INSTANTANEOUS GRATIFICATION

By Christopher Harris and David Laibson*

This Draft: June 6, 2008

Abstract. Extending Barro (1999) and Luttmer & Mariotti (2003), we introduce a new model of time preferences: the *instantaneous-gratification* model. This model applies tractably to a much wider range of settings than existing models. It applies to complete and incomplete-market settings and it works with generic utility functions. It works in settings with linear policy rules and in settings in which equilibrium cannot be supported by linear rules. The instantaneous-gratification model also generates a unique equilibrium, even in infinite-horizon applications, thereby resolving the multiplicity problem hitherto associated with dynamically inconsistent models. Finally, it simultaneously features a single welfare criterion and a behavioral tendency towards overconsumption.

JEL classification: C6, C73, D91, E21.

Keywords: quasi-hyperbolic discounting, time preference, dynamic inconsistency, continuous time, consumption, savings, buffer stock, Euler Relation, dynamic games, altruistic growth.

*We thank Drew Fudenberg, George Mailath, Andrew Postlewaite and seminar participants at European Summer Symposium on Economic Theory (Gerzensee), Harvard University, Princeton University, Stanford University, Universita Bocconi, University of California Berkeley, University of Chicago, University College London, University of Minneapolis, University of Pennsylvania, Universitat Pompeu Fabra, Wharton, and Yale University. Laura Serban and Evdokia Nikolova provided outstanding research assistance. Harris acknowledges financial support from the British Academy. Laibson acknowledges financial support from the National Science Foundation (SBR-9510985), the Olin Foundation, the National Institute on Aging (R01-AG-1665) and the MacArthur Foundation. Christopher Harris: Cambridge University, King's College, Cambridge, United Kingdom, CB2 1ST. David Laibson: Harvard University, Department of Economics, Littauer M-14, Cambridge, MA, 02138.

1. Introduction

The discrete-time quasi-hyperbolic discount function: $\{1, \beta\,\delta, \beta\,\delta^2, \beta\,\delta^3, ...\}$ is used to model high rates of short-run discounting.[1] With $\beta < 1$, this present-biased discount function generates a gap between a high short-run discount rate $(-\ln\beta\,\delta)$ and a low long-run rate $(-\ln\delta)$. The quasi-hyperbolic discount function has been used to study a range of behaviors, including consumption, procrastination, addiction and search.[2]

Extending the work of Barro (1999) and Luttmer and Mariotti (2003) on continuous-time models of non-exponential time preferences, the current paper shows how to operationalize quasi-hyperbolic time preferences in continuous time. Our model — which we call the instantaneous-gratification model or IG model for short — applies tractably to a much wider range of settings than existing models. For example, it applies to incomplete-market settings in which labor income is either uncollaterizable (with the associated liquidity constraints) or uninsurable; and it works with an economically rich class of utility functions which includes, but is much larger than, the class of utility functions with constant relative risk aversion. In particular, we do not have to restrict our analysis to linear policy rules or to settings in which such rules support an equilibrium.

We develop the IG model in two steps. In the first step, following Barro and Luttmer-Mariotti, we assume that the present is valued discretely more than the future, mirroring the one-time drop in valuation implied by the discrete-time quasi-hyperbolic discount function. However, we assume that the transition from the present to the future occurs with a constant hazard rate λ. This assumption reduces the Bellman equation to a pair of stationary differential equations that characterize the current- and continuation-value functions. We call the model obtained after the first step the present-future model or PF model for short.

In the second step, we let the hazard rate λ go to ∞. This brings us to the IG model. The Bellman equation for the IG model is even simpler than that of the PF model: it is a single ordinary differential equation for the continuation-value function. Moreover it turns out to be identical to that of a related optimization problem with a wealth-contingent utility function. The IG model, which is dynamically *inconsistent*, therefore

[1] See Phelps and Pollak (1968) and Laibson (1997). Strotz (1956) first formalized the idea that the short-run discount rate is greater than the long-run discount rate. Loewenstein and Prelec (1992) axiomatize a true hyperboloid.

[2] For some examples, see Akerlof (1991), O'Donoghue and Rabin (1999a, 1999b), Angeletos, Laibson, Repetto, Tobacman and Weinberg (2001), DellaVigna and Malmendier (2004), and Della Vigna and Paserman (2005).

has the same value function as the related optimization problem, which is dynamically *consistent*. The IG model is not, however, observationally equivalent to this optimization problem: the IG model and the optimization problem share the same dynamics, the same long-run discount rate and the same value function, but they have different instantaneous utility functions and different equilibrium policies. In a word: the equivalence applies to the *value* functions but not to the *policy* functions.[3]

The IG model carves out a tractable niche between dynamically inconsistent models and dynamically consistent models. On the one hand, it features dynamically inconsistent behavior and rational expectations. So, at each moment, the individual acts strategically with regard to her future preferences. On the other hand, the fact that the IG value function coincides with the value function of the related optimization problem implies that the IG model inherits many standard regularity properties.[4]

For example, the value-function-equivalence result implies that the IG model has a unique equilibrium. This uniqueness result is surprising, since the quasi-hyperbolic model is a dynamic game. Indeed, Krusell and Smith (2000) have shown that Markov-perfect equilibria are *not* unique in a deterministic discrete-time setting. In contrast, we provide two uniqueness results. First, we prove uniqueness in the case in which asset returns are stochastic. Second, we show that the unique equilibrium of the stochastic IG model converges to an equilibrium of the corresponding deterministic model as the noise in the asset returns goes to zero. In other words, we are able to select a unique equilibrium of the deterministic IG model by using a natural variant of standard equilibrium-refinement procedures.[5]

Similarly, we can give a detailed characterization of the consumption function in the IG model. When the expected rate of return is below a key threshold, the equilibrium consumption function displays a discontinuity at the liquidity constraint. Consequently, consumption will fall discontinuously when a consumer spends down her assets and hits the liquidity constraint. This intuitive prediction is not possible in a dynamically-consistent

[3] See Laibson (1996) and Barro (1999) for the two cases in which observational equivalence of the policy functions also holds: (1) log utility, time-varying interest interest, and no liquidity constraints; or (2) constant relative risk aversion, fixed interest rates, and no liquidity constraints.

[4] In discrete-time quasi-hyperbolic models, standard regularity properties (including differentiability and uniqueness) obtain provided that β is close enough to 1 (Harris and Laibson 2002).

[5] Our uniqueness result even offers something new in settings in which linear policy rules support an equilibrium: it tells us that if one can find an equilibrium in linear policy rules — say by the method of undetermined coefficients — then that equilibrium is unique, not just in the set of equilibria in linear policy rules, but even in the set of *all* policy rules, linear or non-linear.

consumption model. In such models, the timepath of consumption is continuous, even at the point at which the consumer hits a liquidity constraint.

Finally, the IG model features a *single* welfare criterion, even though the model involves dynamically inconsistent behavioral choices. Because the present is valued discretely more than the future, the current self has an incentive to overconsume; but the discretely higher value of the present only lasts for an instant, so this overvaluation does not affect the welfare criterion. Hence, the model simultaneously features a single welfare function and a behavioral tendency toward overconsumption.

In summary, the IG model is generalizable with regard to both institutions and consumption preferences, supports a unique equilibrium, makes new predictions about the consumption function, and identifies a single sensible welfare criterion. In Section 2 we present the PF model of time preferences and formulate some of its properties. In Section 3 we present the consumption problem that we use as our application. In Section 4 we describe the IG model. In Section 5 we show that the IG model has the same Bellman equation as a related dynamically-consistent optimization problem. We use this partial-equivalence result to prove equilibrium existence and uniqueness. We also use it to derive a unique equilibrium of the limiting version of our model in which the return on the financial asset becomes deterministic. In Section 6, we characterize the equilibrium consumption function. In Section 7, we explain how our results can be generalized from the class of utility functions with constant relative risk aversion to the class of utility functions with *bounded* (not constant) relative prudence (cf. Kimball, 1990). Section 8 concludes.

2. The Present-Future Model of Time Preferences

2.1. A Stochastic Discount Function. In the *discrete-time* formulation of quasi-hyperbolic time preferences, it is natural to divide time into two intervals: the present — consisting of only the current period — and the future. All periods, present and future, are discounted *exponentially* with the discount factor $0 < \delta < 1$. Future periods are further discounted with the factor $0 < \beta \leq 1$. Combining these pieces, the present period (i.e. $t = 0$) receives full weight, and future periods (i.e. $t \geq 1$) are given weight $\beta\,\delta^t$.

This model can be generalized in two ways. First, the present could last for an arbitrary length of time, instead of lasting for exactly one period. Second, the duration of the present could be stochastic, instead of being deterministic. Both of these generalizations have natural continuous-time analogues.

Consider an economic self born at time $s_0 = 0$. Call this self 'self 0'. The lifetime of self 0 is divided into two intervals: a 'present', which lasts from s_0 to $s_0 + \tau_0$; and a 'future', which lasts from $s_0 + \tau_0$ to ∞. Think of the present as the interval during which control is exercised by self 0, and of the future as the interval during which control is exercised by subsequent selves. The length τ_0 of the present is stochastic, and is distributed exponentially with hazard rate $\lambda \in [0, \infty)$.

When the present of self 0 ends at $s_0 + \tau_0$, a new self is born and takes control of decision-making. Call this new arrival 'self 1'. The preferences of self 1, like those of self 0, can be divided into two intervals. Self 1 has a present that lasts from $s_1 = s_0 + \tau_0$ to $s_1 + \tau_1$, and a future that lasts from $s_1 + \tau_1$ to ∞. Extending this idea, we assume that at each juncture of present and future a new self is born, yielding a sequence of selves born at dates $\{s_0, s_1, s_2, ...\}$, with respective presents of duration $\{\tau_0, \tau_1, \tau_2, ...\}$, where $s_{n+1} = s_n + \tau_n$. Figure 1 provides a visual representation.

We assume that all selves discount exponentially with discount factor $0 < \delta < 1$. Furthermore, each self values her future discretely less than her present, discounting it by the additional factor $0 < \beta \leq 1$. More explicitly, we assume that self n applies the discount factor $D_n(t)$ to the utility flow at time $s_n + t$, where

$$D_n(t) = \left\{ \begin{array}{lll} \delta^t & \text{if} & t \in [0, \tau_n) \\ \beta\,\delta^t & \text{if} & t \in [\tau_n, \infty) \end{array} \right\}. \tag{1}$$

In other words, her discount function D_n decays exponentially at rate $\gamma = -\ln \delta$ up to time τ_n, drops discontinuously at τ_n to a fraction β of its level just prior to τ_n, and decays exponentially at rate γ thereafter.[6] Figure 2 plots a single realization of this discount function, with $\tau_n = 3.4$.

This continuous-time discount function nests classical exponential discounting: either set $\lambda = 0$, so that the future never arrives; or set $\beta = 1$, so that there is no distinction between present and future. It is similar to some of the deterministic discount functions used in Barro (1999) and Luttmer and Mariotti (2003). However, we assume that τ_n is stochastic. Among other things, this ensures that the expectation of the discount function is smooth.

When $\lambda \to \infty$, the discount function D_n converges to the deterministic function D_∞

[6] The lengths $\{\tau_0, \tau_1, \tau_2, ...\}$ of the present intervals are i.i.d.

Figure 1: Sequential generations of autonomous selves.

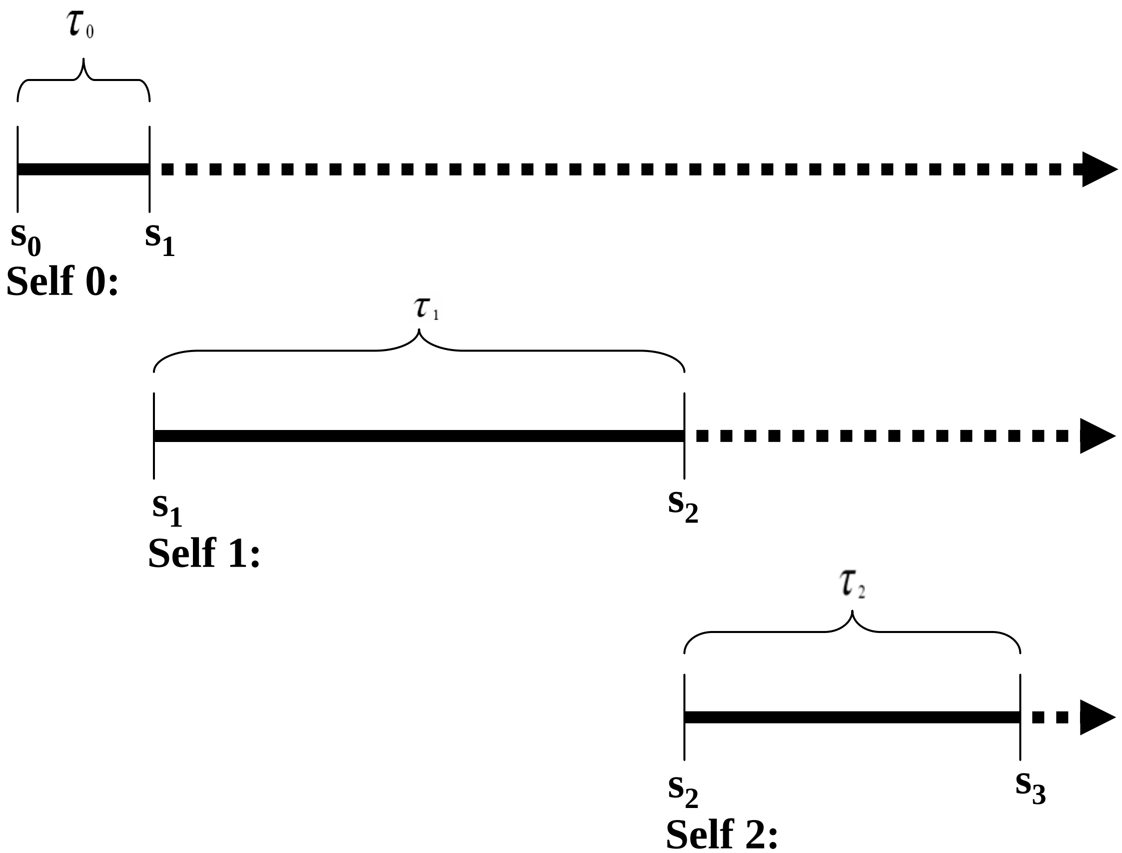

The span of control (solid line) of self n lasts from its time of birth ($t = s_n$) to the time of birth of self n+1 ($t = s_{n+1}$). The length of this control period, $s_{n+1} - s_n$, is the stochastic variable τ_n, which has an exponential distribution.

Figure 2: Realization of discount function ($\beta = 0.7$, $\gamma = 0.1$)

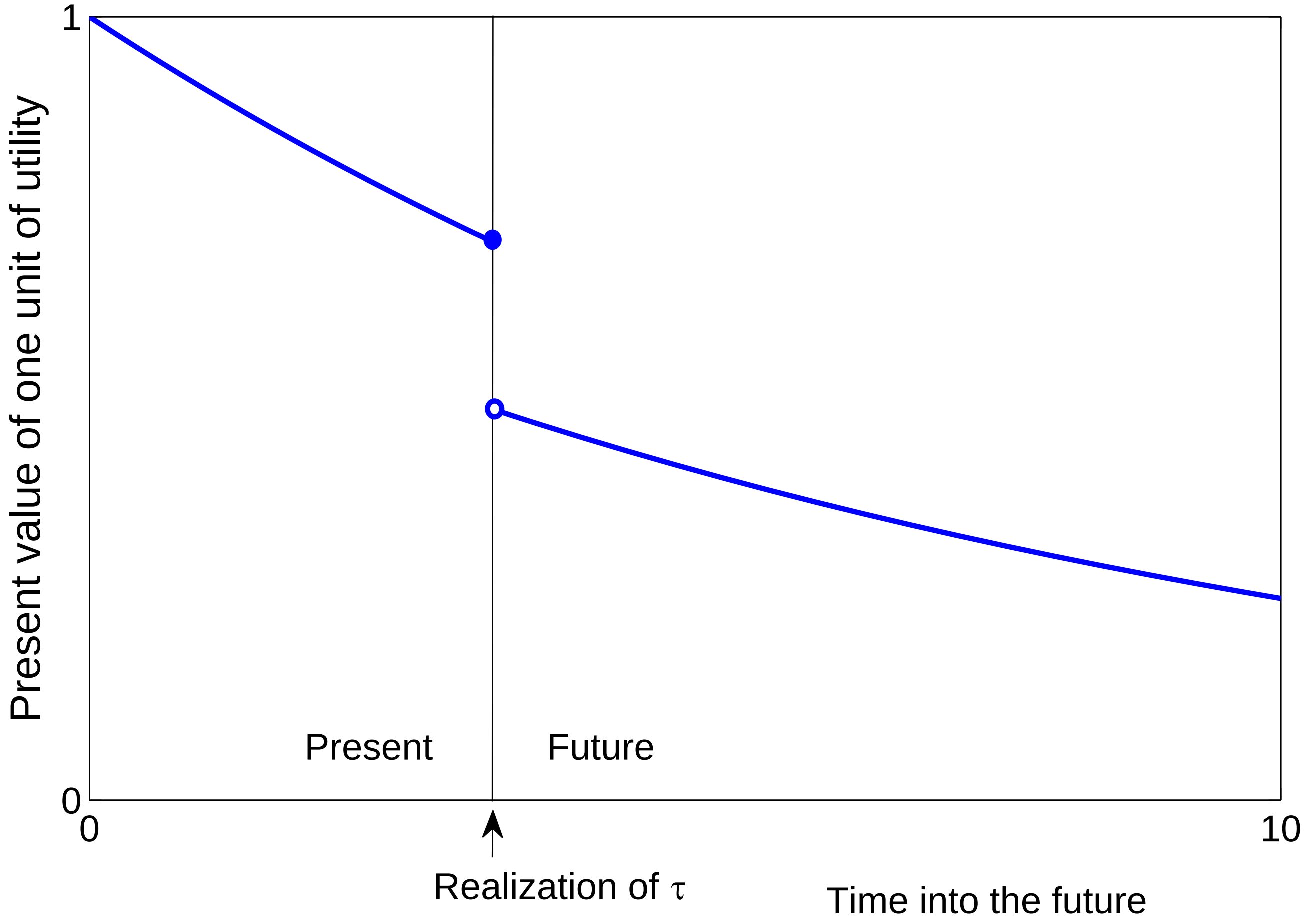

The discount function represents the present value of one unit of future utility. The discount function discretely drops when the present ends and the future begins. This present-to-future transition occurs at a stochastic time. Figure 2 shows a particular realization of this transition.

given by the formula

$$D_\infty(t) = \left\{ \begin{array}{lll} 1 & \text{if} & t = 0 \\ \beta\,\delta^t & \text{if} & t \in (0, \infty) \end{array} \right\}.$$

Characterizing this limiting case is the main focus of the current paper.[7]

2.2. A Reinterpretation Using a Deterministic Discount Function. The arguments in this paper are consistent with a second interpretation of the time preferences described above: one can assume that a new self is born every *instant*; that the present of each self lasts only an instant; and that each self has a *deterministic* discount function $\overline{D}$ equal to the expected value of the *stochastic* discount function D_n described above.[8] More precisely, each self uses the discount function $\overline{D}$ given by the formula

$$\overline{D}(t) = \mathrm{E}\big[D_n(t)\big] = e^{-\lambda\,t}\,\delta^t + (1 - e^{-\lambda\,t})\,\beta\,\delta^t.$$

$\overline{D}(t)$ is the sum of two terms. The first term is the probability $e^{-\lambda\,t}$ that the drop in D_n does not occur before time t, times the discount factor δ^t that applies prior to the drop. The second term is the probability $1 - e^{-\lambda\,t}$ with which the drop in D_n occurs after time t, times the discount factor $\beta\,\delta^t$ that applies after the drop. $\overline{D}(t)$ can also be written in the form

$$(1 - \beta)\,e^{-(\gamma+\lambda)\,t} + \beta\,e^{-\gamma\,t},$$

where $\gamma = -\ln(\delta) > 0$ is the long-run discount rate. Written this way, $\overline{D}(t)$ is seen to be a convex combination of the short-run exponential discount factor $e^{-(\gamma+\lambda)\,t}$, with weight $1 - \beta$, and the long-run exponential discount factor $e^{-\gamma\,t}$, with weight β.

The instantaneous discount rate associated with the deterministic discount function $\overline{D}$ is

$$-\frac{\overline{D}'(t)}{\overline{D}(t)} = \gamma + \frac{\lambda\,e^{-\lambda\,t}\,(1-\beta)\,\delta^t}{\overline{D}(t)}.$$

It too is the sum of two terms. The first term is the long-run (exponential) discount rate γ. The second term is the expected drop in D at time t, namely $\lambda\,e^{-\lambda\,t}\,(1-\beta)\,\delta^t$, divided by the level of $\overline{D}$ at time t. Indeed: $\lambda\,e^{-\lambda\,t}$ is the flow probability with which the drop in D occurs at time t; and $(1-\beta)\,\delta^t$ is the size of the drop in D if the drop occurs at t.

[7] Notice that, because $\tau_n \to 0$ as $\lambda \to \infty$, the expectation of the discount function — which is smooth when λ is finite — has a discontinuity when $\lambda = \infty$. This will not cause any problems for us.

[8] See footnote 12 for a development of this line of argument.

Notice that the instantaneous discount rate decreases from $\gamma + \lambda (1 - \beta)$ at $t = 0$ to γ at $t = \infty$. Figure 3 plots $\overline{D}$ for $\lambda \in \{0, 0.1, 1, \infty\}$.

2.3. Comparison of the Stochastic and Deterministic Discount Functions. The stochastic and deterministic discount functions differ in one important respect: the stochastic discount function assumes a present of non-infinitesimal duration $\tau_n > 0$, whereas the deterministic discount function assumes a present of infinitesimal duration dt. Hence the stochastic discount function assumes a countable number of non-infinitesimal selves, while the deterministic discount function assumes a continuum of infinitesimal selves.

The two formulations are however equivalent, in the critical sense that they generate the same equilibrium behavior. To see why, note that the current self in the stochastic formulation is dynamically consistent during her period of control between time s_n and the stochastic endpoint $s_n + \tau_n$. It therefore makes no difference whether we regard her as a non-infinitesimal agent, who decides how to behave at the outset of her control interval, or as a continuum of infinitesimal agents, each of which makes a decision during its instant of control.

The stochastic formulation has two advantages over the deterministic one. First, it can be set up using only standard mathematical tools. Second, when the stochastic formulation is used, we can derive the IG model in a single step.[9] We therefore focus on the stochastic formulation.

3. Application to a Consumption Problem

Two important qualitative features of consumers' planning problems are liquidity constraints and labor-income uncertainty (cf. Deaton 1991, and Carroll 1992, 1997). We include liquidity constraints, since they make an important difference to the analysis; but we exclude labor-income uncertainty, since it complicates the notation and does not affect our conclusions.

[9] In the analysis using the stochastic discount function, we let λ go to infinity. In doing so, we simultaneously pass from non-infinitesimal to infinitesimal selves and from the finite-λ discount function to the infinite-λ discount function that is the ultimate focus of the paper. By contrast, in order to set up the deterministic discount function, we would first have to formalize the idea of an infinitesimal self. This would involve taking the limit as the span of control of a non-infinitesimal self goes to zero. We would then have to let λ go to infinity, in order to pass from the finite-λ discount function to the infinite-λ discount function.

Figure 3: Expectation of discount function $\beta = 0.7$, $\gamma = 0.1$, $\lambda \in \{ 0, 0.1, 1, \infty \}$

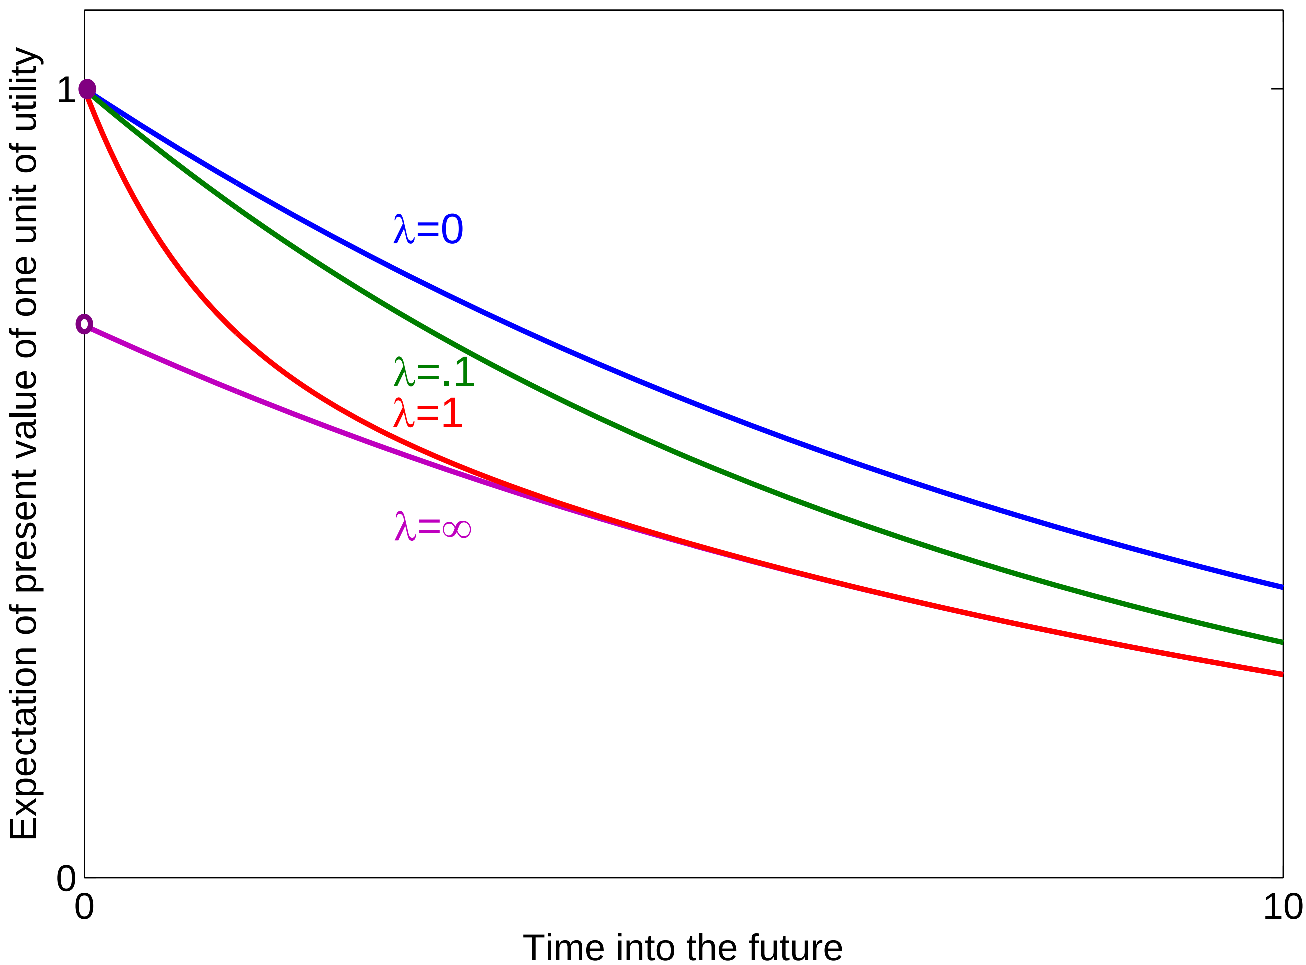

The expectation of the discount function represents the expected present value of one unit of future utility. The expectation integrates over the stochastic present-to-future transition time.

3.1. The Dynamics. At any given point in time $t \in [0, \infty)$, the consumer has stock of (financial) wealth $x \in [0, \infty)$ and receives a flow of labor income $y \in (0, \infty)$. If $x > 0$ then she can choose any consumption level $c \in (0, \infty)$: wealth is a stock and consumption is a flow, so any finite consumption level is achievable provided that it is not maintained for too long. If $x = 0$ then she can only choose a consumption level $c \in (0, y]$: she has no wealth and she cannot borrow, so she cannot consume more than her labor income. In a word, she is is liquidity constrained.

Whatever the consumer does not consume is invested in an asset, the returns on which are distributed normally with mean $\mu\, dt$ and variance $\sigma^2\, dt$, where $\mu \in (-\infty, \infty)$ and $\sigma \in (0, \infty)$. The change in her wealth at time t is therefore

$$dx = (\mu\, x + y - c)\, dt + \sigma\, x\, dz,$$

where z is a standard Wiener process.[10]

3.2. Equilibrium. Recall that the consumer is modeled as a sequence of autonomous selves (see Figure 1). Each self controls consumption in the present and cares about — but does not directly control — consumption in the future.

Consider self n, who takes over control at time s_n and passes control to self $n+1$ at time $s_n + \tau_n$. Recall that this transition occurs with hazard rate rate λ. In order to find the value $w(x_{s_n})$ to self n at time s_n of consumption over the time interval $[s_n, \infty)$, we first find the value $v(x_{s_n+\tau_n})$ to her at time $s_n + \tau_n$ of consumption over the time interval $[s_n + \tau_n, \infty)$. This is given by the formula

$$v(x_{s_n+\tau_n}) = \mathrm{E}_{s_n+\tau_n}\left[\int_{s_n+\tau_n}^{\infty} e^{-\gamma(t-(s_n+\tau_n))}\, u(\widetilde{c}(x_t))\, dt\right],$$

where: $\widetilde{c} : [0, \infty) \to (0, \infty)$ is the consumption function used by selves $n+1$, $n+2$ and so on; $u : (0, \infty) \to \mathbb{R}$ is the instantaneous utility function; $\gamma = -\ln(\delta) > 0$ is the long-run discount rate; $\mathrm{E}_{s_n+\tau_n}$ denotes expectations conditional on the information available at time $s_n + \tau_n$; and $x : [s_n + \tau_n, \infty) \to [0, \infty)$ is the stochastic timepath of wealth generated

[10] We could generalize this framework by adding a stochastic source of labor income. For example, we could assume that — in addition to her basic flow of labor income y — the agent sporadically receives lump-sum bonuses. To preserve stationarity, such bonuses would need to arrive with a constant hazard rate and be drawn from a fixed distribution. We could even allow for non-stationary labor income, at the expense of an extra state variable. We do not pursue these generalizations, since they would not qualitatively change the analysis that follows.

by $\widetilde{c}$ over the time interval $[s_n + \tau_n, \infty)$. We can then find $w(x_{s_n})$. It is given by the formula

$$w(x_{s_n}) = \mathrm{E}_{s_n}\left[\int_{s_n}^{s_n+\tau_n} e^{-\gamma(t-s_n)}\, u(c(x_t))\, dt + \beta\, v(x_{s_n+\tau_n})\right],$$

where: $c : [0,\infty) \to (0,\infty)$ is the consumption function used by self n; E_{s_n} denotes expectations conditional on the information available at time s_n; and $x : [s_n, s_n + \tau_n] \to [0,\infty)$ is the stochastic timepath of wealth generated by c over the time interval $[s_n, s_n + \tau_n]$.

We call v the continuation-value function and w the current-value function. A standard application of stochastic calculus implies that v satisfies the differential equation

$$0 = \tfrac{1}{2}\,\sigma^2\, x^2\, v'' + (\mu\, x + y - \widetilde{c})\, v' - \gamma\, v + u(\widetilde{c}) \tag{2}$$

for $x \in [0,\infty)$, where we have suppressed the dependence of v and $\widetilde{c}$ on x. This equation is derived by decomposing the value function: the expected instantaneous change in the value function (namely $\tfrac{1}{2}\,\sigma^2\, x^2\, v'' + (\mu\, x + y - \widetilde{c})\, v'$) minus the instantaneous loss of value due to discounting (namely $\gamma\, v$) plus the instantaneous utility flow (namely $u(\widetilde{c})$) must be zero. Similarly, w satisfies the differential equation

$$0 = \tfrac{1}{2}\,\sigma^2\, x^2\, w'' + (\mu\, x + y - c)\, w' + \lambda\,(\beta\, v - w) - \gamma\, w + u(c) \tag{3}$$

for $x \in [0,\infty)$, where we have suppressed the dependence of w and c on x. This equation for w is very similar to the equation for v. The only differences are: (i) the equation for v takes the continuation policy $\widetilde{c}$ as its consumption argument rather than the current policy c; and (ii) the equation for w includes the term $\lambda\,(\beta\, v - w)$, which reflects the hazard rate λ of making the transition from the present, valued by the current-value function w, to the future, valued by β times the continuation-value function v.

Finally, if self s behaves optimally — taking the behavior of her future selves as given — then c satisfies the instantaneous optimality condition

$$\left\{\begin{array}{ll} u'(c) = w' & \text{if either (i) } x > 0 \text{ or (ii) } x = 0 \text{ and } w' \geq u'(y) \\ c = y & \text{if } x = 0 \text{ and } w' < u'(y) \end{array}\right\}. \tag{4}$$

In other words, if either (i) $x > 0$ (in which case there is no constraint on consumption) or (ii) $x = 0$ (in which case there is a constraint on consumption) and the marginal current value of wealth w' is at least $u'(y)$ (so that the constraint is not binding), then consumption c is chosen so as to equate the marginal utility of consumption $u'(c)$ to w'. Similarly, if $x = 0$ and w' is less than $u'(y)$ (which means that the constraint on consumption is binding), then $c = y$.

We have a stationary equilibrium if and only if the consumption function c chosen by the current self turns out to be the same as the consumption function $\widetilde{c}$ used by the future selves, in other words if and only if we have a fixed point in consumption functions.

We therefore have the following characterization of equilibrium in the PF model:[11]

Theorem 1. *The consumption function c is a stationary Markov-perfect equilibrium of the PF model if and only if there is a continuation-value function v and a current-value function w such that (c, v, w) together satisfy the following system of differential equations for $x \in [0, \infty)$:*

$$0 = \tfrac{1}{2}\,\sigma^2\,x^2\,v'' + (\mu\,x + y - c)\,v' - \gamma\,v + u(c) \tag{5}$$

$$0 = \tfrac{1}{2}\,\sigma^2\,x^2\,w'' + (\mu\,x + y - c)\,w' + \lambda\,(\beta\,v - w) - \gamma\,w + u(c) \tag{6}$$

$$\left\{ \begin{array}{ll} u'(c) = w' & \text{if either (i) } x > 0 \text{ or (ii) } x = 0 \text{ and } w' \geq u'(y) \\ c = y & \text{if } x = 0 \text{ and } w' < u'(y) \end{array} \right\}. \tag{7}$$

We refer to this system as the Bellman system of the PF consumer.[12]

[11] Pathological equilibria with utility of $-\infty$ can also be constructed: if all future selves choose strategies that imply $-\infty$ discounted utility then the current self does not care if she also chooses a strategy that supports $-\infty$ discounted utility. Such equilibria can however be eliminated by specifying a more refined objective for the consumer.

[12] In the model with the deterministic discount function $\overline{D} : [0, \infty) \to [0, 1]$, the consumption function $c : [0, \infty) \to (0, \infty)$ is a stationary Markov-perfect equilibrium if and only if there is a value function $V : [0, \infty)^2 \to \mathbb{R}$ such that (c, V) together satisfy

$$0 = \tfrac{1}{2}\,\sigma^2\,x^2\,\frac{\partial^2 V}{\partial x^2}(t, x) + (\mu\,x + y - c(x))\,\frac{\partial V}{\partial x}(t, x) + \overline{D}(t)\,u(c(x)) + \frac{\partial V}{\partial t}(t, x)$$

for $(t, x) \in [0, \infty)^2$ and

$$\left\{ \begin{array}{ll} u'(c(x)) = \frac{\partial V}{\partial x}(0, x) & \text{if either (i) } x > 0 \text{ or (ii) } x = 0 \text{ and } \frac{\partial V}{\partial x}(0, x) \geq u'(y) \\ c(x) = y & \text{if } x = 0 \text{ and } \frac{\partial V}{\partial x}(0, x) < u'(y) \end{array} \right\}$$

for all $x \in [0, \infty)$. Here $V(t, x)$ is the value at time 0 of consumption over the interval $[t, \infty)$, and $c(x)$ is the optimal consumption at time 0 given the marginal value of wealth at time 0, namely $\frac{\partial V}{\partial x}(0, x)$. We refer to this pair of equations as the Bellman system of the $\overline{D}$ consumer. It is valid for general $\overline{D}$. However, if $\overline{D}(t) = e^{-\lambda\,t}\,\delta^t + (1 - e^{-\lambda\,t})\,\beta\,\delta^t$ as in the text, then it is easy to show that they have a solution (c, V) in which $V(t, x)$ takes the form $e^{-\lambda\,t}\,\delta^t\,w(x) + (1 - e^{-\lambda\,t})\,\beta\,\delta^t\,v(x)$ if and only if (c, v, w) satisfies the Bellman system of the basic consumer. In other words: any solution of the Bellman system

In general, the PF model can be expected to have a finite number of equilibria. If λ is close to 0 — a dynamically consistent limit case — equilibrium is unique. Likewise, if β is close to 1 — another dynamically consistent limit case — equilibrium is again unique. Much more interestingly, if λ is close to ∞ — a dynamically *inconsistent* limit case — equilibrium is unique. This is the case that we study next.

4. The Instantaneous-Gratification Model

Psychological considerations suggest that the present — in other words, the interval $[s, s+\tau)$ during which consumption is *not* down-weighted by β — is short. This is the same as saying that λ is large, since the future's arrival rate is λ. We are therefore led to consider the limiting case $\lambda \to \infty$, which serves as an approximation of situations in which the duration of the present (namely τ) is short. We refer to the limiting case as the *instantaneous-gratification* model, or IG model for short.

There are two possible ways of characterizing equilibrium in the IG model. The first approach is analogous to the approach that we used to characterize equilibrium in the PF model: we derive the Bellman system of the IG consumer directly from a careful analysis of her objective. The second approach is to derive the Bellman system of the IG consumer indirectly, by taking the limit of the Bellman system of the PF consumer as $\lambda \to \infty$.

The first approach has the advantage that it generates intuitive insights into the logic of the IG model. The second approach has the advantage that it is more compelling from a mathematical point of view.[13] We shall describe both approaches. However, the reader who is convinced by the first can certainly skip the second!

---

of the basic consumer generates a solution of the Bellman system of the $\overline{D}$ consumer; but the possibility remains that the Bellman system of the $\overline{D}$ consumer has other solutions that do not have this form.

[13] An equilibrium c of the PF model can be represented in one of two ways: it can be represented as a fixed point of the mapping taking the consumption function c_F employed by future selves into the consumption functions c_P that are optimal for the present self; or it can be represented as the first component c of a solution (c, v, w) of the Bellman system of the PF consumer. The first representation is traditionally regarded as the primary definition, and the second representation is traditionally regarded as a characterization. The second of our two approaches reverses this traditional point of view. It effectively identifies the PF model with the Bellman system of the PF consumer, and it identifies an equilibrium of the PF model with the first component c of a solution (c, v, w) of the Bellman system of the PF consumer. It then identifies the IG model with the limit of the Bellman system of the PF consumer, and it identifies an equilibrium of the IG model with the first component c of a solution (c, v) of Bellman system of the IG consumer. This approach is especially compelling in the light of the fact that the Bellman system of the IG consumer has a unique solution: not only do we have upper hemicontinuity (in the sense that the limit of any sequence of equilibria of the PF model is an equilibrium of the IG model); but we even have lower hemicontinuity (in the sense any equilibrium of the IG model is the limit of a sequence of equilibria of the PF model).

4.1. The First Approach. In formulating the objective of the IG consumer, it is important to bear in mind that her span of control is an instant. Changes in her behavior therefore have only an infinitesimal effect on her objective. Careful track must therefore be kept of such infinitesimal effects.

Consider self s, and suppose that all future selves use the consumption function $\widetilde{c} : [0, \infty) \to (0, \infty)$. Then the continuation-value function of self s is exactly the same as the continuation-value function of the PF consumer, namely v. In particular, v satisfies the differential equation

$$0 = \tfrac{1}{2}\,\sigma^2\,x^2\,v'' + (\mu\,x + y - \widetilde{c})\,v' - \gamma\,v + u(\widetilde{c}) \tag{8}$$

for $x \in [0, \infty)$, where we have suppressed the dependence of v and $\widetilde{c}$ on x.

Suppose further that self s has wealth x, and that she chooses the consumption level $c \in (0, \infty)$. Then the current value of self s is

$$w(x) = \mathrm{E}_s\big[u(c)\,dt + \beta\,v(x + dx)\big],$$

where $dx = (\mu\,x + y - c)\,dt + \sigma\,x\,dz$. Now, Itô's Lemma implies that

$$v(x + dx) = v(x) + v'(x)\,dx + \tfrac{1}{2}\,v''(x)\,(dx)^2$$

and $(dx)^2 = \sigma^2\,x^2\,dt$. Moreover, $\mathrm{E}_s\big[dx\big] = (\mu\,x + y - c)\,dt$. Hence

$$\begin{aligned} w(x) &= \mathrm{E}_s\big[u(c)\,dt + \beta\,\big(v(x) + v'(x)\,dx + \tfrac{1}{2}\,v''(x)\,(dx)^2\big)\big] \\ &= \beta\,v(x) + \big(u(c) + \beta\,v'(x)\,(\mu\,x + y - c) + \tfrac{1}{2}\,\beta\,v''(x)\,\sigma^2\,x^2\big)\,dt. \end{aligned}$$

That is, there are two contributions to the current value of self s: the non-infinitesimal contribution $\beta\,v$, and the infinitesimal contribution

$$\begin{aligned} &\big(u(c) + \beta\,v'\,(\mu\,x + y - c) + \tfrac{1}{2}\,\beta\,v''\,\sigma^2\,x^2\big)\,dt \\ &\qquad = (u(c) - \beta\,v'\,c)\,dt + \big(\beta\,v'\,(\mu\,x + y) + \tfrac{1}{2}\,\beta\,v''\,\sigma^2\,x^2\big)\,dt, \end{aligned}$$

where we have suppressed the dependence of v on x. Now, the non-infinitesimal contribution does not depend on c. Moreover the infinitesimal contribution only depends on c via the term $(u(c) - \beta\,v'\,c)\,dt$. Hence, in order to maximize her current value, self s need only choose c to maximize

$$u(c) - \beta\, v'\, c.$$

Bearing in mind that self s is free to choose any $c \in (0, \infty)$ when $x > 0$, and that she must choose $c \in (0, y]$ when $x = 0$, it follows that c must satisfy the optimality condition

$$\left\{ \begin{array}{ll} u'(c) = \beta\, v' & \text{if either (i) } x > 0 \text{ or (ii) } x = 0 \text{ and } \beta\, v' \geq u'(y) \\ c = y & \text{if } x = 0 \text{ and } \beta\, v' < u'(y) \end{array} \right\}. \tag{9}$$

In other words, consumption c is chosen so that $u'(c) = \beta\, v'$ if either (i) $x > 0$ (so the liquidity constraint is absent) or (ii) $x = 0$ but $\beta\, v' \geq u'(y)$ (so the liquidity constraint is present but not binding). Similarly, $c = y$ if $x = 0$ and $\beta\, v' < u'(y)$ (so the liquidity constraint is present and binding).

Finally, we have a stationary equilibrium if and only if, whatever the initial wealth x of the current self, her consumption level c is the same as the consumption level that would be chosen by any future self who found herself with wealth x, namely $\widetilde{c}(x)$.

We therefore have the following characterization of equilibrium in the IG model:

Theorem 2. *The consumption function c is a stationary Markov-perfect equilibrium of the IG model if and only if there is a continuation-value function v such that (c, v) together satisfy the following system of differential equations for $x \in [0, \infty)$:*

$$0 = \tfrac{1}{2}\, \sigma^2\, x^2\, v'' + (\mu\, x + y - c)\, v' - \gamma\, v + u(c) \tag{10}$$

$$\left\{ \begin{array}{ll} u'(c) = \beta\, v' & \textit{if either (i) } x > 0 \textit{ or (ii) } x = 0 \textit{ and } \beta\, v' \geq u'(y) \\ c = y & \textit{if } x = 0 \textit{ and } \beta\, v' < u'(y) \end{array} \right\}. \tag{11}$$

We refer to this system as the Bellman system of the IG consumer.

Notice that, in the special case in which $\beta = 1$, the Bellman system of the IG consumer reduces to the Bellman system of an exponential consumer who has utility function u and discount rate γ.

4.2. The Second Approach. In the preceding subsection we formulated the objective of the IG consumer and then found the Bellman system characterizing equilibrium. In the current subsection, we give an alternative derivation: we show that the Bellman system of the IG consumer can be obtained by taking the limit of the Bellman system of the PF consumer as $\lambda \to \infty$.

Suppose that the triple $(c_\lambda, v_\lambda, w_\lambda)$ solves the Bellman system of the PF consumer. In other words, the following equations hold for $x \in [0, \infty)$:

$$0 = \tfrac{1}{2}\,\sigma^2\,x^2\,v''_\lambda + (\mu\,x + y - c_\lambda)\,v'_\lambda - \gamma\,v_\lambda + u(c_\lambda) \tag{12}$$

$$0 = \tfrac{1}{2}\,\sigma^2\,x^2\,w''_\lambda + (\mu\,x + y - c_\lambda)\,w'_\lambda + \lambda\,(\beta\,v_\lambda - w_\lambda) - \gamma\,w_\lambda + u(c_\lambda) \tag{13}$$

$$\left\{\begin{array}{ll} u'(c_\lambda) = w'_\lambda & \text{if either (i) } x > 0 \text{ or (ii) } x = 0 \text{ and } w'_\lambda \geq u'(y) \\ c_\lambda = y & \text{if } x = 0 \text{ and } w'_\lambda < u'(y) \end{array}\right\}. \tag{14}$$

Suppose further that $(c_\lambda, v_\lambda, w_\lambda) \to (c, v, w)$ as $\lambda \to \infty$. Then the equations characterizing (c, v, w) can be derived as follows.

Note first that equation (12) — which is effectively an equation for the continuation-value function — does not depend directly on λ. Indeed, equation (12) only applies after the transition to the future has taken place, so it is not affected in any way by the arrival rate of the future. Letting $\lambda \to \infty$ therefore preserves the form of this equation, yielding the first equation in Theorem 2 above, namely

$$0 = \tfrac{1}{2}\,\sigma^2\,x^2\,v'' + (\mu\,x + y - c)\,v' - \gamma\,v + u(c). \tag{15}$$

In other words, just as v_λ was the expected present discounted value obtained when consumption was chosen according to the exogenously given consumption function c_λ, so v is the expected present discounted value obtained when consumption is chosen according to the exogenously given consumption function c.

On the other hand, equation (13) — which is effectively an equation for the current-value function — does depend directly on λ. This equation can, however, be rearranged to give

$$w_\lambda - \beta\,v_\lambda = \tfrac{1}{\lambda}\left(\tfrac{1}{2}\,\sigma^2\,x^2\,w''_\lambda + (\mu\,x + y - c_\lambda)\,w'_\lambda - \gamma\,w_\lambda + u(c_\lambda)\right).$$

Letting $\lambda \to \infty$ then yields

$$w - \beta\,v = 0. \tag{16}$$

This reflects the fact that, as $\lambda \to \infty$, the discount function drops essentially immediately to a fraction β of its initial value, and that the current-value function w is therefore β times the continuation-value function v.

Thirdly, equation (14) — which is effectively an equation for the consumption function — does not depend directly on λ. Letting $\lambda \to \infty$ therefore preserves the form of this equation, and we have

$$\left\{ \begin{array}{ll} u'(c) = w' & \text{if either (i) } x > 0 \text{ or (ii) } x = 0 \text{ and } w' \geq u'(y) \\ c = y & \text{if } x = 0 \text{ and } w' < u'(y) \end{array} \right\}. \tag{17}$$

In other words, just as c_λ was the optimal consumption function when the current-value function was w_λ, so c is the optimal consumption function when the current-value function is w. Finally, substituting for w in equation (17) from equation (16), we obtain the second equation in Theorem 2.

5. Existence, Uniqueness and Value-Function Equivalence

The single most important property of the IG model is uniqueness: the IG model resolves the multiplicity problem that has plagued the literature on dynamically inconsistent preferences. The current section explains why the IG model yields a unique equilibrium: the Bellman system of the IG consumer coincides with the Bellman system of a related dynamically consistent optimization problem.

The section also discusses a number of other issues, including the existence of a natural equilibrium refinement that picks out a unique equilibrium in the deterministic model.

5.1. Assumptions. Before proceeding, we make the following simple assumptions:[14]

A1 $u(c) = \left\{ \begin{array}{ll} \frac{1}{1-\rho}\left(c^{1-\rho} - 1\right) & \text{if } \rho \neq 1 \\ \ln(c) & \text{if } \rho = 1 \end{array} \right\}$;

A2 $1 - \beta < \rho$;

A3 $\mu < \frac{1}{1-\rho}\gamma + \frac{1}{2}\rho\sigma^2$ if $\rho < 1$.

These assumptions can be weakened considerably, in an economically interesting way, without affecting any of our results.[15] See Section 7 below.

Assumption A1 is standard. It means that u has constant relative risk aversion ρ. Assumption A2 means that the dynamic inconsistency of the IG consumer (as measured by $1-\beta$) is less than the coefficient of relative risk aversion (namely ρ).[16] This assumption

[14]It may also be helpful to recall the mathematically natural restrictions that we have already placed on the parameters, namely $\beta \in (0, 1]$, $\gamma \in (0, \infty)$, $\mu \in (-\infty, \infty)$ and $\sigma \in (0, \infty)$.

[15]To be explicit, Theorems 3-5 and 7-9 all continue to hold as stated under the weaker assumptions. The only exception is Theorem 6, the statement of which needs to be changed slightly to reflect the more general context. See Theorem 10.

[16]The case $1 - \beta > \rho$ can also be analyzed. In that case, the consumer consumes all her wealth immediately and we have $v = \frac{1}{\gamma}u(y)$.

would be satisfied in a standard calibration: empirical estimates of the coefficient of relative risk aversion ρ typically lie between $\frac{1}{2}$ and 5; and the short-run discount factor β is typically thought to lie between $\frac{1}{2}$ and 1.[17] Assumption A3 is adapted from a standard assumption.[18] It is designed to ensure that the expected payoff of the consumer is finite. To see how it works, note first that the consumer can always choose to consume her labor income $y > 0$. Her expected payoff is therefore bounded below (by $\frac{1}{\gamma}\,u(y)$). The main challenge is therefore to ensure that her expected payoff is bounded above. Moreover this challenge only arises if $\rho \leq 1$, i.e. if the consumer's utility function is unbounded above. In that case, we have to ensure that wealth does not grow too fast. Notice that the right-hand side of the inequality for μ is increasing in ρ, and goes to ∞ as ρ goes to 1. In other words; the slower utility grows with consumption, the less stringent the restriction on μ becomes; and no restriction on μ is necessary if $\rho \geq 1$.

5.2. Value-Function Equivalence. Armed with Assumptions A1-A3, we can establish the following result:

Theorem 3 [Value-Function Equivalence]**.** *There exist strictly increasing and concave utility functions $\widehat{u}_+ : (0,\infty) \to \mathbb{R}$ and $\widehat{u}_0 : (0,y] \to \mathbb{R}$ such that, if we*

1. *define a wealth-dependent utility function $\widehat{u}$ by the formula*

$$\widehat{u}(\widehat{c},x) = \left\{ \begin{array}{ll} \widehat{u}_+(\widehat{c}) & \text{if } \widehat{c} \in (0,\infty) \text{ and } x > 0 \\ \widehat{u}_0(\widehat{c}) & \text{if } \widehat{c} \in (0,y] \text{ and } x = 0 \end{array} \right\};$$

2. *introduce a new consumer, whom we shall refer to as the $\widehat{u}$ consumer, who:*
   (a) *discounts the future exponentially at rate γ,*
   (b) *faces the same wealth dynamics as the IG consumer and*
   (c) *has the utility function $\widehat{u}$;*

then the set of possible value functions of the IG consumer coincides with the set of possible value functions of the $\widehat{u}$ consumer.

[17] See Laibson et al (1998) and Ainslie (1992).

[18] In the model with $y = 0$ and $\beta = 1$, it is standard to assume that $\gamma > (1-\rho)\,(\mu - \frac{1}{2}\,\rho\,\sigma^2)$. This ensures that the average propensity to consume is strictly positive. In the present model, for reasons explained below, we only need to consider the case $\rho < 1$. In that case, the standard inequality can be rearranged to give $\mu < \frac{1}{1-\rho}\,\gamma + \frac{1}{2}\,\rho\,\sigma^2$.

Using this result we can reduce the study of the problem of the IG consumer, which is game-theoretic, to the study of the problem of the $\widehat{u}$ consumer, which is decision-theoretic (i.e. non-strategic). There is, however, an important caveat: while the set of possible *value* functions of the IG consumer coincides with the set of possible *value* functions of the $\widehat{u}$ consumer, it is not the case that the set of possible *consumption* functions of the IG consumer coincides with the set of possible *consumption* functions of the $\widehat{u}$ consumer. In particular, value-function equivalence does not translate into observational equivalence.

Proof. We begin by eliminating c from the Bellman *system* of the IG consumer to yield what we call the Bellman *equation* of the IG consumer. Second, we eliminate $\widehat{c}$ from the Bellman *system* of the $\widehat{u}$ consumer to arrive at what we call the Bellman *equation* of the $\widehat{u}$ consumer. Third, we note that if we put

$$\widehat{u}_+(\widehat{c}) = \frac{\psi}{\beta}\, u\Big(\frac{\widehat{c}}{\psi}\Big) + \frac{\psi - 1}{\beta} \text{ for } \widehat{c} \in (0, \infty)$$

and

$$\widehat{u}_0(\widehat{c}) = \left\{ \begin{array}{ll} \widehat{u}_+(\widehat{c}) & \text{for } \widehat{c} \in (0, \psi\, y] \\ \widehat{u}_+(\psi\, y) + (\widehat{c} - \psi\, y)\, \widehat{u}'_+(\psi\, y) & \text{for } \widehat{c} \in [\psi\, y, y] \end{array} \right\},$$

where

$$\psi = \frac{\rho - (1 - \beta)}{\rho},$$

then the Bellman equation of the IG consumer is identical to the Bellman equation of the $\widehat{u}$ consumer. The two equations must therefore have the same set of solutions.

More precisely, for all $\alpha > 0$: let $f_+(\alpha)$ be the unique solution of the equation $u'(c) = \alpha$; and let $f_0(\alpha)$ be the unique c satisfying

$$\left\{ \begin{array}{ll} u'(c) = \alpha & \text{if } \alpha \geq u'(y) \\ c = y & \text{if } \alpha < u'(y) \end{array} \right\}.$$

Put $h_+(\alpha) = u(f_+(\beta\, \alpha)) - \alpha\, f_+(\beta\, \alpha)$; put $h_0(\alpha) = u(f_0(\beta\, \alpha)) - \alpha\, f_0(\beta\, \alpha)$; and put

$$h(\alpha, x) = \left\{ \begin{array}{ll} h_+(\alpha) & \text{if } x > 0 \\ h_0(\alpha) & \text{if } x = 0 \end{array} \right\}.$$

Then the Bellman *system* of the IG consumer reduces to the *single* differential equation

$$0 = \tfrac{1}{2}\,\sigma^2\,x^2\,v'' + (\mu\,x + y)\,v' - \gamma\,v + h(v', x) \tag{18}$$

for $x \in [0, \infty)$. We shall refer to this equation as the Bellman *equation* of the IG consumer.

Next, let $\widehat{u}_+$, $\widehat{u}_0$ and $\widehat{u}$ be given exactly as above. Then standard considerations show that, on $x \in [0, \infty)$, the Bellman system of the $\widehat{u}$ consumer takes the form

$$0 = \tfrac{1}{2}\,\sigma^2\,x^2\,\widehat{v}'' + (\mu\,x + y - \widehat{c})\,\widehat{v}' - \gamma\,\widehat{v} + \widehat{u}(\widehat{c}, x)$$
$$\left\{ \begin{array}{ll} \frac{\partial \widehat{u}}{\partial \widehat{c}}(\widehat{c}, x) = \widehat{v}' & \text{if either (i) } x > 0 \text{ or (ii) } x = 0 \text{ and } \widehat{v}' \geq \widehat{u}_0'(y) \\ \widehat{c} = y & \text{if } x = 0 \text{ and } \widehat{v}' < \widehat{u}_0'(y) \end{array} \right\}.$$

Let $\widehat{f}_+(\alpha)$ be the unique solution of $\widehat{u}_+'(\widehat{c}) = \alpha$; and let $\widehat{f}_0(\alpha)$ be any $\widehat{c}$ satisfying

$$\left\{ \begin{array}{ll} \widehat{u}_+'(\widehat{c}) = \alpha & \text{if } \alpha > \widehat{u}_+'(\psi\,y) \\ \widehat{c} \in [\psi\,y, y] & \text{if } \alpha = \widehat{u}_+'(\psi\,y) \\ \widehat{c} = y & \text{if } \alpha < \widehat{u}_+'(\psi\,y) \end{array} \right\}.$$

Put $\widehat{h}_+(\alpha) = u(\widehat{f}_+(\alpha)) - \alpha\,\widehat{f}_+(\alpha)$; put $\widehat{h}_0(\alpha) = u(\widehat{f}_0(\alpha)) - \alpha\,\widehat{f}_0(\alpha)$; and put

$$\widehat{h}(\alpha, x) = \left\{ \begin{array}{ll} \widehat{h}_+(\alpha) & \text{if } x > 0 \\ \widehat{h}_0(\alpha) & \text{if } x = 0 \end{array} \right\}.$$

Then the Bellman *system* of the $\widehat{u}$ consumer for $x \in [0, \infty)$ reduces to the *single* equation

$$0 = \tfrac{1}{2}\,\sigma^2\,x^2\,\widehat{v}'' + (\mu\,x + y)\,\widehat{v}' - \gamma\,\widehat{v} + \widehat{h}(\widehat{v}', x). \tag{19}$$

We shall refer to this equation as the Bellman *equation* of the $\widehat{u}$ consumer.

Now, it is easy to see that equations (18) and (19) will be identical iff the functions h and $\widehat{h}$ are the same. Moreover, as can be shown by direct calculation, this is indeed the case for the given choice of $\widehat{u}_+$, $\widehat{u}_0$ and $\widehat{u}$. ■

Figure 4 depicts a portion of the graphs of u, $\widehat{u}_+$ and $\widehat{u}_0$ in the case in which $\beta = \frac{2}{3}$, $\rho = 2$ and $y = 1$. It illustrates several important features of u, $\widehat{u}_+$ and $\widehat{u}_0$. First, we have $\widehat{u}_+(\widehat{c}) < u(\widehat{c})$ for all $\widehat{c} > 0$ and $\widehat{u}_0(\widehat{c}) < u(\widehat{c})$ for all $0 < \widehat{c} < y$. This makes sense: the $\widehat{u}$ consumer optimizes fully while the IG consumer does not. Hence the $\widehat{u}$ consumer must be

Figure 4: graph of u, $\hat{u}_+$ and $\hat{u}_0$

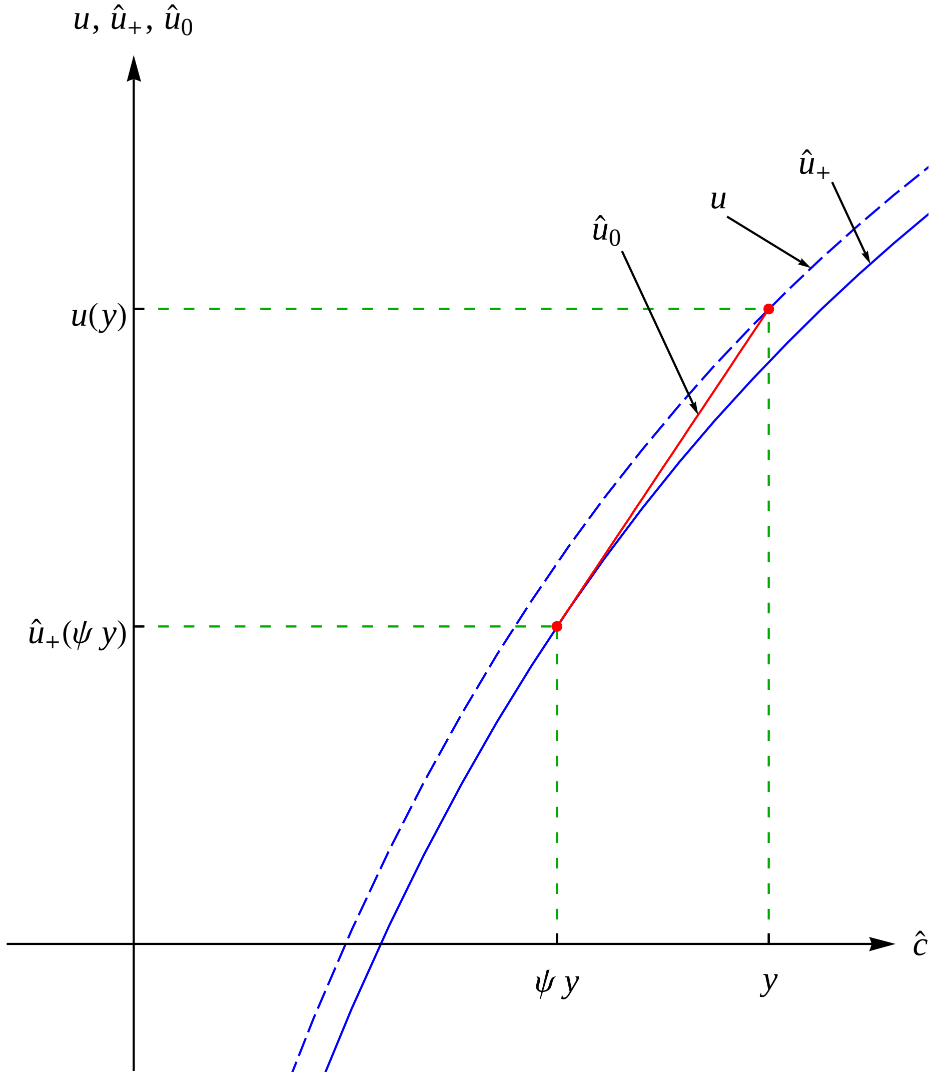

suitably handicapped in order to prevent her from achieving a higher value than the IG consumer. Second, we have $\widehat{u}_0(y) = u(y)$. Once again this makes sense: in the liquidity constrained case, both the $\widehat{u}$ consumer and the IG consumer consume their labor income y forever. So we must have $\widehat{u}_0(y) = u(y)$ if they are both to obtain the same value. Third, the graph of $\widehat{u}_0$ coincides with that of $\widehat{u}_+$ for $\widehat{c} \in (0, \psi\, y]$, and coincides with the tangent to the graph of $\widehat{u}_+$ at $\psi\, y$ for $\widehat{c} \in (\psi\, y, y]$. Notice that the graphs are truncated from below at a utility level of 0.4 and from the right at a consumption level of 1.1. Notice too that the axes intersect at the point $(0.5, -0.5)$, and not at $(0, 0)$.

One important consequence of the Value-Function-Equivalence Theorem is the existence and uniqueness of equilibrium in the IG model:

Theorem 4 [Existence and Uniqueness]**.** *The IG model has a unique equilibrium.*

Proof. Standard considerations show that, if $\widehat{u}_+$ and $\widehat{u}_0$ are chosen as in the proof of Theorem 3, then the Bellman equation of the $\widehat{u}$ consumer has a unique solution $\widehat{v}$. It follows that the Bellman equation of the IG consumer has a unique solution v (and that $v = \widehat{v}$). Finally, the consumption function c associated with v is given by the formula

$$c = \left\{ \begin{array}{ll} f_+(\beta\, v') & \text{if } x > 0 \\ f_0(\beta\, v') & \text{if } x = 0 \end{array} \right\},$$

where f_+ and f_0 are as in the proof of Theorem 3 above. ■

5.3. The Deterministic Case: A Refinement. Until now we have assumed that the standard deviation of asset returns is strictly positive ($\sigma > 0$). In other words, we have been studying the *stochastic* IG model. In the present subsection, we investigate the *deterministic* model ($\sigma = 0$). By viewing the deterministic IG model as a limiting case of the stochastic IG model, we can pinpoint a unique value function for the deterministic IG model. More precisely, we have the following theorem. The proof, which follows standard lines, is omitted.

Theorem 5. *Let v_σ be the value function of the stochastic IG model. Then:*

1. *there is a continuous function $v : [0, \infty) \to \mathbb{R}$ such that $v_\sigma \to v$ uniformly on compact subsets of $[0, \infty)$ as $\sigma \to 0$;*

2. *there exists a consumption function c such that (c,v) together satisfy the following system of differential equations for $x \in [0,\infty)$:*

$$0 = (\mu\,x + y - c)\,v' - \gamma\,v + u(c) \tag{20}$$

$$\left\{\begin{array}{ll} u'(c) = \beta\,v' & \text{if either (i) } x > 0 \text{ or (ii) } x = 0 \text{ and } \beta\,v' \geq u'(y) \\ c = y & \text{if } x = 0 \text{ and } \beta\,v' < u'(y) \end{array}\right\}. \tag{21}$$

We refer to this system as the Bellman system of the deterministic IG consumer. ■

The Bellman system of the deterministic IG consumer is significantly simpler than the Bellman system of the stochastic IG consumer: the former can be transformed into an autonomous first-order differential equation, whereas the latter is a second-order non-autonomous differential equation which cannot be transformed into a simpler form.[19]

In other words, by letting $\sigma \to 0$ in the stochastic IG model, we select a unique sensible equilibrium of the deterministic IG model. Krusell and Smith (2000) consider a deterministic discrete-time hyperbolic consumption model. They show that not only that there are multiple equilibria in their model, but even that equilibrium is indeterminate (i.e. there is a continuum of equilibria). Our results provide a refinement that eliminates not only the indeterminacy but even the multiplicity.

5.4. Consumption Function of the $\widehat{u}$ Consumer. Finally, the consumption function c of the IG consumer and the consumption function $\widehat{c}$ of the $\widehat{u}$ consumer are both determined by v' and x. There is therefore a close relationship between them:

Theorem 6 [Proportionality]**.** *Put $\psi = \frac{\rho-(1-\beta)}{\rho}$. Then $0 < \psi \leq 1$. Moreover*

$$\left\{\begin{array}{ll} c = \frac{1}{\psi}\,\widehat{c} & \text{if either (i) } x > 0 \text{ or (ii) } x = 0 \text{ and } \beta\,v' > u'(y) \\ c = \widehat{c} = y & \text{if } x = 0 \text{ and } \beta\,v' < u'(y) \end{array}\right\}.$$

In other words, there are two possible cases. If the liquidity constraint does not bind (i.e. if $\beta\,v'(0) > u'(y)$), then there is a fixed scalar $\psi \leq 1$ such that $c(x) = \frac{1}{\psi}\,\widehat{c}(x)$ for all $x \geq 0$. On the other hand, if the liquidity constraint does bind (i.e. if $\beta\,v'(0) < u'(y)$), then $c(x) = \frac{1}{\psi}\,\widehat{c}(x)$ for all $x > 0$ but $c(0) = \widehat{c}(0) = y$. Either way, if we restrict attention

[19] See an earlier draft of this paper, Harris and Laibson (2006), for a complete characterization of the value and policy functions of the deterministic case.

to the interior of the wealth space, then the consumption function of the IG consumer is a scalar multiple of the utility function of the $\widehat{u}$ consumer.

Proof. Suppose first that $x > 0$. Then c is determined by the equation $u'(c) = \beta\, v'$ and $\widehat{c}$ is determined by the equation $\widehat{u}'_+(\widehat{c}) = v'$. Moreover it follows from the formula for $\widehat{u}_+$ given in the proof of Theorem 3 that $\widehat{u}'_+(\widehat{c}) = \frac{1}{\beta}\, u'(\frac{\widehat{c}}{\psi})$. Hence $u'(c) = \beta\, v' = \beta\, \widehat{u}'_+(\widehat{c}) = u'(\frac{\widehat{c}}{\psi})$, and therefore $c = \frac{1}{\psi}\, \widehat{c}$. Suppose next that $x = 0$. Then c is determined by the equation $u'(c) = \beta\, v'$ if $\beta\, v' > u'(y)$ and $\widehat{c}$ is determined by the equation $\widehat{u}'_0(\widehat{c}) = v'$ if $v' > \widehat{u}'_0(y)$. Similarly, we have $c = y$ if $\beta\, v' < u'(y)$ and $\widehat{c} = y$ if $v' < \widehat{u}'_0(y)$. Now, it follows from the formulae for $\widehat{u}_0$ and $\widehat{u}_+$ given in the proof of Theorem 3 that $\widehat{u}'_0(y) = \widehat{u}'_+(\psi\, y) = \frac{1}{\beta}\, u'(y)$. Hence $\beta\, v' > u'(y)$ iff $v' > \widehat{u}'_0(y)$ and, in this case, we have $c = \frac{1}{\psi}\, \widehat{c}$ as in the previous paragraph. Similarly, $\beta\, v' < u'(y)$ iff $v' < \widehat{u}'_0(y)$ and, in this case, we have $c = \widehat{c} = y$. ■

6. The Consumption Function of the IG Consumer

In this section, we give a detailed description of the consumption function of the IG consumer, i.e. the function that maps the IG consumer's (financial) wealth to her consumption. Three key properties emerge.

First, the consumption function is continuous in the interior of the wealth space. This is a consequence of Brownian motion in the stochastic process for wealth, which smooths out the value function in the interior of the wealth space and thereby also eliminates discontinuities in the consumption function.

Second, the consumption function may have a discontinuity at the point where the liquidity constraint binds. If the expected rate of return is low enough, the consumption function jumps up when the liquidity constraint ceases to bind. This is a consequence of the consumer's propensity to value immediate rewards discretely more than delayed rewards.

Third, the consumption function may have one region in which it is downward sloping. Such downward sloping intervals will exist if the expected rate of return takes an intermediate value (defined below). However, this non-monotonicity vanishes when a bond is introduced and the investor can take both long and short positions in this asset.

These properties contrast with the properties of the continuous-time *exponential* model, which has an (everywhere) continuous, monotonic consumption function. They also contrast with the properties of the *discrete-time* quasi-hyperbolic model, the consumption function of which may have *several* downward sloping regions and a *countable* number of downward jumps (cf. Laibson 1997, Morris and Postlewaite 1997, Krusell and Smith 2000, Harris and Laibson 2001, Morris 2002).[20]

6.1. Comparative Statics on μ. In order to keep our description of the behavior of the consumption function as simple as possible, we hold the parameters β, γ, σ and ρ fixed and vary μ. (Recall that β is the supplementary discount factor applied to the future, γ is the long-run rate of discounting, μ is the expectation of the return on wealth, σ is the standard deviation of the return on wealth and ρ is the coefficient of relative risk aversion.) There are then three cases to consider. Indeed, recalling Assumption A3, put

$$\overline{\mu} = \left\{ \begin{array}{cc} \frac{1}{1-\rho}\,\gamma + \frac{1}{2}\,\rho\,\sigma^2 & \text{if } \rho < 1 \\ \infty & \text{if } \rho \geq 1 \end{array} \right\}.$$

Then $\overline{\mu} > \gamma$ and there exists $\mu_1 \in (\gamma, \overline{\mu})$ such that the form of the consumption function depends on whether μ lies in the interval $(-\infty, \gamma)$, the interval (γ, μ_1) or the interval $(\mu_1, \overline{\mu})$. We refer to these three cases as the case in which μ is low, the case in which μ is intermediate and the case in which μ is high.

6.2. The Case in which μ is Low. The most interesting case of our model is that in which $\mu < \gamma$. In this case, the expected returns on the asset are not sufficiently attractive to induce the IG consumer to save when her wealth is zero, and the liquidity constraint binds. More precisely, let $\overline{c} \in (y, \infty)$ be the unique solution of the equation

$$u'(\overline{c}) = \beta\, \frac{u(\overline{c}) - u(y)}{\overline{c} - y}. \tag{22}$$

Then:

Theorem 7. *If $\mu < \gamma$ then: $c(0) = y$; $c(0+) = \overline{c} > y$; and $c' > 0$ on $(0, \infty)$.*

In other words: when the IG consumer has no wealth, she consumes all of her labor income; if she acquires even a little wealth, then her consumption jumps up from y to $\overline{c}$; and her consumption increases steadily with further increases in her wealth. In particular, her consumption function is strictly increasing.

Proof. See Appendix A. ■

[20] These multiple downward sloping regions and jumps are not eliminated when a bond is added to the discrete-time model.

Equation (22) can be understood as follows. Let us refer to the moment at which wealth runs out as the 'crunch'. Suppose that the consumption level of the pre-crunch self is $\overline{c}$. Then the cost to the pre-crunch self of putting aside an extra dx units of wealth is $u'(\overline{c})\,dx$. On the other hand, if the post-crunch self receives a windfall consisting of an extra dx units of wealth, then she can raise her consumption level from y to $\overline{c}$ for a length of time $dt = dx\,/\,(\overline{c} - y)$. The benefit to the post-crunch self of this increase in consumption is $(u(\overline{c}) - u(y))\,dt$, and the benefit to the pre-crunch self is $\beta\,(u(\overline{c}) - u(y))\,dt$. The pre-crunch self is therefore indifferent between putting aside the extra dx units of wealth and not putting them aside if and only if

$$u'(\overline{c})\,dx = \beta\,(u(\overline{c}) - u(y))\,dt.$$

Substituting for dt and dividing through by dx, we obtain equation (22).

Figure 5 shows three consumption functions, obtained by fixing $\beta = \frac{2}{3}$, $\gamma = 0.05$, $\sigma = 0.17$, $\rho = 2$ and $y = 1$ (which implies in particular that $\overline{c} = \frac{3}{2}\,y$) and allowing μ to vary over the set $\{0.04, 0.06, 0.10\}$. The top consumption function corresponds to $\mu = 0.04$ and illustrates the low-μ case; the middle consumption function corresponds to $\mu = 0.06$ and illustrates the intermediate-μ case; and the bottom consumption function corresponds to $\mu = 0.10$ and illustrates the high-μ case. As Theorem 7 leads us to expect, for the top consumption function: the liquidity constraint is binding, i.e. $c(0) = y = 1$; there is an upward jump in consumption at $x = 0$, from $c(0) = 1$ to $c(0+) = \overline{c} = \frac{3}{2}$; and consumption rises steadily thereafter.

6.3. The Case in which μ is High. The other polar case of our model is that in which $\mu > \mu_1$. In this case, the expected returns on the asset are sufficiently attractive to induce the IG consumer to save even when her wealth is zero, and the liquidity constraint does not bind. More precisely:

Theorem 8. *If $\mu > \mu_1$ then: $c(0) < y$; $c(0+) = c(0)$; and $c' > 0$ on $[0, \infty)$.*

In other words: even when the IG consumer has no wealth, she still chooses to save out of her labor income; acquiring a little wealth does not lead to a jump in her consumption; and her consumption increases steadily with further increases in her wealth. In particular, her consumption function is strictly increasing.

Proof. See Appendix A. ■

Figure 5: Consumption function

$\beta = 2/3$, $\gamma = 0.05$, $\sigma = 0.17$, $\rho = 2$, $y = 1$

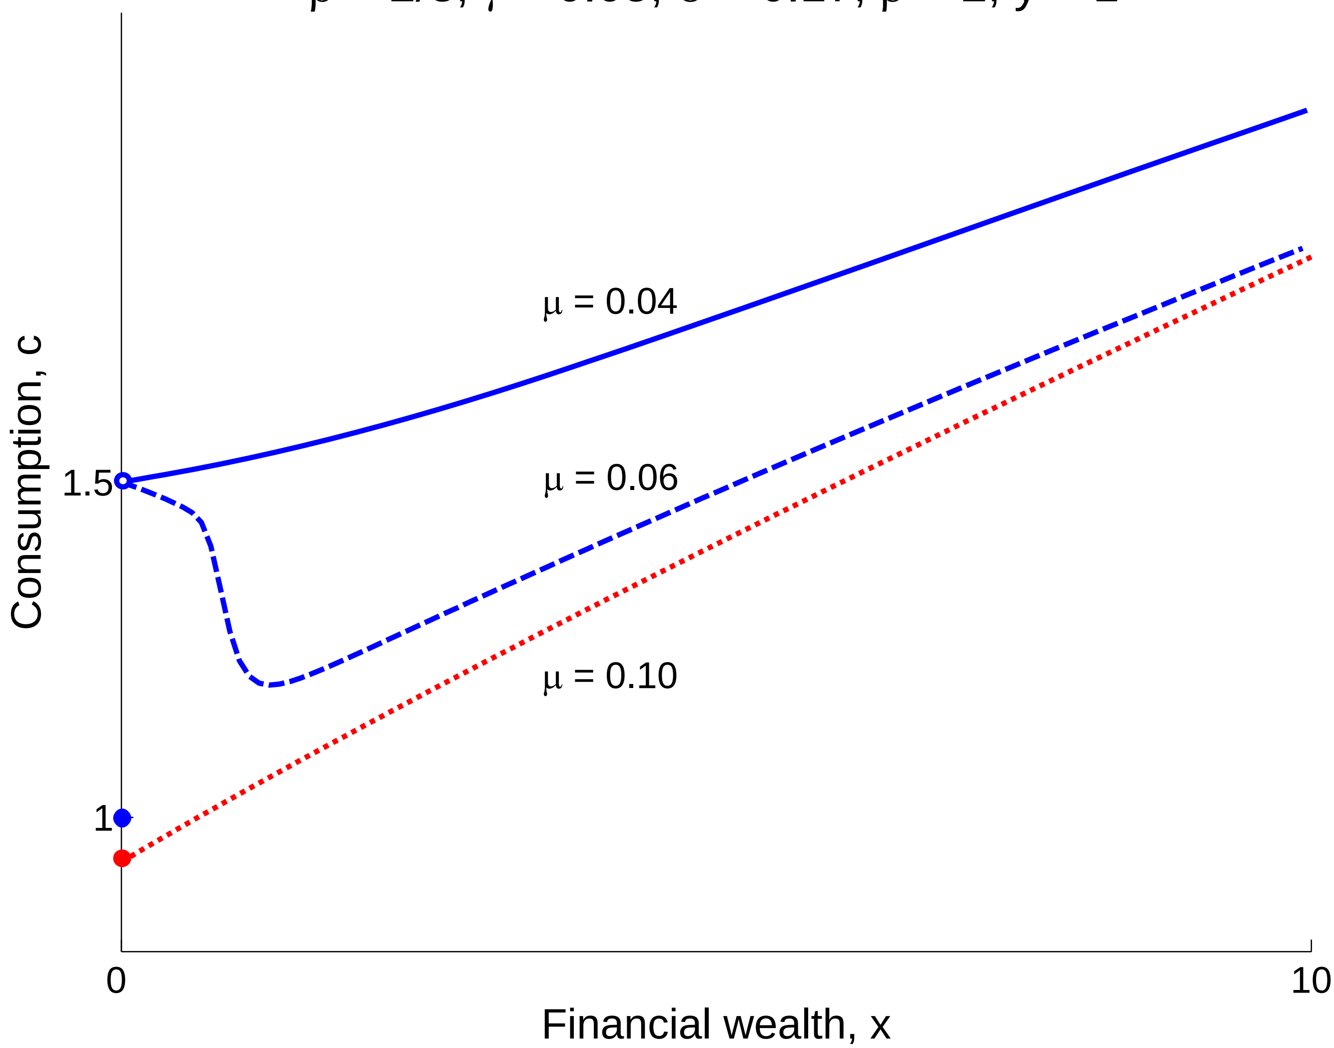

As Theorem 8 leads us to expect, for the bottom consumption function in Figure 5: the liquidity constraint is not binding, i.e. $c(0+) = c(0) < 1$; and consumption rises smoothly over the entire wealth space.

6.4. The Case in which μ is Intermediate. The remaining case of our model is that in which $\gamma < \mu < \mu_1$. Loosely speaking: when wealth is low, this case looks like the case in which μ is low; and when wealth is high, it looks like the case in which μ is high. However, the most striking feature is the behavior of the consumption function during the transition between the two regimes.

Theorem 9. *If $\gamma < \mu < \mu_1$ then: $c(0) = y$; $c(0+) = \overline{c} > y$; there exists $\overline{x} \in (0, \infty)$ such that $c' < 0$ on $(0, \overline{x})$, and $c' > 0$ on $(\overline{x}, \infty)$.*

In other words: when the IG consumer has no wealth, she consumes all of her labor income; if she acquires even a little wealth, then her consumption jumps up from y to $\overline{c}$; as her wealth increases from 0 to $\overline{x}$, her consumption *decreases*; but, once her wealth reaches $\overline{x}$, her consumption increases steadily with further increases in her wealth.

Proof. See Appendix A. ■

As Theorem 8 leads us to expect, for the middle consumption function in Figure 5: the liquidity constraint is binding, i.e. $c(0) = y = 1$; there is an upward jump in consumption at $x = 0$, from $c(0) = 1$ to $c(0+) = \overline{c} = \frac{3}{2}$; consumption then falls for a while before bottoming out and rising thereafter.

Comparing Theorem 9 with Theorems 7 and 8, a simple pattern emerges. The strategic interaction between the current self and future selves induces a form of positive feedback: the higher the marginal propensity to consume of tomorrow's self, the smaller the willingness of the current self to save, and therefore the higher her own marginal propensity to consume. By the same token: the higher the marginal propensity to save of tomorrow's self, the greater the willingness of the current self to save, and therefore the higher her own marginal propensity to save.

There are therefore two possible regimes: a high-consumption regime and a low-consumption regime. When μ is low, the consumer finds herself in the high-consumption regime irrespective of her wealth. When μ is intermediate, the consumer finds herself in the high-consumption regime when her wealth is low, and in the low-consumption regime when her wealth is high. So, naturally, her consumption needs to decrease as her wealth increases in order to effect the transition between the two regimes. Finally, when μ is high, the consumer finds herself in the low-consumption regime irrespective of her wealth.

The non-monotonic consumption function in the intermediate-μ case is not a robust feature of our model. Specifically, we can show that this non-monotonicity vanishes when we introduce a risk-free bond into the economy and allow investors to take long or short positions in the bond. Intuitively, taking a large short position in the bond enables the consumer to take large gambles, enabling her to concavify her value function. This eliminates the regions of non-monotonicity of the consumption function, since a globally concave value function has a slope that is monotonically falling in wealth, and the consumer equates her marginal utility of consumption to β times the slope of her value function.

7. Extension to the Case of General Preferences

In Section 5, we demonstrated that equilibrium exists and is unique in the IG model, and we showed that these results could be extended — by means of a refinement argument — to the deterministic IG model. In Section 6, we presented an exhaustive analysis of the consumption function of the IG model. All of this was done for the special case in which agents have constant relative risk aversion. Although the case of constant relative risk aversion is a leading case for economists, it is a knife-edge case in the space of concave preferences. There is therefore considerable interest in exploring more general preferences. Indeed, it turns out that the crucial economic assumption underpinning our results is not constant relative risk aversion at all, but rather bounds on relative prudence (Kimball, 1990). In other words: the crucial economic concept is not risk aversion but prudence; and we can replace the stringent assumption of constant relative prudence with the much weaker assumption of bounded relative prudence.

More explicitly, Theorems 3-5 and 7-9 all continue to hold as stated under the following generalization of Assumptions A1-A3:

B0 $u'(c) > 0$ and $u''(c) < 0$ for all $c > 0$.

B1 There exist $1 < \underline{\pi} \le \overline{\pi} < \infty$ such that $\underline{\pi} \le \dfrac{-c\,u'''(c)}{u''(c)} \le \overline{\pi}$ for all $c > 0$.

B2 $1 - \beta < \dfrac{\underline{\rho}}{1 + \overline{\pi} - \underline{\pi}}$, where $\underline{\rho} = \underline{\pi} - 1$.

B3 $\mu < \dfrac{1}{1 - \widehat{\underline{\rho}}}\gamma + \frac{1}{2}\,\widehat{\underline{\rho}}\,\sigma^2$ if $\widehat{\underline{\rho}} < 1$, where $\widehat{\underline{\rho}} = \dfrac{\underline{\rho} - (1-\beta)}{\overline{\rho} - (1-\beta)\,(\underline{\pi} - \overline{\rho})}\,\overline{\rho}$ and $\overline{\rho} = \overline{\pi} - 1$.

These assumptions are much more general than Assumptions A1-A3 above and, as can be seen, they are only marginally more complicated to state.

Assumption B0 simply requires that the utility function is strictly increasing and strictly concave on its domain. Assumption B1 requires that the coefficient of relative prudence

$$\pi(c) = \frac{-c\, u'''(c)}{u''(c)}$$

is bounded, and that its lower bound is strictly greater than 1. This assumption implies corresponding bounds for the coefficient of relative risk aversion

$$\rho(c) = \frac{-c\, u''(c)}{u'(c)},$$

namely that $\underline{\rho} \leq \rho(c) \leq \overline{\rho}$, where $\underline{\rho} = \underline{\pi} - 1 > 0$ and $\overline{\rho} = \overline{\pi} - 1 < \infty$ as above.[21]

Assumption B2 is a direct generalization of Assumption A2. It requires that the dynamic inconsistency of the IG consumer (as measured by $1 - \beta$) must be smaller: (i) the lower the minimum possible coefficient of relative risk aversion (as measured by $\underline{\rho}$); and (ii) the larger the fluctuations in the coefficient of relative prudence (as measured by $\overline{\pi} - \underline{\pi}$). Notice that Assumption B2 reduces to Assumption A2 when the consumer has constant relative risk aversion ρ. For in that case we have $\underline{\rho} = \rho$ and $\overline{\pi} = \underline{\pi} = \rho + 1$.

Assumption B3 is likewise a direct generalization of Assumption A3. Indeed, Assumption A3 requires that

$$\mu < \frac{1}{1-\rho}\,\gamma + \tfrac{1}{2}\,\rho\,\sigma^2 \text{ if } \rho < 1.$$

Given that the right-hand side of this inequality is increasing in ρ, the natural way to generalize Assumption A3 is to require that

$$\mu < \frac{1}{1-\underline{\widehat{\rho}}}\,\gamma + \frac{1}{2}\,\underline{\widehat{\rho}}\,\sigma^2 \text{ if } \underline{\widehat{\rho}} < 1,$$

where $\underline{\widehat{\rho}}$ is a lower bound for the coefficient of relative risk aversion $\widehat{\rho}$ of $\widehat{u}_+$. But it can be shown that

$$\widehat{\rho}(\widehat{c}) = \frac{-\widehat{c}\,\widehat{u}''_+(\widehat{c})}{\widehat{u}'_+(\widehat{c})} = \frac{\rho(c) - (1-\beta)}{\rho(c) - (1-\beta)\,(\pi(c) - \rho(c))}\,\rho(c) \geq \frac{\underline{\rho} - (1-\beta)}{\overline{\rho} - (1-\beta)\,(\underline{\pi} - \overline{\rho})}\,\overline{\rho}.$$

[21] This is not a trivial assertion: in the present context, where neither relative risk aversion nor relative prudence are constant, there is no pointwise relationship between π and ρ.

The latter expression can therefore serve as our lower bound. Notice that Assumption B3 reduces to Assumption A3 when the consumer has constant relative risk aversion ρ. For in that case we have $\overline{\rho} = \underline{\rho} = \rho$ and $\underline{\pi} = \rho + 1$. Hence $\underline{\widehat{\rho}} = \rho$.

The crucial step in extending our results is to prove the required generalization of the Value-Function Equivalence Theorem (Theorem 3).[22] In order to do this, we have to find functions $\widehat{u}_+ : (0, \infty) \to \mathbb{R}$ and $\widehat{u}_0 : (0, y] \to \mathbb{R}$ such that, if define a wealth-dependent utility function $\widehat{u}$ by the formula

$$\widehat{u}(\widehat{c}, x) = \left\{ \begin{array}{ll} \widehat{u}_+(\widehat{c}) & \text{if } \widehat{c} \in (0, \infty) \text{ and } x > 0 \\ \widehat{u}_0(\widehat{c}) & \text{if } \widehat{c} \in (0, y] \text{ and } x = 0 \end{array} \right\},$$

then the non-linearity $\widehat{h}$ in the Bellman equation of the $\widehat{u}$ consumer (namely equation (18)) is identical to the non-linearity h in the Bellman equation of the IG consumer (namely equation (19)).

The following observations can be used to construct the function $\widehat{u}_+$. Suppose that we are given a function $h_+ : (0, \infty) \to \mathbb{R}$ such that:

H0 $h'_+(\alpha) < 0$ and $h''_+(\alpha) > 0$ for all $\alpha > 0$; and

H1 there exist $0 < \underline{\theta} \leq \overline{\theta} < \infty$ such that $\underline{\theta} \leq \dfrac{-\alpha\, h''_+(\alpha)}{h'_+(\alpha)} \leq \overline{\theta}$ for all $\alpha > 0$.

Suppose further that we define a function $\widehat{u}_+ : (0, \infty) \to \mathbb{R}$ by the formula

$$\widehat{u}_+(\widehat{c}) = \min_{\alpha \in (0,\infty)} h_+(\alpha) + \widehat{c}\, \alpha.$$

Then it can be verified that

U0 $\widehat{u}'_+(\widehat{c}) > 0$ and $\widehat{u}''_+(\widehat{c}) < 0$ for all $\widehat{c} > 0$; and

U1 $\overline{\theta}^{-1} \leq \dfrac{-\widehat{c}\, \widehat{u}''_+(\widehat{c})}{\widehat{u}'_+(\widehat{c})} \leq \underline{\theta}^{-1}$ for all $\widehat{c} > 0$.

Most important of all, it can be verified that

$$h_+(\alpha) = \max_{\widehat{c} \in (0,\infty)} \widehat{u}_+(\widehat{c}) - \alpha\, \widehat{c}$$

[22] See the previous draft of this paper, Harris and Laibson (2006), for a complete proof.

for all $\alpha > 0$. All we need to do, then, is to ensure that h_+ satisfies H0 and H1. This is achieved by Assumptions B0, B1 and B2.[23]

Now consider the function $\widehat{u}_0$. We have

$$h_0(\alpha) = \left\{ \begin{array}{ll} u(y) - \alpha\, y & \text{for } \alpha \in (-\infty, \frac{1}{\beta}\, u'(y)] \\ h_+(\alpha) & \text{for } \alpha \in [\frac{1}{\beta}\, u'(y), \infty) \end{array} \right\}.$$

We can therefore define a function $\widehat{u}_0 : (0, y] \to \mathbb{R}$ by the formula

$$\widehat{u}_0(\widehat{c}) = \min_{\alpha \in (-\infty,\infty)} h_0(\alpha) + \widehat{c}\,\alpha.$$

Moreover it can be verified that

$$\widehat{u}_0(\widehat{c}) = \left\{ \begin{array}{ll} \widehat{u}_+(\widehat{c}) & \text{for } \widehat{c} \in (0, \psi(y)\, y] \\ \widehat{u}_+(\psi(y)\, y) + (\widehat{c} - \psi(y)\, y)\, \widehat{u}'_+(\psi(y)\, y) & \text{for } \widehat{c} \in [\psi(y)\, y, y] \end{array} \right\},$$

where

$$\psi(y) = \frac{\rho(y) - (1-\beta)}{\rho(y)}.$$

Most important of all, it can be verified that for all $\alpha \in (-\infty, \infty)$

$$h_0(\alpha) = \max_{\widehat{c} \in (0,y]} \widehat{u}_0(\widehat{c}) - \alpha\, \widehat{c}.$$

Once the required generalization of the Value-Function Equivalence Theorem is proved, the proofs of the Theorems 3-5 and 7-9 require little alteration. Theorem 6 no longer holds as stated, but can be replaced with the following theorem:

Theorem 10 [Overconsumption]. *Put* $\psi(c) = \frac{\rho(c) - (1-\beta)}{\rho(c)}$. *Then* $0 < \psi(c) \leq 1$. *Moreover*

$$\left\{ \begin{array}{ll} c = \frac{1}{\psi(c)}\, \widehat{c} & \textit{if either (i) } x > 0 \textit{ or (ii) } x = 0 \textit{ and } \beta\, v' > u'(y) \\ c = \widehat{c} = y & \textit{if } x = 0 \textit{ and } \beta\, v' < u'(y) \end{array} \right\}. \blacksquare$$

Notice that Assumption B3 is not needed for the proof of the Value-Function Equivalence Theorem: it is an integrability assumption that applies equally well to either the IG model or the $\widehat{u}$ model. On the contrary, one of the benefits of the Value-Function

[23] This is an instance of convex duality. The basic reference for convex duality is Rockafellar (1970).

Equivalence Theorem is that, by reducing the IG model to the $\widehat{u}$ model, it enables us to identify the integrability assumption needed for the IG model.

8. Conclusions

We have described a continuous-time model of quasi-hyperbolic discounting that extends the analysis of Barro (1999) and Luttmer and Mariotti (2003). Unlike these previous models, our model allows for a general class of preferences, includes liquidity constraints and places no restrictions on equilibrium policy functions.

Our paper studies a psychological limit case. We take the phrase 'instantaneous gratification' literally, analyzing a case in which individuals prefer gratification in the present instant discretely more than consumption in the momentarily delayed future. In this setting, equilibrium is unique, resolving multiplicity problems in quasi-hyperbolic models.

9. References

Ainslie, George. 1992 *Picoeconomics.* Cambridge: Cambridge University Press.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg (2001) "Hyperbolic discounting, wealth accumulation, and consumption." Forthcoming, *Journal of Economic Perspectives.*

Akerlof, George A. 1991. "Procrastination and Obedience." *American Economic Review. (Papers and Proceedings),* pp. 1-19.

Barro, Robert. 1999. "Laibson Meets Ramsey in the Neoclassical Growth Model." *Quarterly Journal of Economics.* 114(4), pp. 1125-52.

Carroll, Christopher D. 1992. "The Buffer Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity.* 2:1992, pp. 61-156.

———. 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics.* 112, pp. 1-57.

Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica.* 59, pp. 1221-48.

Della Vigna, Stefano and Ulrike Malmendier. 2004. "Contract Design and Self Control: Theory and Evidence," *Quarterly Journal of Economics*, May, 119, pp. 353-402.

Della Vigna, Stefano and M. Daniele Paserman. 2005. "Job Search and Impatience," *Journal of Labor Economics*, July, Vol. 23, pp. 527-588.

Harris, Christopher, and David Laibson. 2001. "Dynamic Choices of Hyperbolic Consumers." *Econometrica.* 69(4), pp. 935-957.

Harris, Christopher J. and David Laibson. 2002. "Hyperbolic discounting and consumption," in eds. Mathias Dewatripont, Lars Peter Hansen, and Stephen Turnovsky, Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress, Volume 1, pp. 258-298.

Harris, Christopher, and David Laibson. 2004. "Instantaneous Gratification," Harvard mimeo.

Kimball, Miles. 1990. "Precautionary Saving in the Small and in Large," *Econometrica*, 58(1), pp. 53-73.

Krusell, Per and Tony Smith. 2000. "Consumption-Savings Decisions with Quasi-Geometric Discounting," *Econometrica.*

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics.* 62: 2, pp. 443-478.

_ _ _ _ _ _ _ _. 1997. "Hyperbolic discount functions and time preference heterogeneity." Harvard mimeo.

_ _ _ _ _ _ _ _, Andrea Repetto, and Jeremy Tobacman. 1998. "Self-Control and Saving for Retirement," *Brookings Papers on Economic Activity* 1, pp.91-196.

Loewenstein, George, and Drazen Prelec. 1992. "Anomalies in Intertemporal Choice: Evidence and an Interpretation." *Quarterly Journal of Economics.* 57, pp. 573-98.

Luttmer, Erzo and Thomas Mariotti. 2003. "Subjective Discount Factors," *Journal of Political Economy.*

Morris, Stephen. 2002. "Continuous Consumption Rules with Nonexponential Discounting," mimeo Yale University.

Morris, Stephen and Andrew Postlewaite. 1997. "Observational Implications of Nonexponential Discounting," mimeo.

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

_ _ _ _ _ _ _ _ and _ _ _ _ _ _ _ _ 1999b. "Incentives for Procrastinators," *Quarterly Journal of Economics*, 114(3), pp. 769-816.

Phelps, E. S., and R. A. Pollak. 1968. "On Second-best National Saving and Game-equilibrium Growth." *Review of Economic Studies.* 35, pp. 185-199.

Rockafellar, R. Tyrrell. 1970. *Convex Analysis.* Princeton University Press: Princeton.

Strotz, Robert H. 1956. "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies.* 23, pp. 165-180.

A. Proof of Theorems 7, 8 and 9

In this appendix, we outline the proof of Theorems 7, 8 and 9.

A.1. Some background information. The first point to note is that the value function is bounded:

Proposition 11. *There exists* $K > 0$ *such that*

$$\tfrac{1}{\gamma}\, u(y) \le v(x) \le K \,\max\{1, u(x)\}$$

for all $x \ge 0$. ■

The lower bound is easy to derive: the $\widehat{u}$ consumer can consume at a high rate until her wealth falls to 0, and then consume at rate y. In this way she obtains a flow utility of at least $\widehat{u}_0(y) = u(y)$. The upper bound is more subtle, and involves two cases. If $\rho > 1$ then $\widehat{u}$ is bounded above. On the other hand, if $\rho \le 1$ then the average propensity to consume of the $\widehat{u}$ consumer, namely $\frac{\widehat{c}}{x}$, converges to the steady-state value

$$\widehat{\eta} = \frac{\gamma - (1-\rho)\,(\mu - \frac{1}{2}\,\rho\,\sigma^2)}{\rho}$$

as x tends to infinity. Moreover

$$v' = \widehat{u}'_+(\widehat{c}) = \tfrac{1}{\beta}\, u'(\tfrac{\widehat{c}}{\psi}) = \tfrac{1}{\beta}\, u'(\tfrac{\widehat{c}}{\psi\, x}\, x) = \tfrac{1}{\beta}\, u'(\tfrac{\widehat{c}}{\psi\, x})\, u'(x),$$

and therefore $\frac{v'}{u'(x)} \to \frac{1}{\beta}\, u'(\frac{\widehat{\eta}}{\psi})$ as $x \uparrow \infty$. The value function v cannot therefore grow any faster than u.

The second point to note is that the value function is smooth:

Proposition 12. *Suppose that* $\beta < 1$. *Then* v *is infinitely differentiable on* $[0, \infty)$. ■

Notice that the discontinuity in $\widehat{u}$ at $x = 0$ (i.e. the fact that $\widehat{u}_+ \neq \widehat{u}_0$) does not translate into a discontinuity in v or any of its derivatives at $x = 0$. On the contrary, Proposition 12 actually depends on this discontinuity: when $\beta = 1$ (and therefore $\widehat{u}_+ = \widehat{u}_0$), v is not smooth at $x = 0$ when $\mu < \gamma$. The discontinuity in $\widehat{u}$ at $x = 0$ does, however, cause a different kind of discontinuity: as we shall see below, there exists $\mu_1 \in (\gamma, \overline{\mu})$ such that $v'(0)$ jumps up from $v'_L = \widehat{u}'_+(\psi\, \overline{c}) < \widehat{u}'_+(y)$ to $v'_R = \widehat{u}'_+(\psi\, y) > \widehat{u}'_+(y)$ at μ_1.

The third point to note is that the shadow value of wealth is always strictly positive:

Proposition 13. *$v' > 0$ on $[0, \infty)$.* ■

This is economically obvious: the $\widehat{u}$ consumer can always consume more in its current span of control.

A.2. An intuitive perspective. Recall that the utility function of the $\widehat{u}$ consumer has two parts: $\widehat{u}(\widehat{c}, x) = \widehat{u}_0(\widehat{c})$ when $x = 0$; and $\widehat{u}(\widehat{c}, x) = \widehat{u}_+(\widehat{c})$ when $x > 0$. Moreover $\widehat{u}_0(\widehat{c}) \geq \widehat{u}_+(\widehat{c})$, with strict inequality when $\widehat{c} \in (\psi\, y, y]$. (See Figure 4.) In other words, the $\widehat{u}$ consumer obtains a utility premium when $x = 0$.

This suggests that, at any given wealth level, the $\widehat{u}$ consumer must choose between two strategies. The first, high-consumption, strategy is to dissave until her wealth runs out, and then enjoy the utility premium that she obtains at $x = 0$. The second, low-consumption, strategy is to save forever in order to take advantage of the higher asset income associated with higher of financial wealth. Which of these two strategies is better will depend on μ. If μ is low, then the high-consumption strategy will be better no matter how large the wealth of the $\widehat{u}$ consumer. Similarly, if μ is high, then the low-consumption strategy will be better no matter how small the wealth of the $\widehat{u}$ consumer. However, if μ is intermediate then the high-consumption strategy will be better when wealth is low (and therefore the utility premium will be enjoyed after a relatively short wait) and the low-consumption strategy will be better when wealth is high (and therefore the prospect of the utility premium is too distant). Moreover consumption can be expected to decrease with wealth over an intermediate range of wealth levels, as the $\widehat{u}$ consumer adjusts from the high-consumption strategy associated with low wealth to the low-consumption strategy associated with high wealth.

A.3. The boundary condition at $x = 0$. The value function v must satisfy two related conditions at $x = 0$. To derive the first of these conditions, note that the Bellman equation of the $\widehat{u}$ consumer takes the form

$$0 = \tfrac{1}{2}\,\sigma^2\, x^2\, v'' + (\mu\, x + y)\, v' - \gamma\, v + \widehat{h}_+(v') \tag{23}$$

for $x > 0$. Letting $x \downarrow 0$ in this equation, taking advantage of Proposition 12 and rearranging yields

$$v(0) = \tfrac{1}{\gamma}\left(y\, v'(0) + \widehat{h}_+(v'(0))\right). \tag{24}$$

The second of these conditions is simply the Bellman equation of the $\widehat{u}$ consumer at $x = 0$ which, on rearrangement, becomes

$$v(0) = \tfrac{1}{\gamma}\left(y\,v'(0) + \widehat{h}_0(v'(0))\right). \tag{25}$$

Figures A1, A2 and A3 illustrate the locus of points $(v'(0), v(0))$ satisfying equation (24), the locus of points $(v'(0), v(0))$ satisfying equation (25) and the locus of points $(v'(0), v(0))$ satisfying both equations.

As Figure A3 shows, there are two possible boundary configurations. First, the $\widehat{u}$ consumer may opt for the utility premium and set $\widehat{c}(0) = y$. In this case $v(0) = \frac{1}{\gamma}\,\widehat{u}_0(y)$, and $v'(0)$ must take on the low value $v'_L = \widehat{u}'_+(\psi\,\overline{c}) < \widehat{u}'_+(y)$ in order to justify the $\widehat{u}$ consumer's high consumption level for small $x > 0$. Second, the $\widehat{u}$ consumer may forgo the utility premium and set $\widehat{c}(0) \leq \psi\,y$. In this case $v(0) \geq \frac{1}{\gamma}\,\widehat{u}_0(y)$, and we must have $v'(0) \geq v'_R = \widehat{u}'_+(\psi\,y) > \widehat{u}'_+(y)$ in order to justify the $\widehat{u}$ consumer's low consumption level at $x = 0$. We refer to these two configurations as the left- and right-hand boundary configurations.

Our next major objective is to show that there exists $\mu_1 \in (\gamma, \overline{\mu})$ such that the left-hand boundary configuration occurs when $\mu < \mu_1$ and the right-hand boundary configuration occurs when $\mu > \mu_1$. To this end, we shall need several supporting results.

A.4. Once convex, always strictly convex. The following result only uses the fact that v satisfies the Bellman equation of the $\widehat{u}$ consumer in the interior of the wealth space.

Proposition 14. *Suppose that $\mu < \gamma$ and that either*

1. *there exists $x_0 \geq 0$ such that $v''(x_0) > 0$, or*
2. *there exists $x_0 > 0$ such that $v''(x_0) \geq 0$.*

Then $v''(x) > 0$ for all $x > x_0$.

Proof. Differentiating the Bellman equation of the $\widehat{u}$ consumer (i.e. equation (23)) with respect to x, we obtain

$$0 = \tfrac{1}{2}\,\sigma^2\,x^2\,v''' + ((\sigma^2 + \mu)\,x + y - \widehat{c})\,v'' - (\gamma - \mu)\,v'. \tag{26}$$

Figure A1: graph of $v(0) = \frac{1}{\gamma}\,(\,y\,v'(0) + \hat{h}_{+}(v'(0))\,)$

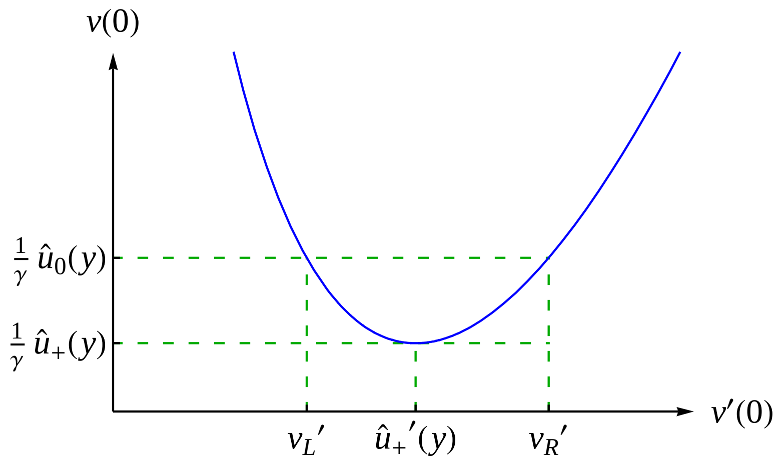

Figure A2: graph of $v(0) = \frac{1}{\gamma}\,(\,y\,v'(0) + \hat{h}_{0}(v'(0))\,)$

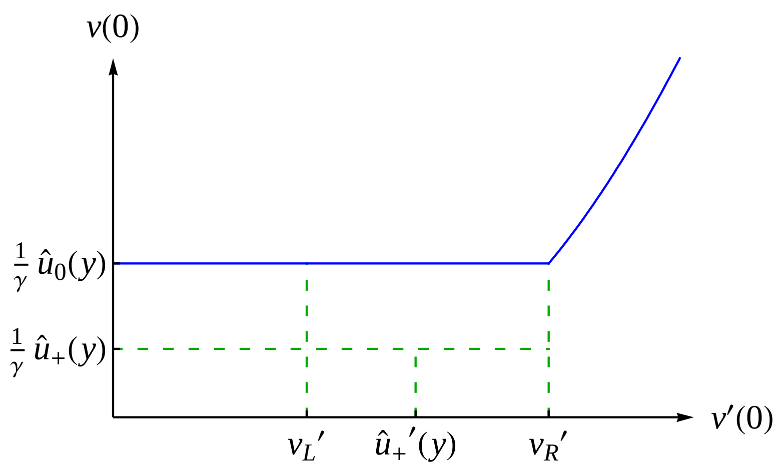

Figure A3: feasible choices of $(v'(0)\,, v(0)\,)$ at $x = 0$

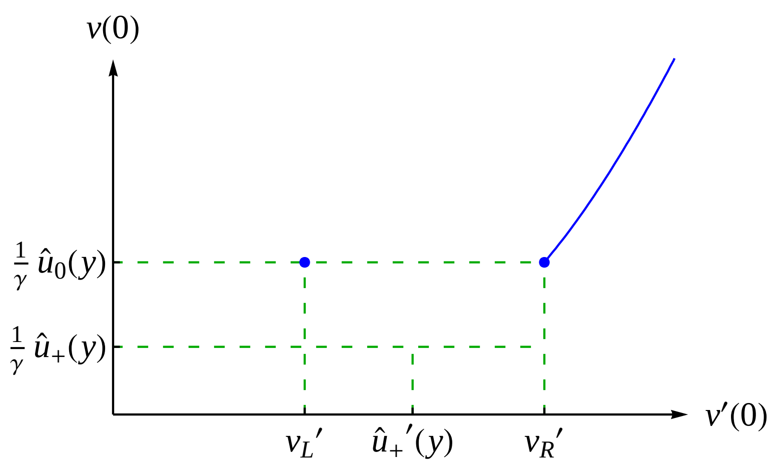

Now suppose for a contradiction that there exist $x_0 \geq 0$ and $x_1 > x_0$ such that $v''(x_0) > 0$ and $v''(x_1) \leq 0$. Let x_2 be the leftmost point in $(x_0, x_1]$ such that $v''(x_2) \leq 0$. Since $v'' > 0$ on $[x_0, x_2)$, we must actually have $v''(x_2) = 0$. Equation (26) then yields

$$v'''(x_2) = \frac{(\gamma - \mu)\, v'(x_2)}{\frac{1}{2}\, \sigma^2\, x_2^2},$$

and the latter expression is strictly positive because $\gamma > \mu$ (by assumption), $v'(x_2) > 0$ (by Proposition 13) and $x_2 > 0$ (by construction). We therefore have $v'' < 0$ to the left of x_2, which is the required contradiction. It remains to consider the case in which there exist $x_0 > 0$ and $x_1 > x_0$ such that $v''(x_0) = 0$ and $v''(x_1) \leq 0$. In that case $v'''(x_0) > 0$, and hence there exists $\widetilde{x}_0 \in (x_0, x_1)$ such that $v''(\widetilde{x}_0) > 0$. We are then back in the previous case. ■

Combining this result with the fact that v satisfies the Bellman equation of the $\widehat{u}$ consumer on the boundary of the wealth space, we obtain the following corollary.

Corollary 15. *Suppose that $\mu < \gamma$. Then $v'' < 0$ for all $x \geq 0$.*

In particular, $\widehat{c}' > 0$ on $(0, \infty)$. This is the case in which the $\widehat{u}$ consumer chooses the high-consumption strategy at all wealth levels.

Proof. Suppose for a contradiction that there exists $x_0 > 0$ such that $v''(x_0) \geq 0$. Then Proposition 14 implies that v is convex on $[x_0, \infty)$. This contradicts Proposition 11, which tells us that v is bounded above by a function that grows at the same rate as u. We therefore have $v'' < 0$ for all $x > 0$. This in turn implies that $v''(0) \leq 0$. Now, letting $x \downarrow 0$ in equation (26) and rearranging, we obtain

$$v''(0) = \frac{(\gamma - \mu)\, v'(0)}{y - \widehat{c}(0+)}.$$

So: either $v'(0) = v'_L$, in which case $y - \widehat{c}(0+) < 0$ and therefore $v''(0) < 0$; or $v'(0) \geq v'_R$, in which case $y - \widehat{c}(0+) > 0$ and therefore $v''(0) > 0$. (Recall that $\gamma - \mu > 0$ by assumption.) In other words, we cannot have $v''(0) = 0$. We conclude that $v''(0) < 0$. This completes the proof. ■

Actually, the proof of Corollary 15 shows more:

Corollary 16. *Suppose that $\mu < \gamma$. Then $v'(0) = v'_L$.* ■

In other words, if $\mu < \gamma$ then the left-hand boundary configuration obtains. In particular, $\widehat{c}(0+) = \psi\,\overline{c} > \widehat{c}(0) = y$.

A.5. Once concave, always strictly concave. The following result likewise only uses the fact that v satisfies the Bellman equation of the $\widehat{u}$ consumer in the interior of the wealth space.

Proposition 17. *Suppose that $\mu > \gamma$ and that either*

1. *there exists $x_0 \geq 0$ such that $v''(x_0) < 0$, or*
2. *there exists $x_0 > 0$ such that $v''(x_0) \leq 0$.*

Then $v''(x) < 0$ for all $x > x_0$.

Proof. The proof is completely analogous to that of Proposition 14. ■

Combining this result with the fact that v satisfies the Bellman equation of the $\widehat{u}$ consumer on the boundary of the wealth space, we obtain the following corollary.

Corollary 18. *Suppose that $\mu > \gamma$. Then: either*

1. *there exists $\overline{x} \in (0, \infty)$ such that $v'' > 0$ on $(0, \overline{x})$, and $v'' < 0$ on $(\overline{x}, \infty)$; or*
2. *$v'' < 0$ for all $x \geq 0$.*

In particular: either there exists $\overline{x} \in (0, \infty)$ such that $\widehat{c}' < 0$ on $(0, \overline{x})$, and $\widehat{c}' > 0$ on $(\overline{x}, \infty)$; or $\widehat{c}' > 0$ on $[0, \infty)$. The first case is the case in which the $\widehat{u}$ consumer chooses the high-consumption strategy at low wealth levels and the low-consumption strategy at high wealth levels. The second case is the case in which the $\widehat{u}$ consumer chooses the low-consumption strategy at all wealth levels.

Proof. Let X_0 be the set of all $x_0 \in [0, \infty)$ such that $v''(x_0) \leq 0$. Proposition 11 implies that we cannot have $v'' > 0$ for all $x \geq 0$, so X_0 is non-empty. Let $\overline{x}$ be the smallest element of X_0. There are then two possibilities: either $\overline{x} > 0$ or $\overline{x} = 0$. If $\overline{x} > 0$, then $v'' > 0$ for all $x \in [0, \overline{x})$. Hence $v''(\overline{x}) = 0$, and Proposition 17 implies that $v'' < 0$ for all $x \in (\overline{x}, \infty)$. On the other hand, if $\overline{x} = 0$ then our construction of $\overline{x}$ yields only that $v''(\overline{x}) \leq 0$. However, as in the proof of Corollary 15, we have

$$v''(0) = \frac{(\gamma - \mu)\, v'(0)}{y - \widehat{c}(0+)}.$$

So: either $v'(0) = v'_L$, in which case $y - \widehat{c}(0+) < 0$ and therefore $v''(0) > 0$; or $v'(0) \geq v'_R$, in which case $y - \widehat{c}(0+) > 0$ and therefore $v''(0) < 0$. (Recall that $\gamma - \mu < 0$ by assumption.) In other words, we cannot have $v''(\overline{x}) = 0$. We conclude that $v'' < 0$ for all $x \geq 0$ when $\overline{x} = 0$. ■

Actually, the proof of Corollary 18 shows slightly more:

Corollary 19. *Suppose that $\mu > \gamma$. Then either*

1. *$v'(0) = v'_L$, in which case $v''(0) > 0$; or*
2. *$v'(0) \geq v'_R$, in which case $v''(0) < 0$.* ■

In other words, if the left-hand boundary configuration obtains then $\widehat{c}(0+) = \psi\,\overline{c} > \widehat{c}(0) = y$. The main surprise in this case is the way in which: (i) the expected initial increase in consumption is confined to a single upward jump in $\widehat{c}$ at $x = 0$; and (ii) the expected decrease in consumption (as the consumer adjusts to the low-consumption strategy) begins immediately to the right of $x = 0$. On the other hand, if the right-hand boundary configuration obtains then $\widehat{c}(0+) = \widehat{c}(0) \leq \psi\, y$.

A.6. The $\widehat{u}_+$ consumer. At this point, we have established that the left-hand boundary configuration obtains when $\mu < \gamma$. However, we have no information as to which of the two possible boundary configurations obtains when $\mu > \gamma$. In order to settle this question, it will be helpful to consider a consumer who

1. discounts the future exponentially at rate γ,
2. faces the same wealth dynamics as the IG consumer and the $\widehat{u}$ consumer and
3. has the wealth-independent utility function $\widehat{u}_+$.

We call this consumer the $\widehat{u}_+$ consumer.

Let the value function of the $\widehat{u}_+$ consumer be $\widehat{v}_+$. Then the relevance of the problem of the $\widehat{u}_+$ consumer can be seen from two closely related observations. First, we have

$$v(0;\mu) = \max\{\tfrac{1}{\gamma}\,\widehat{u}_0(y), \widehat{v}_+(0;\mu)\},$$

where we have made the dependence of v and $\widehat{v}_+$ on the parameter μ explicit. Second, the left-hand boundary configuration obtains if and only if $\widehat{v}_+(0;\mu) < \frac{1}{\gamma}\,\widehat{u}_0(y)$.

We shall need four propositions about $\widehat{v}_+$.

Proposition 20. $\widehat{v}_+(0;\mu) = \frac{1}{\gamma}\,\widehat{u}_+(y)$ *for all* $\mu < \gamma$.

Proof. The proof is based on ideas similar to those used in Section A.4. If $\widehat{v}_+(0;\overline{\mu}) > \frac{1}{\gamma}\,\widehat{u}_+(y)$, then we may set $\eta = \gamma\,\widehat{v}_+(0;\overline{\mu})$ and consider a new problem in which the $\widehat{u}_+$ consumer has the additional option of enjoying the utility level η when $x = 0$ (in which case x will remain unchanged at 0). If $\mu < \gamma$, then the $\widehat{u}_+$ consumer will strictly prefer this option. So $\widehat{v}_+(0;\mu) < \frac{1}{\gamma}\,\eta$ for all $\mu < \gamma$. But $\widehat{v}_+(0;\overline{\mu}) = \frac{1}{\gamma}\,\eta$ by construction. So $\overline{\mu} \geq \gamma$. ■

Proposition 21. $\widehat{v}_+(0;\mu) \to \frac{1}{\gamma}\,\widehat{u}_+(\infty)$ *as* $\mu \uparrow \overline{\mu}$.

Proof. This is obvious. ■

Proposition 22. $\widehat{v}_+(0;\mu)$ *is continuous in* μ.

Proof. This is an instance of the so-called Theorem of the Maximum. We omit the proof. ■

Proposition 23. $\frac{\partial\widehat{v}_+}{\partial\mu}(0;\mu) > 0$ *whenever* $\widehat{v}_+(0;\mu) > \frac{1}{\gamma}\,\widehat{u}_+(y)$.

Proof. The Bellman equation of the $\widehat{u}_+$ consumer takes the form

$$0 = \max_{\widetilde{c}\in(0,\infty)} \left\{ \tfrac{1}{2}\,\sigma^2\,x^2\,\widehat{v}_+'' + (\mu\,x + y - \widetilde{c})\,\widehat{v}_+' - \gamma\,\widehat{v}_+ + \widehat{u}_+(\widetilde{c}) \right\}$$

if $x > 0$ and

$$0 = \max_{\widetilde{c}\in(0,y]} \left\{ (y - \widetilde{c})\,\widehat{v}_+' - \gamma\,\widehat{v}_+ + \widehat{u}_+(\widetilde{c}) \right\}$$

if $x = 0$. Differentiating these equations with respect to μ, putting $a = \frac{\partial\widehat{v}_+}{\partial\mu}$ and denoting the optimal choice of consumption by $\widehat{c}_+$ yields

$$0 = \tfrac{1}{2}\,\sigma^2\,x^2\,a'' + (\mu\,x + y - \widehat{c}_+)\,a' - \gamma\,a + x\,\widehat{v}_+'$$

for all $x \geq 0$. In other words, a is the expected present discounted value (along the equilibrium path) of the flow payoff $x\,\widehat{v}_+'$. (This reflects the fact that a marginal increase in μ when wealth is x generates a marginal increase in income of x, and the fact that this increase in income is valued at the marginal value of wealth, namely $\widehat{v}_+'$.) Now $x\,\widehat{v}_+'(x;\mu) \geq 0$, with strict inequality if $x > 0$. Hence $a(x;\mu) \geq 0$, with strict inequality if either $x > 0$, or $x = 0$ and $\widehat{v}_+(0;\mu) > \frac{1}{\gamma}\,\widehat{u}_+(y)$. Indeed, in the latter case $\widehat{c}_+(0;\mu) < y$, and $\widehat{c}_+(x;\mu)$ is continuous in x at $x = 0$. The dynamics therefore move directly into the interior. ■

Combining Propositions 20-23, and noting that $\widehat{u}_+(y) < \widehat{u}_0(y) < \widehat{u}_+(\infty)$, we see that there is a unique $\mu_1 \in (\gamma, \overline{\mu})$ such that: (i) $\widehat{v}_+(0;\mu) < \frac{1}{\gamma}\,\widehat{u}_0(y)$ for $\mu < \mu_1$; (ii) $\widehat{v}_+(0;\mu_1) = \frac{1}{\gamma}\,\widehat{u}_0(y)$; and (iii) $\widehat{v}_+(0;\mu) > \frac{1}{\gamma}\,\widehat{u}_0(y)$ for $\mu > \mu_1$. Finally, the picture is completed by noting that $\widehat{v}_+(0;\mu) > \frac{1}{\gamma}\,\widehat{u}_+(y)$ for all $\mu > \gamma$, but we shall not prove this here.

A.7. From $\widehat{c}$ to c. At this point we have shown that there exists $\mu_1 \in (\gamma, \overline{\mu})$ such that:

1. If $\mu < \gamma$ then the left-hand boundary configuration holds, i.e. $v'(0) = v'_L = \widehat{u}'_+(\psi\,\overline{c}) < \widehat{u}'_+(y)$. Moreover $v''(0) < 0$, $\widehat{c}(0+) = \psi\,\overline{c} > \widehat{c}(0) = y$ and $\widehat{c}' > 0$ on $(0,\infty)$.

2. If $\gamma < \mu < \mu_1$ then the left-hand boundary configuration still holds, i.e. we still have $v'(0) = v'_L = \widehat{u}'_+(\psi\,\overline{c}) < \widehat{u}'_+(y)$. We also still have $\widehat{c}(0+) = \psi\,\overline{c} > \widehat{c}(0) = y$. However, we now have $v''(0) > 0$. Moreover, there exists $\overline{x} \in (0,\infty)$ such that $\widehat{c}' < 0$ on $(0,\overline{x})$ and $\widehat{c}' > 0$ on $(\overline{x},\infty)$.

3. If $\mu > \mu_1$ then the right-hand boundary configuration holds, i.e. $v'(0) = v'_R = \widehat{u}'_+(\psi\, y) > \widehat{u}'_+(y)$. Moreover $v''(0) < 0$, $\widehat{c}(0+) = \widehat{c}(0) \le \psi\, y$ and $\widehat{c}' > 0$ on $[0,\infty)$.

Now, when the left-hand boundary configuration holds, it is easy to check that $c(0) = y$. On the other hand, when either (i) $x > 0$ or (ii) $x = 0$ and the right-hand boundary configuration holds, we have $u'(c) = \beta\, v' = \beta\,\widehat{u}'_+(\widehat{c})$. So c is increasing (decreasing) iff $\widehat{c}$ is increasing (decreasing). These results therefore translate directly into the results of Theorems 7, 8 and 9.

A.8. The borderline cases. We have said relatively little about the cases $\mu = \gamma$ and $\mu = \mu_1$. The case $\mu = \gamma$ has several interesting features. First, letting $\mu \uparrow \gamma$, we see that $v'' \le 0$ and $\widehat{c}' \ge 0$ on $[0,\infty)$. Second, again letting $\mu \uparrow \gamma$, we obtain $\widehat{v}_+(0) = \frac{1}{\gamma}\,\widehat{u}_+(y) < \frac{1}{\gamma}\,\widehat{u}_0(y)$. It follows that the left-hand boundary configuration obtains. This in turn implies that $\widehat{c}(0+) = \psi\,\overline{c} > y$. Letting $x \downarrow 0$ in equation (26) then yields $0 = (y - \widehat{c}(0+))\, v''(0)$. It follows that $v''(0) = 0$. Third, by considering higher-order analogues of equation (26), one can go on to show that $v^{(n)}(0) = 0$ for all $n \geq 3$ as well. In other words, the only non-zero coefficients in the Taylor expansion for v at $x = 0$ are $v(0)$ and $v'(0)$. At first sight this would seem to suggest that v is linear. However, this would contradict Proposition 11. The resolution lies in the fact that v is not analytic at 0. Rather, v' (and therefore $\widehat{c}$) are so called 'flat functions'. This terminology turns out to be apt: simulations show that v' and $\widehat{c}$ are nearly constant for a significant interval of wealth starting at $x = 0$.

The case $\mu = \mu_1$ involves a number of subtleties. First, even though μ_1 is the point at which we switch from the left- to the right-hand boundary configuration, only the right-hand boundary configuration can occur when $\mu = \mu_1$. This is because $v'(0)$ is essentially the limit $v'(0+)$, and as such is determined by behavior in the interior of the wealth space. Moreover, in the interior of the wealth space, the low-consumption strategy is the preferred strategy of the $\widehat{u}$ consumer. The $\widehat{u}$ consumer does, however, have two equally good options at $x = 0$: since $v'(0) = \widehat{u}'_+(\psi\, y)$ and $\widehat{u}_0$ has slope $\widehat{u}'_+(\psi\, y)$ on $[\psi\, y, y]$, she is indifferent between $\widehat{c}(0) = \psi\, y$ and $\widehat{c}(0) = y$. (She is in fact indifferent among all $\widehat{c}(0) \in [\psi\, y, y]$, but the intermediate options should be seen as the result of strictly randomizing between $\psi\, y$ and y. Moreover they all lead to the same outcome as $\psi\, y$: the dynamics move immediately into the interior of the wealth space.) If she chooses $\widehat{c}(0) = \psi\, y$, then she embarks immediately on the low-consumption strategy. If she chooses $\widehat{c}(0) = y$, then she remains forever with wealth 0. Either way, she ends up with the payoff $v(0) = \frac{1}{\gamma}\, \widehat{u}_0(y)$. Second, as $\mu \uparrow \mu_1$, the length of the interval over which $\widehat{c}$ decreases — which is always an open interval with left-hand endpoint 0 — converges to 0. (So, in effect, $\widehat{c}$ jumps up from y to $\psi\, \overline{c}$ at 0 and then decreases very rapidly back down to something close to $\psi\, y$.) In other words, a boundary layer develops near $x = 0$.