ELSEVIER European Economic Review 42 (1998) 861–871

EUROPEAN
ECONOMIC
REVIEW

Life-cycle consumption and hyperbolic discount functions

David Laibson[a,b,*]

[a] *Harvard University, Cambridge, MA 02138, USA*
[b] *National Bureau of Economic Research, Cambridge, MA 02139, USA*

Abstract

Studies of animal and human behavior suggest that discount functions are approximately hyperbolic (Ainslie, Picoeconomics, 1992). Hyperbolic models explain a wide range of empirical anomalies including (1) missing precautionary savings effects, (2) incongruence between the elasticity of intertemporal substitution and the inverse of the coefficient of relative risk aversion, (3) consumption discontinuities at retirement, (4) variation in patience over the life cycle, (5) consumer self-reports of 'undersaving', (6) disproportionate retirement accumulation in illiquid assets, (7) asset-specific marginal propensities to consume, and (8) declining national savings rates in developed countries. Hyperbolic models also provide a new framework for analyzing the welfare gains that come from pro-saving policies like social security and tax-deferred retirement accounts.

JEL classification: D91; E21

Keywords: Hyperbolic discounting; Dynamic inconsistency; Consumption; Savings; Time preference; Buffer stock; Euler equation

---

[*] Correspondence address: Harvard University, Cambridge, MA 02138, USA. E-mail: dlaibson@harvard.edu.

PII S0014-2921(97)00132-3

1. Introduction

Consumers seem to be of two minds about intertemporal consumption decisions. When sacrifices are distant, patience predominates. 'I want/plan/intend to start exercising next month'. But, when next month actually roles around, the designated sacrifice is often avoided. Think about the success rates of our resolutions to *punctually* revise papers, reply to correspondence, complete referee reports, etc.

Such patterns of 'preference reversals', have been studied by psychologists since the 1960s.[1] To model these preferences, psychologists have adopted the hyperbolic discount function: events τ periods away are discounted with factor $(1+\alpha\tau)^{-\gamma/\alpha}$.[2] Such discount functions imply a monotonically falling discount rate. This discount structure sets up a conflict between today's preferences and the preferences which will be held in the future, implying that preferences are dynamically inconsistent. From today's perspective, the discount rate between two far-off periods, t and $t+1$, is a long-term low discount rate. However, from the time t perspective, the discount rate between t and $t+1$ is a short-term high discount rate. Over a dozen papers have tested this discount structure against the standard exponential framework. The hyperbolic curve has won every competition.[3]

Such preferences have economically important implications for consumers' life-cycle savings decisions. This paper reviews these predicted effects, and relates them to known stylized facts about actual consumption choices. Section 2 formally describes the theoretical framework. Section 3 discusses the positive implications of this framework. Section 4 discusses the normative implications, including policy issues, and discusses directions for future research.

2. Model

Following Strotz (1956) and Pollak (1968), I model an individual as a composite of autonomous temporal selves. These selves interact as players in a finite-horizon dynamic game. I adopt subgame perfection as the equilibrium concept.

The selves are indexed by their respective periods of control ($t = 0, 1, 2, \dots, T$), over a consumption decision. During its period of control, self t receives stochastic labor income y_t, and chooses a consumption level c_t. Cash-on-hand,

---

[1] See Ainslie (1992) for a review.

[2] For an axiomatic derivation and short history of this functional form, see Loewenstein and Prelec (1992).

[3] See Ainslie (1992) and Kirby (1997) for recent reviews.

X_t, evolves according to

$$X_t = (X_{t-1} - c_{t-1})R + y_t, \tag{1}$$

where R is an exogenous real after-tax gross interest rate.

I assume that consumption must be less than or equal to cash on hand in every period:

$$c_t \le X_t \quad \forall t. \tag{2}$$

This implies that consumers may not borrow.[4]

In this game early selves do not have a technology to commit the actions of later selves. But, some real-world institutions – like the pension system in the US – do serve as effective commitment devices, protecting assets from pre-retirement consumption. In Section 3, I discuss such extensions.

It only remains to specify the payoffs of the different selves, or 'players' of this game. Player t receives payoff

$$U_t(c_0, c_1, \dots, c_T) = E_t\left[u(c_t) + \beta \sum_{i=1}^{T-t} \delta^i u(c_{t+i}) \right], \tag{3}$$

where δ and β are discount parameters, and $u(c)$ is a member of the class of CRRA utility functions with relative risk aversion coefficient $\rho \in (0, \infty)$:

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

The model described above unifies two literatures. First, the credit market assumptions are taken from standard liquidity constraint and buffer stock models (e.g., Deaton, 1991; Carroll, 1992). Second, the discount structure captures the qualitative properties of the hyperbolic discount function defined above.

Hyperbolic discount *functions* imply discount *rates* that decline with the passage of time. When $0 < \beta < 1$ the discount structure in Eq. (3) mimics this qualitative property, while maintaining most of the analytical tractibility of the exponential case. I call the discount structure in Eq. (3) 'quasi-hyperbolic'; it implies a discrete-time discount function with values $\{1, \beta\delta, \beta\delta^2, \beta\delta^3, \dots\}$. Fig. 1 graphs the exponential discount function (assuming $\delta = 0.97$), the generalized hyperbolic discount function (assuming $\alpha = 10^5$ and $\gamma = 5 \times 10^3$), and the quasi-hyperbolic discount function ($\beta = 0.6$ and $\delta = 0.99$). The points of the

---

[4] Eq. (2) is not binding if (1) the infimum of the support of the income process is zero and (2) all debts must be repaid by the end of life: $X_{T+1} \ge 0$. Conditions (1) and (2) are common assumptions in buffer stock models, like Carroll (1992).

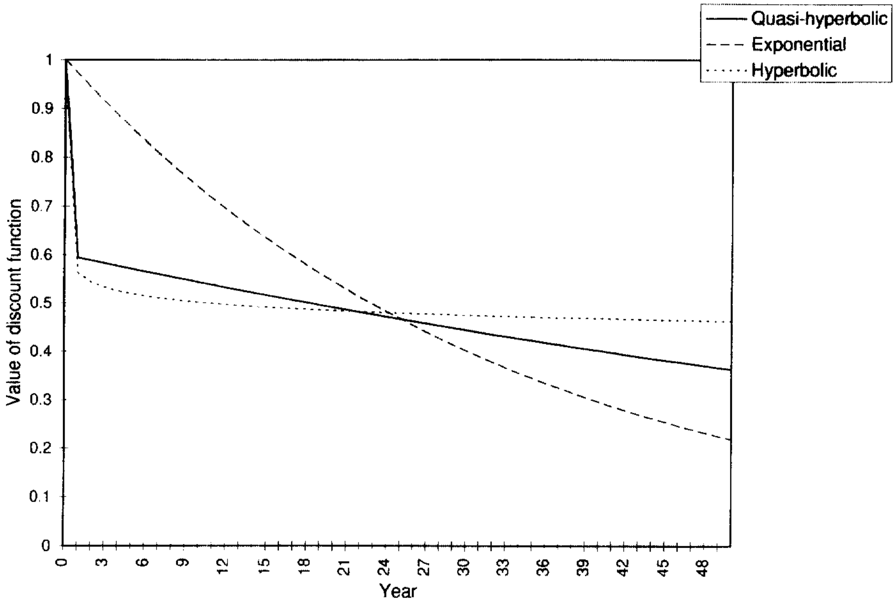

Fig. 1. Discount functions.

discrete-time quasi-hyperbolic function have been connected to generate a smooth curve.

The preferences in Eq. (3) were first analyzed by Phelps and Pollak (1968).[5] Their game is one of imperfect intergenerational altruism, so the players are non-overlapping generations of a dynasty. I assume that the different players are temporally distinct selves of a single person. My set up also differs because I assume the horizon is finite and income is stochastic.[6]

3. Positive analysis

The hyperbolic model predicts a wide range of well-documented empirical regularities. The Euler equation for the hyperbolic economy provides a useful starting point for this discussion.

---

[5] Akerlof (1991) analyzes similar preferences with $\delta = 1$. O'Donoghue and Rabin (1997) also analyze the Phelps and Pollak preferences. Both of these papers consider the problem of a decision-maker who must decide when to do a single task.

[6] Because Phelps and Pollak assume an infinite horizon, their framework admits a continuum of equilibria (Laibson, 1994).

In standard exponential consumption problems (i.e., $\beta = 1$ above), the Euler equation characterizes the equilibrium path:

$$u'(c_t) = E_t R \delta u'(c_{t+1}).$$

In the hyperbolic economy (i.e., $\beta < 1$), this equation generalizes to[7]

$$u'(c_t) = E_t R\left[\left(\frac{\partial c_{t+1}(X)}{\partial X}\right)\beta\delta + \left(1 - \frac{\partial c_{t+1}(X)}{\partial X}\right)\delta\right]u'(c_{t+1}). \tag{4}$$

This generalized Euler equation mirrors the standard Euler equation except for the new discounting term, which I will refer to as the *effective discount factor*:

$$\left[\left(\frac{\partial c_{t+1}(X)}{\partial X}\right)\beta\delta + \left(1 - \frac{\partial c_{t+1}(X)}{\partial X}\right)\delta\right].$$

The effective discount factor varies linearly with next period's marginal propensity to consume (MPC). When $\beta = 1$ this bracketed term is equal to δ (the exponential discounting case). When $\beta < 1$ this bracketed term is a weighted average of the one-period discount factor $\beta\delta$ and the discount factor that applies in all future periods δ. The respective weights are tomorrow's marginal propensity to consume $((\partial c_{t+1}(X))/\partial X)$, and $(1 - (\partial c_{t+1}(X))/\partial X)$.

Intuitively, the effective discount factor is constructed by weighting actual future discount factors according to the marginal propensity to consume in those respective future periods. The effective discount factor weights the discount factor that applies between today and tomorrow – $\beta\delta$ – by the fraction of marginal savings that get consumed tomorrow; all remaining weight is placed on δ – the discount factor that applies between all other future periods. When the MPC is close to one (zero), most weight is placed on $\beta\delta$ (δ) implying a low (high) effective discount factor. Life-cycle variation in the MPC, and hence in the effective discount factor, drives much of the analysis below.

The hyperbolic model replicates the general properties of the buffer stock framework, notably consumption–income tracking. The logic of the buffer stock model carries over to the hyperbolic case: consumers build up small wealth stocks to buffer high-frequency income shocks, but those consumers are

---

[7] This generalization is discussed in Laibson (1997b). The derivation of the generalized Euler equation requires smooth equilibrium consumption functions, a property that is conjectured to arise when the distribution of labor income is made sufficiently smooth and variable. The income variability serves to effectively convexify the continuation payoff functions in this game. The conjecture on the smoothness of the consumption functions has been validated with numerical simulations, although it has not been analytically proven. The generalized Euler equation has also been validated with numerical simulations. For the case of no labor income uncertainty and perfect credit markets, the generalized Euler equation always holds (Laibson, 1996).

insufficiently patient to undertake low-frequency life-time smoothing. However, the hyperbolic model goes beyond the buffer stock predictions, explaining a wide range of additional anomalous empirical regularities.

3.1. *Missing precautionary savings effects*

Since the work of Hall (1978), it has become standard practice to work with linearized versions of the Euler equation. Consider the second-order linearization of the standard Euler equation:[8]

$$\Delta \ln c_{t+1} = \text{constant} + \frac{\rho}{2}\,\mathrm{var}_t\,(\Delta \ln c_{t+1}) + \varepsilon_{t+1}. \tag{5}$$

Empirical estimates of this equation yield an anomalously low value for the coefficient on the conditional variance term (Dynan, 1993). Hence, in the data high expected consumption variability does not predict a high rate of consumption growth.

However, this is exactly what one would expect in the hyperbolic economy. In buffer stock models – exponential or hyperbolic – a high conditional variance of consumption growth is usually associated with a high expected value of the one-period ahead marginal propensity to consume. In hyperbolic models, a high expected MPC implies a low effective discount factor in the generalized Euler equation. This yields a low expected slope of the consumption path, offsetting the precautionary savings effect. Simulations reported in Laibson (1997b) confirm this intuition.

3.2. $EIS < \frac{1}{\rho}$

In the standard exponential economy, the elasticity of intertemporal substitution (EIS) is equal to the inverse of the coefficient of relative risk aversion. For exposition, consider the case of no income uncertainty. Manipulating the standard Euler Equation reveals that

$$\mathrm{EIS} \equiv \frac{\partial \log\left((c_{t+1})/c_t\right)}{\partial r} = \frac{1}{\rho}. \tag{6}$$

This well-known theoretical relationship is contradicted by the available empirical evidence (e.g., Hall, 1988; Carroll and Summers, 1991). Most empirical analysis suggests that the EIS is smaller than the inverse of the coefficient of relative risk aversion.

---

[8] This equation can be derived from the classical Euler equation by assuming that $\Delta \ln c_{t+1}$ is conditionally normally distributed.

The hyperbolic economy predicts this empirical regularity. Using the generalized Euler equation, the EIS is equal to

$$\frac{\partial \log((c_{t+1})/c_t)}{\partial r} = \frac{1}{\rho} + \frac{1}{\rho}\left[\frac{1}{((\partial c_{t+1}(X))/\partial X)(\beta - 1) + 1}\right] \times \frac{\partial((\partial c_{t+1}(X))/\partial X)}{\partial r}(\beta - 1). \tag{7}$$

Note that when $\beta = 1$ the EIS is equal to $1/\rho$, which is the standard case described above. However, when $\beta < 1$ the one-to-one link between the elasticity of substitution and ρ is broken. When $\rho > 1$,

$$\frac{\partial((\partial c_{t+1}(X))/\partial X)}{\partial r} < 0,$$

implying that EIS $< 1/\rho$.

For intuition, note that when the coefficient of relative risk aversion is greater than unity, the income effect dominates the substitution effect. So increases in the interest rate lead to a higher consumption rate and MPC, lowering the effective discount factor in the generalized Euler equation. This change in the effective discount factor implies a lower rate of consumption growth, partially offsetting the effect of a higher interest rate. Laibson (1996) evaluates the magnitude of these effects. For example, when $\rho = 5$, $\delta = 0.99$, and $\beta = 0.5$, then EIS $= 0.11$.

3.3. Consumption discontinuities at retirement

Numerous papers have documented a sharp drop in consumption at retirement, particularly among low wealth households.[9] For example, Bernheim et al. (1997) report that US households in the bottom wealth quartile experience a 30% decline in consumption immediately following retirement. Bernheim et al. demonstrate that these changes cannot be explained by work-related expenses or by substitution between consumption and leisure.

This drop is predicted by the hyperbolic model. Consider a household with relatively low wealth stocks and a high marginal propensity to consume out of cash on hand. Consider the extreme case in which the MPC is close to unity, implying that the effective discount factor is approximately equal to $\beta\delta = 0.6$. Hence, the equilibrium consumption path will drop sharply at retirement (assuming that income drops too), even if the consumer is not strictly liquidity constrained.

---

[9] See Bernheim et al. (1997) for a review.

3.4. Patience heterogeneity: Covariation with wealth and age

Age, income, and wealth, correlate with measured levels of patience (Green et al., 1996). The hyperbolic model predicts that this pattern will arise endogenously, *even when all consumers have identical deep preference parameters.* This effect works through life-cycle variation in the MPC. In buffer stock models, young consumers who expect rising income paths, and/or consumers with low levels of cash on hand, are most likely to have high MPCs, implying a low effective discount factor. By contrast, older consumers who can expect falling income paths, and/or consumers with high levels of cash on hand are most likely to have low MPCs, implying a high effective discount factor. Hence, the effective discount factor varies over the life cycle and is positively correlated with age and financial wealth. Laibson (1997b) demonstrates these relationships in a simulation framework.

3.5. Undersaving

The typical baby boomer household saves 5% of its household income, but has a target saving rate of 15% (Bernheim, 1995). Standard consumption models cannot explain systematic gaps between intentions and actions. However, the hyperbolic model predicts such gaps (Laibson, 1996). Early selves want later selves to act patiently by choosing a high savings rate. But, later selves want to splurge during their period of control. To formally analyze the gap between aspirations and actions in the hyperbolic economy, Laibson (1996) compares equilibrium savings rates to savings rates which would be chosen if early selves could perfectly commit the actions of later selves. The calibrated hyperbolic model predicts gaps of equal magnitude to the gaps reported in consumer surveys.

3.6. Accumulation in illiquid assets

Perfect commitment – proposed as a benchmark above – is not possible in practice. But, partial commitment can be achieved through a host of illiquid assets, like defined benefit pensions, defined contribution pensions (e.g., 401(k)'s), home equity, and public retirement income programs (e.g., social security). Introducing these instruments to the hyperbolic model enables the consumer to partially protect accumulated assets from unplanned splurges, thereby increasing the slope of the equilibrium consumption path (Laibson, 1996, 1997a). Hyperbolic agents choose to create self-imposed liquidity constraints, explaining the anomalously high proportion of the household portfolio that is held in illiquid form and the finding that even precautionary savings are held in illiquid assets (Carroll et al., 1997).

3.7. *Asset-specific MPCs*

Thaler (1990) argues that consumers have different MPCs for different categories of assets. For example, he presents evidence that an unexpected increase in the value of an equity portfolio will have a negligible effect on consumption, while an unexpected job-related bonus will be immediately consumed. The hyperbolic model, augmented with partially illiquid assets, predicts such differences. Hyperbolic consumers choose to invest most of their assets in relatively illiquid instruments, effectively creating endogenous liquidity constraints. In this setting, a windfall to the consumer's liquid wealth will generate an immediate splurge, whereas a windfall to the illiquid portfolios will be saved, even if it only takes a period to move the money out of the illiquid instrument. The current self never wants to initiate that transformation process – from illiquidity to liquidity – because the resulting splurge will not be instantaneous. Laibson (1997a) formalizes these effects in a two-asset model, including an asset whose sale must be initiated one period before the sale proceeds are received.

3.8. *Declining national savings rates*

Industrialized countries have experienced a decline in their net national savings rates since the 1970s. At the same time, these economies have experienced a rapid expansion in their credit markets. In a hyperbolic economy, the introduction of instantaneous consumer credit – which occurred in most industrial economies in the 1980s[10] – eliminates the commitment properties of illiquid assets, resulting in lower equilibrium savings rates. Laibson (1997a) analyzes the equilibrium path of a closed economy in which instantaneous credit has just been introduced. In the new steady state the capital–output ratio falls by 10% and the real interest rate rises by 1.3 percentage points,[11] values which mirror the US experience during this period. Most importantly, the hyperbolic model predicts that the introduction of instantaneous credit lowers welfare for *all* selves, even taking into account the higher consumption that occurs at the beginning of the transition path. Such normative issues are the focus of the next section.

4. Normative analysis and future research

The hyperbolic model provides a framework for answering a broad range of normative questions. First, the hyperbolic model enables economists to

---

[10] See the discussion in Laibson (1997a).

[11] This assumes a β value of 0.6.

formalize and quantify the welfare costs of undersaving. In the hyperbolic model, undersaving occurs if commitment is not available. Laibson (1996), shows that the welfare cost of inadequate commitment is significant – equal to at least $\frac{9}{10}$ of 1 yr of output. These are Pareto-losses, in the sense that all current and future selves are worse off when commitment is not available.

The existence of such large welfare losses raises questions of policy design. What policies make desirable commitments easier to implement? I have already noted that illiquidity can sometimes be welfare improving. However, illiquidity cuts both ways. For example, the US social security system compels consumers to 'save' for retirement, but does not allow those consumers to use their accumulating claims to buffer income shocks during their working life. Optimal policies would create retirement instruments which trade off the need for commitment with the need for flexibility.

Future work should evaluate these tradeoffs using simulation models which capture life-cycle uncertainty arising from shocks to consumers' income, health, etc. Do existing instruments, like 401(k)'s, optimally address these tradeoffs? For example, should workers be allowed to borrow from their retirement plans? What are the optimal penalties for early withdrawal? Should consumers be given the choice to participate in the social security system?

The hyperbolic model can be readily calibrated and used to evaluate the behavior of sophisticated consumers who have to struggle to master their propensities to overspend. The hyperbolic model can provide quantitative answers to the policy questions identified above. However, the model is only easy to apply because it side steps many important and conceptually difficult consumer phenomena, including mental accounts (Thaler and Shefrin, 1981), bounded rationality (Bernheim, 1995), and impulsivity (Loewenstein, 1996). The hyperbolic model will undoubtedly be improved when it incorporates these central behavioral features, sacrificing parsimony for realism.

Acknowledgements

This work has been supported financially by the National Science Foundation (SBR-95-10985) and the MacArthur Foundation.

References

Ainslie, G.W., 1992. Picoeconomics. Cambridge University Press, Cambridge.

Akerlof, G.A., 1991. Procrastination and obedience. American Economic Review, AEA Papers and Proceedings, pp 1–19.

Bernheim, B.D., 1995. Do households appreciate their financial vulnerabilites? an analysis of actions, perceptions, and public policy. Tax Policy and Economic Growth. American Council for Capital Formation, Washington D.C., pp 1–30.

Bernheim, B.D., Skinner, J., Weinberg, S., 1997. What accounts for the variation in retirement saving among US households? Mimeo., Stanford University.

Carroll, C.D., 1992. The buffer stock theory of saving: some macroeconomic evidence. Brookings Papers on Economic Activity (2) 61–156.

Carroll, C.D., Summers, L.H., 1991. Consumption growth parallels income growth: some new evidence. In: Bernheim, B.D., Shoven, J.B. (Eds.), National Saving and Economic Performance. Chicago University Press for NBER, Chicago, IL.

Carroll, C.D., Dynan, K., Krane, S., 1977. Unemployment risk and precautionary wealth: evidence from households' balance sheets. Mimeo.

Deaton, A., 1991. Saving and liquidity constraints. Econometrica 59, 1221–1248.

Dynan, K.E., 1993. How prudent are consumers? Journal of Political Economy 101 (6), 1104–1113.

Green, L., Myerson, J., Lichtman, D., Rosen, S., Fry, A., 1996. Temporal discounting in choise between delayed rewards: the role of age and income. Psychology and Aging 11 (1), 79–84.

Hall, R.E., 1978. Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence. Journal of Political Economy 96, 971–987.

Hall, R.E., 1988. Intertemporal substitution in consumption. Journal of Political Economy 96, 339–357.

Kirby, K.N., 1997. Bidding on the future: evidence against normative discounting of delayed rewards. Journal of Experimental Psychology (General) 126 (1) 57–71.

Laibson, D.I., 1994. Self-control and savings. MIT Thesis, Cambridge, MA, USA.

Laibson, D.I., 1996. Hyperbolic discount functions, undersaving, and savings policy. NBER Working Paper No. 5635, Cambridge, MA.

Laibson, D.I., 1997a. Golden eggs and hyperbolic discounting. Quarterly Journal of Economics 443–477.

Laibson, D.I., 1997b. Hyperbolic discount functions and time preference heterogeneity. Mimeo., Harvard University, Cambridge, MA.

Loewenstein, G., Prelec, D., 1992. Anomalies in intertemporal choice: evidence and an interpretation. Quarterly Journal of Economics 57 (2), 573–598.

Loewenstein, G., 1996. Out of control: visceral influences on behavior. Organizational Behavior and Human Decision Processes 65 (3) 272–292.

O'Donoghue, T., Rabin, R., 1997. Doing it now or doing it later. Mimeo., University of California, Berkeley.

Phelps, E.S., Pollak, R.A., 1968. On second-best national saving and game-equilibrium growth. Review of Economic Studies 35, 185–199.

Pollak, R.A., 1968. Consistent planning. Review of Economic Studies 35, 201–208.

Strotz, R.H., 1956. Myopia and inconsistency in dynamic utility maximization. Review of Economic Studies 23, 165–180.

Thaler, R.H., Shefrin, H.M., 1981. An economic theory of self-control. Journal of Political Economy 89 (2), 392–406.

Thaler, R.H., 1990. Saving, fungibility, and mental accounts. Journal of Economic Perspectives 4, 193–205.