

Subjective Discounting in an Exchange Economy*

Erzo G.J. Luttmer
University of Minnesota and CEPR

Thomas Mariotti
London School of Economics,
GREMAQ, Université de Toulouse 1, and CEPR

December 15, 2002

Abstract

This paper describes the equilibrium of a discrete-time exchange economy in which consumers with arbitrary subjective discount factors and homothetic period utility functions follow linear Markov consumption and portfolio strategies. Explicit expressions are given for state prices and consumption-wealth ratios. We provide an analytically convenient continuous-time approximation and show how subjective rates of time preference affect risk-free rates but not instantaneous risk-return trade-offs. Hyperbolic discount factors can be a source of return volatility, but they cannot be used to address asset pricing puzzles related to high-frequency Sharpe ratios.

1. INTRODUCTION

Many explicit dynamic equilibrium models in macroeconomics and finance are based on the assumption that consumers have time and state separable preferences and that they discount future utilities at a constant rate. This implies that consumer choices are dynamically consistent.

Psychologists have questioned the validity of the assumption of exponential discounting on the basis of experimental evidence (Chung and Herrnstein (1967), Ainslie (1975), Kirby and Herrnstein (1995)). These studies suggest instead that subjective discount functions are approximately hyperbolic. According to this literature, events

*We thank the participants in numerous seminars, and the Editor, Fernando Alvarez, for helpful comments. The FMG and STICERD at the London School of Economics provided financial support. An earlier version of this paper circulated as CEPR discussion paper 2503.

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 functions 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 of 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 on consumption (and possibly in some long-lived securities as well.) These contracts can be perfectly enforced, so that consumers are allowed to borrow up to the present value of their endowments. In contrast, there are no enforceable contracts that allow consumers to commit in advance to a particular sequence of consumption choices.

Following Strotz (1956) 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 all have the same constant relative risk aversion (CRRA) preferences. Endowment processes may differ across consumers. The only restriction we impose on subjective discount factors is that the utility of aggregate endowments is finite.

We consider competitive equilibria in which consumers follow intrapersonal strategies that give rise to consumption and portfolio choices that are proportional to wealth. These strategies are required to constitute a Markov perfect equilibrium of the intrapersonal game. If there is a long-lived asset in positive net supply, then there is one and only one such competitive equilibrium. The fact that individual behavior is the outcome of a game implies that standard derivations of a transversality condition

¹Beyond the consumption-savings problem, the "multiple selves" methodology has been applied to a broad set of self-control issues. See for instance Benabou and Tirole (2002), Carrillo and Mariotti (2000) and O'Donoghue and Rabin (1999, 2001).

do not apply. We provide conditions on preferences and endowments that nevertheless rule out bubbles on long-lived assets. These conditions also ensure existence of equilibrium in the absence of a long-lived asset. We obtain explicit expressions for the consumption and portfolio strategies, as well as for the equilibrium prices of state-contingent claims.

In two special cases, equilibrium prices will be exactly as if consumers discount utility geometrically, even when they do not. This happens when endowments are such that expected utility growth is constant, or when period utility functions are logarithmic. In both cases, agents consume a constant fraction of wealth in each period, irrespective of the shape of the subjective discount function. This observational equivalence result generalizes a result of Barro (1999) to the case of an economy with uncertainty.

To gain more insight into the properties of equilibrium state prices, we examine a limit economy obtained by letting the length of a period go to zero. In this limit economy, the standard consumption capital asset pricing model (C-CAPM) of Breeden (1979) applies, irrespective of how consumers discount utility: instantaneous risk premia are given by instantaneous covariances of returns with marginal utility. The risk-free rate process, however, is affected by time-inconsistency. It is the sum of two terms. One is the expected instantaneous growth rate of marginal utility, as is the case when consumers are time-consistent. The other is a weighted average of consumer subjective rates of time preference at different horizons. The weights are proportional to discounted expected future period utilities. This second term introduces a new source of risk-free rate dynamics: variation in expected future endowment growth shifts weights across subjective discount rates at different horizons, and this in turn affects the current risk-free rate.

The fact that Breeden's formula for instantaneous risk premia continues to hold implies that the high-frequency Sharpe ratios of returns in an economy with non-geometric subjective discount functions are the same as those in an economy with time-consistent consumers. Many models of preferences are not able to generate the high Sharpe ratios found in many data sets (Hansen and Jagannathan (1991)). This is a puzzle that can therefore not be addressed using only non-standard assumptions about subjective rates of time preference.

Although ratios of risk premia over return volatilities do not depend on how consumers discount utility, the return volatilities themselves typically do. In particular, we give conditions under which certain forms of hyperbolic discounting can increase the volatility of aggregate wealth. Intuitively, the high subjective discount rates for nearby utilities implied by hyperbolic discount functions must be matched by low subjective discount rates for utilities in the distant future, or else interest rates would be too high compared to what is observed in the data. These low long-run subjective discount rates can make wealth more sensitive to information about long-run

endowment growth rates.

Related Literature One of the subjective discount functions we consider is the “quasi-hyperbolic” discount function introduced by Phelps and Pollak (1968) in a model of imperfect intergenerational altruism. It 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 (2001) 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 consecutive dates, but also on the marginal propensity to consume out of wealth. Krusell and Smith (1999) consider an economy with quasi-hyperbolic discount factors 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. In contrast, we incorporate the high short-run rates of time preference suggested by the psychological literature, and emphasize that one can use low long-run rates of time preference to account for the low risk-free rates found in the data.

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 examine the combined implications of hyperbolic discounting and monotone dynamics in expected utility for the dynamics of the risk-free rate and the volatility of aggregate wealth. Section 6 contains concluding remarks.

2. AN EXCHANGE ECONOMY

2.1. Preferences and Endowments

We consider a discrete-time, infinite-horizon economy. Time is labelled by $t = 0, 1, 2, \dots$ and uncertainty is described by a probability space endowed with a non-decreasing 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.

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 or “stock” 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.

Following Peleg and Yaari (1973), we view the representative consumer as composed of a sequence of autonomous decision makers, indexed by time. We refer to the decision maker at date t as the “date- t consumer.” This date- t consumer evaluates current and future consumption according to a utility function:

$$U_t(\{c_{t+n}\}_{n=0}^{\infty}) = \mathbb{E}_t \left[\sum_{n=0}^{\infty} \delta_n u(c_{t+n}) \right], \quad (1)$$

where $\delta_0 = 1$ and $\delta_n > 0$. These 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 is given by:

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

for some $\gamma > 0$, $\gamma \neq 1$. We take $\gamma = 1$ to mean $u(c) = \ln(c)$. It is easy to modify (1) to incorporate habits or subsistence levels of consumption.

2.2. Examples of Subjective Discount Factors

It is convenient to set:

$$\delta_n = \delta(n\tau),$$

where δ is a positive discount function defined on $[0, \infty)$, and $\tau > 0$ is the length of a period. The discrete-time subjective rate of time preference can be defined as $-\ln(\delta_{n+1}/\delta_n)/\tau$. For smooth discount functions, this converges to $\rho(t) = -D\delta(t)/\delta(t)$ as τ goes to zero.

Experimental studies by psychologists (see Ainslie (1975), Kirby and Herrnstein (1995), 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. O’Donoghue and Rabin (1999, 2001) refer to this phenomenon as “present bias.” A discount function that captures present bias is:

$$\delta(t) = (1 + \zeta t)^{-\xi/\zeta} \exp(-\rho t). \quad (2)$$

This combines an exponential discount function with the “generalized hyperbolic” discount function for which Loewenstein and Prelec (1992) have suggested an axiomatic justification. It generalizes the hyperbolic discount function proposed by Ainslie (1975) to interpret experiments that indicate reversals over time of preferences for rewards at different horizons.

The resulting subjective discount rate is given by $\rho(t) = \rho + \xi/(1 + \zeta t)$. This converges to $\rho + \xi$ as the horizon goes to zero, and it declines monotonically to ρ as

the horizon goes to infinity. Increasing ρ increases the discount rate at all horizons, and increasing ξ raises the subjective discount rate more at short horizons than at long horizons. 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 . At $\xi = 0$, there is no difference between short-run and long-run discount rates, and one obtains the standard exponential discount function $\exp(-\rho t)$. Similarly, letting ζ go to zero yields the exponential discount function $\exp(-(\rho + \xi)t)$.

A particularly convenient limiting discount function is obtained by letting the speed parameter ζ and the discrepancy ξ between short and long-run discount rates go to infinity at the same time. Specifically, take some $\delta_+ < 1$, set $\xi = -\ln(\delta_+)\zeta/\ln(\zeta)$, and let ζ go to infinity. Then $\delta(t)$ converges to:

$$\delta(t) = \delta_+ \exp(-\rho t) \tag{3}$$

for all $t > 0$. The corresponding discrete-time discount factors can be written as $\delta_n = \delta_+ \beta^n$, for some positive β and all $n > 0$. These are the discount factors used by Phelps and Pollak (1968), and more recently Laibson (1994, 1997), Harris and Laibson (2001), and many others. Laibson (1994) argued that these discount factors can provide a good approximation to hyperbolic discount factors and called this “quasi-hyperbolic discounting.”

2.3. Markets

Our main assumption is that markets are sequentially complete. One-period ahead state-contingent claims are traded at every date and in every state. The cost at date t of a portfolio of state contingent claims that delivers b_{t+1} units of consumption at date $t + 1$ is $E_t[\pi_{t+1}b_{t+1}]/\pi_t$ units of consumption, where $\{\pi_t\}_{t=0}^\infty$ is a strictly positive sequence of probability-weighted state-contingent prices. Stocks are traded at all dates t , at ex-dividend prices equal to s_t units of consumption per unit of stock. We restrict attention to equilibria in which stock prices are non-negative.

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

$$\pi_t c_t + E_t[\pi_{t+1}b_{t+1}] + \pi_t s_t k_t \leq \pi_t [n_t + b_t + (s_t + d_t)k_{t-1}], \tag{4}$$

and subject to the present-value borrowing constraint:

$$E_{t+1} \left[\sum_{s=t+1}^{\infty} \pi_s n_s \right] + \pi_{t+1} [b_{t+1} + (s_{t+1} + d_{t+1})k_t] \geq 0. \tag{5}$$

There are no contingent claims outstanding at the initial date, and so $b_0 = 0$. The constraint (5) restricts the portfolio of b_{t+1} and k_t chosen by the date- t consumer to be such that wealth at the beginning of period $t + 1$ is non-negative in every state.

In view of the fact that markets are complete, and since there are no constraints on portfolios other than that current consumption and next-period wealth must remain non-negative, we focus on equilibria in which prices satisfy the no-arbitrage condition:

$$\pi_t s_t = \mathbb{E}_t [\pi_{t+1}(s_{t+1} + d_{t+1})]. \quad (6)$$

A violation of this condition would imply that the date- t consumer could choose to consume an unlimited amount.²

The only way in which the portfolio choices of the date- t consumer affect the set of feasible choices at later dates is through their effect on date- $t + 1$ wealth. It will be convenient to write w_t for the consumer's wealth at the beginning of date t :

$$\pi_t w_t = \mathbb{E}_t \left[\sum_{s=t}^{\infty} \pi_s n_s \right] + \pi_t [b_t + (s_t + d_t)k_{t-1}]. \quad (7)$$

Given the no-arbitrage condition (6), it then follows that the set of budget-feasible consumption choices defined by (4)-(5) is equivalent to the set of sequences $\{c_t\}_{t=0}^{\infty}$ that satisfy, for some sequence $\{w_{t+1}\}_{t=0}^{\infty}$:

$$\begin{aligned} \pi_t c_t + \mathbb{E}_t[\pi_{t+1} w_{t+1}] &\leq \pi_t w_t, \\ c_t, w_{t+1} &\geq 0, \end{aligned} \quad (8)$$

with initial wealth w_0 defined by (7). Note that the stochastic process of endowments $\{e_t\}_{t=0}^{\infty}$ only affects the consumer's set of feasible consumption choices via initial wealth, unlike in models with borrowing constraints that are tighter than (5), such as in Harris and Laibson (2001). For an individual consumer there is therefore no commitment value to changing his or her endowment process using, say, pension investments that are not directly tradable.

2.4. Intrapersonal and Competitive Equilibrium

In any period, the consumer can in no way restrict his future actions, other than through the amount of wealth transferred to the next period. Behavior of the infinitely-lived consumer is therefore not the outcome of a single utility maximization, but of a strategic interaction among the sequence of date- t consumers who make choices at successive dates.

This results in an intrapersonal game in which each date- t consumer chooses his current consumption c_t and next-period state-contingent wealth w_{t+1} , taking as given a sequence of prices $\{\pi_t, s_t\}_{t=0}^{\infty}$ and the strategies of the date- s consumers for all $s > t$. A strategy for the date- t consumer in the intrapersonal game is a mapping

²For certain period utility functions u , it may be possible to construct equilibria in which such arbitrage opportunities are not exploited. We will not consider this possibility here.

(C_t, W_{t+1}) that specifies, for any history h_t of the game up to date t , a consumption level $c_t = C_t(h_t)$, and amounts of state-contingent wealth $w_{t+1} = W_{t+1}(h_t)$, such that the budget constraint (8) is satisfied given w_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.

A competitive equilibrium of the representative agent economy is given by a strategy profile $\{(C_t, W_{t+1})\}_{t=0}^{\infty}$ 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 condition (6);
- (ii) The present value of endowments is finite at prices $\{\pi_t\}_{t=0}^{\infty}$;
- (iii) $\{(C_t, W_{t+1})\}_{t=0}^{\infty}$ is an intrapersonal equilibrium at prices $\{\pi_t, s_t\}_{t=0}^{\infty}$;
- (iv) Goods, state-contingent claims, and stock markets clear at every date and in every state:

$$C_t(h_t) = e_t,$$

$$\pi_{t+1}W_{t+1}(h_t) = E_{t+1} \left[\sum_{s=t+1}^{\infty} \pi_s n_s \right] + \pi_{t+1}(s_{t+1} + d_{t+1})k_{-1},$$

where h_t is the date- t history of an intrapersonal equilibrium path implied by $\{(C_t, W_{t+1})\}_{t=0}^{\infty}$.

The price-taking assumption entails that each date- t consumer uses the prevailing prices when evaluating the payoff of a deviation from the intrapersonal equilibrium. Markets need not clear following a deviation in the intrapersonal game.

3. EQUILIBRIUM

Because consumers are assumed to be infinitely lived, the intrapersonal game that determines behavior may have many subgame perfect equilibria. We focus here on Markov perfect equilibria in which the consumption and portfolio choices of the consumer at any date depend only on current wealth and on exogenous variables.³ This rules out bootstrap strategies that might otherwise be used to mitigate the consumer's self-control problem (see Laibson (1994)).

³This is also the perspective adopted by Phelps and Pollak (1968) and Harris and Laibson (2001).

3.1. Intrapersonal Equilibrium

Fix prices, and consider a Markov perfect equilibrium for the intrapersonal game. With slight abuse of notation, let $C_t(w)$ and $W_{t+1}(w)$ denote the consumption and wealth choices of a date- t consumer who starts with date- t wealth w . Given these Markov strategies, let $C_{t,t+n}(w)$ be the implied level of consumption chosen by the date- $t+n$ consumer in the subgame in which the date- t consumer starts with date- t wealth w . Define:

$$F_t(w) = \mathbb{E}_t \left[\sum_{n=0}^{\infty} \delta_n u(C_{t,t+n}(w)) \right], \quad V_t(w) = \mathbb{E}_t \left[\sum_{n=0}^{\infty} \delta_{n+1} u(C_{t,t+n}(w)) \right]. \quad (9)$$

The current value $F_t(w)$ is the expected utility for the date- t consumer in the subgame in which the date- t consumer starts with wealth w . The continuation value $V_t(w)$ is the utility perceived by the date- $t-1$ consumer for this same subgame. In any Markov equilibrium, $C_t(w)$ and $W_{t+1}(w)$ must be a best response for the date- t consumer when his initial wealth is w . That is, we must have:

$$F_t(w) = \max_{c_t, w_{t+1} \geq 0} \{u(c_t) + \mathbb{E}_t [V_{t+1}(w_{t+1})] : \pi_t c_t + \mathbb{E}_t [\pi_{t+1} w_{t+1}] \leq \pi_t w\} \quad (10)$$

for all possible levels of initial date- t wealth w .⁴

Starting from initial wealth w_0 defined by (7), write $w_{t+1} = W_{t+1}(w_t)$ and $c_t = C_t(w_t)$ for all t . Under regularity conditions that will be verified below, the sequence of date- t first-order and envelope conditions for (10) can then be written as:

$$\frac{\pi_{t+1}}{\pi_t} = \frac{DV_{t+1}(w_{t+1})}{Du(c_t)}, \quad (11)$$

$$DF_t(w_t) = Du(c_t). \quad (12)$$

Combining the first-order condition at date t with the envelope condition at date $t+1$, we obtain:

$$\frac{\pi_{t+1}}{\pi_t} = \frac{DV_{t+1}(w_{t+1}) Du(c_{t+1})}{DF_{t+1}(w_{t+1}) Du(c_t)}. \quad (13)$$

Clearly, if $\delta_n = \beta^n$, then (9) implies $V_{t+1}(w_{t+1}) = \beta F_{t+1}(w_{t+1})$, and (13) yields the standard first-order condition. More generally, (13) shows that the standard geometric subjective discount factor must be replaced by a ratio of marginal utilities of wealth based on the continuation value function V_{t+1} and next period's current value function F_{t+1} .

The homotheticity of preferences suggests a Markov perfect equilibrium of the intrapersonal game that is linear in wealth:

$$C_t(w) = \phi_t w, \quad W_{t+1}(w) = \psi_{t+1} w, \quad (14)$$

⁴This equation is related to the “quasi-Bellman equation” developed by Harris and Laibson (2001) for quasi-hyperbolic subjective discount functions.

for some time and state dependent coefficients ϕ_t and ψ_{t+1} that do not depend on wealth. The strategies (14) imply that $F_t(w)$ and $V_t(w)$ as defined in (9) are both proportional to $w^{1-\gamma}/(1-\gamma)$. The maximization in (10) therefore implies a date- t best response that is linear in date- t wealth, and the right-hand side of (10) will again be proportional to $w^{1-\gamma}/(1-\gamma)$. The Markov strategies (14) constitute a subgame perfect equilibrium of the intrapersonal game if the coefficients $\{\phi_t, \psi_{t+1}\}_{t=0}^\infty$ are such that (10) is satisfied at every date and in every state for F_t and V_{t+1} constructed from these coefficients according to (9).

3.2. Competitive Equilibrium

Except for the special case of logarithmic preferences, it is hard to explicitly solve for a Markov equilibrium of the intrapersonal game at arbitrary prices. But to construct a competitive equilibrium, we only need to solve for an intrapersonal equilibrium at market clearing prices.

Equilibrium State Prices The homotheticity of the period utility function implies that $F_t(w)/[C_t(w)]^{1-\gamma}$ and $V_t(w)/[C_t(w)]^{1-\gamma}$ only depend on w via the consumption growth rates $C_{t,t+n}(w)/C_t(w)$. But the linearity of the strategies (14) implies that $C_{t,t+n}(w)/C_t(w)$ does not depend on w , while market clearing requires that this ratio is equal to e_{t+n}/e_t . Therefore, at market clearing prices, the value functions (9) can be written as:

$$F_t(w) = \frac{\Gamma_t(\phi_t w)^{1-\gamma}}{1-\gamma}, \quad V_t(w) = \frac{\Delta_t(\phi_t w)^{1-\gamma}}{1-\gamma}, \quad (15)$$

where the coefficients Γ_t and Δ_t are given by:

$$\Gamma_t = \mathbb{E}_t \left[\sum_{n=0}^{\infty} \delta_n \left(\frac{e_{t+n}}{e_t} \right)^{1-\gamma} \right], \quad \Delta_t = \mathbb{E}_t \left[\sum_{n=0}^{\infty} \delta_{n+1} \left(\frac{e_{t+n}}{e_t} \right)^{1-\gamma} \right]. \quad (16)$$

The envelope condition (12) combined with $c_t = \phi_t w_t$ and the expression for $F_t(w)$ in (15) yields $\phi_t = 1/\Gamma_t$. Together with $w_{t+1} = \psi_{t+1} w_t$ and the fact that market clearing requires that $c_{t+1}/c_t = e_{t+1}/e_t$ we then obtain the following coefficients for the Markov strategies:

$$\phi_t = \frac{1}{\Gamma_t}, \quad \psi_{t+1} = \frac{e_{t+1}\Gamma_{t+1}}{e_t\Gamma_t}. \quad (17)$$

This expresses the coefficients of the linear Markov strategies (14) in terms of the endowment process of the economy. To construct an explicit formula for equilibrium state prices, we can use (13), (15)-(16) and market clearing to obtain:

$$\frac{\pi_{t+1}}{\pi_t} = \frac{\Delta_{t+1}}{\Gamma_{t+1}} \left(\frac{e_{t+1}}{e_t} \right)^{-\gamma}. \quad (18)$$

In any equilibrium in which consumers follow Markov strategies that are linear in wealth, state prices must be of this form.

Remaining Equilibrium Conditions Clearly, the construction leading up to (17)-(18) requires that Γ_t and Δ_t are finite. That is, utility of the aggregate endowment process must be finite at every date and in every state. If this is the case, then the strategies defined by (14) and (17) are such that consumption grows at the same rate as the aggregate endowments. The only remaining requirement for a competitive equilibrium is that the levels of the consumption and endowment processes are the same as well.

By iterating on (8) and using the market clearing requirement $e_T = w_T/\Gamma_T$ one can verify that initial wealth must be equal to the present value of the aggregate endowments, plus $\lim_{T \rightarrow \infty} E_0[\pi_T e_T \Gamma_T]$. Similarly, by iterating on (6) and using the definition of wealth (7) one obtains that initial wealth must also be equal to the present value of the aggregate endowments, plus $\lim_{T \rightarrow \infty} E_0[\pi_T s_T]k_{-1}$. Our construction yields an equilibrium if these two limit terms are equal. If the stock is in positive net supply, this can be guaranteed by assigning a high enough value to s_0 ; in particular, there may have to be a bubble on the stock. On the other hand, if $k_{-1} = 0$, then an assumption is needed to ensure that $E_0[\pi_T e_T \Gamma_T]$ goes to zero as T gets large. Using (18) one can verify that this is equivalent to requiring that:

$$\lim_{T \rightarrow \infty} E_0 \left[\prod_{t=1}^T \frac{\Delta_t}{\Gamma_{t-1}} \left(\frac{e_t}{e_{t-1}} \right)^{1-\gamma} \right] = 0. \quad (19)$$

This is not something that follows simply from assuming that utility of the aggregate endowments is finite. Either of the following two assumptions is sufficient to ensure that (19) holds.

Condition A *There is an $\alpha > 0$ such that $\alpha\Gamma_t \in (0, 1]$ at all dates and in all states.*

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

$$E_t \left[\sum_{n=0}^{\infty} \beta^n \left(\frac{e_{t+n}}{e_t} \right)^{1-\gamma} \right] < \infty \quad (20)$$

at every date and in every state.

The proof of the following proposition is given in Appendix A.⁵

Proposition 1 *If there is a long-lived asset in positive net supply, or if at least one of Conditions A or B holds, then there is an equilibrium in which consumption and*

⁵The proposition does not cover the case of quasi-hyperbolic preferences with $\delta_+ > 1$, but a separate proof can be given in this case.

portfolio choices are proportional to wealth. The strategies are given by (14) and (16)-(17), and equilibrium state prices are given by (18).

It is possible to construct examples in which there must be a bubble on the long-lived asset.⁶ In such examples, the consumption-wealth ratio goes to zero fast enough so that consumers do not exhaust their present-value budget constraints. This is consistent with equilibrium only if there is a long-lived asset that trades above the present value of its dividends. For time-consistent consumers, single-agent optimization implies that consumers exhaust their present-value budget constraints, and this is the argument that is traditionally used to derive a transversality condition that can then in turn be used to rule out bubbles on long-lived assets (see Scheinkman (1977), Brock (1979) and Obstfeld and Rogoff (1983)). Here, the requirement that consumers play a subgame-perfect equilibrium does not guarantee that they exhaust their present-value budget constraints, and so a transversality condition need not apply.

The possible need for Conditions A or B or bubbles is a consequence of our focus on Markov strategies. If we allow for non-Markov strategies, then there is always a competitive equilibrium in which there are no bubbles on long-lived assets. The equilibrium consumption and portfolio choices are affine functions of wealth, and the formula for state prices (18) continues to hold.

Logarithmic Preferences The above derivations do not apply directly to the case of $\gamma = 1$, although (16)-(18) do hold. For logarithmic preferences, the intrapersonal game can be solved explicitly. The solution is $\phi_t = 1 - \bar{\beta}$ and $\psi_{t+1} = \bar{\beta}\pi_t/\pi_{t+1}$ where $\bar{\beta} = 1 - 1/\sum_{n=0}^{\infty} \delta_n$, given any state prices for which wealth is finite. One can use this to solve for equilibrium state prices, even when consumers are heterogeneous in terms of their subjective discount functions. The ratio of consumption (wealth) to aggregate consumption (wealth) converges to one for the consumer who is the most patient on average (the one with the highest $\bar{\beta}$.) In an economy with both time-consistent and time-inconsistent consumers, this could easily be a time-inconsistent consumer with high subjective rates of time preference at nearby horizons—as long as this consumer’s long-run subjective rates of time preference are sufficiently low.

3.3. Implications

Effective Subjective Discount Factors As can be seen from the equation for state prices (18), the two variables that determine state prices in this economy are endowment growth and the ratio $\Delta_{t+1}/\Gamma_{t+1}$. If consumers discount geometrically, the ratio $\Delta_{t+1}/\Gamma_{t+1}$ is constant and equal to the subjective discount factor β . For general

⁶ An example is an economy with $e_t = (t!)^{(1+\varepsilon)/(1-\gamma)}$ and $\delta_n = (\beta^n/n!)^{1+\varepsilon}$. This yields a positive limit in (19) if $\varepsilon > 0$, even though expected utility is finite.

subjective discount factors, $\Delta_{t+1}/\Gamma_{t+1}$ is an “effective subjective discount factor” that can be expressed as a weighted average of δ_{n+1}/δ_n :

$$\frac{\Delta_{t+1}}{\Gamma_{t+1}} = \sum_{n=0}^{\infty} \omega_{n,t+1} \left(\frac{\delta_{n+1}}{\delta_n} \right), \quad (21)$$

where the weights $\omega_{n,t+1}$ are given by:

$$\omega_{n,t+1} = \frac{\mathbb{E}_{t+1} [\delta_n (e_{t+1+n}/e_{t+1})^{1-\gamma}]}{\mathbb{E}_{t+1} [\sum_{n=0}^{\infty} \delta_n (e_{t+1+n}/e_{t+1})^{1-\gamma}]}.$$

These weights are proportional to the expected utility of date- $t+1+n$ consumption from the perspective of the date- $t+1$ 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:”

$$\frac{\pi_{t+1}}{\pi_t} = (\phi_{t+1} \delta_+ \beta + (1 - \phi_{t+1}) \beta) \left(\frac{e_{t+1}}{e_t} \right)^{-\gamma} \quad (22)$$

of Harris and Laibson (2001). More generally, consider the properties of (21) if subjective discount rates are relatively high at nearby horizons, and low at distant horizons. The discount rates implied by $\Delta_{t+1}/\Gamma_{t+1}$ 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 $\Delta_{t+1}/\Gamma_{t+1}$. If 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 $\Delta_{t+1}/\Gamma_{t+1}$.

Identification The distinguishing characteristic of non-geometric discounting is the fact that the usual geometric subjective discount factor is replaced by $\Delta_{t+1}/\Gamma_{t+1}$ in (18). This implies that it is not possible to differentiate non-geometric discount factors from geometric ones if $\Delta_{t+1}/\Gamma_{t+1}$ happens to be constant. This will be the case if the conditionally expected utility ratio $\mathbb{E}_t [(e_{t+1}/e_t)^{1-\gamma}]$ is constant. For any subjective discount function, one can then construct an alternative economy with a geometric subjective discount factor given by $\bar{\beta} = \Delta_{t+1}/\Gamma_{t+1}$. 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 a sum of expected utility ratios, discounted using the geometric subjective discount factor $\bar{\beta}$. 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+1} and Δ_{t+1} are simply sums of subjective discount factors and thus $\Delta_{t+1}/\Gamma_{t+1}$ is obviously constant across time.⁷

4. A CONTINUOUS-TIME APPROXIMATION

In the continuous-time version of the consumption capital asset pricing model studied by Breeden (1979), instantaneous expected returns in excess of the risk-free rate are equal to the instantaneous covariance of returns with marginal utility. The risk-free rate is equal to a constant subjective rate of time preference plus the expected instantaneous growth rate of marginal utility. In this section we describe how Breeden's results change when consumers are time-inconsistent. This will allow us to highlight properties of short-horizon returns in an economy with time-inconsistent consumers that are not apparent from the discrete-time state price formula (18).

4.1. Limit Properties of Discrete-Time Economies

Suppose that there exists an underlying continuous-time endowment process $\{e_t\}_{t \geq 0}$ that evolves according to a diffusion:

$$de_t = e_t (\mu_e(x_t)dt + \sigma_e(x_t)^\top dW_t), \quad (23)$$

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

$$dx_t = \mu_x(x_t)dt + \sigma_x(x_t)^\top dW_t. \quad (24)$$

Suppose also that the subjective discount function δ is defined for all $t \in [0, \infty)$ and normalized so that $\delta(0) = 1$.

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

$$\Gamma_t(\tau) = E_t \left[\sum_{n=0}^{\infty} \delta(n\tau) \left(\frac{e_{t+n\tau}}{e_t} \right)^{1-\gamma} \right], \quad \Delta_t(\tau) = E_t \left[\sum_{n=0}^{\infty} \delta((n+1)\tau) \left(\frac{e_{t+n\tau}}{e_t} \right)^{1-\gamma} \right]. \quad (25)$$

When t is an integer multiple of τ , these definitions correspond to those given in (16). In equilibrium, the ratio of consumption per unit of time over wealth in this

⁷Barro (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).

discrete-time economy is equal to $1/\tau\Gamma_t(\tau)$, for $t = 0, \tau, 2\tau, \dots$, as in (17). From (18), the one-period ahead state prices are given by:

$$\frac{\pi_{t+\tau}(\tau)}{\pi_t(\tau)} = \frac{\Delta_{t+\tau}(\tau)}{\Gamma_{t+\tau}(\tau)} \left(\frac{e_{t+\tau}}{e_t} \right)^{-\gamma}, \quad (26)$$

again for $t = 0, \tau, 2\tau, \dots$.

We now construct a continuous-time state-price process from (26) and examine the properties of its sample path as τ goes to zero. We adopt the normalization $\pi_0(\tau) = e_0^{-\gamma}$. Define $\pi_t(\tau)$ to be the product of $\pi_0(\tau)$ and the sequence of one-period ahead state prices (26) up to period $\lfloor t/\tau \rfloor - 1$, where $\lfloor t/\tau \rfloor$ is the integer part of t/τ . A somewhat indirect but revealing way to write this is:

$$\pi_t(\tau) = \exp \left(\sum_{n=1}^{\lfloor t/\tau \rfloor} \left(\frac{1}{\tau} \ln \left(1 - \left(\frac{\Gamma_{n\tau}(\tau) - \Delta_{n\tau}(\tau)}{\tau\Gamma_{n\tau}(\tau)} \right) \tau \right) \right) \tau \right) e_{\tau\lfloor t/\tau \rfloor}^{-\gamma}. \quad (27)$$

For every $\tau > 0$, this defines a continuous-time process indexed by t that coincides with (26) for any t equal to an integer multiple of τ .

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

$$\pi_t = \exp \left(- \int_0^t \frac{\Phi(x_v)}{\Gamma(x_v)} dv \right) e_t^{-\gamma}, \quad (28)$$

where:

$$\Phi(x_t) = -\mathbb{E}_t \left[\int_{[0, \infty)} \left(\frac{e_{t+v}}{e_t} \right)^{1-\gamma} d\delta(v) \right] \quad (29)$$

and:

$$\Gamma(x_t) = \mathbb{E}_t \left[\int_0^\infty \delta(v) \left(\frac{e_{t+v}}{e_t} \right)^{1-\gamma} dv \right]. \quad (30)$$

Appendix B gives a precise statement and proof of this result. The expressions (28)-(30) are as might be expected given (25)-(27). In particular, note that $\tau\Gamma_t(\tau)$ can be interpreted as an integral of discounted future expected utility ratios against time that converges to $\Gamma(x_t)$, and that the difference $\Gamma_t(\tau) - \Delta_t(\tau)$ can be interpreted as an integral of future expected utility ratios against δ that converges to $\Phi(x_t)$ as τ goes to zero.

Effective Subjective Discount Rates The ratio $\Phi(x_t)/\Gamma(x_t)$ will be referred to as the “effective subjective discount rate.” It replaces the constant subjective discount rate that would appear in (28) if consumers discounted utility geometrically.

If δ is sufficiently smooth so that one can write $\delta(t) = \exp(-\int_0^t \rho(v) dv)$, then the effective subjective discount rate simplifies to a weighted average of the subjective discount rates $\rho(v)$:

$$\frac{\Phi(x_t)}{\Gamma(x_t)} = \int_0^\infty \rho(v) \omega(x_t, v) dv, \quad (31)$$

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

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

as in (21). Of course, (31) will be constant if δ is exponential.

Alternatively, consider the quasi-hyperbolic discount function (3). This is clearly non-smooth: an arbitrarily small positive delay results in a discrete drop in the discount function. This drop shows up in the effective subjective discount rate via:

$$\frac{\Phi(x_t)}{\Gamma(x_t)} = \rho + \frac{1 - \delta_+}{\Gamma(x_t)}. \quad (32)$$

The associated wealth-consumption ratio can be written as:

$$\Gamma(x_t) = \delta_+ \mathbb{E}_t \left[\int_0^\infty \exp(-\rho v) \left(\frac{e_{t+v}}{e_t} \right)^{1-\gamma} dv \right]. \quad (33)$$

As expected, (32)-(33) implies that the effective subjective discount rate is increasing in ρ and the instantaneous discount $1 - \delta_+$.

Remark Proposition 2 describes the limiting properties of a sequence of discrete-time economies as the period length goes to zero. It is also possible to analyze the continuous-time economy directly. Following Barro’s (1999) analysis of the deterministic Cass-Koopmans growth model, 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. In Luttmer and Mariotti (2000a) we prove that this also yields (28) when the discount function for an economy with a commitment period of length τ is given by $\delta_\tau(t) = 1$ if $t \in [0, \tau)$ and $\delta_\tau(t) = \delta(t)$ if $t \in [\tau, \infty)$. If $\delta_\tau(t) = \delta(t)$ for all t instead, then any discontinuity of δ at zero would not be reflected in the limit, and the domain of integration in (29) would be $(0, \infty)$ instead of $[0, \infty)$.

4.2. Interest Rates and Risk Premia

Let R_t be the date- t cumulative return on some asset. That is, one unit of consumption invested at date t yields R_T/R_t units of consumption at date $T > t$ if the asset is held from t to T , and any dividends are reinvested. Suppose that we have an asset with cumulative returns that satisfy:

$$dR_t = R_t (\mu_R(x_t)dt + \sigma_R(x_t)^\top dW_t).$$

The state prices (28) form an Ito process that can be written as:

$$d\pi_t = -\pi_t (r(x_t)dt + \lambda(x_t)^\top dW_t).$$

for some $r(x_t)$ and $\lambda(x_t)$ that are given below. To avoid arbitrage opportunities, cumulative returns must satisfy $\pi_t R_t = E_t[\pi_T R_T]$ for all $t \leq T$. Thus $\pi_t R_t$ has no drift, and an application of Ito's lemma therefore implies the well-known relation:

$$\mu_R(x_t) - r(x_t) = \sigma_R(x_t)^\top \lambda(x_t). \quad (34)$$

If $\sigma_R(x_t) = 0$, then the cumulative return R_t is instantaneously risk-free, and so $r(x_t)$ can be interpreted as the risk-free rate. The coefficient $\lambda(x_t)$ is usually referred to as the "market price of risk." Using (28) and Ito's lemma, we obtain:

$$r(x_t) = \frac{\Phi(x_t)}{\Gamma(x_t)} + \gamma \mu_e(x_t) - \frac{1}{2} \gamma (1 + \gamma) \sigma_e(x_t)^\top \sigma_e(x_t), \quad (35)$$

$$\lambda(x_t) = \gamma \sigma_e(x_t). \quad (36)$$

Equations (34)-(36) show how the equilibrium risk-free rate and risk premia are determined by preferences and endowments.

The key implication of (35)-(36) is that the shape of the subjective discount function δ influences state prices only via the risk-free rate. In turn, the risk-free rate is only affected by the shape of δ through its effect on the effective subjective discount rate $\Phi(x_t)/\Gamma(x_t)$. The other determinants of the risk-free rate are the usual intertemporal substitution and precautionary savings effects represented by the second and third terms in (35). The dynamics of the risk-free rate will depend on the interaction of these standard effects with the variation in the effective subjective discount rate that arises when discounting is not exponential.

By contrast, the shape of the subjective discount function δ has no effect on the market price of risk $\lambda(x_t)$. Recall that the instantaneous Sharpe ratio for the cumulative return R_t is defined by $[\mu_R(x_t) - r(x_t)]/[\sigma_R(x_t)^\top \sigma_R(x_t)]^{1/2}$. From (34), this is equal to $\sigma_R(x_t)^\top \lambda(x_t)/[\sigma_R(x_t)^\top \sigma_R(x_t)]^{1/2}$. As pointed out by Hansen and Jagannathan (1991), the absolute value of this ratio is a lower bound for $[\lambda(x_t)^\top \lambda(x_t)]^{1/2}$, and thus, because of (36), for $\gamma[\sigma_e(x_t)^\top \sigma_e(x_t)]^{1/2}$. Given several returns processes,

one can tighten this lower bound by using the portfolio with the highest Sharpe ratio. If we can use means and variances of monthly returns to approximate instantaneous Sharpe ratios, then the estimates reported in Hansen and Jagannathan (1991) indicate that extreme levels of risk aversion are required to reconcile return data and data on aggregate consumption in the US. Because the market price of risk does not depend on the shape of δ , this is a puzzle that cannot be understood by modifying standard assumptions about how consumers discount utility.

It should be emphasized, though, that this does not mean that risk premia and return volatilities do not depend on the specification of the subjective discount function. The dividends of an asset are discounted by the product of $\exp(-\int_0^t \Phi(x_s)/\Gamma(x_s)ds)$ and $e_t^{-\gamma}$, and the price of the asset depends on how both these factors correlate with dividends. For an important example, consider the infinitely-lived asset that generates dividends equal to aggregate endowments. The price of this asset is simply aggregate wealth $e_t\Gamma(x_t)$. The return on aggregate wealth consist of an instantaneously risk-free dividend yield $1/\Gamma(x_t)$, and risky capital gains that arise from changes in $e_t\Gamma(x_t)$. Ito's lemma applied to $e_t\Gamma(x_t)$ therefore implies that:

$$\sigma_R(x_t)^\top = \sigma_e(x_t)^\top + \frac{D\Gamma(x_t)\sigma_x(x_t)^\top}{\Gamma(x_t)} \quad (37)$$

when R_t is taken to be the return on aggregate wealth. But the wealth-consumption ratio $\Gamma(x_t)$ is simply the expected value of $(e_{t+v}/e_t)^{1-\gamma}$ for $v > 0$, discounted by $\delta(v)$. This means that $D\Gamma(x_t)/\Gamma(x_t)$, and thus the volatility of the return on aggregate wealth and, by (34), the risk premium on aggregate wealth will depend on the shape of the subjective discount function δ .

5. PRESENT BIAS, INTEREST RATE DYNAMICS, AND VOLATILITY

Equilibrium prices depend on the shape of the subjective discount function via its impact on the effective subjective discount rate $\Phi(x_t)/\Gamma(x_t)$. We want to examine in more detail how present bias affects this discount rate, as well as the volatility of the consumption-wealth ratio $1/\Gamma(x_t)$. From (29) and (30) it is clear that the dependence of expected utility growth on the current state and the horizon is going to be important. For general endowment processes, this dependence can be quite complicated, and this makes it hard to determine the effects of changing the shape of the subjective discount function. In the following, we identify a class of endowment processes for which the interaction between utility growth and present bias can be examined analytically. We then present a parametric example for which we give a complete characterization of the effects of present bias on interest rates and consumption-wealth ratios.

5.1. Monotone Utility Dynamics

Recall from (29) and (30) that $\Phi(x)$ and $\Gamma(x)$ are integrals of the “expected utility ratio:”

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

against, respectively, $-\delta(t)$ and $\delta(t) dt$. Note that this ratio is negatively related to expected utility growth if $\gamma > 1$. Suppose now that $\{x_t\}_{t \geq 0}$ is a stationary scalar diffusion, and consider the following restriction on endowments and utility.

Condition M *The expected utility ratio satisfies:*

$$D_x G(x, t) \frac{\partial}{\partial t} \left(\frac{D_x G(x, t)}{G(x, t)} \right) \geq 0.$$

Since $D_x G(x, 0) = 0$, Condition M implies that $D_x G(x, t)$ has the same sign for all $t > 0$. Condition M therefore says that a change in the state has an impact on the log of the expected utility ratio that increases monotonically as the horizon increases. Roughly, this property can arise if a change in x_0 has an immediate impact on e_0 , small or no effects on endowments in the long-run, and effects on intermediate e_t that decline monotonically in t .

Using Condition M, we can sign the correlation between the expected utility ratio $G(x, t)$ and effective subjective discount rate $\Phi(x)/\Gamma(x)$, and we can determine the effect on $D\Gamma(x)/\Gamma(x)$ of certain parametric changes in δ . Recall that for smooth discount functions, $\Phi(x)/\Gamma(x)$ is a weighted average of the subjective discount rate $\rho(t) = -D\delta(t)/\delta(t)$, with weights $\omega(x, t)$ that are proportional to $\delta(t) G(x, t)$. We can therefore write:

$$\frac{\partial}{\partial x} \left(\frac{\Phi(x)}{\Gamma(x)} \right) = \int_0^\infty \rho(t) \left(\frac{D_x \omega(x, t)}{\omega(x, t)} \right) \omega(x, t) dt.$$

Note that the right-hand side of this equation can be interpreted as a covariance. This leads to the following lemma.

Lemma 1 *Suppose that δ is smooth, with subjective discount rates $\rho(t)$ that are decreasing in t . Then Condition M implies:*

$$D_x G(x, t) \frac{\partial}{\partial x} \left(\frac{\Phi(x)}{\Gamma(x)} \right) \leq 0.$$

We say that subjective discount rates exhibit “monotone present bias” when $\rho(t)$ is a decreasing function of t . Lemma 1 states that the effective subjective discount rate $\Phi(x)/\Gamma(x)$ and the expected utility ratio $G(x, t)$ vary with x in opposite directions when subjective discount rates exhibit monotone present bias. For example, if $\gamma > 1$

and if a high value of x implies high future endowment growth, then a high value of x will tend to imply a low value of expected utility ratio $G(x, t)$. Thus Lemma 1 implies that high expected endowment growth tends to go together with high effective subjective discount rates when $\gamma > 1$.

Next, consider varying the subjective discount function parametrically. Write δ_α for a subjective discount function indexed by some scalar parameter α , and let $\Gamma_\alpha(x)$ be the corresponding wealth-consumption ratio. Observe that one can write:

$$\frac{\partial}{\partial \alpha} \left(\frac{D\Gamma_\alpha(x)}{\Gamma_\alpha(x)} \right) = \int_0^\infty \left(\frac{D_x G(x, t)}{G(x, t)} \right) \left(\frac{D_\alpha \omega_\alpha(x, t)}{\omega_\alpha(x, t)} \right) \omega_\alpha(x, t) dt,$$

where the weights $\omega_\alpha(x, t)$ are proportional to $\delta_\alpha(t) G(x, t)$. Note again that the right-hand side of this equation can be interpreted as a covariance.

Lemma 2 *Suppose that δ_α is smooth, with subjective discount rates $\rho_\alpha(t)$ that are increasing in α for all t . Then Condition M implies:*

$$\frac{\partial}{\partial \alpha} \left| \frac{D_x \Gamma_\alpha(x)}{\Gamma_\alpha(x)} \right| \leq 0.$$

For example, if $\delta_\alpha(t) = \exp(-\alpha t)$, this simply says that geometrically discounted expected utility ratios become more sensitive to the state as one lowers the discount rate. For general subjective discount functions, this conclusion continues to hold if one lowers subjective discount rates at all horizons.

It should be emphasized that Condition M is only a sufficient condition. The conclusions of Lemmas 1 and 2 apply more generally if endowments are such that any oscillations over time in $D_x G(x, t)/G(x, t)$ are relatively small.

5.2. Parametric Examples

Consider the endowment process $e_t = \exp(\eta t - x_t)$, where x_t is a square-root process:

$$dx_t = \kappa(\mu - x_t) dt + \sigma \sqrt{x_t} dW_t, \quad (38)$$

for positive κ , μ and σ (see Feller (1951), or Cox, Ingersoll and Ross (1985)). It can be verified that Condition M is satisfied for these endowments.⁸ Since x_t is non-negative and mean-reverting, this example implies that logarithmic endowments fluctuate below some linear trend. The example can be extended to allow for a stochastic trend while preserving Condition M. For high values of x , endowments are far below trend. Expected endowment growth is then high and uncertain. For $\gamma > 1$, this implies a low expected utility ratio $G(x, t)$.

⁸This follows from the fact that $D_x G(x, t)/G(x, t) = -xA((1 + Be^{-\kappa t})^{-1} - (1 + B)^{-1})$, where $A = 2\kappa/\sigma^2$ and $B = (1 - \gamma)/(1 - \gamma - A)$.

We now consider the effects of present bias using two alternative subjective discount functions. In both cases, we adjust the subject discount function so that the risk-free rate remains unchanged on average. Very high short-run subjective discount rates have been suggested in the literature. Only if long-run subjective discount rates are correspondingly low will the implied risk-free rate be in the range observed in most data sets.

Quasi-Hyperbolic Discounting Recall that the effective subjective discount rate for quasi-hyperbolic discount function is increasing in both the long-run discount rate ρ and the instantaneous discount $1 - \delta_+$. Making consumers present-biased amounts to increasing the discount $1 - \delta_+$. By (33), this has no effect on $D\Gamma(x_t)/\Gamma(x_t)$. But to keep the risk-free rate the same on average, we have to lower ρ . Lemma 2 applied to $\delta_\rho(t) = \exp(-\rho t)$ implies that this makes the wealth-consumption ratio more volatile.

A higher value of x implies higher future endowment growth. Suppose $\gamma > 1$, so that $G(x, t)$ and therefore $\Gamma(x)$ is decreasing in x . This means that the wealth-consumption ratio and current consumption move together. The added volatility of the wealth-consumption ratio that comes about from making consumers present-biased therefore increases the risk premium on aggregate wealth (by (34) and (37)).

Consider the risk-free rate. Equation (32) implies that increasing the present bias discount $1 - \delta_+$ together with the fact that $\Gamma(x)$ is decreasing in x when $\gamma > 1$ implies that the effective subjective discount rate $\Phi(x)/\Gamma(x)$ is increasing in x . Higher future endowment growth therefore makes consumers more impatient. It follows that the effective subjective discount rate and the intertemporal substitution effect $\gamma\kappa((1 + \sigma^2/(2\kappa))x - \mu)$ are positively correlated. The precautionary savings effect is given by $-\gamma(1 + \gamma)\sigma^2x/2$, and thus negatively correlated with both the effective subjective discount rate and the intertemporal substitution effect. If the precautionary savings effect dominates the intertemporal substitution effect, then making consumers present-biased will reduce the variability of the risk-free rate.

Generalized Hyperbolic Discounting Clearly, the discount function (2) exhibits monotone present bias. Therefore Lemma 1 applies, and the effective subjective discount rate will move together with the intertemporal substitution effect if $\gamma > 1$. We therefore obtain the same co-movements among the three determinants of the risk-free rate as in the case of quasi-hyperbolic preferences.

Note that $-D\delta(t)/\delta(t)$ is increasing in ρ and ξ , and decreasing in ζ (since a higher value of ζ speeds up the transition from $\rho + \xi$ to ρ .) To make consumers more present-biased while keeping the level of the risk-free rate the same, one can simultaneously increase ξ and ζ , and lower ρ . Lemma 2 implies that the increase in ξ tends to make the wealth-consumption ratio less volatile, while the increase in ζ and the decrease in ρ will make it more volatile. This suggest an ambiguous effect of present bias.

But if we take ζ to be large and ξ on the same order as $\zeta/\ln(\zeta)$, then the hyperbolic discount function is well approximated by a quasi-hyperbolic discount function, and our results for the latter should apply. Numerical experiments confirm that this is indeed the case.

Implications Our examples suggest that if endowments are such that expected utility ratios exhibit monotone dynamics, and if subjective discount functions exhibit monotone present bias, then the effective subjective discount rate tends to move together with the intertemporal substitution term that determines the risk-free rate. This can make the risk-free rate more or less volatile, depending on the magnitude of the precautionary savings effect. At the same time, these examples indicate that the wealth-consumption ratio becomes unambiguously more volatile as a result of present bias.

In US data, aggregate consumption growth does not seem to exhibit much predictability, and this implies that the effective impatience is going to be close to constant. Nevertheless, there are two ways in which our results can potentially shed light on asset pricing puzzles. First, participation in financial markets may be limited, and the consumption processes of those who do participate may exhibit significant predictability. Second, our results continue to apply when the period utility function is replaced by $[(c_t - \underline{c}_t)^{1-\gamma} - 1]/(1 - \gamma)$, where \underline{c}_t is a subsistence or habit level of consumption. For example, the habit persistence model of Campbell and Cochrane (1999) generates a net-of-habit consumption process $c_t - \underline{c}_t$ that has stochastic properties that are very similar to those of $-x_t$, where x_t is given by (38). An interesting question for further research is whether present bias moves the implications of these models closer to the data.

6. CONCLUDING REMARKS

In an infinite-horizon exchange economy in which consumers cannot commit to future choices and in which period utilities are expected to grow at a constant rate, price and consumption data can be interpreted as resulting from the optimal choices of consumers whose subjective rates of time preference are constant. A similar observational equivalence applies for essentially arbitrary endowment processes if period utility functions happen to be logarithmic.

Inferences about the shape of subjective discount functions can be made when consumers can make irreversible commitments regarding future consumption levels, as in Laibson (1997), or when consumers face binding borrowing constraints, as in the buffer-stock savings model analyzed by Harris and Laibson (1999). In this paper, we abstract from market frictions or other commitment devices. Instead, we examine the observable implications of alternative assumptions about subjective discount

functions when there is serial dependence in utility.

An important feature of the continuous-time approximation we present is the fact that the instantaneous market price of risk does not depend on the subjective discount function of consumers. In other words, short-horizon Sharpe ratios are not affected by how consumers discount utility. Subjective rates of time preference do influence the dynamics of the instantaneous risk-free rate. In turn, this affects how the dividends of long-lived assets are discounted, and this has implications for the volatility and risk premia on such assets. For example, present bias can, under certain conditions, make aggregate wealth more volatile, even when it reduces the volatility of the risk-free rate.

For the period utility functions we consider, serial dependence in utility is the result of serial dependence in endowment growth. Serial dependence in utility can also be generated by habit persistence and by the consumption of durable goods. If preferences are homothetic, linear equilibria of the type derived in this paper can again be constructed. How these aspects of consumer preferences interact with hyperbolic subjective rates of time preference is an interesting subject for further research. Linear equilibria can also be constructed for economies in which utility is no longer time and state separable, as long as preferences remain homothetic. In Luttmer and Mariotti (2000b) we extend the work of Epstein and Zin (1989, 1991) on non-expected utility to a very general class of non-recursive homothetic preferences.

A EQUILIBRIUM

To prove Proposition 1, we proceed in two steps. We first give a precise statement and proof of the first-order condition (11). Next, we check that Conditions A and B imply the existence of an equilibrium in which there is no long-lived asset that is available in positive net supply.

Proof of the First-Order Condition (11). Given current wealth $w_t > 0$, the date- t consumer's decision problem can be written as:

$$\max_{(c_t, w_{t+1}) \in \mathbb{R}_+ \times L_{t+1}^+} \left\{ u(c_t) + \mathbb{E}_t [A_{t+1} u(w_{t+1})] : c_t + \mathbb{E}_t \left[\frac{\pi_{t+1}}{\pi_t} w_{t+1} \right] \leq w_t \right\}, \quad (39)$$

where L_{t+1}^+ is the set of nonnegative \mathcal{F}_{t+1} -measurable random variables, $A_{t+1} \in L_{t+1}^+$ is an almost surely positive random variable that is equal to $\Delta_{t+1}/\Gamma_{t+1}^{1-\gamma}$ in equilibrium, and $\pi_{t+1}/\pi_t \in L_{t+1}^+$ represents the relative price of next-period consumption.

Lemma 3 *Suppose that (39) has a solution (c_t^*, w_{t+1}^*) at which expected discounted utility is finite. Then (11) holds almost surely, i.e., $\pi_{t+1}/\pi_t = A_{t+1} Du(w_{t+1}^*)/Du(c_t^*)$.*

Proof. Since A_{t+1} is almost surely positive and the marginal utility Du is infinite at zero, c_t^* and w_{t+1}^* must be both positive at the maximum. Furthermore, given

that u is strictly increasing, the budget constraint must be binding at the maximum. Perturb w_{t+1}^* to $w_{t+1} = (1 + \varepsilon) w_{t+1}^*$ for some ε close enough to 0 and take $c_t = w_t - \mathbb{E}_t[(\pi_{t+1}/\pi_t)w_{t+1}]$. Clearly, the mapping $\varepsilon \mapsto \mathbb{E}_t[A_{t+1}u((1 + \varepsilon) w_{t+1}^*)]$ is differentiable on \mathbb{R}_{++} . The optimality of (c_t^*, w_{t+1}^*) then requires that:

$$\left. \frac{\partial}{\partial \varepsilon} \left(u \left(w_t - (1 + \varepsilon) \mathbb{E}_t \left[\frac{\pi_{t+1}}{\pi_t} w_{t+1}^* \right] \right) + \mathbb{E}_t[A_{t+1}u((1 + \varepsilon) w_{t+1}^*)] \right) \right|_{\varepsilon=0} = 0,$$

or, equivalently:

$$\mathbb{E}_t \left[\left(-Du(c_t^*) \frac{\pi_{t+1}}{\pi_t} + A_{t+1}Du(w_{t+1}^*) \right) w_{t+1}^* \right] = 0. \quad (40)$$

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

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

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 on \mathbb{R}_{++} together with the dominated convergence theorem therefore implies that the right-derivative of the mapping $\varepsilon \mapsto \mathbb{E}_t[A_{t+1}u((1 + \varepsilon \iota_B) w_{t+1}^*)]$ at $\varepsilon = 0$ is well-defined and given by $\mathbb{E}_t[A_{t+1} \iota_B w_{t+1}^* Du(w_{t+1}^*)]$. The optimality of (c_t^*, w_{t+1}^*) then requires that:

$$\mathbb{E}_t \left[\left(-Du(c_t^*) \frac{\pi_{t+1}}{\pi_t} + A_{t+1}Du(w_{t+1}^*) \right) w_{t+1}^* \iota_B \right] \leq 0. \quad (41)$$

Note that because B is an arbitrary element of \mathcal{F}_{t+1} , equations (40) and (41) imply that $(-Du(c_t^*) \pi_{t+1}/\pi_t + A_{t+1}Du(w_{t+1}^*)) w_{t+1}^* = 0$ almost surely. Since w_{t+1}^* is always positive, the result follows. \blacksquare

Proof of Proposition 1. By assumption, Γ_t , Δ_t and $A_{t+1} = \Delta_{t+1}/\Gamma_{t+1}^{1-\gamma}$ are well-defined and almost surely finite for any $t \geq 0$. By construction, the consumption and wealth choices $(\phi_t w_t, \psi_{t+1} w_t)$ satisfy the first-order condition (11) at market-clearing prices $\pi_{t+1}/\pi_t = (\Delta_{t+1}/\Gamma_{t+1})(e_{t+1}/e_t)^{-\gamma}$, where ϕ_t and ψ_{t+1} are given by (17). It then follows from the concavity of the objective function in (39) that $(\phi_t w_t, \psi_{t+1} w_t)$ is optimal from the perspective of the date- t consumer. In the text, we have only imposed that consumption choices grow at the same rate as aggregate endowments. Goods market clearing at all dates is therefore equivalent to the condition $w_0 = e_0 \Gamma_0$. Iteration on (8) together with goods market clearing implies that:

$$\pi_t w_t = \mathbb{E}_t \left[\sum_{n=0}^{\infty} \pi_{t+n} e_{t+n} \right] + \lim_{N \rightarrow \infty} \mathbb{E}_t [\pi_{t+N} e_{t+N} \Gamma_{t+N}] \quad (42)$$

for all $t \geq 0$. Note that Γ_t and Δ_{t+1} are related via $\Gamma_t = 1 + \mathbb{E}_t [(e_{t+1}/e_t)^{1-\gamma} \Delta_{t+1}]$. Iterating on this identity, using (18) to eliminate Δ_{t+1} , one can verify that requiring (42) at date 0 is in fact equivalent to the goods market clearing condition $w_0 = e_0 \Gamma_0$. Together with the fact that stock prices must be non-negative, (6) implies that $\pi_t s_t = \mathbb{E}_t [\sum_{s=t+1}^{\infty} \pi_s d_s] + \pi_t z_t$, where $\{z_t\}_{t=0}^{\infty}$ is a non-negative sequence of random variables that satisfies:

$$\pi_t z_t = \mathbb{E}_t [\pi_{t+1} z_{t+1}] \quad (43)$$

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.” The stock market and markets for contingent claims clear when $k_t = k_{-1}$ and $b_{t+1} = 0$ for all t . This implies that, in equilibrium, $\pi_t w_t = \mathbb{E}_t [\sum_{s=t}^{\infty} \pi_s e_s] + \pi_t z_t k_{t-1}$. It therefore follows from (7) and (42) that we have an equilibrium if and only if for some sequence $\{z_t\}_{t=0}^{\infty}$:

$$\pi_t z_t k_{-1} = \lim_{N \rightarrow \infty} \mathbb{E}_t [\pi_{t+N} e_{t+N} \Gamma_{t+N}]. \quad (44)$$

If $k_{-1} > 0$, then (44) uniquely determines the bubble process $\{z_t\}_{t=0}^{\infty}$. Note that the right-hand side of (44) is by construction a non-negative martingale, and thus (43) holds. If instead $k_{-1} = 0$, then any non-negative $\{z_t\}_{t=0}^{\infty}$ that satisfies (43) will be consistent with equilibrium provided the right-hand side of (44) is zero. Using (18), one can see that the right-hand side of (44) is zero if and only if (19) holds. Thus, to conclude the proof, we need only to check that Conditions A and B imply (19). Consider first Condition A. Using (16) one can write:

$$\begin{aligned} \mathbb{E}_0 \left[\prod_{t=1}^T \frac{\Delta_t}{\Gamma_{t-1}} \left(\frac{e_t}{e_{t-1}} \right)^{1-\gamma} \right] &= \mathbb{E}_0 \left[\prod_{t=1}^{T-1} \frac{\Delta_t}{\Gamma_{t-1}} \left(\frac{e_t}{e_{t-1}} \right)^{1-\gamma} \mathbb{E}_{T-1} \left[\frac{\Delta_T}{\Gamma_{T-1}} \left(\frac{e_T}{e_{T-1}} \right)^{1-\gamma} \right] \right] \\ &= \mathbb{E}_0 \left[\prod_{t=1}^{T-1} \frac{\Delta_t}{\Gamma_{t-1}} \left(\frac{e_t}{e_{t-1}} \right)^{1-\gamma} \left(1 - \frac{1}{\Gamma_{T-1}} \right) \right] \\ &\leq \mathbb{E}_0 \left[\prod_{t=1}^{T-1} \frac{\Delta_t}{\Gamma_{t-1}} \left(\frac{e_t}{e_{t-1}} \right)^{1-\gamma} \right] (1 - \kappa) \\ &\quad \vdots \\ &\leq (1 - \kappa)^T, \end{aligned}$$

from which (19) follows. Under Condition B we can write $\Delta_t/\Gamma_t \leq \beta$ for any $t \geq 0$ since Δ_t/Γ_t is an average of $\delta_{n+1}/\delta_n \leq \beta$. Also, $\delta_n \leq \beta^n$ for any $n \geq 0$ and therefore: $\Gamma_t \leq \mathbb{E}_t [\sum_{n=0}^{\infty} \beta^n (e_{t+n}/e_t)^{1-\gamma}]$. This yields:

$$\mathbb{E}_0 \left[\left(\prod_{t=1}^T \frac{\Delta_t}{\Gamma_t} \left(\frac{e_t}{e_{t-1}} \right)^{1-\gamma} \right) \Gamma_T \right] \leq \mathbb{E}_0 \left[\beta^T \left(\frac{e_T}{e_0} \right)^{1-\gamma} \Gamma_T \right] \leq \mathbb{E}_0 \left[\sum_{n=T}^{\infty} \beta^n \left(\frac{e_n}{e_0} \right)^{1-\gamma} \right].$$

Condition B implies that the right-hand side of this inequality converges to zero as T goes to infinity. Hence the result. \blacksquare

B CONTINUOUS-TIME APPROXIMATION

In the following, $\{x_t\}_{t \geq 0}$ denotes a continuous-time Markov process defined on the same probability space $(\Omega, \mathcal{F}, \mathbb{P})$ 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, for any $x \in X$ and $t \geq 0$, $G(x, t)$ is defined as in the text. The following assumptions will be maintained in the remainder of this Appendix.

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

Assumption 2 *There exists a function $M : X \rightarrow \mathbb{R}_+$ 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 δ is non-increasing and integrable over \mathbb{R}_+ , $\lim_{t \rightarrow \infty} \delta(t) = 0$. Thus δ 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)). By Fubini's theorem, we may rewrite (29) and (30) as $\Phi(x) = \int_{[0, \infty)} G(x, t) d\mu_\delta(t)$ and $\Gamma(x) = \int_0^\infty \delta(t) G(x, t) dt$, respectively, 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 $\delta > 0$ on some set of positive Lebesgue measure. Assumption 3 ensures that Γ is continuous on X . It follows that Γ is bounded away from zero on compact subsets of X . Also, Φ is positive and bounded above by M . For any $\tau > 0$ and $x \in X$, define $\Gamma(x, \tau) = \sum_{n=0}^\infty \delta(n\tau) G(x, n\tau)$ and $\Delta(x, \tau) = \sum_{n=0}^\infty \delta((n+1)\tau) G(x, n\tau)$, 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 4 *For any $x \in X$, $\lim_{\tau \rightarrow 0^+} \tau \Gamma(x, \tau) = \Gamma(x)$.*

Proof. For any $\tau > 0$, $t \geq 0$, and $x \in X$, let $G_\tau(x, t) = G(x, \tau \lceil t/\tau \rceil)$ and $\delta_\tau(t) = \delta(\tau \lceil t/\tau \rceil)$, where $\lceil t/\tau \rceil$ 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 1-2 imply that $\delta_\tau(t) G_\tau(x, t) \leq M(x)$ for any $t \in [0, 1)$ and $\tau > 0$, as well as $\delta_\tau(t) G_\tau(x, t) \leq \delta(t-1) M(x)$ for all $t \geq 1$ and $\tau \in (0, 1]$. Thus $\delta_\tau G_\tau(x, \cdot)$ is

dominated for all $\tau \in (0, 1]$ by an integrable function of t . Moreover, the continuity of $G(x, \cdot)$ implies that $\lim_{\tau \rightarrow 0^+} G_\tau(x, t) = G(x, t)$ and the left-continuity of δ implies that $\lim_{\tau \rightarrow 0^+} \delta_\tau(t) = \delta(t)$. The result then follows by Lebesgue's dominated convergence theorem. \blacksquare

Lemma 5 *As τ goes to 0^+ , $\tau \Gamma(\cdot, \tau)$ converges uniformly to Γ on every compact subset of X .*

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

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

for any $x, y \in K$ such that $\|x - y\| < \eta$, where the second inequality follows from the fact that $\int_0^\infty \delta_\tau(v) dv = \sum_{n=0}^\infty \delta(n\tau)\tau \leq \tau + \int_0^\infty \delta(v) dv$ as δ 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 $\tau \Gamma(\cdot, \tau)$ to Γ as τ goes to 0^+ is therefore uniform on K . \blacksquare

Lemma 6 *For any $x \in X$, $\lim_{\tau \rightarrow 0^+} \Phi(x, \tau) = \Phi(x)$.*

Proof. Using the notation of Lemma 4, 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, $\lim_{\tau \rightarrow 0^+} G_\tau(x, t) = G(x, t)$ and $G_\tau(x, t) \leq M(x)$ for any $\tau > 0$, $t \geq 0$ and $x \in X$. Since μ_δ is a probability measure, the result follows immediately from Lebesgue's dominated convergence theorem. \blacksquare

We can now state and prove our main result.

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

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

$$\pi_t^\omega(\tau) = \exp\left(\sum_{n=1}^{\lfloor t/\tau \rfloor} \ln\left(\frac{\Delta(x_{n\tau}^\omega, \tau)}{\Gamma(x_{n\tau}^\omega, \tau)}\right)\right)(e_t^\omega)^{-\gamma} = \exp\left(\int_0^t \iota_{t,\tau}(v) Q(x_{\tau \lceil v/\tau \rceil}^\omega, \tau) dv\right)(e_t^\omega)^{-\gamma},$$

where $\iota_{t,\tau}(v) = 1$ if $v/\tau \in [1, \lceil t/\tau \rceil]$ 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 $\pi_t^\omega(\tau) = \pi_{\tau \lceil t/\tau \rceil}^\omega(\tau)$. By Lemmas 4 and 6, $\lim_{\tau \rightarrow 0^+} Q(x, \tau) = -\Phi(x)/\Gamma(x)$ for any $x \in X$. Next, recall that $\lim_{\tau \rightarrow 0^+} \tau \Gamma(x, \tau) = \Gamma(x)$ uniformly on compact subsets K of X . Hence, since x_v^ω is a continuous function of v , $\lim_{\tau \rightarrow 0^+} \tau \Gamma(x_{\tau \lceil v/\tau \rceil}^\omega, \tau) = \Gamma(x_v^\omega)$ uniformly on $[0, t]$. Therefore $\tau \Gamma(x_{\tau \lceil v/\tau \rceil}^\omega, \tau)$ 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 $Q(x_{\tau \lceil v/\tau \rceil}^\omega, \tau)$ 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 almost every $\omega \in \Omega$. ■

REFERENCES

- [1] Ainslie, G. (1975): "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control," *Psychological Bulletin*, Vol. 82, No. 4 (July), 463-496.
- [2] Ainslie, G. (1992): *Picoeconomics*, Cambridge, Cambridge University Press.
- [3] Barro, R.J. (1999): "Ramsey Meets Laibson in the Neoclassical Growth Model," *Quarterly Journal of Economics*, Vol. 114, No. 4 (November), 1125-1152.
- [4] Benabou, R., and J. Tirole (2002): "Self-Confidence and Personal Motivation," *Quarterly Journal of Economics*, Vol. 117, No. 3 (August), 871-915.
- [5] Breeden, D. (1979): "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," *Journal of Financial Economics*, Vol. 7, No. 3 (September), 265-296.
- [6] Brock, W.A. (1979): "An Integration of Stochastic Growth Theory and the Theory of Finance, Part I: The Growth Model," in *General Equilibrium, Growth and Trade*, edited by J. Green and J.A. Scheinkman. New York, Academic Press.
- [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. 107, No. 2 (April), 205-251.
- [8] Carrillo, J.D., and T. Mariotti (2000): "Strategic Ignorance as a Self-Disciplining Device," *Review of Economic Studies*, Vol. 67, No. 3 (July), 529-544.

- [9] Chung, S.H., and R.J. Herrnstein (1967): "Choice and Delay of Reinforcement," *Journal of the Experimental Analysis of Behavior*, Vol. 10, No. 1 (January), 67-74.
- [10] Cox, J.C., J.E. Ingersoll, and S.A. Ross (1985): "A Theory of the Term Structure of Interest Rates," *Econometrica*, Vol. 53, No. 2 (March), 385-406.
- [11] Epstein, L.G. and S.E. Zin (1989): "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, Vol. 57, No. 4 (July), 937-969.
- [12] Epstein, L.G. and S.E. Zin (1991): "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis," *Journal of Political Economy*, Vol. 9, No. 2 (April), 263-286.
- [13] Feller, W. (1951): "Two Singular Diffusion Problems," *Annals of Mathematics*, Vol. 54, No. 1 (July), 173-182.
- [14] Hansen, L.P. and R. Jagannathan (1991): "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, Vol. 99, No. 2 (April), 225-262.
- [15] Harris, C., and D. Laibson (2001): "Dynamic Choices of Hyperbolic Consumers," *Econometrica*, Vol. 69, No. 4 (July), 935-957.
- [16] Kirby, K.N., and R.J. Herrnstein (1995): "Preference Reversals Due to Myopic Discounting of Delayed Rewards," *Psychological Science*, Vol. 6, No. 2 (March), 83-89.
- [17] Krusell, P., and A. Smith (1999): "Consumption and Savings Decisions with Quasi-Geometric Discounting," mimeo, Carnegie Mellon University.
- [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. 112, No. 2 (May), 443-477.
- [21] Lang, S. (1993): *Real and Functional Analysis*, New York, Springer-Verlag.
- [22] Loewenstein, G., and D. Prelec (1992): "Anomalies in Intertemporal Choice: Evidence and an Interpretation," *Quarterly Journal of Economics*, Vol. 107, No. 3 (April), 573-597.

- [23] Luttmer, E.G.J. and T. Mariotti (2000a): "Subjective Discount Factors," CEPR Discussion Paper 2503.
- [24] Luttmer, E.G.J. and T. Mariotti (2000b): "Non-Recursive Homothetic Preferences," in progress.
- [25] Obstfeld, M. and K. Rogoff (1983): "Speculative Hyperinflations in Maximizing Models: Can We Rule Them Out?," *Journal of Political Economy*, Vol. 91, No. 4 (August), 675-687.
- [26] O'Donoghue, T., and M. Rabin (1999): "Doing it Now or Later," *American Economic Review*, Vol. 89, No. 1 (March), 103-124.
- [27] O'Donoghue, T., and M. Rabin (2001): "Choice and Procrastination," *Quarterly Journal of Economics*, Vol. 116, No. 1 (February), 121-160.
- [28] Peleg, B., and M. Yaari (1973): "On the Existence of a Consistent Course of Actions when Tastes are Changing," *Review of Economic Studies*, Vol. 40, No. 3 (July), 391-401.
- [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] Pollak, R.A. (1968): "Consistent Planning," *Review of Economic Studies*, Vol. 35, No. 2 (April), 201-208.
- [31] Scheinkman, J.A. (1977): "Notes on Asset Pricing," mimeo, University of Chicago.
- [32] Strotz, R.H. (1956): "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, Vol. 23, No. 3, 165-180.
- [33] Thaler, R. (1981): "Some Empirical Evidence on Dynamic Inconsistency," *Economics Letters*, Vol. 8, No. 3, 201-207.