W as well as their coefficients indeed have the forms (B.7), (B.8) and (B.9) for all dates, once we use the formula (B.10).

With that in mind, define, for all t and all A_t

$$\hat{V}(A_t,t)\equiv\int_t^N e^{-\rho(s-t)}\ln(c_s(A_s))ds \text{ and } \hat{W}(A_t,t)\equiv\int_t^N e^{-\rho|T-s|}\ln(c_s(A_s))ds$$

where the family of functions $\{c_s(A_s)\}$ is given by (B.10). By totally differentiating with respect to t, it is easy to check that:

$$\rho\hat{V}(A_t,t)=\ln(c_t)+\dot{A}_t\hat{V}_A(A_t,t)+\hat{V}_t(A_t,t) \tag{B.11}$$

and

$$e^{-\rho|T-t|}\ln(c_t)+\dot{A}\hat{W}_A(A_t,t)+\hat{W}_t(A_t,t)=0. \tag{B.12}$$

Optimality of the plan (B.10) will require us to equate (V,W) with $(\hat{V},\hat{W})$, and that will verify both the conjectured functional forms in (B.7) and (B.8) as well as the differential equations (B.9) for the endogenous coefficients (a_t,b_t,p_t,q_t).

To derive the boundary values, set $t=N$ in (B.5) and (B.6), from which we see that $V(A_N,N)=\ln c_N$ and $W(A_N,N)=e^{-\rho(N-T)}\ln c_N$. To derive the ODEs for a_t and p_t, set $V=\hat{V}$ in (B.11) so that

$$\begin{aligned}\rho V(A_t,t)&=\ln(c_t)+\dot{A}_tV_A(A_t,t)+V_t(A_t,t)\\&=\ln(A_t)+\ln(\lambda_t)+\frac{\dot{A}_ta_t}{A_t}+\dot{a}_t\ln(A_t)+\dot{b}_t\\&=(1+\dot{a}_t)\ln(A_t)+\ln(\lambda_t)+\dot{b}_t+\frac{a_t\dot{A}_t}{A_t}.\end{aligned}$$

Using the law of motion $\dot{A}_t = rA_t - c_t$ in the equality above, we see after some trivial simplification that

$$\rho V(A_t, t) = (1 + \dot{a}_t)\ln(A_t) + \ln(\lambda_t) + \dot{b}_t + a_t(r - \lambda_t).$$

A similar calculation can be performed by setting $W = \hat{W}$ in (B.12) to obtain:

$$\begin{aligned}
0 &= e^{-\rho|T-t|}\ln(c_t) + \dot{A}_t W_A(A_t, t) + W_t(A_t, t) \\
&= e^{-\rho|T-t|}\ln(\lambda_t[A_t]) + \frac{p_t \dot{A}_t}{A_t} + \dot{p}_t \ln(A_t) + \dot{q}_t \\
&= (\dot{p}_t + e^{-\rho|T-t|})\ln(A_t) + e^{-\rho|T-t|}\ln(\lambda_t) + \dot{q}_t + \frac{p_t(\dot{A}_t)}{A_t} \\
&= (\dot{p}_t + e^{-\rho|T-t|})\ln(A_t) + e^{-\rho|T-t|}\ln(\lambda_t) + \dot{q}_t + p_t(r - \lambda_t),
\end{aligned}$$

Now use the formula for V in (B.7) and for W in (B.8) and equate coefficients to obtain the desired ODEs for $\{a_t, p_t\}$.

In the process we also obtain differential equations for $\{b_t, q_t\}$:

$$\begin{aligned}
\rho b_t &= \ln(\lambda_t) + a_t(r - \lambda_t) + \dot{b}_t \\
\dot{q}_t &= p_t(\lambda_t - r) - e^{-\rho|T-t|}\ln(\lambda_t)
\end{aligned} \tag{B.13}$$

from which it is readily verified that these coefficients satisfy the properties stated in the initial conjecture. All that remains is to provide explicit and unique solutions for a_t and p_t, which we can easily do given (B.9) and the boundary conditions. ■

B.3. **Proof of Proposition 3.** Let $D = e^{-\rho(N-T)}$, $x(t) = e^{-\rho(T-t)}$, $f(x) = x^2/e^{-\rho T}$, and

$$g(x) = \frac{(\rho-1)e^{-\rho(N-t)} + 1 + 1 - e^{-\rho(T-t)}}{e^{-\rho t}[(\rho-1)e^{-\rho(N-t)} + 1]} = xe^{-\rho T} + \frac{1-x}{[(\rho-1)D + x^{-1}]e^{-\rho T}}$$

Then

$$\zeta_t = \frac{\alpha + (1-\alpha)f(x(t))}{\alpha + (1-\alpha)g(x(t))} \equiv h(x(t)).$$

It is straightforward to verify that $g''(x) < 0$. Furthermore, $h'(x) = 0$ implies that $f'(x)/g'(x) = h(x)$, which in turn implies that $g'(x) > 0$ whenever $h'(x) = 0$. Making use of all of the above, we see that $h''(x) > 0$ whenever $h'(x) = 0$. This indicates that $h(x)$ has at most one stationary point, and if that exists, $h(x)$ is decreasing in x to the left of that point and increasing to its right.

More tedious calculations show that at $x = e^{-\rho T}$ (i.e., at $t = 0$), we have $f(x) < g(x)$, while at $x = 1$ (i.e., at $t = T$), $f(x) = g(x)$. Therefore $h(x(0)) < h(x(T)) = 1$. Furthermore, $f'(x)g(x) - f(x)g'(x) > 0$ at $x = e^{-\rho T}$. Hence,

$$h'(x) = \frac{1-\alpha}{[\alpha + (1-\alpha)g(x)]^2} \times \\ \{f'(x)g(x) - f(x)g'(x) + \alpha[f'(x) - g'(x) - f'(x)g(x) + f(x)g'(x)]\}$$

is positive at $\alpha = 0$, and possibly turns negative for larger α's if $f'(x) - g'(x) - f'(x)g(x)$
$+ f(x)g'(x) < 0$. The initial slope of $h(x)$ determines whether $h(x)$ is U-shaped or always increasing. If the initial slope is negative, then it is U-shaped. Otherwise, it is always increasing.

Since $x = x(t)$ is an increasing function of t, the same property holds for ζ_t. ∎

B.4. **Proof of Proposition 4.** The proof is completed in the following steps:

(i) Divide the full interval $[0, N]$ into sub-intervals of length Δ.

(ii) Solve the problem facing each agent at the initialization t of an interval, who controls consumption over $[t, t + \Delta)$, given some starting asset A, and under the restriction that total asset holdings at the end of the period must equal some provisionally exogenous value $\hat{A}$, feasible given A. The solution is an optimal commitment plan as in Proposition 2.

(iii) Next, for each t and initial asset A, solve for the equilibrium choice of $\hat{A}$, under the conjecture that current and future agents behave as in (ii). This problem is in effect a game played by finitely many agents each taking a single action, and can be solved via backward induction.

(iv) Compute the limiting strategy profile as $\Delta \to 0$.

To this end, fix a t at the start of any sub-interval of length Δ, as well as "starting" and "ending" assets as described in (ii) above, and consider the problem faced by the agent controlling the interval $[t, t + \Delta]$. This problem is no different from our planning problem, except that (t, Δ) takes the place of $(0, N)$, and there is an asset promise of $\hat{A}$. But the latter only means that the effective present value of the asset at the agent's disposal is $B \equiv A - e^{-r\Delta}\hat{A} \geqslant 0$.[27]

It follows that for each of these "mini-problems," the solution is exactly as given by the planning problem: for any $t \in \{0, \Delta, 2\Delta, \ldots, N-\Delta\}$, and for $s \in [t, t+\Delta]$,

$$c_t(A_s, s; \hat{A}) = \left[\frac{\alpha e^{-\rho(s-t)} + (1-\alpha)e^{-\rho|T-s|}}{\alpha e^{-\rho(s-t)} a_s^t + (1-\alpha) p_s^t}\right] [A_s - e^{-r(t+\Delta-s)}\hat{A}] \tag{B.14}$$

where $\{a_s^t, p_s^t\}$ solve equations analogous to (8) and (9) for the planning problem:

$$a_s^t = \rho^{-1}\left[(\rho-1)e^{-\rho(t+\Delta-s)} + 1\right] \tag{B.15}$$

and

$$p_s^t = \begin{cases} \rho^{-1}\left[e^{-\rho(s-T)} + (\rho-1)e^{-\rho|(t+\Delta)-T|}\right] \text{ for } & s > T \\ \rho^{-1}\left[-e^{-\rho(T-s)} + (\rho+1)e^{-\rho|T-(t+\Delta)|}\right] & s < T, \end{cases} \tag{B.16}$$

These expressions provide a full characterization of the optimal consumption plan for the agent controlling $[t, t+\Delta)$, and aiming for $\hat{A}$ at the end of this period.

There are also forward and backward values generated by the mini-plan, and these are exactly as in the planning problem, except that they are defined on $A - e^{-r\Delta}\hat{A}$:

$$V_t(A, \hat{A}) = a_t^t \ln(A - e^{-r\Delta}\hat{A}) + b^t, \tag{B.17}$$

and

$$W_t(A, \hat{A}) = p_t^t \ln(A - e^{-r\Delta}\hat{A}) + q^t, \tag{B.18}$$

where the additive constants b^t and q^t can be solved just as in the planning problem (see (B.13)), but we will not need to do so here.

Now we turn to the discrete game in which each of the players also choose the target $\hat{A}$ at the end of their control, anticipating a continuation payoff from that choice. We proceed by induction. At $N - \Delta$, define

$$V^*_{N-\Delta}(A) = \int_{N-\Delta}^{N} e^{-\rho(s-[N-\Delta])} \ln c_s ds$$

and

$$W^*_{N-\Delta}(A) = \int_{N-\Delta}^{N} e^{-\rho|T-s|} \ln c_s ds$$

[27] Because $\hat{A}$ is feasible given A, B must be nonnegative.

to be the values generated by the planning problem at the very last decision stage, starting at A with a continuation asset of 0. It is easy to see that

$$V^*_{N-\Delta}(A) = \gamma_{N-\Delta}\ln(A) + \text{constant}$$

and

$$W^*_{N-\Delta}(A) = \nu_{N-\Delta}\ln(A) + \text{constant},$$

where

$$\gamma_{N-\Delta} = a^{N-\Delta}_{N-\Delta} = \frac{1}{\rho}[(\rho-1)e^{-\Delta\rho} + 1]$$

and

$$\nu_{N-\Delta} = p^{N-\Delta}_{N-\Delta} = \frac{1}{\rho}[e^{-\rho|N-\Delta-T|} + (\rho-1)e^{-\rho|N-T|}],$$

using (B.15) and (B.16). (The exact value of the constants need not concern us.)

Inductively, suppose that at decision date $t+\Delta$,

$$V^*_{t+\Delta}(A) = \gamma_{t+\Delta}\ln(A) + \text{constant} \tag{B.19}$$

and

$$W^*_{t+\Delta}(A) = \nu_{t+\Delta}\ln(A) + \text{constant}, \tag{B.20}$$

where $\gamma_{t+\Delta}$ and $\nu_{t+\Delta}$ satisfy the difference equations

$$\begin{aligned} \gamma_{t+\Delta} &= a^{t+\Delta}_{t+\Delta} + e^{-\rho\Delta}\gamma_{t+2\Delta} \\ \nu_{t+\Delta} &= p^{t+\Delta}_{t+\Delta} + \nu_{t+2\Delta}. \end{aligned} \tag{B.21}$$

The agent at date t chooses $\hat{A}$, for each initial asset A, to maximize

$$\alpha V_t(A,\hat{A}) + (1-\alpha)W_t(A,\hat{A}) + \alpha e^{-\rho\Delta}V^*_{t+\Delta}(\hat{A}) + (1-\alpha)W^*_{t+\Delta}(\hat{A}),$$

which, using (B.17), (B.18), (B.19) and (B.20), is equivalent to maximizing

$$[\alpha a^t_t + (1-\alpha)p^t_t]\ln[A - e^{-r\Delta}\hat{A}] + \alpha e^{-\rho\Delta}\gamma_{t+\Delta}\ln(A) + (1-\alpha)\nu_{t+\Delta}\ln(A)$$

by choosing $\hat{A}$. It is easy to see that the solution is given by

$$\sigma_t(A) = e^{r\Delta}\Big[\frac{\alpha e^{-\rho\Delta}\gamma_{t+\Delta} + (1-\alpha)\nu_{t+\Delta}}{\alpha(a^t_t + e^{-\rho\Delta}\gamma_{t+\Delta}) + (1-\alpha)(p^t_t + \nu_{t+\Delta})}\Big]A$$

where $\{a_t^t, p_t^t\}$ solve (B.15) and (B.16). From this expression, we can obtain solutions for $V_t^*(A), W_t^*(A)$ that solve (B.19)–(B.21) at date t. In particular, using the solutions

$$\begin{aligned} \gamma_t &= a_t^t + e^{-\rho\Delta}\gamma_{t+\Delta} \\ \nu_t &= p_t^t + \nu_{t+\Delta} \end{aligned} \tag{B.22}$$

we can write $\sigma_t(A) = e^{r\Delta}\mu_t A$, where

$$\mu_t = \left[\frac{\alpha(\gamma_t - a_t^t) + (1-\alpha)(\nu_t - p_t^t)}{\alpha(\gamma_t) + (1-\alpha)\nu_t}\right].$$

Now the inductive step is complete.

To finish the proof of the Proposition, we substitute the equilibrium rule $\sigma_t(A)$ for $\hat{A}$ in (B.14), which yields an instantaneous equilibrium consumption policy at date t:

$$c_t^*(A) = c_t(A, \sigma_t(A)) = \lambda_t[A - e^{-r\Delta}(e^{r\Delta}\mu_t A)] = \lambda_t(1-\mu_t)A,$$

where

$$\lambda_t \equiv \frac{\alpha + (1-\alpha)e^{-\rho|T-t|}}{\alpha a_t^t + (1-\alpha)p_t^t}.$$

Now we pass to the limit as $\Delta \to 0$. Equations (B.15) and (B.16) tell us that $a_t^t = 1 + \mathcal{O}(\Delta^2)$ and $p_t^t = e^{-\rho|T-t|} + \mathcal{O}(\Delta^2)$, so that by (B.22),

$$\begin{aligned} \gamma_t &= 1 + (1-\rho\Delta)\gamma_{t+\Delta} + \mathcal{O}(\Delta^2) \\ \nu_t &= e^{-\rho|T-t|} + \nu_{t+\Delta} + \mathcal{O}(\Delta^2). \end{aligned}$$

Thus, as $\Delta \to 0$, γ_t and ν_t satisfy the ODEs $\rho\dot{\gamma}_t = 1 + \gamma_t$ and $\dot{\nu}_t = -e^{-\rho|T-t|}$. Furthermore, it is easily verified that $\lim_{\Delta\to 0}\gamma_{N-\Delta} = 1$ and $\lim_{\Delta\to 0}\nu_{N-\Delta} = e^{-\rho|T-N|}$. These limiting boundary conditions, combined with the ODEs for γ_t and ν_t above, are precisely those that characterized the coefficients a_t and p_t in equations (8) and (9) of Proposition 2 and 4. Therefore $\lim_{\Delta\to 0}(\gamma_t, \nu_t) = (a_t, p_t)$, and so

$$\begin{aligned}\lambda_t(1-\mu_t)A &= \left[\frac{\alpha+(1-\alpha)e^{-\rho|T-t|}}{\alpha a_t^t+(1-\alpha)p_t^t}\right]\left[\frac{\alpha a_t^t+(1-\alpha)p_t^t}{\alpha(\gamma_t)+(1-\alpha)\nu_t}\right]A \\ &\to \left[\frac{\alpha+(1-\alpha)e^{-\rho|T-t|}}{\alpha+(1-\alpha)e^{-\rho|T-t|}}\right]\left[\frac{\alpha+(1-\alpha)e^{-\rho|T-t|}}{\alpha a_t+(1-\alpha)p_t}\right]A \\ &= \left[\frac{\alpha+(1-\alpha)e^{-\rho|T-t|}}{\alpha a_t+(1-\alpha)p_t}\right]A\end{aligned}$$

as $\Delta \to 0$, as required. ■

B.5. **Proof of Proposition 6.** Consider any situation in which the receiver believes the sender to be using the effort schedule $\{a_t\}$, while in fact the sender uses the schedule $\{a_t + \Delta_t\}$. Then, if the receiver obtains signals $\mathbf{z}^{\mathbf{t}} = (z_t)_{t\in\mathbf{t}}$ at the finite set of dates $\mathbf{t}$, she believes that the signals were generated according to

$$z_t = \theta + a_t + \epsilon_t \text{ for each } t \in \mathbf{t},$$

whereas in fact they were generated according to the true process

$$z_t = \theta + a_t + \Delta_t + \epsilon_t \text{ for each } t \in \mathbf{t}.$$

Writing $y_t \equiv z_t - a_t$, the receiver will form the point update

$$\theta(|\mathbf{t}|) \equiv \frac{\mu_0\tau_0 + \tau_\epsilon \sum_{t\in\mathbf{t}} y_t}{\tau_0 + |\mathbf{t}|\tau_\epsilon}, \tag{B.23}$$

which *only* depends on parameters $y_t = \theta + \epsilon_t$ and on $|\mathbf{t}|$, but not on $\mathbf{t}$. But from the sender's (correct) perspective, this translates into the random variable

$$\frac{\mu_0\tau_0 + \tau_\epsilon \sum_{t\in\mathbf{t}}[y_t + \Delta_t]}{\tau_0 + |\mathbf{t}|\tau_\epsilon} = \theta(|\mathbf{t}|) + \frac{\tau_\epsilon \sum_{t\in\mathbf{t}} \Delta_t}{\tau_0 + |\mathbf{t}|\tau_\epsilon}, \tag{B.24}$$

where — owing to the possible variation of Δ_t with t — this "augmented" random variable could well depend on the dates in $\mathbf{t}$. Thus a random payoff of $U(\theta(|\mathbf{t}|) + \frac{\tau_\epsilon \sum_{t\in\mathbf{t}} \Delta_t}{\tau_0+|\mathbf{t}|\tau_\epsilon})$ is delivered to the sender, where this value is simply $U(\mu_0)$ when $\mathbf{t}$ is empty. We now proceed to calculate the sender's *expected* payoff.

Let $p^{|\mathbf{t}|}$ denote the probability of $|\mathbf{t}|$ signals to be retained at some finite vector of dates $\mathbf{t}$. It includes the probability that other signals are generated at other dates but are forgotten. Since Poisson arrivals are conditionally independent, it is clear that $p^{|\mathbf{t}|}$ depends on the vector $\mathbf{t}$ only through the integer $|\mathbf{t}| \geq 0$. Then expected payoff is given by

$$\sum_{|\mathbf{t}|=0}^{\infty} p^{|\mathbf{t}|} \int_{\mathbf{t}} \left[\prod_{t \in \mathbf{t}} e^{-\rho_b (T-t)} \right] \mathbb{E}U \left(\theta(|\mathbf{t}|) + \frac{\tau_\epsilon \sum_{t \in \mathbf{t}} \Delta_t}{\tau_0 + |\mathbf{t}|\tau_\epsilon} \right) d\mathbf{t}, \tag{B.25}$$

where, conditional on the event that the remembered set of signals is $\mathbf{t}$, we take the expectation $\mathbb{E}$ of payoff over noisy signal realizations embodied in $\theta(|\mathbf{t}|)$.[28]

We also note that the derivative of $\mathbb{E}U \left(\theta(|\mathbf{t}|) + \frac{\tau_\epsilon \sum_{t \in \mathbf{t}} \Delta_t}{\tau_0 + |\mathbf{t}|\tau_\epsilon} \right)$ with respect to Δ_t at any t, and evaluated at $\Delta_t = 0$, is independent of t. It is given by

$$\frac{\partial}{\partial \Delta_t} \mathbb{E}U \left(\theta(|\mathbf{t}|) + \frac{\tau_\epsilon \sum_{t \in \mathbf{t}} \Delta_t}{\tau_0 + |\mathbf{t}|\tau_\epsilon} \right)_{|\Delta_t = 0} = \frac{\tau_\epsilon}{\tau_0 + |\mathbf{t}|\tau_\epsilon} \mathbb{E}U' \left(\theta(|\mathbf{t}|) \right). \tag{B.26}$$

Equations (B.25) and (B.26) immediately show that a schedule *not* as described in the statement of the Proposition cannot be an equilibrium. For then there is a set of positive measure around some $t < \bar{T}$ such that $a_s > 0$ on that set, and a set of positive measure around some $t' > t$ for which $a_s < \bar{a}$ on that set. Using (B.26), we see that a small surprise reallocation of effort from t to t' must then be a profitable deviation by the sender; a contradiction.

Finally, we argue that when U is concave, the described schedule in the statement of the Proposition *is* an equilibrium. Consider any deviation to an alternative effort schedule a'_t, which generates $\Delta_t \equiv a'_t - a_t$ for every t. Because $a_t = 0$ for $t \in [0, \bar{T})$ and $a_t = \bar{a}$ for $t \in [\bar{T}, T]$, it must be that $\Delta(t) \geq 0$ for $t \in [0, \bar{T})$ and $\Delta_t \leq 0$ for $t \in [\bar{T}, T]$. It follows that for every finite set of dates $\mathbf{t}$,

$$\begin{aligned}
&\left[\prod_{t \in \mathbf{t}} e^{-\rho_b (T-t)} \right] \left[\mathbb{E}U \left(\theta(|\mathbf{t}|) + \frac{\tau_\epsilon \sum_{t \in \mathbf{t}} \Delta_t}{\tau_0 + |\mathbf{t}|\tau_\epsilon} \right) - \mathbb{E}U \left(\theta(|\mathbf{t}|) \right) \right] \\
&\leq e^{-\rho_b (T-\bar{T})|\mathbf{t}|} \left[\mathbb{E}U \left(\theta(|\mathbf{t}|) + \frac{\tau_\epsilon \sum_{t \in \mathbf{t}} \Delta_t}{\tau_0 + |\mathbf{t}|\tau_\epsilon} \right) - \mathbb{E}U \left(\theta(|\mathbf{t}|) \right) \right] \\
&\leq e^{-\rho_b (T-\bar{T})|\mathbf{t}|} \left[\frac{\tau_\epsilon}{\tau_0 + |\mathbf{t}|\tau_\epsilon} \right] \mathbb{E}U' \left(\theta(|\mathbf{t}|) \right) \Delta_t, \qquad \text{(B.27)}
\end{aligned}$$

where the first inequality comes from replacing t by $\bar{T}$ in $e^{-\rho_b (T-t)}$, and the second inequality comes from (a) viewing $\mathbb{E}U$ as a concave function of Δ_t, (b) the fact

[28] The empty product has value 1 by convention, with $U(\theta(\mathbf{y}^{\mathbf{t}} + \mathbf{\Delta}^{\mathbf{t}})) = U(\mu_0)$ in that case.

that $f(x) - f(x_0) \leq f'(x_0)(x - x_0)$ for any concave function and any x and x_0, and (c) applying (B.26).

Let G be the payoff gain from the contemplated deviation. Then:

$$\begin{aligned} G &= \sum_{|\mathbf{t}|=0}^{\infty} p^{|\mathbf{t}|} \int_{\mathbf{t}} \left[\prod_{t \in \mathbf{t}} e^{-\rho_b (T-t)} \right] \left[\mathbb{E}U \left(\theta(|\mathbf{t}|) + \frac{\tau_\epsilon \sum_{t \in \mathbf{t}} \Delta_t}{\tau_0 + |\mathbf{t}|\tau_\epsilon} \right) - \mathbb{E}U \left(\theta(|\mathbf{t}|) \right) \right] d\mathbf{t} \\ &\leq \sum_{|\mathbf{t}|=0}^{\infty} p^{|\mathbf{t}|} \int_{\mathbf{t}} e^{-\rho_b (T-\bar{T})|\mathbf{t}|} \left[\frac{\tau_\epsilon}{\tau_0 + |\mathbf{t}|\tau_\epsilon} \right] \mathbb{E}U' \left(\theta(|\mathbf{t}|) \right) \Delta_t \, d\mathbf{t} \\ &= \left[\sum_{|\mathbf{t}|=0}^{\infty} p^{|\mathbf{t}|} e^{-\rho_b (T-\bar{T})|\mathbf{t}|} \left[\frac{\tau_\epsilon}{\tau_0 + |\mathbf{t}|\tau_\epsilon} \right] \mathbb{E}U' \left(\theta(|\mathbf{t}|) \right) \right] \cdot \int_t \Delta_t \, dt \\ &\leq 0, \end{aligned}$$

where the first inequality uses (B.27) and the last inequality uses the fact that no effort schedule can exceed the total budget.

When $\rho_b = 0$, the replacement of t by T in (B.27) is no longer necessary, and so the above arguments apply to show that any effort schedule that satisfies (17) with equality constitutes an equilibrium. ■

B.6. **Proof of Proposition 7.** . Recall that the payoff gain from a contemplated deviation to $\{\Delta_t\}$ is given by

$$G = \sum_{|\mathbf{t}|=0}^{\infty} p^{|\mathbf{t}|} \int_{\mathbf{t}} \left[\prod_{t \in \mathbf{t}} e^{-\rho_b (T-t)} \right] \left[\mathbb{E}U \left(\theta(|\mathbf{t}|) + \frac{\tau_\epsilon \sum_{t \in \mathbf{t}} \Delta_t}{\tau_0 + |\mathbf{t}|\tau_\epsilon} \right) - \mathbb{E}U \left(\theta(|\mathbf{t}|) \right) \right] d\mathbf{t}.$$

Given the additive time-separability of costs, we may specialize to an infinitesimal deviation entertained *only* at some t for an infinitesimal interval dt. As such, all terms above are zero, except for the probability of just an additional *remembered* signal arriving during dt, which occurs with probability $\lambda e^{-\rho_b (T-t)} dt$. Using (B.26), the gain evaluated at $\Delta_t = 0$ is therefore

$$G_t \equiv \left[\lambda e^{-\rho_b (T-t)} dt \right] \left[\frac{\tau_\epsilon}{\tau_0 + |\mathbf{t}|\tau_\epsilon} \right] \mathbb{E}U' \left(\theta(|\mathbf{t}|) \right) = c'(a_t) dt, \tag{B.28}$$

which is set equal to the marginal cost $c'(a_t)dt$ of that deviation as a necessary condition of any equilibrium. From (B.28), it is clear that if $\rho_b > 0$, then this marginal gain G_t is rising as a function of t, and so $c'(a_t)$ must also rise. It follows that $a(t)$ is strictly increasing in any equilibrium with $\rho_b > 0$. On the other hand, if $\rho_b = 0$, $a(t)$ is constant over time.

To show existence, define an effort schedule $\{a_t^*\}$ by (B.28). Consider any deviation to $a_t = a_t^* + \Delta_t$. Then

$$\begin{aligned}
G &= \sum_{|\mathbf{t}|=0}^{\infty} p^{|\mathbf{t}|} \int_{\mathbf{t}} \left[\prod_{t\in\mathbf{t}} e^{-\rho_b(T-t)}\right] \left[\mathbb{E}U\left(\theta(|\mathbf{t}|) + \frac{\tau_\epsilon \sum_{t\in\mathbf{t}} \Delta_t}{\tau_0 + |\mathbf{t}|\tau_\epsilon}\right) - \mathbb{E}U\left(\theta(|\mathbf{t}|)\right)\right] d\mathbf{t} \\
&\leq \sum_{|\mathbf{t}|=0}^{\infty} p^{|\mathbf{t}|} \int_{\mathbf{t}} \left[\prod_{t\in\mathbf{t}} e^{-\rho_b(T-t)}\right] \left[\frac{\tau_\epsilon}{\tau_0 + |\mathbf{t}|\tau_\epsilon}\right] \mathbb{E}U'\left(\theta(|\mathbf{t}|)\right) \Delta_t \, d\mathbf{t} \\
&= \sum_{|\mathbf{t}|=0}^{\infty} p^{|\mathbf{t}|} \int_{\mathbf{t}} c'(a_t^*)\Delta_t \, d\mathbf{t} \\
&= \int_t c'(a_t^*)\Delta_t \, dt \\
&\leq \int_t [c(a_t^* + \Delta) - c(a_t^*)] \, dt,
\end{aligned}$$

where the first inequality follows from viewing $\mathbb{E}U$ as a concave function of Δ_t, just as in (B.27), the equality after that from the first order condition (B.28), the equality after because $\sum_{|\mathbf{t}|=0}^{\infty} p^{|\mathbf{t}|} = 1$ and the integral need only be taken over t rather than $\mathbf{t}$, and the final inequality from the convexity of the cost function. This calculation shows that no gain is possible under any deviation from $\{a_t^*\}$, and the proof is complete. ■

References

Ainslie, G. (1991), "Derivation of 'Rational' Economic Behavior from Hyperbolic Discount Curves," *American Economic Review* **81**, 334–340.

Ariely, D. (1998), "Combining Experiences over Time: The Effects of Duration, Intensity Changes, and On-line Measurements on Retrospective pain Evaluations," *Journal of Behavior Decision Making* **11**, 19–45.

Ariely, D. and G. Loewenstein (2000), "When Does Duration Matter in Judgment and Decision Making," *Journal of Experimental Psychology* **129**, 508–523 .

ASHRAF, N., N. GONS AND D. KARLAN (2003), "A Review of Commitment Savings Products in Developing Countries," *Asian Development Bank Economics and Research Department Working Paper Series* **45**.

BANDERA, A. (1977), *Social Learning Theory*, Englewood Cliffs, NJ: Prentice-Hall.

BECKER, G. AND C. MULLIGAN (1997), "The Endogenous Determination of Time Preference," *Quarterly Journal of Economics* **112**, 729–758.

BECKER, G., M. GROSSMAN AND K. M. MURPHY, (1994), "An empirical analysis of cigarette addiction," *American Economic Review* **84**, 396-418.

BERGSTROM, T. (1999), "Systems of Benevolent Utility Functions," *Journal of Public Economic Theory* **1**, 71–100.

BERNHEIM, D. AND D. RAY (1987), "Economic Growth with Integenerational Altruism," *Review of Economic Studies* **54**, 227–41.

BERNHEIM, D., RAY, D. AND S. YELTEKIN (2015), "Poverty and Self-Control," *Econometrica* **83**, 1877–1911.

BISIN, A. AND T. VERDIER (1998), "On the Cultural Transmission of Preferences for Social Status," *Journal of Public Economics* **70**, 75–97.

BISIN, A. AND T. VERDIER (2000), "Beyond the Melting Pot: Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits," *Quarterly Journal of Economics*, **115**, 955–988.

BJÖRK, T AND M. KHAPKO AND A. MURGOCI (2017), "On time-inconsistent stochastic control in continuous time," *Finance Stoch.*, **21**, 331–360.

BREGMAN, G. AND M. KILLEN (1999), "Adolescents' and Young Adults' Reasoning About Career Choice and the Role of Parental Influence," *Journal of Research on Adolescence* **9**, 253–75.

BRIM, O. (1966), "Socialization through the Life Cycle," in *Socialization after Childhood* (O. Brim and S. Wheeler (eds), New York: John Wiley and Sons.

CAPLIN, A. AND J. LEAHY (2004), "The Social Discount Rate," *Journal of Political Economy* **112**, 1257–1268.

Chapman, G. (2000), “Preferences for Improving and Declining Sequences of Health Outcomes,” *Journal of Behavioral Decision Making*, **13**, 203–218.

Doepke, M. and F. Zilibotti (2017), “Parenting with Style: Altruism and Paternalism in Intergenerational Preference Transmission,” *Econometrica* **85(5)**, 1331-1371.

Dryler, H. (1998), “Parental Role Models, Gender and Educational Choice,” *British Journal of Sociology* **49**, 375–398.

Fredrickson, B.L., and Kahneman, D. (1993), “Duration Neglect in Retrospective Evaluations of Affective Episodes,” *Journal of Personality and Social Psychology* **65**, 44–55.

Galperti, S. and B. Strulovici (2017), “A Theory of Intergenerational Altruism,” *Econometrica* **85**, 1175-1218.

Gilboa, I. A. Postlewaite and L. Samuelson (2016), “Memorable Consumption,” *Journal of Economic Theory* **165**, 414-455.

Gollier, C., and M. Weitzman (2010), “How Should the Distant Future Be Discounted When Discount Rates Are Uncertain?” *Economics Letters* **107**, 350–353.

Gruber, J., and B. Köszegi (2001), “Is Addiction Rational: Theory and Evidence,” *Quarterly Journal of Economics*, **116**, 1261-1303.

Harris, C. and D. Laibson (2001), “Dynamic Choices of Hyperbolic Consumers,” *Econometrica* **269**, 935–957.

Harrison, G. W., Morten I. Lau and Melonie B. Williams (2002), “Estimating Individual Discount Rates in Denmark: A Field Experiment,” *American Economic Review* **92**, 1606-1617.

Hartman, S. and O. Harris (1992), “The Role of Parental Influence in Leadership,” *Journal of Social Psychology* **132**, 153–167.

Hayashi, F. (1997), *Understanding Saving: Evidence from the United States and Japan*, Cambridge, MA: The MIT Press.

Henretta, J. (1984), "Parental Status and Child's Home Ownership," *American Sociological Review* **49**, 131–140.

Hess, R. and J. Torney (1967), *The Development of Political Attitudes in Children*, Chicago: Aldine Publishing Company.

Holmström, B. (1999), "Managerial Incentive Problems: A Dynamic Perspective," *The Review of Economic Studies* **66:1**, 169–182.

Hori, H. and S. Kanaya (1989), "Utility Functionals with Nonpaternalistic Intergenerational Altruism," *Journal of Economic Theory* **49**, 241–265.

Jackson, M., and L. Yariv (2015), "Collective Dynamic Choice: The Necessity of Time Inconsistency," *American Economic Journal: Microeconomics.***7:4**, 150-178.

Laibson, D. (1994), "Self-Control and Saving," Mimeo., Department of Economics, Harvard University.

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

Laibson, D. (1998), "Life-Cycle Consumption and Hyperbolic Discount Functions," *European Economic Review* **42**, 861–871.

Loewenstein, G. and D. Prelec, (1992) "Anomalies in Intertemporal Choice: Evidence and an Interpretation," *Quarterly Journal of Economics* **107**, 573–97.

Loewenstein, G. and D. Prelec, (1993) "Preferences for Sequences of Outcomes," *Psychological Review* **100**, 91–108.

Loewenstein, G. and N. Sicherman, (1991) "Do Workers Prefer Increasing Wage Profiles," *Journal of Labor Economics* **9**, 67-84.

Loewenstein, G. and R. Thaler (1989), "Anomalies: Intertemporal Choice," *Journal of Economic Perspectives* **3**, 181–193.

Moschis, G. (1987), *Consumer Socialization: A Life Cycle Perspective*, Lexington, MA: D.C. Heath and Company.

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

Pearce, D. (2008), "Nonpaternalistic Sympathy and the Inefficiency of Consistent Intertemporal Plans," in *Foundations in Microeconomic Theory* **89**, ed. by M. O. Jackson and A. McLennan. Berlin: Springer, 213–231.

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

Ray, D. (1987), "Nonpaternalistic Intergenerational Altruism," *Journal of Economic Theory* **41**, 112–132.

Ray, D. and R. Wang (2001), "On Some Implications of Backward Discounting," mimeo., New York University.

Ray, D. ⓡ R. Vohra (2020), "Games of Love and Hate," *Journal of Political Economy* **128**, 1789–1825.

Read, D. and N.L. Read (2004), "Time Discounting Over the Lifespan," *Organization Behaviour and Human Decision Processes* **94**, 22-32.

Salili, F. (1994), "Age, Sex, and Cultural Differences in the Meaning and Dimensions of Achievement," *Personality and Social Psychology Bulletin* **20**, 635–648.

Sozou, P.D. and R.M. Seymour (2003), "Augmented Discounting: Interaction between aging and time-preference," *Proceedings of the Royal Society of London B* **270**, 1047-1053.

Strotz, R. (1956), "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies* **23**, 165–180.

Varey, C.A., and D. Kahneman (1992), "Experiences Extended across Time: Evaluation of Moments and Episodes," *Journal of Behavioral Decision Making* **5**, 169–185.

Vásquez, J. and M. Weretka (2020), "Affective Empathy in Non-Cooperative Games," *Games and Economic Behavior* **121**, 548–556.

Ward, S. (1974), "Consumer Socialization," *Journal of Consumer Research* **1** (September), 1–16.

Weinberg, B. (2001), “An Incentive Model of the Effect of Parental Income on Children,” *Journal of Political Economy* **109**(2), 226–280.

World Bank (2015a), “Real Interest Rate,” *World Development Indicators, World DataBank.*

World Bank (2015b), “Adjusted Savings: Gross Savings (% of GNI),” *World Development Indicators, World DataBank.*

Online Appendix

Debraj Ray, Nikhil Vellodi, Ruqu Wang

Appendix C. Discussion on the Profile of Patience

In this section, we elaborate on the specification of weights placed by the agent on her various selves. As mentioned, our results go through relatively unscathed with even more general schemes, and in fact some interesting insights can be drawn from such considerations.

Recall that our date t agent places weight α on her current self, β on her retirement self at date T, and uniform weight ω on all other future selves. Consider the following variations.

C.1. **Weights Only On Future Selves Between Current Age and Retirement.** Suppose our agent spreads uniform weight over all her "intermediate" future selves, i.e. weight ω is placed on all selves in (t, T), in addition to the weights on her current and retirement selves, which remain α and β respectively. Make exactly the same assumption A.1 in the main text. Then very little changes in Proposition 1. Intuitively, a weight on future selves beyond T, as in the main text, only serves to enhance the importance of the retirement self, thereby raising the effective value of β but otherwise causing no difference at all to the pre-retirement results. This intuition is fully borne out. The instantaneous rate of impatience $i(t, s)$ for $s \in [t, T)$ now reads

$$i(t,s) = \Big[\frac{(\rho_f \alpha - \omega)e^{-\rho_f(s-t)} - (\rho_b \beta - \omega)e^{-\rho_b(T-s)}}{(\alpha - \frac{\omega}{\rho_f})e^{-\rho_f(s-t)} + (\beta - \frac{\omega}{\rho_b})e^{-\rho_b(T-s)} + \frac{\omega}{\rho_f} + \frac{\omega}{\rho_b}}\Big]$$

and in particular,

$$i(t,t) = \Big[\frac{(\rho_f \alpha - \omega) - (\rho_b \beta - \omega)e^{-\rho_b(T-t)}}{(\alpha - \frac{\omega}{\rho_f}) + (\beta - \frac{\omega}{\rho_b})e^{-\rho_b(T-t)} + \frac{\omega}{\rho_f} + \frac{\omega}{\rho_b}}\Big]$$

Just as before, our assumptions guarantee that the numerator of $i(t, s)$ is decreasing in s while the denominator is increasing, and the same is true of $i(t, t)$. Thus, the results regarding how $i(t, s)$ and $i(t, t)$ fall in the approach to retirement remain entirely unchanged.

Part 5 of the Proposition is now sharper. There is still a jump in impatience as t crosses T, and moreover there will be an immediate reversion to standard geometric discounting post-retirement, at the rate ρ_f.

C.2. **Time-Varying Weights.** We might also think it reasonable to allow the weights α, β, ω to depend on time. We use here the formulation that weights are placed on the current self, the retirement self and all future selves. (The other case in which weights are placed on all "intermediate" future selves works in exactly the same way.) We can without any loss of generality choose a particular normalization. So we presume that at any date $t < T$, our individual places non-negative weights $\{\alpha(t), \beta(t), \omega(t)\}$ on the three sets of selves, and that these sum to 1; i.e., $\alpha(t) + \beta(t) + (N - t)\omega(t) = 1$. We impose the following restrictions:

[A.2] $\min\{\rho_f, \rho_b\} \min\{\alpha(t), \beta(t)\} > 2\omega(t)$ for all $t < T$.

[A.3] $\alpha(t)$ and $\beta(t)$ are non-decreasing.

Proposition 1 in the main text remains valid under these assumptions. As a matter of fact, observe that the time-varying nature of the weights only enters the picture when t is changing, which means that we need only reconsider Parts 4 and 5. To see why $i(t,t)$ is decreasing before retirement under these assumptions, observe that:

$$\begin{aligned} i(t,t) &= \left[\frac{\rho_f \alpha(t) - \rho_b \beta(t) e^{-\rho_b(T-t} - \omega(t)[1 - e^{-\rho_b(T-t}]}{\rho_b \alpha(t) + \rho_b \beta(t) e^{-\rho_b(T-t} + \omega(t)[1 - e^{-\rho_b(T-t}]}\right] \\ &= \left[\frac{(\rho_f \alpha(t) - \omega(t)) - (\rho_b \beta(t) - \omega(t)) e^{-\rho_b(T-t)}}{(\alpha(t) + \frac{\omega(t)}{\rho_b}) + (\beta(t) - \frac{\omega(t)}{\rho_b}) e^{-\rho_b(T-t)}}\right] \end{aligned} \tag{C.1}$$

By inspecting the formula (C.1) above and recalling the normalization, notice that $i(t,t)$ increases most rapidly when all the extra weight is handed over to $\alpha(t)$ as t increases. That is, it suffices to examine the case in which $\beta(t)$ and $\omega(t)$ are constant in t, and $\alpha'(t) = \omega(t)$. Differentiating $i(t,t)$ with respect to t and imposing these restrictions on weights, it can be shown that $i(t,t)$ is decreasing in t if and only if

$$(\rho_b \alpha(t) - \omega)(\rho_b \beta - \omega) e^{-\rho_b(T-t)} \geq \omega^2$$

from which it is clear that Condition A.2 is sufficient for this expression to hold.

For $t > T$, follows the analysis of the main text, or is exponential and constant if no weight is additionally placed on selves beyond T.

C.3. **Weights on Past Selves.** Yet another alternative is to place a uniform weight *all* selves, past or future, in addition to the privileged weights α and β on the current and retirement selves. Now, at any date, the "effective weight" on the current self t is α *plus* the discounted effect of all weights on "ancestral selves." As time passes, these ancestral selves will accumulate, effectively lending weight to α. Meanwhile, the effective weight on all other selves remains steady up to retirement. We therefore have a special case of the previous section on changing weights, where β and ω are unchanged, but where the cumulated weight $\alpha(t)$ on the current self can be viewed as

$$\alpha(t) = \alpha + \omega \int_0^t e^{-\rho_f(t-i)} di \tag{C.2}$$

It is possible to embed this model in the case of time-varying weights. First make Assumption A.2 on the original weights; i.e. assume $\min\{\rho_f, \rho_b\} \min\{\alpha, \beta\} > 2\omega$. Next, renormalize all weights by dividing by the sum $\alpha(t) + \beta + (N - t)\omega$, where $\alpha(t)$ is defined by (C.2). Then it is easy to see that A.2 and A.3 are both satisfied for the renormalized set of weights, and the earlier analysis applies.[29]

C.4. **No Retirement Self.** Finally, consider the case in which the agent places uniform weight on all future selves, without the presence of a privileged future retirement self. This is a model that might seem to natural to some: after all, if the agent is forward-looking enough to consider one future self, why not all consider all such selves equally? We claim, however, that such equal treatment is at odds with the basic premise in our paper, which places special emphasis on a stock-taking self in the future. Thus our weights are fundamentally "bi-modal" in nature, with privilege accorded to the shadow parent, as well as the current self.

That said, we consider briefly this alternative model. Indeed, the predictions in Proposition 1 generally not hold. Consider for instance, the expressions for $i(t,s)$

[29]It is easy to check that $\alpha'(t) < \omega$ for all t, so that $\alpha(t) + \beta + \omega(N - t)$ falls with t. Therefore, if we renormalize, the new weights will satisfy [A.3], while [A.2] is unaffected by the renormalization and continues to apply.

and $i(t,t)$:

$$i(t,s)\Big[\frac{(\rho_f\alpha-\omega)e^{-\rho_f(s-t)}+\omega e^{-\rho_b(N-s)}}{(\alpha-\frac{\omega}{\rho_f})e^{-\rho_f(s-t)}-\frac{\omega}{\rho_b}e^{-\rho_b(N-s)}+\frac{\omega}{\rho_f}+\frac{\omega}{\rho_b}}\Big]$$

and

$$i(t,t)=\Big[\frac{(\rho_f\alpha-\omega)+\omega e^{-\rho_b(N-t)}}{(\alpha+\frac{\omega}{\rho_b})-\frac{\omega}{\rho_b}e^{-\rho_b(N-t)}}\Big].$$

Looking first at $i(t,s)$, it is clear that both the denominator and numerator are always positive, and hence negative discounting cannot occur. Further, note that the denominator of the expression for $i(t,s)$ above is decreasing in s, and hence $i(t,s)$ is *increasing* if and only if

$$e^{-\rho_f(s-t)}e^{\rho_b(N-s)}\geq\frac{\rho_f(\rho_f\alpha-\omega)}{\rho_b\omega}$$

For s close to t, this expression will hold for N sufficiently large, or alternatively for ω close enough to its upper bound $\rho_f\alpha$. Turning to $i(t,t)$, it is immediate that $i(t,t)$ is now in fact *increasing* in t. Previously, the conflict was largely between the current and retirement self. Given that the latter is initially situated far away, its influence was initially discounted. Over time, the conflict was felt to a greater degree as the influence of this future self became a proximate reality. Without such a focal self, the passage of time simply erodes the influence of future selves as they are passed, and thus the rate of impatience simply increases.