No. 2503

SUBJECTIVE DISCOUNT FACTORS

Erzo G J Luttmer and Thomas Mariotti

FINANCIAL ECONOMICS

Centre for Economic Policy Research

ISSN 0265-8003

SUBJECTIVE DISCOUNT FACTORS

Erzo G J Luttmer, London School of Economics and CEPR
Thomas Mariotti, London School of Economics

Discussion Paper No. 2503
July 2000

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: http://www.cepr.org

CEPR Discussion Paper No. 2503

July 2000

ABSTRACT

Subjective Discount Factors*

This Paper describes the equilibrium of a discrete-time exchange economy in which consumers with arbitrary subjective discount factors and quasi-homothetic period utility functions follow linear Markov consumption and portfolio strategies. Explicit expressions are given for state prices and consumption–wealth ratios. If utility is logarithmic or endowment growth is i.i.d., then this economy is observationally equivalent to one in which consumers discount geometrically. We provide analytically convenient continuous-time approximations and examine the effects of non-geometric subjective discount factors in an economy in which log endowments are subject to temporary and permanent shocks that are governed by a Feller (1951) square-root process. Hyperbolic and quasi-hyperbolic discount factors can significantly increase the volatility of aggregate wealth and raise the expected excess return on aggregate wealth.

JEL Classification: D50, D91 and G12
Keywords: asset pricing, consumption–wealth ratios, general equilibrium, hyperbolic discounting and volatility

Erzo G J Luttmer
Thomas Mariotti
Department of Economics
London School of Economics
Houghton Street
London WC2A 2AE
UK
Tel: (44 20) 7955 7897
(44 20) 7955 7517
Fax: (44 20) 7831 1840 (both)
Email: e.luttmer@lse.ac.uk
t.mariotti@lse.ac.uk

* This is a work in progress. We thank Per Krusell, Rohit Rahi, and seminar participants at the Bank of England, CEPR, the Universities of Minnesota and Siena, and the Stockholm School of Economics for comments. The FMG and STICERD at the London School of Economics provided financial support. This Paper is produced as part of a CEPR research network on The Industrial Organisation of Banking and Financial Markets in Europe, funded by the European Commission under the Training and Mobility of Researchers Programme (Contract No ERBFMRX-CT98-0222).

Submitted 12 June 2000

NON-TECHNICAL SUMMARY

This Paper examines an economy populated by consumers whose preferences are time-inconsistent. These are consumers who may, for instance, prefer one pound today to two pounds tomorrow, while at the same time preferring two pounds a year and a day from today to one pound a year from today. More generally, these consumers discount future utilities using subjective rates of time preference that may not be constant. An important implication of these preferences is that the course of action preferred by an agent today need not coincide with the one he or she would like to implement tomorrow. As a result, self-control and the degree to which agents can commit to future choices become central issues for decision-making. In the absence of a perfect commitment technology, individual choices of consumers with these types of preferences can be viewed as the outcome of an ‘intra-personal game’ in which at any date the current incarnation of an individual consumer plays a game against future incarnations of the same consumer.

If consumers can perfectly commit to a sequence of consumption choices, then the term structure of interest rates will reflect the rates at which consumers discount utility at different horizons. Given that the term structure of interest rates is typically upward sloping, this would imply that consumers discount utility at lower rates over the short term than they do over the long term. This is in contrast to evidence collected by psychologists that suggests that subjective rates of time preference decrease as a function of the horizon over which utility is evaluated.

In this Paper, we consider an exchange economy in which markets are frictionless and consumers cannot commit themselves in advance to future consumption and portfolio choices.

A first contribution of this Paper is to show how to construct a competitive equilibrium in which consumer choices are the outcome of such a game. Our construction relies on a class of preferences that generates linear decision rules. We obtain explicit formulas for the prices of state-contingent claims and for equilibrium consumption–wealth ratios. Because consumer choices result from the strategic interaction between current and future incarnations of the same consumer, standard arguments that rule out bubbles on long-lived assets do not apply. We provide an example in which the presence of a long-lived asset that is in positive net supply is required to ensure existence of a competitive equilibrium. We give conditions on endowment growth and subjective discount rates that rule out bubbles. These conditions cover most interesting applications.

In addition, we derive analytically convenient expressions for interest rates and risk premia by letting the length of a period go to zero. In the resulting limit economy, we show that the ratio between the instantaneous expected return and volatility is not affected by rates at which consumers discount utility.

Short-term interest rates in this economy depend on a utility-weighted average of subjective rates of time-preference. If utility is expected to grow at a constant rate, then the relevant utility weights are constant and the economy is observationally equivalent to one in which consumers are time-consistent. On the other hand, fluctuations in the growth rate of consumption affect these utility weights. This causes interest rates to move in ways that differ from the standard model. In particular, changes in expected future growth rates of consumption affect the current short-term risk-free rate. In standard versions of economies with time-consistent preferences, short-term interest rates would only depend on the one-period ahead growth rate of consumption.

Psychologists have argued that, as in the example mentioned above, subjective rates of time preference are higher at nearby horizons than at distant horizons. To reconcile the low level of interest rates observed in the data with the high short-run subjective discount rates emphasized by psychologists we need to assume that long-run subjective discount rates are quite low. For the preferences we consider, wealth–consumption ratios and expected discounted utility growth are proportional. Low discount rates for distant utilities tend to make expected discounted utility (and therefore wealth) more sensitive to shocks that predict changes in future endowment growth rates. This raises the volatility of aggregate wealth. In turn, the increased volatility of wealth raises the risk premium on aggregate wealth.

Using numerical simulations we show that if consumers discount utility in this way, aggregate wealth can be significantly more volatile relative to consumption than is the case when consumers discount utility at a constant rate. This may be part of an explanation for the empirical evidence that suggests that important components of aggregate wealth, such as the stock market, are too volatile.

1. Introduction

Most explicit dynamic equilibrium models in macroeconomics and finance are based on the assumption that consumers have time and state separable preferences and discount future utilities exponentially. This implies that subjective discount rates are constant, and thus that consumer choices are dynamically consistent.

Psychologists have questioned the validity of the assumption of exponential discounting on the basis of experimental evidence (Herrnstein (1961), Ainslie (1975, 1992)). These studies suggest instead that subjective discount functions are approximately hyperbolic. According to this literature, events in the near future tend to be discounted at a higher rate than events that occur in the long-run. This creates a conflict between an individual agent's preferences at different points in time. The course of action preferred today by a hyperbolic agent does not coincide with the one he knows he would like to implement tomorrow. As a result, self-control and the degree to which agents are able to commit to future choices become central issues for decision making.

If consumers can perfectly commit to a sequence of consumption choices, then standard consumer theory applies, whether subjective discount factors are geometric or not. Consider for example an exchange economy with time-invariant period utility functions, constant aggregate endowments, and date-zero markets for consumption at all future dates and in all future states. In this economy, the term structure of interest rates coincides with the term structure of the representative consumer's rates of time preference. As Loewenstein and Prelec (1992) have suggested, the experimental evidence would then lead one to expect higher yields on short-maturity bonds than on long-maturity bonds. Empirical studies on the term structure of interest rates indicate that on average the opposite is true.

In this paper, we assume instead that consumers have no means through which they can commit to future consumption choices. Under this assumption, we examine an exchange economy with a sequence of markets. In every period, consumers can trade in a complete set of one-period state contingent claims (and possibly in some long-lived securities as well.) These contracts can be perfectly enforced, so that consumers can be allowed to borrow up to the present value of their endowments. In contrast, there are no enforceable contracts that allow consumers to commit to a particular sequence of consumption choices.

The only restriction we impose on subjective discount factors is that the utility of aggregate endowments is finite. Following Pollak (1968), Phelps and Pollak (1968) and many others since, we take individual consumption and portfolio choices to be the outcome of an "intrapersonal game" in which the same individual consumer is represented by a different player at every date.[1] Consumers are assumed to have the same constant relative risk aversion (CRRA) preferences. Endowments processes may differ across consumers.

We consider competitive equilibria in which consumers all follow intrapersonal

[1] Beyond the consumption-savings problem, the "multiple selves" methodology has been applied to a broad set of self-control issues. See for instance Akerlof (1991), Benabou and Tirole (1999), Carrillo and Mariotti (1997), O'Donoghue and Rabin (1996, 1997).

strategies that give rise to consumption and portfolio choices that are linear in wealth. We give conditions under which there exists a competitive equilibrium in which these types of strategies constitute a Markov perfect equilibrium in the intrapersonal game. The linearity of the individual decision rules implies that standard aggregation results apply. We obtain explicit expressions for the consumption and portfolio strategies, as well as for the equilibrium prices of state contingent claims. The fact that individual behavior is the outcome of a game implies that standard derivations of a transversality condition do not apply. We provide conditions on preferences and endowments that nevertheless rule out bubbles on long-lived assets.

An important first question is to what extent subjective rates of time preference can be inferred from market data. We generalize the observational equivalence result of Barro (1999) to the case of an economy with uncertainty. Subjective discount functions cannot be inferred from state prices or consumption-wealth ratios if the period utility function is logarithmic, or if the conditionally expected growth rate of utility is constant. In these circumstances, the representative agent consumes a constant fraction of wealth in each period, irrespective of the shape of the subjective discount function. It follows that state prices will be the same as if consumers were discounting utilities exponentially.[2]

Although we obtain explicit expressions for state prices, the effect of non-geometric discounting on interest rates and risk premia is not easy to analyze in a discrete-time economy. We derive formulas for the risk-free rate, the market price of risk, and the risk premium on aggregate wealth in a limit economy obtained by letting the length of a period go to zero. We show how interest rates depend on a utility-weighted average of subjective discount rates, and establish that the market price of risk in the limit economy does not depend on the subjective discount function.

This continuous-time limit allows us to further examine the impact of time-inconsistency on asset prices in an economy in which log endowment growth is governed by a Feller (1951) square-root process. Log-endowments may be either trend-stationary or difference-stationary. The square-root process generates persistence in the volatility of endowment growth, and possibly in mean endowment growth as well. As a result of this persistence, geometric and non-geometric discount factors are no longer observationally equivalent.

In this environment, we show that the type of discount functions suggested by experimental studies make aggregate wealth more volatile. The intuition is that the high subjective discount rates for near-future utilities typically generated by hyperbolic discount functions must be matched by low or even significantly negative subjective discount rates for long-run utilities, or else interest rates would be too high compared to what is observed in the data. For the type of endowment processes we consider, these low or negative discount rates at long horizons make wealth more sensitive to new information, thereby increasing volatility. Since the market price of risk is unaffected by consumer rates of time preference, this raises the risk premium

[2]In the presence of market frictions such as transaction costs or borrowing constraints, consumption will typically not be a constant fraction of wealth, even if endowment growth is *i.i.d* or if preferences are logarithmic. This allows one to distinguish between geometric and non-geometric subjective discount factors. See Harris and Laibson (1999).

on aggregate wealth. We provide two numerical examples to illustrate that these effects can be quantitatively significant, with the risk premium on aggregate wealth increasing by several percentage points per annum.[3]

These results suggest that variable subjective rates of time preference may be part of an explanation for the high volatility and risk premia observed in historical data on stock returns (LeRoy and Porter (1981), Shiller (1981), Mehra and Prescott (1985)).

Related Literature Strotz (1956) and Phelps and Pollak (1968) are early authors who considered additively time-separable preferences with non-geometric discounting. The discrete-time "quasi-hyperbolic" discount function introduced by Phelps and Pollak in a model of imperfect intergenerational altruism was later used by Laibson (1994) to capture the qualitative features of hyperbolic discounting for an individual consumer. Laibson (1997) shows that a partially illiquid asset may be used as a commitment device by consumers with time-inconsistent preferences. Harris and Laibson (1999) study the dynamic choices of a quasi-hyperbolic consumer facing a constant risk-free interest rate and subject to borrowing constraints. They derive an Euler equation that depends not only on the level of consumption at two dates, but also on the marginal propensity to consume out of wealth. Krusell and Smith (1999) consider a version of the economy in Mehra and Prescott (1985) with quasi-hyperbolic discount factors in which observational equivalence does not obtain, and argue that consumers must have negative rates of time-preference for nearby utilities in order to account for the low level of interest rates observed in US data.

Gul and Pesendorfer (1999a) have proposed an axiomatic derivation of preferences for commitment that does not rely on non-geometric discounting. The recursive structure of the resulting "dynamic self-control" preferences allows them to apply standard dynamic programming techniques to determine optimal consumer choices, thereby avoiding the multiplicity of equilibria typically associated with the multiple selves interpretation of individual behavior. They study the impact of such preferences for competitive equilibrium in Gul and Pesendorfer (1999b), and argue that an increase in the cost of self-control of consumers can raise the equity premium.

Outline of the Paper The economy is described in Section 2. In Section 3, we analyze the intrapersonal game faced by a typical consumer and derive competitive equilibrium prices. Section 4 derives expressions for interest rates and risk premia in a limiting economy obtained by letting the length of a period go to zero. In Section 5, we assess the quantitative impact of time-inconsistent preferences when endowment growth exhibits serial dependence. Section 6 contains concluding remarks. Readers mainly interested in our quantitative experiments can focus on Section 3.3 together with equations (10), (15)-(16) and (23), compare (23) with the continuous-time results (32)-(34), and then continue with Section 5.

[3] Habit persistence also makes the shape of the subjective discount function matter for asset prices by generating predictability in period utilities. We have experimented with the preferences proposed by Campbell and Cochrane (1999), assuming, as they do, that log-endowment growth is *i.i.d.* Quasi-hyperbolic subjective discount factors in this economy can raise the risk premium on aggregate wealth by as much as 3% per annum.

2. An Exchange Economy

2.1. Environment

We consider a discrete-time, infinite-horizon economy. Time is labelled by $t = 0, 1, 2, ...$

Information Uncertainty is described by a probability space endowed with a filtration $\{\mathcal{F}_t\}_{t=0}^{\infty}$. For each t, we denote by $E_t[\cdot]$ the conditional expectation operator with respect to $\mathcal{F}_t$. Throughout, random variables indexed by t are taken to be $\mathcal{F}_t$-measurable.

Endowments There is a single good available for consumption in every period. The representative consumer's non-negative endowments of this good are denoted by $\{n_t\}_{t=0}^{\infty}$. One interpretation is that the consumer supplies labor inelastically and has access to a linear technology that converts labor into consumption goods. There are also $k_{-1} \geq 0$ units of a long-lived asset that produces non-negative dividends $\{d_t\}_{t=0}^{\infty}$. Aggregate endowments are denoted by $e_t = n_t + d_t k_{-1}$. We assume that aggregate endowments are strictly positive at all dates, with probability one.

Preferences Following Strotz (1956), we view the individual consumer as composed of a sequence of autonomous temporal incarnations, hereafter called date-t consumers, indexed by their period of control over consumption and portfolio decisions, $t = 0, 1, 2, \ldots$ The date-t consumer evaluates current and future consumption according to a utility function:

$$U_t(\{c_s\}_{s=t}^{\infty}) = E_t\left[\sum_{n=0}^{\infty} \delta_n u(c_{t+n})\right], \tag{1}$$

where $\delta_0 = 1$ and $\delta_n \geq 0$.[4] The sum of the δ_n may or may not be finite. The consumer's preferences are time-inconsistent if δ_{n+1}/δ_n is not a constant function of n. Throughout, we assume that the period utility function $u(\cdot)$ is given by:

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}$$

for some $\gamma > 0$, $\gamma \neq 1$. We take $\gamma = 1$ to mean $u(x) = \ln(x)$.

It is possible to modify (1) to incorporate exogenous subsistence levels $\{\underline{c}_t\}_{t=0}^{\infty}$ by letting the period utility function in period t be:

$$\frac{(c - \underline{c}_t)^{1-\gamma}}{1-\gamma} \tag{2}$$

instead of $u(c)$. If we assume that $n_t > \underline{c}_t$ then consumers can achieve subsistence without participating in markets. In this case, all results presented below will continue

[4] Most of what follows in this section generalizes easily to the case in which the date-t consumer uses stochastic subjective discount factors $\{\delta_{t,t+n}\}_{n=0}^{\infty}$ with $\delta_{t,t+n}$ in $\mathcal{F}_{t+n}$.

to apply, without further conditions, if throughout we replace c_t, n_t and e_t by $c_t - \underline{c}_t$, $n_t - \underline{c}_t$ and $e_t - \underline{c}_t$, respectively. The assumption $n_t > \underline{c}_t$ implies that $e_t - \underline{c}_t > d_t k_{-1}$. Therefore net-of-subsistence endowments will be positive, and the present value of the dividends generated by the long-lived asset is guaranteed to be finite if the present value of net-of-subsistence endowments is finite.

Example 1 Standard geometric discounting is equivalent to $\delta_n = \beta^n$ for some $\beta > 0$. The implied rate of time preference is constant and equal to $-\ln(\beta)$.

Example 2 Experimental studies by psychologists (see Herrnstein (1961), and, for a survey, Ainslie (1992)) and economists (Thaler (1981)) suggest that rates of time preference tend to decline as a function of the horizon over which utility is discounted. Loewenstein and Prelec (1992) have proposed an axiomatic justification for the following "generalized hyperbolic" discount function:

$$\delta_n = (1 + \zeta n)^{-\xi/\zeta},$$

where ζ and ξ are both positive. This generalizes the hyperbolic discount function proposed by Ainslie (1975) to account for the reversal over time of preferences for rewards at different horizons. The corresponding discount rate $\xi\,(1 + \zeta n)^{-1}$ declines hyperbolically with the horizon n. The limit case, as $\zeta \to 0$, is the geometric discount function $\delta_n = \exp(-\xi n)$.

Example 3 Phelps and Pollak (1968), and more recently Laibson (1997) and Harris and Laibson (1999), among others, have considered what is now called "quasi-hyperbolic" discounting, following Laibson (1994):

$$\delta_n = \begin{cases} (\delta\beta)^n & \text{if} \quad n \leq N \\ \delta^N \beta^n & \text{if} \quad n > N \end{cases}.$$

Laibson (1994), considering the case $N = 1$ and $\delta < 1$, has argued that this provides a good approximation to the hyperbolic discount function given in the previous example. More generally, the parameter N can be used to construct a kink in the subjective discount function at different horizons. This may be useful when considering alternative assumptions about the length of a period.

2.2. Markets

Our main assumption is that markets are complete.

Assets One-period ahead state-contingent claims are traded at every date and in every state. For any date t, we denote by b_{t+1} the portfolio of such claims purchased by the date-t consumer and maturing at date $t+1$. There are no claims outstanding at date 0: $b_0 = 0$. The long-lived asset or "stock" is traded at every date t at an ex-dividend price s_t.

The Budget Set At date t, the date-t consumer can choose non-negative consumption c_t, a stock portfolio k_t, and a portfolio of state-contingent claims b_{t+1}, subject to the period-t budget constraint:

$$\pi_t c_t + E_t[\pi_{t+1} b_{t+1}] + \pi_t s_t k_t \leq \pi_t n_t + \pi_t b_t + \pi_t (s_t + d_t) k_{t-1}. \tag{3}$$

The $\{\pi_t\}_{t=0}^{\infty}$ are strictly positive probability-weighted state contingent prices. We will usually take $\pi_0 = 1$, making consumption at date 0 the numeraire. In addition, at every date t, the consumer faces the borrowing constraint:

$$\pi_{t+1}[b_{t+1} + (s_{t+1} + d_{t+1})k_t] \geq -E_{t+1}\left[\sum_{s=t+1}^{\infty} \pi_s n_s\right]. \tag{4}$$

Part of the definition of a competitive equilibrium will be that the right-hand side of (4) is finite. We shall assume that consumers who hold a positive amount of the stock can freely dispose of it if the stock price happens to be negative. In any equilibrium, therefore, s_t must be non-negative.

Arbitrage Since we assume that there are one-period claims for every contingency, it must be the case that:

$$\pi_t s_t = E_t\left[\pi_{t+1}(s_{t+1} + d_{t+1})\right]$$

for all $t \geq 0$, or else the date-t consumer would be able to construct an arbitrage.[5] Together with the fact that stock prices must be non-negative, this implies that:

$$\pi_t s_t = E_t\left[\sum_{s=t+1}^{\infty} \pi_s d_s\right] + \pi_t z_t, \tag{5}$$

where $\{z_t\}_{t=0}^{\infty}$ is a non-negative sequence of random variables that satisfies:

$$\pi_t z_t = E_t\left[\pi_{t+1} z_{t+1}\right] \tag{6}$$

for all $t \geq 0$. This says that the value of the stock must be equal to the present value of dividends, plus a non-negative "bubble." In any equilibrium, the price of the stock must be finite, and thus the present value of the dividends must be finite in equilibrium.[6] Using (5) and (6), together with the definition of e_t, we can define the consumer's wealth w_t at date t to be:

$$\pi_t w_t \equiv E_t\left[\sum_{s=t}^{\infty} \pi_s e_s\right] + \pi_t(b_t + z_t k_{t-1}). \tag{7}$$

[5] *A priori*, it might be possible to construct subgame-perfect equilibria of the intrapersonal game in which such arbitrage opportunities are not exploited. We will not consider this possibility here.

[6] An arbitrage strategy that attempts to exploit a situation in which s_t exceeds the present value of future dividends by selling the stock short and using part of the proceeds to finance dividend payments will at some date violate the borrowing constraint (4). The inequality in (4) will be strict in equilibrium and we will give conditions under which $z_t = 0$ if $k_{-1} > 0$.

In the absence of arbitrage opportunities, the portfolio decision in each period amounts to a choice of the amount of state contingent wealth available in the next period. The set of budget-feasible consumption choices defined by (3)-(4) is therefore equivalent to the set of sequences $\{c_t\}_{t=0}^{\infty}$ that for some sequence $\{w_{t+1}\}_{t=0}^{\infty}$ satisfy:

$$\begin{aligned} \pi_t c_t + E_t[\pi_{t+1} w_{t+1}] &\leq \pi_t w_t, \\ c_t, w_{t+1} &\geq 0 \end{aligned} \tag{8}$$

for all $t \geq 0$, where w_0 is given by (7). Note that w_0 is completely determined by state prices if $k_{-1} = 0$. If $k_{-1} > 0$, then the initial value of the bubble on the stock also affects the consumer's initial wealth. Note also that the stochastic process of endowments $\{e_t\}_{t=0}^{\infty}$ only affects the consumer's set of feasible consumption choices via w_0, unlike in models with borrowing constraints that are tighter than (4), such as Harris and Laibson (1999). This means that there is no commitment value to changing the process of $\{e_t\}_{t=0}^{\infty}$, using, for instance, pension commitments that are not directly tradable.

2.3. Intrapersonal and Competitive Equilibrium

Throughout, we maintain the assumption that consumers are price takers. In the absence of any commitment technology, the consumer's individual behavior is the outcome of a strategic interaction between his successive temporal incarnations, and not of a single optimization problem as in the standard theory of consumer demand. In the resulting intrapersonal game, each date-t consumer chooses his current consumption and a portfolio of assets, taking as given a sequence of prices $\{\pi_t, s_t\}_{t=0}^{\infty}$ and the strategies of his successors.

The Intrapersonal Game Given a sequence of prices $\{\pi_t, s_t\}_{t=0}^{\infty}$, a strategy for the date-t consumer in the intrapersonal game is a mapping $\langle C_t(\cdot), W_{t+1}(\cdot)\rangle$ that specifies, for any history h_t of the game up to date t:

(i) A consumption level at date t, $c_t = C_t(h_t)$;

(ii) A state contingent wealth level at date $t+1$, $w_{t+1} = W_{t+1}(h_t)$,

such that the budget constraints (8) are satisfied given the wealth level w_t at date t. The history h_t consists of all events observed by the date-t consumer, including the realizations of endowments and prices, as well as past consumption and wealth choices. Given a price sequence $\{\pi_t, s_t\}_{t=0}^{\infty}$, an intrapersonal equilibrium is a subgame-perfect equilibrium of the intrapersonal game played by the sequence of date-t consumers.

Competitive Equilibrium A competitive equilibrium of the representative agent economy is given by a strategy profile $\{\langle C_t(\cdot), W_{t+1}(\cdot)\rangle\}_{t=0}^{\infty}$ in the intrapersonal game and a price sequence $\{\pi_t, s_t\}_{t=0}^{\infty}$ such that

(i) $\{\pi_t, s_t\}_{t=0}^{\infty}$ satisfies the arbitrage conditions (5)-(6);

(ii) Initial wealth w_0 is finite at prices $\{\pi_t, s_t\}_{t=0}^{\infty}$;

(iii) $\{\langle C_t(\cdot), W_{t+1}(\cdot)\rangle\}_{t=0}^{\infty}$ is an intrapersonal equilibrium at prices $\{\pi_t, s_t\}_{t=0}^{\infty}$;

(iv) On any intrapersonal equilibrium path, the goods market clear at any date.

It follows from the price-taking assumption that when evaluating the payoff of a deviation from the intrapersonal equilibrium, each date-t consumer takes into account the impact of prevailing prices on the continuation equilibrium. In particular, markets need not clear following a deviation.

3. Equilibrium

For an infinitely-lived consumer, the coordination problem between his temporal incarnations may lead to the existence of multiple intrapersonal equilibria. In particular, bootstrap strategies may be used very naturally in a multi-self context to mitigate the consumer's self-control problem (Laibson (1994)). We study these equilibria in Luttmer and Mariotti (1999). As a benchmark, we focus here on intrapersonal Markov equilibria in which the consumption and portfolio choices of the consumer at any date will depend only on his wealth and on exogenous variables.[7] In the following, we first characterize the linear intrapersonal Markov equilibria. We then construct a competitive equilibrium for the representative agent economy.[8]

3.1. Intrapersonal Equilibrium

We shall construct a subgame perfect equilibrium for the intrapersonal game in which a date-t consumer who starts with wealth w chooses consumption $c_t(w)$ and next-period wealth $w_{t+1}(w)$ according to the following linear rule:

$$\begin{aligned} c_t(w) &= \phi_t w \\ w_{t+1}(w) &= \psi_{t+1} w \end{aligned} \tag{9}$$

for some time and state dependent coefficients ϕ_t and ψ_{t+1} that do not depend on wealth. These strategies are Markov strategies in the sense that the choices of the date-t consumer are only influenced by past actions through the state variable wealth.

Given the strategies (9), let $c_{t,t+n}(w)$ denote consumption at date $t+n$ when wealth at the beginning of date t is given by w. Define:

$$V_t(w) = E_t\left[\sum_{n=0}^{\infty} \delta_{n+1} u(c_{t,t+n}(w))\right]. \tag{10}$$

[7] This is also the perspective adopted by Phelps and Pollak (1968) and Harris and Laibson (1999). See Maskin and Tirole (1998) for a thorough investigation of the Markov perfect equilibrium concept.

[8] The derivations we present in this section only cover the case $\gamma \neq 1$. Analogous derivations can be performed for the case of logarithmic preferences. They show that our formulas for wealth-consumption ratios and state prices ((21) and (23) below) continue to apply.

(Note that the first subjective discount factor in $V_t(w)$ is δ_1.) Then (9) yields:

$$V_t(w) = A_t \left(\frac{w^{1-\gamma}}{1-\gamma} \right),$$

where:

$$A_t = E_t \left[\sum_{n=0}^{\infty} \delta_{n+1} \phi_{t+n}^{1-\gamma} \left(\prod_{k=1}^{n} \psi_{t+k} \right)^{1-\gamma} \right]. \tag{11}$$

To construct a subgame perfect equilibrium we need to determine the best response of the date-t consumer when the date-$t+n$ consumers consume a fraction ϕ_{t+n} of their wealth and let wealth grow at a rate ψ_{t+n}, for all $n \geq 1$. Given wealth w_t, this best response is the solution to:

$$\max_{c_t, w_{t+1} \geq 0} \left\{ \frac{c_t^{1-\gamma}}{1-\gamma} + E_t \left[V_{t+1}(w_{t+1}) \right] : \pi_t c_t + E_t[\pi_{t+1} w_{t+1}] \leq \pi_t w_t \right\}.$$

The first-order conditions for this maximization can be written as:[9]

$$\frac{\pi_{t+1}}{\pi_t} = A_{t+1} \left(\frac{w_{t+1}}{c_t} \right)^{-\gamma}. \tag{12}$$

Together with the budget constraint this yields optimal choices for c_t and w_{t+1}. It is not difficult to see that these choices are again of the form (9). The coefficients ϕ_t and ψ_{t+1} that describe the best response of the date-t consumer are determined by:

$$\phi_t = \left(1 + E_t \left[\left(\frac{\pi_{t+1}}{\pi_t} \right)^{1-1/\gamma} A_{t+1}^{1/\gamma} \right] \right)^{-1} \tag{13}$$

$$\psi_{t+1} = \left(1 + E_t \left[\left(\frac{\pi_{t+1}}{\pi_t} \right)^{1-1/\gamma} A_{t+1}^{1/\gamma} \right] \right)^{-1} \left(\frac{A_{t+1}}{\pi_{t+1}/\pi_t} \right)^{1/\gamma}. \tag{14}$$

A subgame perfect equilibrium for the intrapersonal game is now given by the strategies (9) and a sequence $\{A_t, \phi_t, \psi_{t+1}\}_{t=0}^{\infty}$ that satisfies (11) and (13)-(14). Explicit solutions for $\{A_t, \phi_t, \psi_{t+1}\}_{t=0}^{\infty}$ are easily obtained in the case of time-consistent preferences or if the period utility function is logarithmic. When $\gamma \neq 1$, the system of equations linking together $\{A_t, \phi_t, \psi_{t+1}\}_{t=0}^{\infty}$ is non-linear. Given arbitrary state prices, typically no explicit solutions are available, although one can show that there exists a unique and constant solution for $\{A_t, \phi_t, \psi_{t+1}\}_{t=0}^{\infty}$ if $\gamma > 1$ and $\{\pi_{t+1}/\pi_t\}_{t=0}^{\infty}$ is *i.i.d.* Nevertheless, explicit solutions can be obtained at equilibrium prices.

3.2. Competitive Equilibrium

State prices $\{\pi_t\}_{t=0}^{\infty}$, together with strategies for the intrapersonal game that satisfy (9), (11) and (13)-(14), form a competitive equilibrium if and only if the implied consumption choices clear goods markets at all dates and in all states.

[9] A precise justification is given in the proof of Proposition 1.

The first step in constructing an equilibrium will be to use the market clearing conditions to express A_t in terms of the endowment process. To simplify the resulting formulas, define:

$$\Gamma_t = E_t\left[\sum_{n=0}^{\infty}\delta_n\left(\frac{e_{t+n}}{e_t}\right)^{1-\gamma}\right] \tag{15}$$

$$\Delta_t = E_t\left[\sum_{n=0}^{\infty}\delta_{n+1}\left(\frac{e_{t+n}}{e_t}\right)^{1-\gamma}\right], \tag{16}$$

and note that Γ_t and Δ_{t+1} are related via:

$$\Gamma_t = 1 + E_t\left[\left(\frac{e_{t+1}}{e_t}\right)^{1-\gamma}\Delta_{t+1}\right]. \tag{17}$$

Of course, we shall need to assume that the subjective discount factors and endowments are such that Γ_t and Δ_t are finite. It is immediate from (15) that this is true if and only if the discounted utility from the aggregate endowments is finite from the perspective of every date-t consumer.

It follows from the linear consumption and portfolio strategies (9) that market clearing at all dates and in all states is equivalent to:

$$e_0 = \phi_0 w_0 \tag{18}$$

together with:

$$\frac{e_{t+1}}{e_t} = \frac{\phi_{t+1}\psi_{t+1}}{\phi_t} \tag{19}$$

for all $t \geq 0$ and in all states. Using (19) and the definitions (11) and (16) of A_t and Δ_t one can verify that A_t must therefore satisfy:

$$A_t = \phi_t^{1-\gamma}E_t\left[\sum_{n=0}^{\infty}\delta_{n+1}\left(\frac{\phi_{t+n}}{\phi_t}\prod_{k=1}^{n}\psi_{t+k}\right)^{1-\gamma}\right] = \phi_t^{1-\gamma}\Delta_t \tag{20}$$

in any equilibrium. At $c_t = e_t$, the first-order condition (12) can be written as $\pi_{t+1}/\pi_t = A_{t+1}\phi_{t+1}^{\gamma}(e_{t+1}/e_t)^{-\gamma}$. Together with (13) this gives:

$$\phi_t = \left(1 + E_t\left[A_{t+1}\phi_{t+1}^{-(1-\gamma)}\left(\frac{e_{t+1}}{e_t}\right)^{1-\gamma}\right]\right)^{-1}.$$

In combination with (17) and (20) this implies that ϕ_t must be given by:

$$\phi_t = \frac{1}{\Gamma_t} \tag{21}$$

in any equilibrium. For these consumption-wealth ratios, portfolio choices ψ_{t+1} must satisfy:

$$\psi_{t+1} = \frac{e_{t+1}\Gamma_{t+1}}{e_t\Gamma_t}. \tag{22}$$

to ensure that consumption grows at the rate required by (19). Note from (20) and (21) that $A_t\phi_t^\gamma = \Delta_t/\Gamma_t$. State prices must therefore be given by:

$$\frac{\pi_{t+1}}{\pi_t} = \frac{\Delta_{t+1}}{\Gamma_{t+1}}\left(\frac{e_{t+1}}{e_t}\right)^{-\gamma}. \tag{23}$$

At this point, we have constructed consumption-wealth ratios (21), portfolio choices (22), and state prices (23), for which the growth rate of consumption is equal to the growth rate of endowments in all states and at all dates $t \geq 1$. It remains to show that the goods market clears at date 0.

It follows from (18), (21), and the definition (7) of wealth at date 0 that market clearing at date 0 requires that:

$$\Gamma_0 = \frac{1}{\pi_0 e_0} E_0\left[\sum_{t=0}^{\infty} \pi_t e_t\right] + \frac{z_0 k_{-1}}{e_0}.$$

Given our construction of state prices, we shall argue that the first term on the right-hand side of this equation is finite and no greater than Γ_0. If $k_{-1} > 0$, this establishes that an equilibrium exists: one can simply adjust $z_0 \geq 0$ to make up the difference between $e_0\Gamma_0$ and the present value of endowments. If $z_0 > 0$ is required, then there is a bubble on the long-lived asset. We shall give conditions that ensure that $e_0\Gamma_0$ is actually equal to the present value of endowments. These conditions then imply that there can, in fact, be no bubble on the long-lived asset. Furthermore, these conditions ensure that an equilibrium also exists if $k_{-1} = 0$.

To see why $e_0\Gamma_0$ cannot be less than the present value of endowments, note that the solution for state prices (23) together with (17) implies:

$$\Gamma_t = 1 + E_t\left[\left(\frac{\pi_{t+1}e_{t+1}}{\pi_t e_t}\right)\Gamma_{t+1}\right].$$

Because state prices and endowments are positive we can thus write:

$$\Gamma_0 = E_0\left[\sum_{t=0}^{\infty} \frac{\pi_t e_t}{\pi_0 e_0}\right] + \lim_{T\to\infty} E_0\left[\left(\frac{\pi_T e_T}{\pi_0 e_0}\right)\Gamma_T\right].$$

Since Γ_t is positive by construction, the limit term on the right-hand side of this equation is non-negative, and thus $e_0\Gamma_0$ must dominate the present value of endowments.

To rule out bubbles and ensure existence of equilibrium in case $k_{-1} = 0$ we therefore need to show that (using (23)):

$$\lim_{T\to\infty} E_0\left[\prod_{t=1}^{T} \frac{\Delta_t}{\Gamma_{t-1}}\left(\frac{e_t}{e_{t-1}}\right)^{1-\gamma}\right] = 0. \tag{24}$$

This does not follow simply from assuming that Δ_t and Γ_t are finite at all dates and in all states. We shall consider two assumptions under which (24) holds. These conditions cover our applications in Section 5.

Condition A *There is a* $\kappa \in (0,1)$ *such that* $\kappa\Gamma_t \in (0,1]$ *at all dates and in all states.*

Condition B *There is some* $\beta > 0$ *such that* $\delta_{n+1}/\delta_n \leq \beta$ *for all* $n \geq 0$*, and for which:*

$$E_t\left[\sum_{n=0}^{\infty}\beta^n\left(\frac{e_{t+n}}{e_t}\right)^{1-\gamma}\right] < \infty \tag{25}$$

at all dates and in all states.

Condition A is equivalent to stating that the equilibrium consumption-wealth ratio is bounded away from zero. This ensures that consumers do not over-accumulate wealth. Condition B is a simple way to translate arguments that rule out bubbles in economies in which consumers discount geometrically to economies with more complicated subjective discount factors. One application of Condition B is an economy with $\gamma > 1$ and $\delta_{n+1} \leq \delta_n$, and with endowments that are expected to grow exponentially in the long run, following every date and state. In this case we can simply take $\beta = 1$. The proof of the following proposition is given in Appendix A.

Proposition 1 *If at least one of Conditions A or B holds, then there is an equilibrium and there is no bubble on the long-lived asset if it is in positive net supply.*

In Luttmer and Mariotti (1999) we consider the $N = 1$ case of the quasi-hyperbolic preferences introduced in Example 3. For this case, we show that there is always a unique linear Markov equilibrium, provided Γ_0 is finite.[10]

A Bubble Example Proposition 1 implies that there can be no bubble if preferences are time-consistent, since in that case Condition B applies. If Conditions A and B do not hold, then a linear Markov equilibrium may not exist when $k_{-1} = 0$, or there may be a bubble on the long-lived asset when $k_{-1} > 0$. As an example, consider an economy with $\gamma \neq 1$ and endowments given by $e_0 > 0$ and:

$$\frac{e_t}{e_{t-1}} = t^{(1+\varepsilon)/(1-\gamma)},$$

for some $\varepsilon > 0$ and for all $t \geq 1$. These endowment growth rates are so large that utility will not be finite for any geometric subjective discount function. No equilibrium would exist if consumers were time consistent, and Condition B must be violated. Suppose instead that the representative consumer's subjective discount factors are given by:

$$\delta_n = \left(\frac{\beta^n}{n!}\right)^{1+\varepsilon},$$

where β is positive and smaller than one. These subjective discount factors decline fast enough to ensure that:

$$\Delta_t = \sum_{n=0}^{\infty}\left(\frac{\beta^{n+1}(t+n)!}{(n+1)!t!}\right)^{1+\varepsilon}$$

[10] If $\delta \leq 1$ in Example 3, this follows directly from Proposition 1 as Condition B is then satisfied. If $\delta > 1$, a separate argument is needed as Condition B may not hold.

is finite.[11] Clearly, Δ_t is increasing and unbounded in t. Recall from (17) that $\Gamma_{t-1} = 1 + (e_t/e_{t-1})^{1-\gamma}\Delta_t$. So Condition A must be violated, and:

$$\begin{aligned}\prod_{t=1}^{T} \frac{\Delta_t}{\Gamma_{t-1}} \left(\frac{e_t}{e_{t-1}}\right)^{1-\gamma} &= \prod_{t=1}^{T} \frac{(e_t/e_{t-1})^{1-\gamma}\,\Delta_t}{1+(e_t/e_{t-1})^{1-\gamma}\,\Delta_t} \\ &\geq \prod_{t=1}^{T} \frac{(e_t/e_{t-1})^{1-\gamma}\,\Delta_0}{1+(e_t/e_{t-1})^{1-\gamma}\,\Delta_0} \\ &= \prod_{t=1}^{T}\left(1-\frac{1}{1+\Delta_0 t^{1+\varepsilon}}\right).\end{aligned}$$

This converges to a positive number as $T \to \infty$, since ε and Δ_0 are both strictly positive. Hence, (24) is violated, and this economy has no linear Markov equilibrium if there is no long-lived asset in positive net supply. If $k_{-1} > 0$, such an equilibrium exists, but there must be a bubble on the long-lived asset.

Average Subjective Discount Factors As can be seen from the equation for state prices (23), the two variables that determine state prices in this economy are endowment growth and the ratio Δ_t/Γ_t. If consumers discount geometrically, the ratio Δ_t/Γ_t is constant and equal to the subjective discount factor β. For general subjective discount factors, the ratio Δ_t/Γ_t can be expressed as a weighted average of δ_{n+1}/δ_n:

$$\frac{\Delta_t}{\Gamma_t} = \sum_{n=0}^{\infty} \omega_{n,t}\left(\frac{\delta_{n+1}}{\delta_n}\right), \tag{26}$$

where the weights $\omega_{n,t}$ are given by:

$$\omega_{n,t} = \frac{E_t\left[\delta_n\,(e_{t+n}/e_t)^{1-\gamma}\right]}{E_t\left[\sum_{n=0}^{\infty}\delta_n\,(e_{t+n}/e_t)^{1-\gamma}\right]}.$$

These weights are proportional to the expected utility of date-$t+n$ consumption from the perspective of the date-t consumer. In the special case of quasi-hyperbolic discounting, $\delta_0 = 1$ and $\delta_n = \delta\beta^n$ for all $n \geq 1$, this yields the "generalized Euler equation:"

$$\frac{\pi_{t+1}}{\pi_t} = \beta\left(1-\frac{1-\delta}{\Gamma_t}\right)\left(\frac{e_{t+1}}{e_t}\right)^{-\gamma}, \tag{27}$$

of Harris and Laibson (1999). More generally, consider the properties of (26) if subjective discount rates are relatively high at nearby horizons, and low at distant horizons. The discount rates implied by Δ_t/Γ_t will then depend on the timing of endowment growth. If $\gamma > 1$, high early endowment growth lowers the weights on δ_{n+1}/δ_n for small values of n and this lowers the discount rate implied by Δ_t/Γ_t. If

[11] Note that $\sum_{n=0}^{\infty}\beta^n\binom{t-1+n}{n} = (1-\beta)^{-t}$. From this one can compute Δ_t at $\varepsilon = 0$. The result then follows for $\varepsilon > 0$ since all but finitely many of the terms in the series that defines Δ_t must be smaller than one.

the same amount of endowment growth is delayed, more weight is put on δ_{n+1}/δ_n for small values of n. Delayed growth therefore increases the discount rate implied by Δ_t/Γ_t.

Identification The characteristic feature of non-geometric discounting is the fact that Δ_t/Γ_t replaces the usual geometric subjective discount factor. This implies that it is not possible to distinguish non-geometric discount factors from geometric ones if Δ_t/Γ_t happens to be constant. This will be the case if conditionally expected utility growth, $E_t\left[(e_{t+1}/e_t)^{1-\gamma}\right]$, is constant. For any subjective discount function, one can then construct an alternative economy with a geometric subjective discount factor given by $\overline{\beta} = \Delta_t/\Gamma_t$. State prices will be the same in both economies. It is not difficult to verify that the wealth-consumption ratio in the alternative economy is again equal to expected utility growth discounted using the geometric subjective discount factor $\overline{\beta}$. Thus consumption-wealth ratios cannot be used to identify properties of the subjective discount function either.

As an example, one can take endowment growth to be *i.i.d.* and the information structure $\{\mathcal{F}_t\}_{t=0}^{\infty}$ such that at any date nothing is known about future endowment growth. Alternatively, one can take preferences to be logarithmic. If $\gamma = 1$, then Γ_t and Δ_t are simply sums of subjective discount factors, and therefore constant across time.[12]

3.3. Some Intuition about State Prices and Marginal Utilities

In constructing our equilibrium, we have already defined a value function $V_t(w)$ in (10). This value function represents the expected utility from date t on, as viewed by the date-$t-1$ consumer, in the date-t subgame in which the date-t consumer starts with wealth w. Now let $F_t(w)$ be expected utility for the date-t consumer in the same subgame:

$$F_t(w) = E_t\left[\sum_{n=0}^{\infty} \delta_n u(c_{t,t+n}(w))\right].$$

Clearly, if $\delta_n = \beta^n$ then $V_t(w) = \beta F_t(w)$. More generally, the two value functions are related via:

$$F_t(w_t) = \max_{c_t, w_{t+1}} \left\{u(c_t) + E_t\left[V_{t+1}(w_{t+1})\right]\right\},$$

where the maximization is subject to the budget constraint (8).[13] The first-order and envelope conditions for this maximization are:

$$\begin{aligned} \frac{\pi_{t+1}}{\pi_t} &= \frac{DV_{t+1}(w_{t+1})}{Du(c_t)}, \\ DF_t(w_t) &= Du(c_t) \end{aligned}$$

[12] Barro (1999) observes that in the standard deterministic Cass-Koopmans growth model one cannot infer from data whether consumers discount geometrically or not if the economy is in steady state, or if preferences are logarithmic. See also Laibson (1996).

[13] This equation is related to the "quasi-Bellman equation" developed by Harris and Laibson (1999) for quasi-hyperbolic subjective discount functions.

in all states and at all dates. Combining the first-order condition at date t with the envelope condition at date $t+1$, one can write:

$$\frac{\pi_{t+1}}{\pi_t} = \frac{DV_{t+1}(w_{t+1})}{DF_{t+1}(w_{t+1})}\frac{Du(c_{t+1})}{Du(c_t)}. \tag{28}$$

That is, the usual (geometric) subjective discount factor is replaced by a ratio of marginal utilities of wealth based on the "one-period ahead" value function V_{t+1}, and next period's "current" value function F_{t+1}. Given that preferences are homothetic, and that consumption and portfolio strategies are linear in wealth, it is not difficult to see that $V_t(w) = \Delta_t\,(w/\Gamma_t)^{1-\gamma}/(1-\gamma)$ and $F_t(w) = \Gamma_t\,(w/\Gamma_t)^{1-\gamma}/(1-\gamma)$ in equilibrium. Combining this with (28) implies that equilibrium state prices are indeed given by (23).

4. Continuous-Time Approximations

In the discrete-time economies we have considered so far, it was relatively easy to construct a competitive equilibrium. However, characterizing the properties of this equilibrium is more difficult. As we show in this section, this task is simplified considerably when we let the length of a period go to zero.

4.1. Limit Properties of Discrete-Time Economies

Suppose that the subjective discount function $\delta(t)$ is defined for all $t \in [0,\infty)$, and normalized so that $\delta(0) = 1$. Suppose also that there exists an underlying continuous-time endowment process $\{e_t\}_{t\geq 0}$ that evolves according to a diffusion:

$$d\ln(e_t) = \mu_e(x_t)dt + \sigma_e(x_t)^\top dB_t, \tag{29}$$

where $\{B_t\}_{t\geq 0}$ is a vector of independent standard Brownian motions, and $\{x_t\}_{t\geq 0}$ is a vector of state variables that satisfies:[14]

$$dx_t = \mu_x(x_t)dt + \sigma_x(x_t)^\top dB_t. \tag{30}$$

Given the discount function $\delta(\cdot)$ and the endowment process $\{e_t\}_{t\geq 0}$, we construct a sequence of discrete-time economies as follows. For any $\tau > 0$, interpreted as the length of a period, consider a discrete-time economy with a sequence of subjective discount factors $\delta_n = \delta(n\tau)$, $n = 0, 1, 2, \ldots$, and endowments in period n given by $\tau e_{n\tau}$. For any $t \geq 0$ and $\tau > 0$, let:

$$\begin{aligned}
\Gamma_t(\tau) &= E_t\left[\sum_{n=0}^{\infty}\delta(n\tau)\left(\frac{e_{t+n\tau}}{e_t}\right)^{1-\gamma}\right], \\
\Delta_t(\tau) &= E_t\left[\sum_{n=0}^{\infty}\delta((n+1)\tau)\left(\frac{e_{t+n\tau}}{e_t}\right)^{1-\gamma}\right]
\end{aligned}$$

[14] Without loss of generality, we take the dimension of B_t to be equal to one plus the dimension of x_t.

denote the analogues of (15) and (16) for the discrete-time economy with period length τ. It then follows from (23) and an appropriate normalization that the state prices for this economy are given by $\{\pi_{n\tau}(\tau)\}_{n=0}^{\infty}$, where for any $t \geq 0$:

$$\pi_t(\tau) = \exp\left(\sum_{n=1}^{[\![t/\tau]\!]} \left(\frac{1}{\tau} \ln\left(\frac{\Delta_{n\tau}(\tau)}{\Gamma_{n\tau}(\tau)}\right)\right)\tau\right) e_{\tau[\![t/\tau]\!]}^{-\gamma}, \tag{31}$$

and $[\![t/\tau]\!]$ denotes the integer part of t/τ. From (21), $1/\Gamma_{n\tau}(\tau)$ is the consumption-wealth ratio based on consumption at date $n\tau$ in the discrete-time economy with period length τ.

We now examine the sample path properties of the discrete-time state price process (31) when τ goes to zero. Note first that the consumption-wealth ratio and the ratio $1/\Delta_{n\tau}(\tau)$ vanish as τ goes to zero. In Appendix B we provide conditions that ensure that, for every $t \geq 0$, the quantities $\tau\,\Delta_t(\tau)$ and $\tau\,\Gamma_t(\tau)$ converge almost surely to $\Gamma(x_t)$ as τ goes to zero, where the function $\Gamma(\cdot)$ is defined by:[15]

$$\Gamma(x_t) = E_t\left[\int_0^{\infty} \delta(v)\left(\frac{e_{t+v}}{e_t}\right)^{1-\gamma} dv\right]. \tag{32}$$

Furthermore, we show that, for every $t \geq 0$, the difference $\Gamma_t(\tau) - \Delta_t(\tau)$ converges almost surely to $\Phi(x_t)$ as τ goes to zero, where the function $\Phi(\cdot)$ is defined by:

$$\Phi(x_t) = -E_t\left[\int_{[0,\infty)} \left(\frac{e_{t+v}}{e_t}\right)^{1-\gamma} d\delta(v)\right], \tag{33}$$

and the integration in (33) is with respect to the measure induced by $\delta(\cdot)$ on $[0,\infty)$. This, together with (31), implies the following proposition.

Proposition 2 *Under regularity conditions, the state prices $\{\pi_t(\tau)\}_{t\geq 0}$ for the discrete-time economy with period length τ converge to:*

$$\{\pi_t\}_{t\geq 0} = \left\{\exp\left(-\int_0^t \frac{\Phi(x_v)}{\Gamma(x_v)}\,dv\right) e_t^{-\gamma}\right\}_{t\geq 0} \tag{34}$$

as τ goes to zero, almost surely.

Appendix B gives a precise statement and proof of this result. Note that the sense in which the process $\{\pi_t(\tau)\}_{t\geq 0}$ converges to $\{\pi_t\}_{t\geq 0}$ as τ goes to zero is quite strong: almost every sample path of state prices in the discrete-time economy with period length τ converges pointwise to the corresponding continuous-time sample path.

[15]Note that since $\{x_t\}_{t\geq 0}$ and $\{e_t\}_{t\geq 0}$ are time-homogenous diffusion processes, they both satisfy the Markov property, hence $\Gamma(\cdot)$ is independent of t. The same applies to the function $\Phi(\cdot)$ defined below.

Interest Rates and Risk Premia From (34), the limiting state price process $\{\pi_t\}_{t\geq 0}$ satisfies:

$$d\pi_t = -\pi_t \left(r(x_t)dt + \sigma_\pi(x_t)^\top dB_t \right), \tag{35}$$

where the instantaneous risk-free rate $r(x_t)$ and the "market price of risk" $\sigma_\pi(x_t)$ are determined by:

$$r(x_t) = \frac{\Phi(x_t)}{\Gamma(x_t)} + \gamma\mu_e(x_t) - \frac{1}{2}\gamma^2\sigma_e(x_t)^\top\sigma_e(x_t), \tag{36}$$

$$\sigma_\pi(x_t) = \gamma\sigma_e(x_t). \tag{37}$$

Observe that the shape of the subjective discount function $\delta(\cdot)$ is reflected in the first term in (36). Clearly, if $\delta(t) = \exp(-\rho t)$, then $\Phi(x_t) = \rho\Gamma(x_t)$ and this term reduces to ρ as expected. The last two terms in (36) are the usual terms that result from endowment growth and curvature of the period utility function. It follows from (37) that the market price of risk $\sigma_\pi(x_t)$ is not affected by the shape of the subjective discount function $\delta(\cdot)$.

Let $R(x_t)$ be the risk premium on aggregate wealth—that is, the instantaneous expected excess return on aggregate wealth. In the limit economy, aggregate wealth at time t is equal to $e_t\Gamma(x_t)$. Standard calculations imply that $R(x_t)$ is given by the negative of the instantaneous covariance between $e_t\Gamma(x_t)$ and π_t, divided by $\pi_t e_t\Gamma(x_t)$. This yields:

$$R(x_t) = \gamma \left(\sigma_e(x_t)^\top + \frac{D\Gamma(x_t)\sigma_x(x_t)^\top}{\Gamma(x_t)} \right) \sigma_e(x_t). \tag{38}$$

The risk premium on aggregate wealth therefore does depend on the discount function $\delta(\cdot)$, via its effect on the wealth-consumption ratio $\Gamma(x_t)$.

If $\delta(\cdot)$ is sufficiently smooth so that one can write $\delta(t) = \exp\big(-\int_0^t \rho(v)\,dv\big)$, then the first term in (36) simplifies to a weighted average of the subjective discount rates $\{\rho(v)\}_{v\geq 0}$:

$$\frac{\Phi(x_t)}{\Gamma(x_t)} = \int_0^\infty \rho(v)\,\omega(x_t, v)\,dv, \tag{39}$$

where the weights $\omega(x_t, v)$ are given by:

$$\omega(x_t, v) = \frac{E_t\left[\delta(v)\,(e_{t+v}/e_t)^{1-\gamma}\right]}{E_t\left[\int_0^\infty \delta(v)\,(e_{t+v}/e_t)^{1-\gamma}\,dv\right]},$$

as in (26). Alternatively, non-smooth subjective discount functions arise if consumers discount any positive delay of utility by a discrete amount, while utility at positive horizons is discounted at an instantaneous rate. Consider for example the discount function $\delta(0) = 1$, and $\delta(t) = \delta\exp(-\rho t)$ for all $t > 0$. The jump in the discount function shows up in the risk-free rate (36) via:

$$\frac{\Phi(x_t)}{\Gamma(x_t)} = \rho + \frac{1-\delta}{\Gamma(x_t)}. \tag{40}$$

Note that $\delta(\cdot)$ can be seen as the limit as τ goes to zero of discrete-time quasi-hyperbolic discount functions with $\delta_0(\tau) = 1$ and $\delta_n(\tau) = \delta\exp(-\rho n\tau)$ for each $n \geq 1$.

One feature of this limit is that the wealth-consumption ratio (32) is proportional to the time-inconsistency parameter δ:

$$\Gamma(x_t) = \delta E_t \left[\int_0^\infty \exp(-\rho v) \left(\frac{e_{t+v}}{e_t} \right)^{1-\gamma} dv \right].$$

This implies that the vector $D\Gamma(x_t)/\Gamma(x_t)$, and therefore the risk-premium on aggregate wealth (38), only depend on ρ and not on δ.

4.2. A Continuous-Time Economy with Partial Commitment

So far we have studied the limiting properties of a sequence of discrete-time economies as the period length goes to zero. However, it is also possible to analyze the continuous-time economy directly. Following Barro (1999), one way to do this is to assume that consumers can commit to a particular consumption strategy for a short period of time, and then let this commitment period go to zero. We now investigate under what conditions this continuous-time approximation is equivalent to the one given in the previous section.

A Sequence of Partial Commitments Assume as above that the endowment and state variable processes $\{e_t\}_{t\geq 0}$ and $\{x_t\}_{t\geq 0}$ evolve as in (29) and (30), and consider an infinitely lived consumer whose preferences at any date t over continuous-time future streams of consumption $\{c_{t+v}\}_{v\geq 0}$ are given by:

$$U_t(\{c_{t+v}\}_{v\geq 0}) = E_t \left[\int_0^\infty \delta(v)\, u(c_{t+v})\, dv \right],$$

where $u(\cdot)$ is a CRRA utility function, as before. Suppose that at any date $t = n\tau$, $n = 0, 1, 2, \ldots$, the date-t consumer can commit to consumption choices for an episode $\tau > 0$. Thus, at any such date t, the date-t consumer commits to a stream of consumption $\{c_{t+v}\}_{v\in[0,\tau)}$, as well as to an amount of wealth $w_{t+\tau}$ to be left to the date-$t+\tau$ consumer. The budget constraint for the date-t consumer is then:

$$E_t \left[\int_0^\tau \pi_{t+v} c_{t+v}\, dv + \pi_{t+\tau} w_{t+\tau} \right] \leq \pi_t w_t, \tag{41}$$

together with the requirement that $w_{t+\tau} \geq 0$.[16] For any date $t = n\tau$, $n = 0, 1, 2, \ldots$, let $\{c_{t,t+v}(w)\}_{v\geq 0}$ and $\{w_{t,t+k\tau}(w)\}_{k=1}^\infty$ be the consumption and portfolio strategies in the subgame in which the date-t consumer starts period t with wealth w. As in the case of the discrete-time economy, we will focus on consumption and portfolio strategies that are linear in current wealth:

$$\begin{aligned} c_{t,t+v}(w) &= \phi_{t+v} w, \\ w_{t,t+v}(w) &= \psi_{t+v} w, \end{aligned}$$

[16] One possible interpretation is that at every date $t = n\tau$ there is a complete set of state-contingent claims for every date $t+v$, $v \in [0, \tau]$. Alternatively, one can assume that there is continuous trading in a certain set of long-lived assets, in such a way that markets are dynamically complete.

where $v \in [0, \tau)$. For any date $t = n\tau$, $n = 1, 2, \ldots$, let $V_t(w)$ be the discounted utility expected by the date-$t - \tau$ consumer from date t on, in the subgame in which the date-t consumer starts with wealth w:

$$V_t(w) = E_t \left[\int_0^\infty \delta(v + \tau)\, u(c_{t,t+v}(w))\, dv \right].$$

As in the discrete-time economy, the linearity of the consumption and wealth strategies implies that $V_t(w) = A_t w^{1-\gamma}/(1-\gamma)$. As a result, the optimal consumption and portfolio choices for the date-t consumer will be linear if those of his successors are linear. One can thus solve for the equilibrium decision rules and state prices exactly as in Section 2. For any $t = n\tau$, $n = 0, 1, 2, \ldots$, the analogues of (15) and (16) can be defined as $\widetilde{\Gamma}_t = \Gamma(x_t)$, and:

$$\widetilde{\Delta}_t(\tau) = E_t \left[\int_0^\infty \delta(v + \tau) \left(\frac{e_{t+v}}{e_t} \right)^{1-\gamma} dv \right].$$

Note that $\widetilde{\Gamma}_t$ is independent of the commitment period τ. Consumption and wealth choices are determined by:

$$\begin{aligned} \widetilde{\phi}_{t+v} &= \frac{e_{t+v}}{e_t \widetilde{\Gamma}_t}, \\ \widetilde{\psi}_{t+\tau} &= \frac{e_{t+\tau} \widetilde{\Gamma}_{t+\tau}}{e_t \widetilde{\Gamma}_t}, \end{aligned}$$

for each $v \in [0, \tau)$, and equilibrium state prices are given by:

$$\begin{aligned} \frac{\widetilde{\pi}_{t+v}(\tau)}{\widetilde{\pi}_t(\tau)} &= \delta(v) \left(\frac{e_{t+v}}{e_t} \right)^{-\gamma}, \\ \frac{\widetilde{\pi}_{t+\tau}(\tau)}{\widetilde{\pi}_t(\tau)} &= \frac{\widetilde{\Delta}_{t+\tau}(\tau)}{\widetilde{\Gamma}_{t+\tau}} \left(\frac{e_{t+\tau}}{e_t} \right)^{-\gamma}, \end{aligned}$$

where $v \in [0, \tau)$ and $t = n\tau$, $n = 1, 2, \ldots$. If we let $\widetilde{\pi}_0(\tau) = e_0^{-\gamma}$, the process for state prices can be written as:

$$\widetilde{\pi}_t(\tau) = \delta(t - \tau[\![t/\tau]\!]) \exp\left(\sum_{k=1}^{[\![t/\tau]\!]} \ln\left(\frac{\widetilde{\Delta}_{k\tau}(\tau)}{\widetilde{\Gamma}_{k\tau}} \right) \right) e_t^{-\gamma}. \tag{42}$$

Observe that these state prices will in general differ from the ones obtained in an economy with geometric discount factors, even if endowment growth is *i.i.d.* or if utility is logarithmic. The factor $\delta(t - \tau[\![t/\tau]\!])$ introduces a periodicity in state prices that allows one to identify $\delta(t)$ over the interval $[0, \tau)$.

Small-τ Limits We now investigate the sample path properties of the state price process (31) when τ goes to zero. Note first that, as τ goes to zero, $\delta(t - \tau[\![t/\tau]\!])$

converges to $\delta(0^+)$, for any $t \geq 0$. If $\delta(0^+) < 1$, this only amounts to a different normalization of the limiting state price process, and we can therefore ignore it. Using the fact that:

$$\sum_{k=1}^{[\![t/\tau]\!]} \ln\left(\frac{\widetilde{\Delta}_{k\tau}(\tau)}{\widetilde{\Gamma}_{k\tau}}\right) = \sum_{k=1}^{[\![t/\tau]\!]} \frac{1}{\tau} \ln\left(1 - \tau\left(\frac{\frac{1}{\tau}\left(\widetilde{\Gamma}_{k\tau} - \widetilde{\Delta}_{k\tau}(\tau)\right)}{\widetilde{\Gamma}_{k\tau}}\right)\right)\tau$$

for all $t \geq 0$ and $\tau > 0$, one can show that the state prices $\{\widetilde{\pi}_t(\tau)\}_{t\geq 0}$ for the continuous-time economy with commitment period τ converge almost surely as τ goes to zero to a state-price process proportional to:

$$\{\widetilde{\pi}_t\}_{t\geq 0} = \left\{\exp\left(-\int_0^t \frac{\widetilde{\Phi}_v}{\widetilde{\Gamma}_v}\, dv\right) e_t^{-\gamma}\right\}_{t\geq 0}, \tag{43}$$

where, for each $t \geq 0$:

$$\widetilde{\Phi}_t = \lim_{\tau\to 0^+} \frac{\widetilde{\Gamma}_t - \widetilde{\Delta}_t(\tau)}{\tau} = -E_t\left[\int_{(0,\infty)} \left(\frac{e_{t+v}}{e_t}\right)^{1-\gamma} d\delta(v)\right], \tag{44}$$

If $\delta(\cdot)$ is continuous at 0, then $\widetilde{\Phi}_t = \Phi(x_t)$ for each $t \geq 0$, almost surely, where $\Phi(\cdot)$ is given by (33). Hence, in that case, the sample paths of the limit state-price process $\{\widetilde{\pi}_t\}_{t\geq 0}$ coincide almost surely with those of $\{\pi_t\}_{t\geq 0}$ obtained as the limit of state prices for discrete-time economies when the period length goes to zero. But in general, (33) and (44) differ by an amount $1 - \delta(0^+)$, so that the two limit state-price processes do not coincide.

This potential discrepancy is an artifact of the commitment technology. The reason is that for any $\tau > 0$ and $n = 0, 1, 2, \ldots$, the presence of a jump in the discount function $\delta(\cdot)$ at 0 does not affect the optimal choice of the date-$n\tau$ consumer who can commit to a consumption stream over the interval of time $[n\tau, (n+1)\tau)$. The discontinuity of the discount function at 0 cannot therefore be reflected in the limit state-price process $\{\widetilde{\pi}_t\}_{t\geq 0}$ obtained as τ goes to zero. A possible way to capture the impact of such sharp changes in the discount function in the neighborhood of the present period would be to consider a sequence of economies indexed by $\tau > 0$, where now τ not only represents the length of the commitment period, but also determines the shape of the subjective discount function. For instance:

$$\delta_\tau(t) = \begin{cases} 1 & \text{if} \quad t \in [0, \tau) \\ \delta(t) & \text{if} \quad t \in [\tau, \infty) \end{cases}, \tag{45}$$

This corresponds to a case in which consumers do not discount over the commitment period. Let $\Gamma_t(\tau)$ be the equilibrium wealth-consumption ratio for such an economy. In Appendix C, we show that:

$$\lim_{\tau\to 0^+} \frac{\Gamma_t(\tau) - \Delta_t(\tau)}{\tau} = -E_t\left[\int_{[0,\infty)} \left(\frac{e_{t+v}}{e_t}\right)^{1-\gamma} d\delta(v)\right] = \Phi(x_t) \tag{46}$$

for each $t \geq 0$, almost surely. It is easy to check that (46) implies that the sample paths of the process obtained as the limit of state-prices for the continuous-time economies when the commitment period length goes to zero coincide almost surely with those of the process constructed as the limit of state-prices for discrete-time economies when the period length goes to zero. This provides an alternative motivation for (34).

5. Time-Inconsistency and Volatility

As emphasized in Section 2, economies with different subjective discount factors are observationally equivalent if conditionally expected utility growth is constant. In this section we shall consider a model of endowment growth that exhibits serial dependence, and show how quasi-hyperbolic and hyperbolic discount factors can make aggregate wealth more volatile than when preferences are time consistent.

We shall assume that consumers discount using a subjective discount function $\delta(\cdot)$ that satisfies $\delta(0) = 1$ and either:

$$\delta(t) = \delta \exp(-\rho t)$$

for any $t > 0$ (the quasi-hyperbolic case), or:

$$\delta(t) = (1 + \zeta t)^{-\xi/\zeta} \exp(-\rho t)$$

for any $t > 0$ (the hyperbolic case.) Taking $\delta = 1$ or $\zeta \to 0$ corresponds to geometric discounting, at rates ρ and $\rho + \xi$, respectively. Note that the second discount function combines a geometric part with the generalized hyperbolic discount function proposed by Loewenstein and Prelec (1992). For this combined discount function, the subjective rate of time preference converges to $\rho + \xi$ as the horizon goes to zero and to ρ as the horizon goes to infinity. The parameter ζ governs the speed at which the subjective rate of time preference changes from its short-run value $\rho + \xi$ to its long-run value ρ. The parameter ρ also determines the consumer's long-term subjective discount rate in the quasi-hyperbolic case. Keeping the value of ρ low, or even negative, while at the same time choosing $\delta < 1$ or $\zeta > 0$ allows us to examine preferences that exhibit both high discount rates in the short run and very low, or even negative discount rates in the long run.[17] Recall from (36) and (38) that the risk-free rate depends on $\Phi(x)/\Gamma(x)$, and that the risk premium on aggregate wealth depends on $D\Gamma(x)/\Gamma(x)$. A key factor in our analysis will be how these ratios depend on the long-run rate of time preference ρ. Our specification of the endowment process will enable us to characterize this dependence explicitly.

[17] Kocherlakota (1990) shows that competitive equilibria may exist in an infinite-horizon economy even if consumers have negative discount rates, and suggests this as part of solution to the equity premium and risk-free rate puzzles. This has sometimes been objected to on intuitive grounds. Since we can allow for strongly positive rates of time preference at short horizons, such objections have less force in our context.

5.1. A Square-Root Endowment Process

We consider an economy in which endowments evolve according to:

$$\ln(e_t) = \ln(e_0) + \eta t + \lambda \int_0^t (\mu - x_s)\, ds + \theta \int_0^t \sqrt{x_s}\, dB_s, \tag{47}$$

$$x_t = x_0 + \kappa \int_0^t (\mu - x_s)\, ds + \sigma \int_0^t \sqrt{x_s}\, dB_s \tag{48}$$

where κ, μ and σ are positive. The $\{e_t\}_{t\geq 0}$ process is a special case of one considered by Heston (1993). The $\{x_t\}_{t\geq 0}$ process is the Feller (1951) square-root process used by Cox, Ingersoll and Ross (1985) and others. This process is stationary when x_0 has the right initial distribution. Its stationary density is proportional to $x^{\nu-1}e^{-\omega x}$, where $\nu = 2\kappa\mu/\sigma^2$ and $\omega = 2\kappa/\sigma^2$. The mean of x_t is μ and its variance is ν/ω^2. For small ν, the stationary density is strongly skewed to the right. De-trended log endowments are stationary if and only if:

$$\frac{\lambda}{\kappa} = \frac{\theta}{\sigma}.$$

If this condition is not met, then log endowments are difference-stationary. One can then interpret $(\lambda/\kappa)x_t$ as a stationary deviation from the stochastic trend:

$$\eta t + \left(\frac{\theta}{\sigma} - \frac{\lambda}{\kappa}\right) \sigma \int_0^t \sqrt{x_s}\, dB_s.$$

In our calibrations we shall take $\lambda < 0$. If x_t is large, then the level of endowments is far below its stochastic trend, and the evolution of both the stochastic trend as well as the gap between trend and actual endowments are more uncertain than usual. Since the stationary density of $\{x_t\}_{t\geq 0}$ is skewed to the right, the endowment process spends relatively little time below its trend.

Note that one can normalize the $\{x_t\}_{t\geq 0}$ process and at the same time re-scale λ and θ without changing the properties of the endowment growth process. If x_t is not observable, then there is no loss of generality in setting $\omega = 1$. The shape parameter of the stationary density of $\{x_t\}_{t\geq 0}$ is then simply $\nu = \mu$.

Expected Utility Growth We shall use the continuous-time approximations of Section 4.2 to obtain interest rates and risk premia. These variables depend on the functions $\Gamma(\cdot)$ and $\Phi(\cdot)$. Both these functions can be expressed as integrals over time of appropriately discounted expected utility growth. Since log-endowment growth is a stationary Markov process, we can express expected utility growth over a period t as a function of the current state x:

$$G(x,t) = E\left[\left(\frac{e_t}{e_0}\right)^{1-\gamma} \middle| x_0 = x\right].$$

A convenient feature of our specification of endowment growth is that $G(x,t)$ can be computed analytically. In Appendix D we show that under certain restrictions on the parameters:

$$G(x,t)=\left(\frac{1+Q}{1+Qe^{-R\kappa t}}\right)^{\nu}\exp\left(-Pt+\left(\frac{1}{2}(R+S)-\left(\frac{R}{1+Qe^{-R\kappa t}}\right)\right)\omega x\right) \quad (49)$$

for all $x, t \geq 0$, where:

$$\begin{aligned} P &= (\gamma-1)\eta+\kappa\nu\left(\left(\frac{\gamma-1}{\omega}\right)\left(\frac{\lambda}{\kappa}\right)+\frac{1}{2}(R-S)\right), \\ Q &= \frac{R-S}{R+S}, \\ R &= \sqrt{1+4\left(\frac{\gamma-1}{\omega}\right)\left(\frac{\theta}{\sigma}-\frac{\lambda}{\kappa}\right)}, \\ S &= 1+2\left(\frac{\gamma-1}{\omega}\right)\left(\frac{\theta}{\sigma}\right). \end{aligned}$$

From hereon, we shall assume that consumers are more risk averse than when utility is logarithmic. That is, $\gamma > 1$. The parameter restrictions under which (49) is valid then correspond to requiring that $p = (\gamma-1)\lambda/(\omega\kappa)$ and $q = (\gamma-1)\theta/(\omega\sigma)$ lie below the solid curve in the Figure 1 (see Equation (61) in Appendix D; essentially, Q must be no smaller than -1.) If this condition is not satisfied then $G(x,t)$ diverges to ∞ at some finite t. If $G(x,t)$ is finite for all t, then both $\Gamma(x)$ and $\Phi(x)$ are finite if and only if $P+\rho > 0$. If log-endowment growth is trend-stationary, then $P = (\gamma-1)\eta$ and $P+\rho > 0$ amounts to the familiar condition $\rho > (1-\gamma)\eta$. If log endowments are difference stationary, then $P < (\gamma-1)\eta$ and larger values of ρ are required to ensure utility is finite.

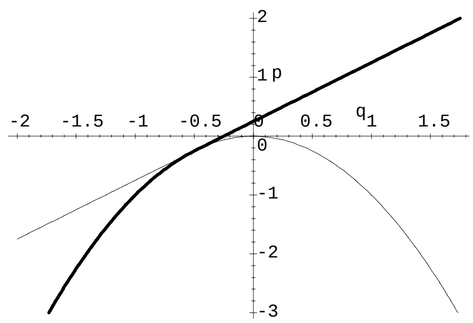

Figure 1: Feasible values of $p = \frac{(\gamma-1)\lambda}{\omega\kappa}$ and $q = \frac{(\gamma-1)\theta}{\omega\sigma}$.

5.2. Interest Rates and Risk Premia

The expressions for the risk-free rate (36), the market price of risk (37), and for the risk premium on aggregate wealth (38) specialize to:

$$r(x) = \frac{\Phi(x)}{\Gamma(x)}+\gamma\left(\eta-\frac{1}{2}\mu\gamma\theta^2-\left(\lambda+\frac{1}{2}\gamma\theta^2\right)(x-\mu)\right),$$

$$\begin{aligned}\sigma_\pi(x) &= \gamma\theta\sqrt{x}, \\ R(x) &= \gamma\theta\left(\theta+\frac{\sigma D\Gamma(x)}{\Gamma(x)}\right)x\end{aligned} \tag{50}$$

for all $x \geq 0$.

The dependence on $\delta(\cdot)$ of the risk-free rate and the risk premium on aggregate wealth is not difficult to characterize in the case of quasi-hyperbolic preferences. As shown in (40), quasi-hyperbolic discount factors imply that $\Gamma(x)$ is proportional to δ and that $\Phi(x)/\Gamma(x) = \rho + (1-\delta)/\Gamma(x)$. Therefore $\Phi(x)/\Gamma(x)$ is increasing in ρ as long as $\delta \leq 1$, and decreasing in δ. Consider making consumers more time-inconsistent by lowering δ while at the same time also lowering ρ to keep the mean of the risk-free rate unchanged. Since $\Gamma(x)$ is proportional to δ, lowering δ does not affect $D\Gamma(x)/\Gamma(x)$. But lowering ρ does.

To see how, let $f(t|x)$ be the density $\delta(t)G(x,t)/\Gamma(x)$ and observe that the derivative of $\ln(G(x,t))$ with respect to x is monotone in t and of one sign. By taking a derivative with respect to ρ of $D\Gamma(x)/\Gamma(x) = \int_0^\infty [\partial \ln(G(x,t))/\partial x]\, f(t|x)dt$ one can verify that this ensures that:[18]

$$\frac{\partial}{\partial\rho}\left|\frac{D\Gamma(x)}{\Gamma(x)}\right| < 0. \tag{51}$$

That is, the log consumption-wealth ratio becomes more volatile as the long-run rate of time preference ρ declines.

The risk premium on aggregate wealth is driven both by the volatility of the consumption-wealth ratio, and by the correlation of this ratio with endowment growth. One may verify that $(R-S)D\Gamma(x) < 0$ unless R happens to be equal to S ($R > S$ corresponds to the area under the parabola in Figure 1.) The expression for the risk premium on aggregate wealth then implies that:

$$\frac{\partial R(x)}{\partial\rho} < 0 \text{ for all } x \Leftrightarrow \theta(R-S) < 0.$$

If $\theta(R-S) < 0$, then θ and $D\Gamma(x)$ are of the same sign, and thus innovations in aggregate endowments and in the wealth-consumption ratio are positively correlated. This implies that the return on aggregate wealth and aggregate endowments are always positively correlated. This ensures that the risk premium on aggregate wealth is always positive and increasing in the volatility of aggregate wealth. It follows that making consumers more time-inconsistent by lowering δ and ρ raises the instantaneous risk premium on aggregate wealth if $\theta(R-S) < 0$.

The case of hyperbolic discount factors is a bit more complicated. Now $\Phi(x)/\Gamma(x)$ satisfies (39) and one can verify that $\Phi(x)/\Gamma(x)$ is increasing in ρ and ξ, and decreasing in ζ. By lowering ρ and raising ξ and ζ one can make the subjective discount factors "more time-inconsistent" (discount nearby utilities at higher rates and distant utilities at lower rates) while keeping the mean of the risk-free rate the same. For

[18] Note that $f(t|x)$ is a density for any ρ, and that $D\Gamma(x)/\Gamma(x)$ depends on ρ only via $f(t|x)$. The derivative of $D\Gamma(x)/\Gamma(x)$ with respect to ρ can then be written as a covariance (interpreting t as a "random variable" with density $f(t|x)$) between $\partial \ln(\delta(t))/\partial\rho = -t$ and $\partial \ln(G(x,t))/\partial x$.

the same reasons as in the case of quasi-hyperbolic discount factors, (51) holds, and thus lowering ρ again makes wealth more volatile. But now $D\Gamma(x)/\Gamma(x)$ also changes with ξ and ζ. The numerical results in the next section illustrate how increases in ξ and ζ can accentuate the increase in volatility of the consumption-wealth ratio that is generated by lowering ρ.

5.3. Quantitative Results[19]

We conduct two sets of experiments, loosely motivated by US data on aggregate consumption, dividends, and asset returns. As an imperfect proxy for the return on aggregate wealth we shall use the return on the value-weighted index of the New York Stock Exchange (NYSE.) Table 1 below provides a guide for specifying the endowment process. It shows means and standard deviations for annual US per capita consumption and dividend growth for two sample periods. Mean growth rates of consumption and dividends vary within a relatively narrow range, but the variances differ significantly. To get some idea of the quantitative implications of our model we shall take $\eta = .0175$ and let the standard deviation of endowments be 3% per annum in Experiment I and 6% per annum in Experiment II. This is in the mid-range of the statistics reported in Table 1.

Table 1: Endowments

		20^{th} Century	Postwar
Consumption:	Mean	1.77	1.89
	St. dev.	3.26	1.22
Dividends:	Mean	1.49	2.23
	St. dev.	14.2	6.00

(Annual growth rates in percent.)

We shall adopt the normalization $\omega = 1$ and take $\kappa = .1$ and $\nu = 1.5$. The parameter κ governs the persistence of the $\{x_t\}_{t\geq 0}$ process and, indirectly, that of interest rates and consumption-wealth ratios. Data on short-term interest rates and NYSE price-dividend ratios suggest that both should be highly persistent. Thus κ should be small. A small value of κ also implies that endowment growth exhibits only low levels of serial correlation, as seems to be the case for US per capita consumption. A value of 1.5 for ν makes the stationary distribution of $\{x_t\}_{t\geq 0}$ highly skewed (see the densities plotted in the background of Figures 4-6), and this may capture business cycle asymmetries of the type documented by Hamilton (1989).

To complete the description of the endowment process, we need to determine λ and θ. In Experiment I we take log endowments to be trend-stationary. This restriction suffices to determine λ and θ. In Experiment II we take the risk-free rate to be constant under geometric discounting. For a given value of γ this also identifies

[19] This section is preliminary and only intended to illustrate some possibilities. A more complete empirical investigation is in progress. All statistics reported in this section are gathered from Campbell (1999), and Campbell and Cochrane (1999). These papers contain the original sources of the data and the precise sample periods used.

λ and θ. In both experiments, we shall take $\gamma = 7.5$. As we shall see, this yields mean Sharpe ratios $E[\sigma_\pi(x_t)]$ that are in the right range for US stock returns. The resulting estimates of λ and θ are negative in both experiments, so that "recessions" (periods of below-trend consumption) are relatively short and uncertain. Moreover, $R > S$ for these estimates so that $\Gamma(x)$ is a decreasing function of the state x. The parameter restrictions imposed so far are summarized in Table 2.

Table 2: Parameters

Experiment:	I	II
γ	7.5	7.5
κ	.10	.10
$\nu = 2\kappa\mu/\sigma^2$	1.5	1.5
$\omega = 2\kappa/\sigma^2$	1.0	1.0
$\frac{\lambda}{\kappa} - \frac{\theta}{\sigma}$	0	—
$\lambda + \frac{1}{2}\gamma\theta^2$	—	0
$E\left[\ln\left(e_{t+1}/e_t\right)\right]$	.0175	.0175
$std\left[\ln\left(e_{t+1}/e_t\right)\right]$	.0300	.0600

Given this specification of endowments and of the period utility function, we shall now describe the implications of various assumptions about subjective rates of time preference. We choose the parameters of $\delta(\cdot)$ so that the mean of the instantaneous risk-free rate is equal to 2.5% per annum in Experiment I and 1% per annum in Experiment II. These values are in the range of the mean for the century-long sample (2.92%) and the postwar sample (.94%) reported in Campbell and Cochrane (1999).

We shall be interested to see what type of subjective discount factors can generate a risk premium on aggregate wealth that is comparable to the equity premium in US data. As one measure of the risk premium on aggregate wealth we report $E[R(x_t)]$. This is the mean of the instantaneous expected excess return on wealth. Although analytically convenient, this is not a direct measure of a discrete-time return. For comparison with data, we also report the mean of the annual log return on aggregate wealth. One can show that the mean of the log return on aggregate wealth over a period τ is given by $(\eta + E\left[\Gamma_t^{-1}\right])\tau$, and the tables below report values for $\tau = 1$. The concavity of the logarithm implies that this return is lower than might be expected from $E[R(x_t)]$ and the level of the instantaneous risk-free rate.

Experiment I Table 3 reports means and standard deviations of returns for various specifications of the subjective discount function. For the geometric discount function, the parameter ρ is determined so that the mean of the risk-free rate is equal to 2.5%. The results for this parameterization are reported in column G of Table 3. The mean risk premium on aggregate wealth implied by these parameters is 3.55% per annum, somewhat below the historical average for the period 1871-1993. Note that ρ has to be negative in order to match the relatively low real interest rates found in the data.

The columns labelled H in Table 3 show results for different hyperbolic discount factors.[20] In each case, ρ is set equal to $-.11$, which is less than 10^{-2} away from

[20] The integrals over t that determine $\Gamma(x)$ and $\Phi(x)$ are calculated numerically using the NAG routine `D01AJF` after the change of variables $t' = t/(1+t)$.

P. Choosing ρ close to the boundary of the region that ensures that wealth is finite maximizes the volatility of the wealth-consumption ratio. For given values of ζ, the values of ξ reported in Table 3 were chosen so that the mean of the risk-free rate is equal to 2.5%.The results in Table 3 show that the risk premium on aggregate wealth increases noticeably as consumers discount more rapidly over short horizons (that is, as ζ gets larger.) For large values of ζ, this risk premium is about 1.1% higher than when consumers discount geometrically.

Table 3: Experiment I

Case:	G	H	H	H	QH	US-XX	US-PW
δ	1	–	–	–	.200	–	–
ρ	$-.080$	$-.110$	$-.110$	$-.110$	$-.107$	–	–
ζ	–	1	10	10^6	–	–	–
ξ	–	.487	3.19	$1.19^{\times 10^5}$	–	–	–
$\lvert E[\sigma_\pi(x_t)]\rvert$	.21	.21	.21	.21	.21	.22	.43
$E[r(x_t)]$	2.50	2.50	2.50	2.50	2.50	1.96	.79
$std[r(x_t)]$	2.97	3.83	4.19	4.66	4.54	8.92	1.76
$E\left[R(x_t)\right]$	3.55	4.18	4.42	4.71	4.64	–	–
$\eta + E\left[1/\Gamma(x_t)\right]$	4.92	5.07	5.11	5.14	5.14	6.82	7.63
$E[\Gamma(x_t)]$	34.4	34.3	34.6	35.1	34.9	21.1	24.7

Returns are real, annual, and measured in percent. The column US-XX is based on 20th century data for the US and US-PW is based on postwar data. Risk-free rate statistics are from Table 2 of Campbell (1999). Other data are from Table 2 of Campbell and Cochrane (1999).

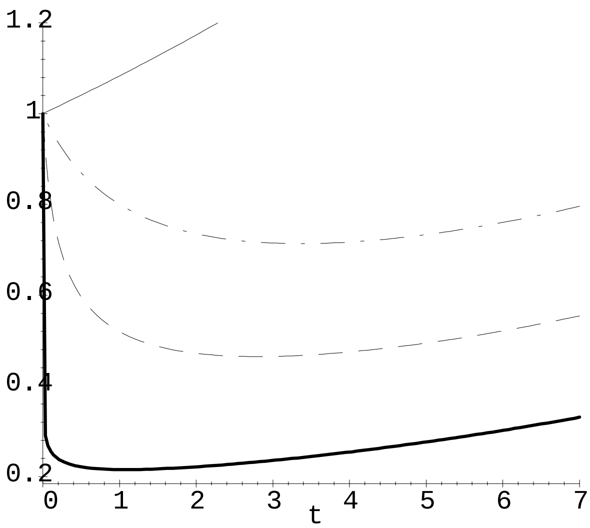

Figure 2: $\delta(t) = \exp(-\rho t)(1+\zeta t)^{-\xi/\zeta}$: $\rho = -.11$ and $\zeta = 1, \xi = .48$ (dash-dot), $\zeta = 10, \xi = 3.19$ (dash), $\zeta = 10^6, \xi = 1.19 \times 10^5$ (solid); the thin curve represents $\delta(t) = \exp(.08t)$.

Figure 2 illustrates the shape of the hyperbolic part of several of the discount factors used in Table 3. As the figure suggests, the generalized hyperbolic discount factor corresponding to $\zeta = 10^6$ should be well approximated by a quasi-hyperbolic discount factor with δ around .2. The results for this case are reported in column QH of Table 3. Again, the parameter ρ was chosen to ensure that the model implies a mean risk-free rate of 2.5% per annum. The risk premium on aggregate wealth reported in the third H column and the QH column of Table 3 are quite similar, as expected.

Experiment II The results for the experiment with volatile endowment growth are reported in Table 4. The non-geometric examples are constructed by setting $\rho = -.10$. The corresponding discount factors are illustrated in Figure 3 below. The impact on the risk premium on aggregate wealth of making consumers more time-inconsistent is much more dramatic in this experiment than in Experiment I. The values of ξ and ζ required to obtain a significant impact on returns are much smaller than in Table 3. One reason is that the change in ρ is larger: from $-.08$ to $-.11$ in Table 3 and from $-.015$ to $-.10$ in Table 4. In turn, this is a consequence of the fact that the precautionary savings term $-\mu\gamma^2\theta^2/2$ in the expression for the risk-free rate is larger in this experiment than in Experiment I. A second reason is the fact that in this experiment the Sharpe ratio is twice as large as in Experiment I. This doubles the impact that changes in the volatility of aggregate wealth have on the expected return on aggregate wealth.

Table 4: Experiment II

Case:	G	H	H	H	QH	US-XX	US-PW
δ	1	–	–	–	.200	–	–
ρ	$-.015$	$-.100$	$-.100$	$-.100$	$-.071$	–	–
ζ	–	.100	.250	.500	–	–	–
ξ	–	.195	.326	.514	–	–	–
$\lvert E[\sigma_\pi(x_t)]\rvert$	.42	.42	.42	.42	.42	.22	.43
$E[r(x_t)]$	1.00	1.00	1.00	1.00	1.00	1.96	.79
$std[r(x_t)]$	0.00	1.00	1.62	2.13	2.62	8.92	1.76
$E\left[R(x_t)\right]$	5.70	7.04	7.87	8.54	9.43	–	–
$\eta + E\left[1/\Gamma(x_t)\right]$	6.60	7.45	7.94	8.31	8.80	6. 82	7.63
$E[\Gamma(x_t)]$	19.8	17.4	17.6	17.0	16.3	21.1	24.7

Returns are real, annual, and measured in percent. The column US-XX is based on 20^{th} century data for the US and US-PW is based on postwar data. Risk-free rate statistics are from Table 2 of Campbell (1999). Other data are from Table 2 of Campbell and Cochrane (1999).

The results in Tables 3 and 4 illustrate the fact that hyperbolic discount factors can lead to more volatile wealth and to higher risk premia on aggregate wealth. If consumers discount near-future utilities at high rates (as they do when discount factors are strongly hyperbolic; note that $\rho+\xi < 0$ in all cases in Tables 3 and 4) then they must also discount distant future utilities at a low rate, or else interest rates would be much higher than found in the data. This means that ρ must be lower (more negative) than would be the case if consumers discounted utility geometrically. For our specification of the endowment process, this makes $D\Gamma(x)/\Gamma(x)$ more negative. As can be seen from (50) this should increase the risk premium on aggregate wealth ($\theta < 0$.) The three hyperbolic cases in Tables 3 and 4 show that the same effect also operates for a given value of ρ, as one makes the discount function more hyperbolic while continuing to match the historical average of the risk-free rate.

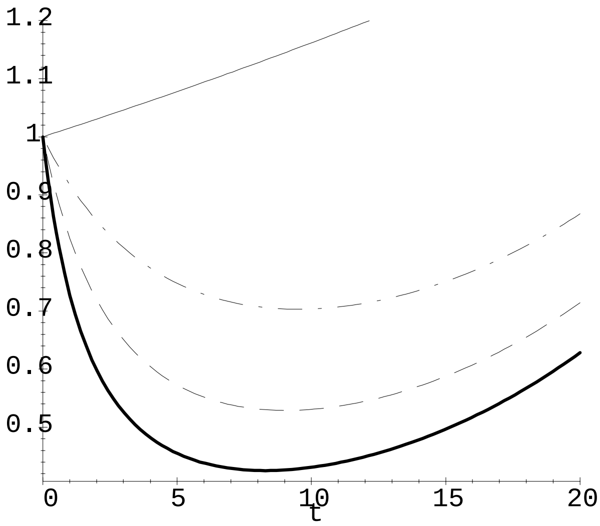

Figure 3: $\delta(t) = \exp(-\rho t)(1+\zeta t)^{-\xi/\zeta}$: $\rho = -.10$ and $\zeta = .1, \xi = .195$ (dash-dot), $\zeta = .25, \xi = .326$ (dash), $\zeta = .5, \xi = .514$ (solid); the thin line represents $\delta(t) = \exp(.015t)$.

These effects are made more explicit in Figures 4-6. These figures report $\log(\Gamma(x))$, $R(x)$ and $r(x)$ for the geometric discount function reported in Table 4, and for the hyperbolic case of $\zeta = .5$. The stationary density of $\{x_t\}_{t\geq 0}$ is plotted in the background. The log wealth-consumption ratio is not far from linear and more steeply downward sloping for the hyperbolic case.

In both Table 3 and Table 4 the risk-free rate becomes more volatile as consumers become more time-inconsistent. This is because $\Phi(x)/\Gamma(x)$ becomes more volatile and $\Phi(x)/\Gamma(x)$ has a non-negative correlation with $\left(\lambda + \gamma\theta^2/2\right)x$. This need not always be the case. One can show that these variables are negatively correlated when p and q in Figure 1 lie in between the parabolas $p + q^2 = 0$ and $p + q^2\gamma/(\gamma - 1) = 0$. In Experiment II, $\lambda + \gamma\theta^2/2$ is zero and p and q lie on the latter parabola. Thus, for values of λ/κ that are a bit less negative than is the case in Experiment II, time-inconsistency reduces (at least initially) the volatility of the risk-free rate, while at the same time increasing the volatility of the wealth-consumption ratio.

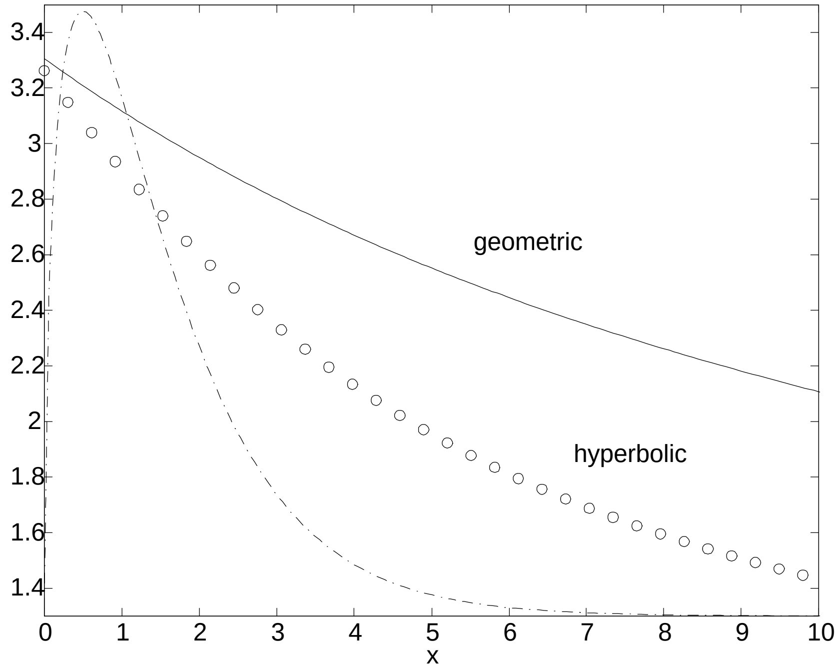

Figure 4: Log Wealth-Consumption Ratios (Experiment II, $\zeta = .5, \xi = .514$.)

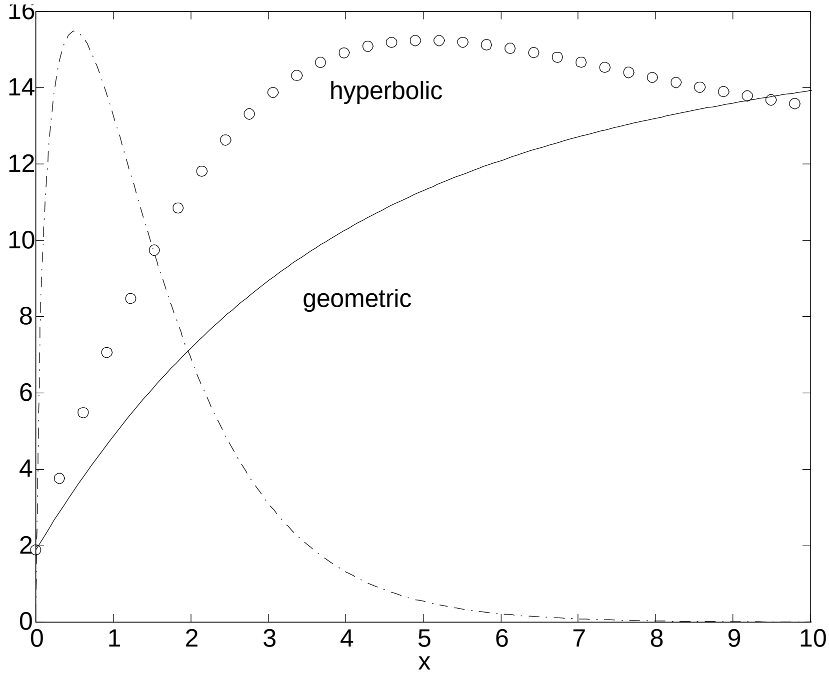

Figure 5: Risk Premia on Aggregate Wealth (Experiment II, $\zeta = .5, \xi = .514$.)

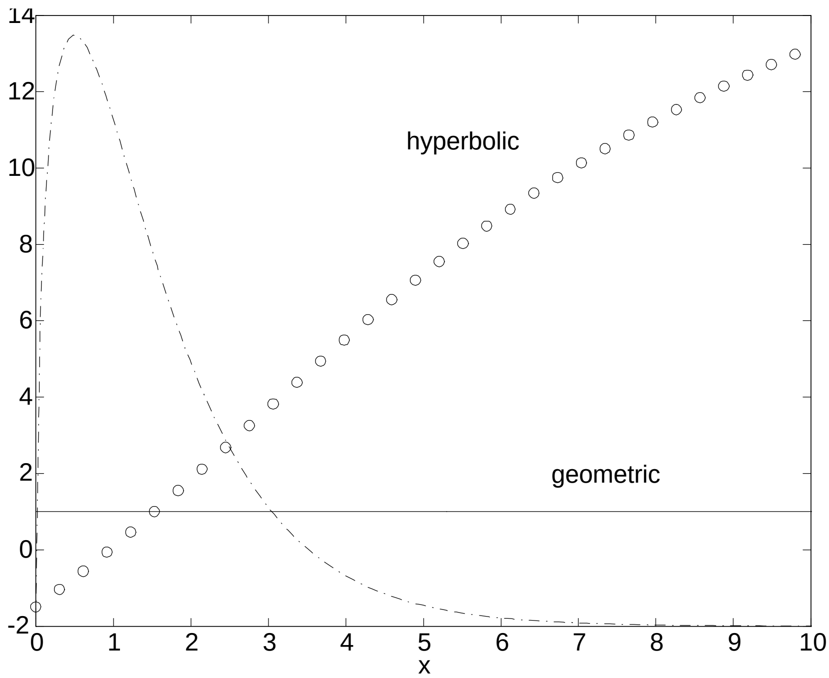

Figure 6: Risk-free Rates (Experiment II, $\zeta = .5, \xi = .514$.)

6. Concluding Remarks

In an infinite-horizon exchange economy in which utility is expected to grow at a constant rate, and in which consumers cannot commit to future choices, price and wealth data can always be interpreted as resulting from the optimal choices of consumers with geometric discount factors. While the fact that expected utility must be finite allows one to rule out some subjective discount functions, there will typically be a large set of subjective discount functions consistent with the same data. If utility is logarithmic, this observational equivalence result applies even for essentially arbitrary endowment processes.

Sources of evidence that may allow one to make inferences about subjective discount factors include situations in which consumers can make only partially reversible commitments, as in Laibson (1997), or in which consumers face binding borrowing constraints, as in the buffer-stock savings model analyzed by Harris and Laibson (1999). In this paper, we have abstracted from commitment devices or market frictions. Instead, we have examined the extent to which serial dependence in utility alone allows one to distinguish among various types of subjective discount factors.

In the presence of serial dependence in endowment growth, implicit discount rates are no longer constant when consumers have hyperbolic or quasi-hyperbolic discount factors. As a result, equilibrium prices will differ from those in an economy in which consumers discount utility geometrically.

We have explored the quantitative implications of hyperbolic and quasi-hyperbolic discounting in an economy that exhibits persistence in the volatility of endowment growth, and possibly in mean endowment growth. We show that when consumers have hyperbolic or quasi-hyperbolic discount factors, aggregate wealth can become significantly more volatile and the risk premium on aggregate wealth can become significantly higher than when consumers discount geometrically.

Standard models tend to generate stock prices that are less volatile than appears to be consistent with the data (LeRoy and Porter (1981), Shiller (1981)). The risk premium on stocks implied by these models also appears to be too low (Mehra and Prescott (1985)). Given that stocks are a non-negligible fraction of aggregate wealth, our results suggest that hyperbolic discounting can contribute to an understanding of these phenomena.

For the period utility functions we consider, these results hinge on the presence of serial dependence in endowment growth. Serial dependence in utility can also be generated by habit persistence and the consumption of durable goods. How these aspects of consumer preferences interact with non-geometric discounting is the subject of ongoing research.

An important feature of the continuous-time approximation we present in this paper is the fact that the instantaneous market price of risk is not affected by how consumers discount utility. Subjective rates of time preference only affect the level and dynamics of the instantaneous risk-free rate. For the endowment process we consider in the quantitative part of this paper it is easy to calculate analytically the term structure of interest rates when consumers are time-consistent. The impact of time-inconsistent preferences on the term structure of interest rates needs to be investigated further.

A Ruling out Bubbles

To prove Proposition 1, we proceed in two steps. We first give a precise statement and proof of the first-order condition (12). Next, we check that Conditions A and B imply the existence of an equilibrium in which there is no-bubble on the long-lived asset.

Proof of the First-Order Condition (12). Given current wealth $w_t > 0$, the date-t consumer's decision problem can be written as:

$$\max_{(c_t, w_{t+1}) \in \mathbb{R}_+ \times L^+_{t+1}} \left\{ u(c_t) + E_t\left[A_{t+1} u(w_{t+1})\right] : c_t + E_t\left[\frac{\pi_{t+1}}{\pi_t} w_{t+1}\right] \leq w_t \right\}, \tag{52}$$

where L^+_{t+1} is the set of nonnegative $\mathcal{F}_{t+1}$–measurable random variables, $A_{t+1} \in L^+_{t+1}$ is a P-almost surely positive random variable to be determined in equilibrium, and $\pi_{t+1}/\pi_t \in L^+_{t+1}$ represents the relative price of next-period consumption.

Lemma 1 *Suppose that (52) has a solution (c^*_t, w^*_{t+1}) at which expected discounted utility is finite. Then (12) holds P-almost surely.*

Proof. Since A_{t+1} is positive P-almost surely, the fact that the marginal utility $Du(\cdot)$ at zero is infinite implies that c^*_t and w^*_{t+1} are both positive. Furthermore, given that $u(\cdot)$ is strictly increasing, we know that the budget constraint must be binding at the optimum. Perturb w^*_{t+1} to $w_{t+1} = \alpha w^*_{t+1}$ for some $\alpha > 0$ small enough and take $c_t = w_t - E_t[(\pi_{t+1}/\pi_t) w_{t+1}]$. By homotheticity of $u(\cdot)$, the mapping $\alpha \mapsto E_t[A_{t+1} u(\alpha w^*_{t+1})]$ is differentiable on $\mathbb{R}_{++}$. The optimality of (c^*_t, w^*_{t+1}) then implies:

$$\frac{\partial}{\partial \alpha} \left(u\left(w_t - \alpha E_t\left[\frac{\pi_{t+1}}{\pi_t} w^*_{t+1}\right]\right) + E_t[A_{t+1} u(\alpha w^*_{t+1})] \right) \Bigg|_{\alpha=1} = 0,$$

or, equivalently:

$$E_t\left[\left(-Du(c^*_t)\frac{\pi_{t+1}}{\pi_t} + A_{t+1} Du(w^*_{t+1})\right) w^*_{t+1}\right] = 0. \tag{53}$$

Alternatively, consider perturbing wealth to $w_{t+1} = (1 + \alpha \iota_B) w^*_{t+1}$ for some $\alpha \geq 0$ and $B \in \mathcal{F}_{t+1}$. For α small enough, $w_t - E_t[(\pi_{t+1}/\pi_t) w_{t+1}]$ is positive. Furthermore, since $u(\cdot)$ is increasing and concave, one has:

$$\left| \frac{A_{t+1}(u((1+\alpha \iota_B) w^*_{t+1}) - u(w^*_{t+1}))}{\alpha} \right| \leq A_{t+1} \iota_B w^*_{t+1} Du(w^*_{t+1}) \leq A_{t+1} w^*_{t+1} Du(w^*_{t+1}).$$

Since $w^*_{t+1} Du(w^*_{t+1}) = (1-\gamma) u(w^*_{t+1})$ and expected utility is finite, the right-hand side of this inequality is integrable. The differentiability of $u(\cdot)$ on $\mathbb{R}_{++}$ together with the dominated convergence theorem therefore imply that the right-derivative of the mapping $\alpha \mapsto E_t[A_{t+1} u((1+\alpha \iota_B) w^*_{t+1})]$ at $\alpha = 0$ is well-defined and given by $E_t[A_{t+1} \iota_B w^*_{t+1} u(w^*_{t+1})]$. The optimality of (c^*_t, w^*_{t+1}) then requires that:

$$E_t\left[\left(-Du(c^*_t)\frac{\pi_{t+1}}{\pi_t} + A_{t+1} Du(w^*_{t+1})\right) w^*_{t+1} \iota_B\right] \leq 0. \tag{54}$$

Note that because B is an arbitrary element of $\mathcal{F}_{t+1}$, equations (53) and (54) imply that $\left(-Du(c_t^*)\,\pi_{t+1}/\pi_t + A_{t+1}Du(w_{t+1}^*)\right) w_{t+1}^* = 0$ P-almost surely. Since w_{t+1}^* is always positive, the result follows. ■

We are now ready to complete the proof of Proposition 1.

Proof of Proposition 1. Conditions A and B ensure that the quantities Γ_t, Δ_t and A_{t+1} are well-defined and P-almost surely finite for any $t \geq 0$. By construction, the consumption and wealth choices $(e_t, e_{t+1}\Gamma_{t+1})$ satisfy the first-order condition (12) at market-clearing prices $\pi_{t+1}/\pi_t = (\Delta_{t+1}/\Gamma_{t+1})(e_{t+1}/e_t)^{-\gamma}$ for any $t \geq 0$. It then follows from the concavity of the objective function in (52) that $(e_t, e_{t+1}\Gamma_{t+1})$ is optimal from the perspective of the date-t consumer, and since Γ_t is finite, that the value of problem (52) is finite. To conclude the proof, we need only to check that Conditions A and B imply (24). Consider first Condition A. Using (17) one can write:

$$\begin{aligned}
E_0\left[\prod_{t=1}^{T}\frac{\Delta_t}{\Gamma_{t-1}}\left(\frac{e_t}{e_{t-1}}\right)^{1-\gamma}\right] &= E_0\left[\prod_{t=1}^{T-1}\frac{\Delta_t}{\Gamma_{t-1}}\left(\frac{e_t}{e_{t-1}}\right)^{1-\gamma} E_{T-1}\left[\frac{\Delta_T}{\Gamma_{T-1}}\left(\frac{e_T}{e_{T-1}}\right)^{1-\gamma}\right]\right] \\
&= E_0\left[\prod_{t=1}^{T-1}\frac{\Delta_t}{\Gamma_{t-1}}\left(\frac{e_t}{e_{t-1}}\right)^{1-\gamma}\left(1-\frac{1}{\Gamma_{T-1}}\right)\right] \\
&\leq E_0\left[\prod_{t=1}^{T-1}\frac{\Delta_t}{\Gamma_{t-1}}\left(\frac{e_t}{e_{t-1}}\right)^{1-\gamma}\right](1-\kappa) \\
&\ \ \vdots \\
&\leq (1-\kappa)^T,
\end{aligned}$$

from which it is immediate that (24) is satisfied. Under Condition B we can write $\Delta_t/\Gamma_t \leq \beta$ since Δ_t/Γ_t is an average of $\delta_{n+1}/\delta_n \leq \beta$. Also, $\delta_n \leq \beta^n$ and therefore:

$$\Gamma_t \leq E_t\left[\sum_{n=0}^{\infty}\beta^n\left(\frac{e_{t+n}}{e_t}\right)^{1-\gamma}\right].$$

This yields:

$$E_0\left[\left(\prod_{t=1}^{T}\frac{\Delta_t}{\Gamma_t}\left(\frac{e_t}{e_{t-1}}\right)^{1-\gamma}\right)\Gamma_T\right] \leq E_0\left[\beta^T\left(\frac{e_T}{e_0}\right)^{1-\gamma}\Gamma_T\right] \leq E_0\left[\sum_{n=T}^{\infty}\beta^n\left(\frac{e_n}{e_0}\right)^{1-\gamma}\right].$$

(25) implies that the RHS of this inequality converges to zero, hence the result. ■

B Continuous-Time Approximation

In the following, $\{x_t\}_{t\geq 0}$ denotes a continuous-time Markov process defined on the same probability space $(\Omega, \mathcal{F}, P)$ as the endowment process $\{e_t\}_{t\geq 0}$, and taking its values in some state space $X \subset \mathbb{R}^N$. Endowments $\{e_t\}_{t\geq 0}$ are positive and we write:

$$G(x,t) = E\left[\left(\frac{e_t}{e_0}\right)^{1-\gamma}\,\middle|\, x_0 = x\right]$$

for any $x \in X$ and $t \geq 0$. Unless stated otherwise, the following assumptions will be maintained in the remainder of this Appendix.

Assumption 1 *The subjective discount function $\delta : \mathbb{R}_+ \to [0,1]$ is non-increasing, left-continuous and positive on a set of positive Lebesgue measure. Moreover, $\delta(\cdot)$ is integrable over $\mathbb{R}_+$ and $\delta(0) = 1$.*

Assumption 2 *There exists a function $M : X \to \mathbb{R}_+$ that is bounded on compact subsets of X and such that $G(x,t) \leq M(x)$ for all $x \in X$ and $t \geq 0$.*

Assumption 3 *The function $G(x,\cdot)$ is continuous for every $x \in X$. Moreover, the family of functions $\{G(\cdot,t)\}_{t\in\mathbb{R}_+}$ is equicontinuous at any $x \in X$.*

Since $\delta(\cdot)$ is non-increasing and integrable over $\mathbb{R}_+$, $\lim_{t\to\infty} \delta(t) = 0$. Thus $\delta(\cdot)$ induces a unique probability measure μ_δ on the Borel sets of $\mathbb{R}_+$ such that $\mu_\delta([s,t)) = \delta(s) - \delta(t)$ for any $t > s \geq 0$ (Lang (1993, Proposition X.1.8)). Note that by Fubini's theorem, we may rewrite (32) and (33) as:

$$\begin{aligned} \Gamma(x) &= \int_0^\infty \delta(t)\, G(x,t)\, dt, \\ \Phi(x) &= \int_{[0,\infty)} G(x,t)\, d\mu_\delta(t) \end{aligned}$$

for any $x \in X$. Assumptions 1-3 ensure that these functions are well defined and finite. Furthermore, $\Gamma(x)$ is positive for all x in X since $G(x,\cdot) > 0$ and $\delta(\cdot) > 0$ on some set of positive Lebesgue measure. Assumption 3 ensures that $\Gamma(\cdot)$ is continuous on X. It follows that $\Gamma(\cdot)$ is bounded away from zero on compact subsets of X. Also, $\Phi(\cdot)$ is positive and bounded above by $M(\cdot)$. For any $\tau > 0$ and $x \in X$, define:

$$\begin{aligned} \Gamma(x,\tau) &= \sum_{n=0}^{\infty} \delta(n\tau)\, G(x,n\tau), \\ \Delta(x,\tau) &= \sum_{n=0}^{\infty} \delta((n+1)\tau)\, G(x,n\tau), \end{aligned}$$

and let $\Phi(x,\tau) = \Gamma(x,\tau) - \Delta(x,\tau)$. Assumptions 1 and 2 ensures that these sums are finite for any $\tau > 0$ and $x \in X$, and that $\Gamma(x,\tau)$ is positive for all $\tau > 0$ small enough. It follows from the monotone convergence theorem that $\Gamma(x_t,\tau) = \Gamma_t(\tau)$ and $\Delta(x_t,\tau) = \Delta_t(\tau)$ for any $t \geq 0$ and $\tau > 0$.

Lemma 2 *For any $x \in X$, $\lim_{\tau\to 0^+} \tau\,\Gamma(x,\tau) = \Gamma(x)$.*

Proof. For any $\tau > 0$, $t \geq 0$, and $x \in X$, let $G_\tau(x,t) = G(x,\tau[\![t/\tau]\!])$ and $\delta_\tau(t) = \delta(\tau[\![t/\tau]\!])$, where $[\![t/\tau]\!]$ is the integer part of t/τ. Observe that:

$$\tau\,\Gamma(x,\tau) = \sum_{n=0}^{\infty} \delta(n\tau)\, G(x,n\tau)\,\tau = \int_0^\infty \delta_\tau(t)\, G_\tau(x,t)\, dt.$$

Assumptions 1 and 2 imply that $\delta_\tau(t)\, G_\tau(x,t) \le M(x)$ for any $t \in [0,1)$ and $\tau > 0$, as well as $\delta_\tau(t)\, G_\tau(x,t) \le \delta(t-1)\, M(x)$ for all $t \ge 1$ and $\tau \in (0,1]$. Thus $\delta_\tau(\cdot)\, G_\tau(x,\cdot)$ is dominated for all $\tau \in (0,1]$ by an integrable function of t. Moreover, the continuity of $G(x,\cdot)$ implies that $\lim_{\tau\to 0^+} G_\tau(x,t) = G(x,t)$ and the left-continuity of $\delta(\cdot)$ implies that $\lim_{\tau\to 0^+} \delta_\tau(t) = \delta(t)$. The result then follows by Lebesgue's dominated convergence theorem. ■

Lemma 3 *The convergence of $\tau\,\Gamma(x,\tau)$ to $\Gamma(x)$ as $\tau \to 0^+$ is uniform on any compact subset of X.*

Proof. Assumption 3 implies that the family $\{G(\cdot,t)\}_{t\in\mathbb{R}_+}$ is uniformly equicontinuous on any compact K of X. Hence, for any $\varepsilon > 0$, there exists $\eta > 0$ such that for all $t \ge 0$ and $\tau \in (0,1]$, $|G_\tau(x,t) - G_\tau(y,t)| \le \varepsilon / \left(1 + \int_0^\infty \delta(v)\, dv\right)$ for any $x, y \in K$ such that $\|x-y\| < \eta$. It follows that for any $\tau \in (0,1]$,

$$|\tau\,\Gamma(x,t) - \tau\,\Gamma(y,t)| \le \int_0^\infty \delta_\tau(v)\, |G_\tau(x,v) - G_\tau(y,v)|\, dv \le \varepsilon$$

for any $x, y \in K$ such that $\|x-y\| < \eta$, where the second inequality follows from the fact that $\int_0^\infty \delta_\tau(v)\, dv = \sum_{n=0}^\infty \delta(n\tau)\,\tau \le \tau + \int_0^\infty \delta(v)\, dv$ as $\delta(\cdot)$ is decreasing over $\mathbb{R}_+$. This implies that the family $\{\tau\Gamma(\cdot,\tau)\}_{\tau\in(0,1]}$ is uniformly equicontinuous on K. By Ascoli's theorem (Lang (1993, Theorem III.3.1 and Corollary III.3.3)), the convergence of $\tau\,\Gamma(x,\tau)$ to $\Gamma(x)$ as $\tau \to 0^+$ is therefore uniform on K. ■

Lemma 4 *For any $x \in X$, $\lim_{\tau\to 0^+} \Phi(x,\tau) = \Phi(x)$.*

Proof. Using the notation of Lemma 2, observe that for any $x \in X$ and $\tau > 0$:

$$\Phi(x,\tau) = \sum_{n=0}^\infty G(x,n\tau)\,(\delta(n\tau) - \delta((n+1)\tau)) = \int_{[0,\infty)} G_\tau(x,t)\, d\mu_\delta(t).$$

As before, $\lim_{\tau\to 0^+} G_\tau(x,t) = G(x,t)$ and $G_\tau(x,t) \le M(x)$ for any $\tau > 0$, $t \ge 0$ and $x \in X$. Since μ_δ is a probability measure, the result follows immediately from Lebesgue's dominated convergence theorem. ■

We can now state and prove our main result.

Proposition 2 *If the Markov processes $\{x_t\}_{t\ge 0}$ and $\{e_t\}_{t\ge 0}$ have P-almost surely continuous sample paths, and if Assumptions 1-3 are satisfied, then for P-almost every $\omega \in \Omega$, $\lim_{\tau\to 0^+} \pi_t^\omega(\tau) = \exp\left(-\int_0^t \frac{\Phi(x_v^\omega)}{\Gamma(x_v^\omega)}\, dv\right)(e_t^\omega)^{-\gamma} = \pi_t^\omega$ for each $t \ge 0$.*

Proof. Consider a continuous sample path $\{x_t^\omega\}_{t\ge 0}$. At any date $t = k\tau$, $k \in \mathbb{N}$, discrete-time state prices are given by:

$$\begin{aligned}
\pi_t^\omega(\tau) &= \exp\left(\sum_{n=1}^{[\![t/\tau]\!]} \ln\left(\frac{\Delta(x_{n\tau}^\omega,\tau)}{\Gamma(x_{n\tau}^\omega,\tau)}\right)\right)(e_t^\omega)^{-\gamma} \\
&= \exp\left(\int_0^t \iota_{t,\tau}(v) Q(x_{\tau[\![v/\tau]\!]}^\omega,\tau)\, dv\right)(e_t^\omega)^{-\gamma},
\end{aligned}$$

where $\iota_{t,\tau}(v) = 1$ if $v/\tau \in [1, [\![t/\tau]\!]]$ and zero otherwise, and:

$$Q(x,\tau) = \frac{1}{\tau} \ln\left(1 - \tau\left(\frac{\Phi(x,\tau)}{\tau\,\Gamma(x,\tau)}\right)\right)$$

for all $\tau > 0$ and $x \in X$. At dates $t \neq k\tau$, $k \in \mathbb{N}$, simply set $\pi_t^\omega(\tau) = \pi^\omega_{\tau[\![t/\tau]\!]}(\tau)$. By Lemmas 2 and 4, $\lim_{\tau\to 0^+} Q(x,\tau) = -\Phi(x)/\Gamma(x)$ for any $x \in X$. Next, recall that $\lim_{\tau\to 0^+} \tau\,\Gamma(x,\tau) = \Gamma(x)$ uniformly on compact subsets K of X. Hence, since x_v^ω is a continuous function of v, $\lim_{\tau\to 0^+} \tau\,\Gamma(x^\omega_{\tau[\![v/\tau]\!]}, \tau) = \Gamma(x_v^\omega)$ uniformly on $[0,t]$. Therefore $\tau\,\Gamma(x^\omega_{\tau[\![v/\tau]\!]}, \tau)$ is bounded away from zero on $[0,t]$ for all $\tau > 0$ small enough. Because μ_δ is a probability measure and $G(x,\cdot) \in [0, M(x)]$, we know that $\Phi(x,\tau) \in [0, M(x)]$ for any $x \in X$ and $\tau > 0$ small enough. Thus $Q(x^\omega_{\tau[\![v/\tau]\!]}, \tau)$ is uniformly bounded on $[0,t]$, for all $\tau > 0$ small enough. The result then follows from Lebesgue's dominated convergence theorem and the continuity of $\{x_t^\omega, e_t^\omega\}_{t\geq 0}$ for P-almost every $\omega \in \Omega$. ■

Remark 1 In some applications, one can ensure that Assumptions 1-3 are satisfied by replacing $\delta(t)$ and $G(x,t)$ by $\delta^*(t) = \exp(-\alpha t)\,\delta(t)$ and $G^*(x,t) = \exp(\alpha t)\,G(x,t)$, for some coefficient $\alpha > 0$. The analogs of $\Gamma(x,\tau)$ and $\Delta(x,\tau)$ are then $\Gamma^*(x,\tau) = \Gamma(x,\tau)$ and $\Delta^*(x,\tau) = \exp(\alpha t)\,\Delta(x,\tau)$, and discrete-time state prices are given by:

$$\pi_t^\omega(\tau) = \exp\left(\alpha\tau[\![t/\tau]\!] + \sum_{n=1}^{[\![t/\tau]\!]} \ln\left(\frac{\Delta^*(x^\omega_{n\tau},\tau)}{\Gamma^*(x^\omega_{n\tau},\tau)}\right)\right)(e_t^\omega)^{-\gamma}$$

for any $\omega \in \Omega$ and $\tau > 0$, and at any dates $t = k\tau$, $k \in \mathbb{N}$. If δ^* and G^* satisfy Assumptions 1-3, then Proposition 2 implies that for P-almost every $\omega \in \Omega$, $\lim_{\tau\to 0^+} \pi_t^\omega(\tau) = \exp\big(\alpha t - \int_0^t \frac{\Phi(x_v^\omega)+\alpha\Gamma(x_v^\omega)}{\Gamma(x_v^\omega)}\,dv\big)(e_t^\omega)^{-\gamma} = \pi_t^\omega$ for each $t \geq 0$.

Remark 2 Proposition 1 also holds for any left-continuous and integrable subjective discount function $\delta : \mathbb{R}_+ \to \mathbb{R}_+$ that is of bounded variation and decreasing in an interval $[T,\infty)$. In that case there exist two decreasing and left-continuous functions δ^+ and δ^- such that $\delta = \delta^+ - \delta^-$ and $\delta^- = 0$ on $[T,\infty)$. These functions induce two finite measures μ_δ^+ and μ_δ^- on the Borel sets of $\mathbb{R}_+$ that satisfy $\mu_\delta^+[s,t) = \delta^+(s) - \delta^+(t)$ and $\mu_\delta^-[s,t) = \delta^-(s) - \delta^-(t)$ for any $t > s \geq 0$. All the arguments above can then be generalized in a straightforward manner, provided that the measure μ_δ that appears in the expression of $\Phi(x)$ is replaced by the signed measure $\mu_\delta^+ - \mu_\delta^-$.

C Partial Commitment

In this Appendix we will maintain the Assumptions 1-3 of Appendix B. We add the following two assumptions.

Assumption 4 *The function $G(x,\cdot)$ is continuously differentiable for each $x \in X$.*

Assumption 5 *There exists a function $\mathcal{M} : X \to \mathbb{R}_+$ that is bounded on compact subsets of X and such that $|D_2G(x,t)| \leq \mathcal{M}(x)$ for each $x \in X$ and $t > 0$.*

For any $x \in X$ and $\tau > 0$, define:

$$\begin{aligned}\widetilde{\Delta}(x,\tau) &= \int_0^\infty \delta(t+\tau)\, G(x,t)\, dt,\\ \widetilde{\Phi}(x) &= \int_{(0,\infty)} G(x,t)\, d\mu_\delta(t),\end{aligned}$$

and let $\widetilde{\Phi}(x,\tau) = \frac{1}{\tau}\Big(\Gamma(x) - \widetilde{\Delta}(x,\tau)\Big)$. Assumption 1-3 ensure that these functions are well-defined and finite. As before, $\Gamma(\cdot)$ is bounded away from zero on compact subsets of X, $\widetilde{\Phi}(\cdot,\cdot)$ is positive, and μ_δ is the probability measure induced by the discount function $\delta(\cdot)$. It follows from Fubini's theorem that $\widetilde{\Delta}(x_t,\tau) = \widetilde{\Delta}_t(\tau)$ and $\widetilde{\Phi}(x_t) = \widetilde{\Phi}_t$ for any $t \geq 0$ and $\tau > 0$.

Lemma 5 *For any $x \in X$, $\lim_{\tau\to 0^+} \widetilde{\Phi}(x,\tau) = \widetilde{\Phi}(x)$.*

Proof. For any $x \in X$ and $\tau > 0$, we can rewrite $\widetilde{\Phi}(x,\tau)$ as:

$$\widetilde{\Phi}(x,\tau) = \frac{1}{\tau}\int_0^\tau \delta(t)\, G(x,t)\, dt + \int_0^\infty \delta(t+\tau)\,\frac{G(x,t+\tau)-G(x,t)}{\tau}\, dt. \tag{55}$$

Note that $\frac{1}{\tau}\int_0^\tau |\delta(t)\, G(x,t) - \delta(0^+)\, G(x,0)|\, dt \;\leq\; \delta(0^+)\max_{t\in[0,\tau]} |\, G(x,t) - \; G(x,0)|$ since $\delta(\cdot)$ is non-increasing. It follows that $\lim_{\tau\to 0^+}\frac{1}{\tau}\int_0^\tau \delta(t)\, G(x,t)\, dt = \delta(0^+)\, G(x,0)$ by continuity of $G(x,\cdot)$. Next, by the mean-value theorem, $|G(x,t+\tau) - G(x,t)| \leq \tau \sup_{\xi\in[0,1]} |D_2G(x,t+\xi\tau)|$. Assumption 5 implies that the integrand in the second term of the right-hand side of (55) is bounded in absolute value by $\delta(t)\mathcal{M}(x)$ for all $t > 0$. Since $\delta(\cdot)$ is integrable, we thus have, by Lebesgue's dominated convergence theorem:

$$\lim_{\tau\to 0^+}\int_0^\infty \delta(t+\tau)\,\frac{G(x,t+\tau)-G(x,t)}{\tau}\, dt = \int_0^\infty \delta(t^+)\, D_2G(x,t)\, dt.$$

This in turn is equal to $\int_0^\infty \delta(t)\, D_2G(x,t)\, dt$ since $\delta(\cdot)$ is monotone and hence continuous except on a set of measure zero. We therefore have:

$$\lim_{\tau\to 0^+}\widetilde{\Phi}(x,\tau) = \delta(0^+)\, G(x,0) + \int_0^\infty \delta(t)\, D_2G(x,t)\, dt. \tag{56}$$

Let $T > 0$. By Assumptions 1 and 4, the function $\delta(\cdot)\, D_2G(x,\cdot)$ is continuous except on a set of measure zero and therefore Riemann integrable on $[0,T]$. Moreover, since $G(x,\cdot)$ is continuously differentiable and $\delta(\cdot)$ is of bounded variation, we have $\int_0^T \delta(t)\, D_2G(x,t)\, dt = \int_0^T \delta(t)\, dG(x,t)$, where the second integral is to be interpreted in the Riemann-Stieltjes sense. From the integration by parts formula for these integrals and $\delta(0) = 1$ we obtain:

$$\int_0^T \delta(t)\, dG(x,t) = -\int_0^T G(x,t)\, d\delta(t) + \delta(T)\, G(x,T) - G(x,0).$$

By the Riesz representation theorem (Lang (1993, Theorem IX.2.7)), the positive linear functional $f(\cdot) \mapsto -\int_0^\infty f(t)\, d\delta(t)$ on the space of continuous real-valued functions $f(\cdot)$ with compact support in $\mathbb{R}_+$ can be represented in a unique way by a measure on the Borel sets of $\mathbb{R}_+$, and it is immediate to check that the corresponding measure is μ_δ. Using the fact that $\lim_{T\to\infty} \delta(T)\, G(x,T) = 0$ and that $\mu_\delta(\{0\}) = 1 - \delta(0^+)$ together with the above integration by parts formula, we get:

$$\int_0^\infty \delta(t)\, D_2 G(x,t)\, dt = \int_{(0,\infty)} G(x,t)\, d\mu_\delta(t) - \delta(0^+)\, G(x,0)$$

by Lebesgue's dominated convergence theorem. The result then follows from (56). ■

From this, we can now deduce (43).

Proposition 3 *If the Markov processes $\{x_t\}_{t\geq 0}$ and $\{e_t\}_{t\geq 0}$ have P-almost surely continuous sample paths, and if Assumptions 1-5 are satisfied, then for P-almost every $\omega \in \Omega$, $\lim_{\tau\to 0^+} \widetilde{\pi}_t^\omega(\tau) = \exp\big(-\int_0^t \frac{\widetilde{\Phi}(x_v^\omega)}{\Gamma(x_v^\omega)}\, dv\big)(e_t^\omega)^{-\gamma} = \widetilde{\pi}_t^\omega$ for each $t \geq 0$.*

Proof. Consider a continuous sample path $\{x_t^\omega\}_{t\geq 0}$. At any date $t \geq 0$, the state prices for the continuous-time economy with commitment period $\tau > 0$ are given by:

$$\begin{aligned}
\widetilde{\pi}_t^\omega(\tau) &= \exp\left(\sum_{n=1}^{[\![t/\tau]\!]} \ln\left(\frac{\widetilde{\Delta}(x_{n\tau}^\omega, \tau)}{\Gamma(x_{n\tau}^\omega)}\right)\right)(e_t^\omega)^{-\gamma} \\
&= \exp\left(\int_0^t \iota_{t,\tau}(v)\, \widetilde{Q}(x_{\tau[\![v/\tau]\!]}^\omega, \tau)\, dv\right)(e_t^\omega)^{-\gamma},
\end{aligned}$$

where $\iota_{t,\tau}(v) = 1$ if $v/\tau \in [1, [\![t/\tau]\!]]$ and zero otherwise, and:

$$\widetilde{Q}(x,\tau) = \frac{1}{\tau} \ln\left(1 - \tau\, \frac{\widetilde{\Phi}(x,\tau)}{\Gamma(x)}\right)$$

for all $\tau > 0$ and $x \in X$. By Lemma 5, $\lim_{\tau\to 0^+} \widetilde{Q}(x,\tau) = -\widetilde{\Phi}(x)/\Gamma(x)$ for any $x \in X$. Note that since x_v^ω is a continuous function of v and $\Gamma(\cdot)$ is bounded away from zero on compact subsets of X, $\Gamma(x_{\tau[\![v/\tau]\!]}^\omega)$ is bounded away from zero for all $\tau > 0$ small enough. Next, observe from (55) that $\widetilde{\Phi}(x,\tau) \leq M(x) + \mathcal{M}(x)$ for all $\tau > 0$ and $x \in X$. Since both $M(\cdot)$ and $\mathcal{M}(\cdot)$ are bounded on compact subsets of X and x_v^ω is a continuous function of v, it follows that $\widetilde{Q}(x_{\tau[\![v/\tau]\!]}^\omega, \tau)$ is uniformly bounded on $[0,t]$ for all $\tau > 0$ small enough. The result then follows from Lebesgue's dominated convergence theorem and the continuity of $\{x_t^\omega, e_t^\omega\}_{t\geq 0}$ for P-almost every $\omega \in \Omega$. ■

Remark 3 Consider now the discount function $\delta_\tau(\cdot)$ in (45). In that case, the analogue of $\widetilde{\Phi}(x,\tau)$ is $\widetilde{\Phi}(x,\tau) + \frac{1}{\tau}\int_0^\tau (1 - \delta(t))\, G(x,t)\, dt$. Arguing along the same lines as in the proof of Lemma 4, it is not difficult to show that this converges to $\Phi(x)$. Finally, the proof of Proposition 3 can easily be adapted to prove that the corresponding state-price process converges to $\{\pi_t\}_{t\geq 0}$ as τ goes to zero.

D A Conditional Expectation

Write $\alpha = \gamma - 1$ and note:

$$G(x, \tau) = E\left[\exp\left(-\alpha\left(\eta\tau + \lambda\int_t^{t+\tau}(\mu - x_s)ds + \theta\int_t^{t+\tau}\sqrt{x_s}dB_s\right)\right)\Big| x_t = x\right].$$

It is convenient to define:

$$H(x, \tau) = \exp\left(\alpha\eta\tau\right)G(x, \tau).$$

Then:

$$\begin{aligned} & \exp\left(-\alpha\left(\lambda\int_0^t(\mu - x_s)ds + \theta\int_0^t\sqrt{x_s}dB_s\right)\right)H(x_t, T - t) \\ = \quad & E\left[\exp\left(-\alpha\left(\lambda\int_0^T(\mu - x_s)ds + \theta\int_0^T\sqrt{x_s}dB_s\right)\right)\Big| x_t = x\right]. \end{aligned}$$

The right-hand side of this equation is a martingale. The drift of the left-hand side must therefore be zero. By Ito's Lemma, this is the case when:

$$D_tH(x, t) = \left(-\alpha\lambda(\mu - x) + \frac{1}{2}\alpha^2\theta^2x\right)H(x, t) + \mathcal{L}H(x, t) - \alpha\theta\sigma xD_xH(x, t), \quad (57)$$

where:

$$\mathcal{L}H(x, t) = \kappa(\mu - x)D_xH(x, t) + \frac{1}{2}\sigma^2xD_{xx}H(x, t).$$

Note that $G(x, 0) = 1$. The boundary condition for $H(x, t)$ at $t = 0$ is therefore:

$$H(x, 0) = 1. \quad (58)$$

The solution of the PDE (57)-(58) is of the form:

$$H(x, t) = a(t)\exp\left(b(t)x\right).$$

The PDE (57)-(58) implies that $a(\cdot)$ and $b(\cdot)$ must satisfy:

$$\frac{Da(t)}{a(t)} + Db(t)\,x = \left(\kappa\mu - (\kappa + \alpha\theta\sigma)\,x\right)b(t) + \frac{1}{2}\sigma^2x\,b^2(t) + \left(-\alpha\lambda\mu + \left(\alpha\lambda + \frac{1}{2}\alpha^2\theta^2\right)x\right)$$

for all $x, t \geq 0$. Therefore:

$$\begin{aligned} Da(t) &= \kappa\mu a(t)\left(b(t) - \alpha\left(\frac{\lambda}{\kappa}\right)\right), & (59) \\ Db(t) &= \left(\alpha\lambda + \frac{1}{2}\alpha^2\theta^2\right) - (\kappa + \alpha\theta\sigma)\,b(t) + \frac{1}{2}\sigma^2b^2(t), & (60) \end{aligned}$$

together with the boundary conditions $a(0) = 1$ and $b(0) = 0$. The differential equation (60) is a Riccati equation that can be solved explicitly. Its solution does not explode in finite time if and only if the $b(t)$-polynomial on the right-hand side of (60) has at least one real root and the largest of these roots is non-negative. If $\alpha > 0$, it can be verified that this is equivalent to:

$$\frac{(\gamma - 1)\lambda}{\omega\kappa} \leq -\left(\frac{(\gamma - 1)\theta}{\omega\sigma}\right)^2 + \left(\frac{(\gamma - 1)\theta}{\omega\sigma} + \frac{1}{2}\right)\max\left\{\frac{(\gamma - 1)\theta}{\omega\sigma} + \frac{1}{2}, 0\right\}, \quad (61)$$

and that the solution for $G(x, t)$ is given by (49).

References

[1] Ainslie, G. (1975): "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control," *Psychological Bulletin*, 82: 463-496.

[2] Ainslie, G. (1992): *Picoeconomics*, Cambridge: Cambridge University Press.

[3] Akerlof, G.A. (1991): "Procrastination and Obedience," *American Economic Review*, 81: 1-19.

[4] Barro, R.J. (1999): "Ramsey Meets Laibson in the Neoclassical Growth Model," *Quarterly Journal of Economics*, Vol. CXIV, Issue 4 (November): 1125-1152.

[5] Benabou, R., and J. Tirole (1999): "Self-Confidence: Intrapersonal Strategies," mimeo, GREMAQ.

[6] Campbell, J.Y. (1999): "Asset Prices, Consumption, and the Business Cycle," in *Handbook of Macroeconomics*, edited by John B. Taylor and Michael Woodford, Amsterdam, North-Holland.

[7] Campbell, J.Y. and J.H. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, Vol. 2, No. 2 (April): 205-251.

[8] Carrillo, J.D., and T. Mariotti (1997): "Strategic Ignorance as a Self-Disciplining Device," forthcoming *Review of Economic Studies*.

[9] Cox, J.C., J.E. Ingersoll, and S.A. Ross (1985): "A Theory of the Term Structure of Interest Rates," *Econometrica*, 53: 385-406.

[10] Feller, W. (1951): "Two Singular Diffusion Problems," *Annals of Mathematics*, 54, 173-182.

[11] Gul, F. and W. Pesendorfer (1999a): "Temptation and Self-Control," forthcoming, *Econometrica*.

[12] Gul, F. and W. Pesendorfer (1999b): "Self-Control and the Theory of Consumption," Princeton University.

[13] Hamilton, J. (1989): "A New Approach to the Economic Analysis of Nonstationary Times Series and the Business Cycle," *Econometrica*, 57, 357-384.

[14] Harris, C., and D. Laibson (1999): "Dynamic Choices of Hyperbolic Consumers," forthcoming, *Econometrica*.

[15] Herrnstein, R.J. (1961): "Relative and Absolute Strengths of Response as a Function of Frequency of Reinforcement," *Journal of the Experimental Analysis of Animal Behavior*, 4: 267-272.

[16] Heston, S.L. (1993): "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *The Review of Financial Studies*, Vol. 6, No. 2: 327-343.

[17] Kocherlakota, N. (1990): "On the 'Discount Factor' in Growth Economies," *Journal of Monetary Economics*, 25: 43-47.

[18] Laibson, D. (1994): "Self-Control and Saving," mimeo, Harvard University.

[19] Laibson, D. (1996): "Hyperbolic Discount Functions, Undersaving, and Savings Policy," NBER working paper 5635.

[20] Laibson, D. (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, Vol. CXII, Issue 2 (May): 443-477.

[21] Lang, S. (1993): *Real and Functional Analysis*, New York: Springer-Verlag.

[22] LeRoy, S. and R. Porter (1981): "The Present Value Relation: Tests Based on Variance Bounds," *Econometrica*, 49, 555-577.

[23] Loewenstein, G., and R. Thaler (1989): "Anomalies: Intertemporal Choices," *Journal of Economic Perspectives*, 3: 181-193.

[24] Loewenstein, G., and D. Prelec (1992): "Anomalies in Intertemporal Choice: Evidence and an Interpretation," Quarterly Journal of Economics, Vol. CVII, Issue 3 (April), 573-597.

[25] Luttmer, E.G.J., and T. Mariotti (1999): "The Equilibrium Implications of Time-Inconsistent Preferences," mimeo, LSE.

[26] Krusell, P., and A. Smith (1999): "Consumption and Savings Decisions with Quasi-Geometric Discounting," mimeo, Carnegie Mellon University.

[27] Maskin, E., and J. Tirole (1998): "Markov Perfect Equilibrium I: Observable Actions," mimeo, GREMAQ.

[28] Mehra, R., and E.C. Prescott (1985): "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15: 145-162.

[29] Phelps, E.S., and R.A. Pollak (1968): "On Second-Best National Saving and Game-Equilibrium Growth," *Review of Economic Studies*, Vol. 35, No. 2 (April): 185-199.

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

[31] O'Donoghue, T., and M. Rabin (1999): "Incentives for Procrastinators," *Quarterly Journal of Economics*, Vol. CXIV, Issue 4 (August): 769-816.

[32] Shiller, R.J. (1981): "Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?" *American Economic Review,* 71, 421-436.

[33] Strotz, R.H. (1956): "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, Vol. 23, No. 3: 165-180.

[34] Thaler, R. (1981): "Some Empirical Evidence on Dynamic Inconsistency," *Economics Letters*, 4: 193-205.