

Summer 2001

Quarterly Review

Dollarization and
the Conquest of Hyperinflation
in Divided Societies (p. 3)

Russell W. Cooper
Hubert Kempf

Looking for Evidence
of Time-Inconsistent Preferences
in Asset Market Data (p. 13)

Narayana R. Kocherlakota

Federal Reserve Bank of Minneapolis

Quarterly Review

Vol. 25, No. 3

ISSN 0271-5287

This publication primarily presents economic research aimed at improving policymaking by the Federal Reserve System and other governmental authorities.

Any views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

Editor: Arthur J. Rolnick

Associate Editors: Patrick J. Kehoe, Warren E. Weber

Economic Advisory Board: V. V. Chari, Patrick J. Kehoe,

Narayana R. Kocherlakota, Erzo G. J. Luttmer

Managing Editor: Kathleen S. Rolfe

Article Editors: Kathleen S. Rolfe, Jenni C. Schoppers

Production Editor: Jenni C. Schoppers

Designer: Phil Swenson

Typesetter: Mary E. Anomalay

Circulation Assistant: Elaine R. Reed

The *Quarterly Review* is published by the Research Department of the Federal Reserve Bank of Minneapolis. Subscriptions are available free of charge.

Quarterly Review articles that are reprints or revisions of papers published elsewhere may not be reprinted without the written permission of the original publisher. All other *Quarterly Review* articles may be reprinted without charge. If you reprint an article, please fully credit the source—the Minneapolis Federal Reserve Bank as well as the *Quarterly Review*—and include with the reprint a version of the standard Federal Reserve disclaimer (italicized above). Also, please send one copy of any publication that includes a reprint to the Minneapolis Fed Research Department.

Electronic files of *Quarterly Review* articles are available through the Minneapolis Fed's home page on the World Wide Web: <http://www.minneapolisfed.org>.

Comments and questions about the *Quarterly Review* may be sent to

Quarterly Review

Research Department

Federal Reserve Bank of Minneapolis

P. O. Box 291

Minneapolis, Minnesota 55480-0291

(Phone 612-204-6455 / Fax 612-204-5515).

Subscription requests may also be sent to the circulation assistant at elaine.reed@mpls.frb.org; editorial comments and questions, to the managing editor at ksr@res.mpls.frb.fed.us.

Looking for Evidence of Time-Inconsistent Preferences in Asset Market Data

Narayana R. Kocherlakota*

Consultant

Research Department

Federal Reserve Bank of Minneapolis

and Carlson Professor of Economics

University of Minnesota

Jan is about to go out to her neighborhood bar. Before drinking anything there, Jan would like to sign a legally binding contract stating that she is allowed to drink only four beers that night. Why does she want to sign such a contract? She knows that after having four beers, she will want to have a fifth, and she wants to prevent herself from doing so.

Jan is exhibiting what economists call *time-inconsistent* preferences: her preferences for beer, at a given date and state, change over time without the arrival of new information. An essential feature of time-inconsistency is the desire for *self-commitment*. People like Jan with time-inconsistent preferences are willing to pay a cost to restrict their future choices.

Until recently, economists have typically assumed that the preferences of most people are consistent over time. In the last five years, however, research into the consequences of time-inconsistency has increased. Much of this has been spurred by the work of Laibson (1997) on consumption and saving.¹ Laibson assumes that people would like to be able to commit to save more at some future date than they think they actually would otherwise save when they get to that date. He then considers the consequences of these kinds of preferences for standard macroeconomic phenomena like the covariation of household consumption and income and the level of household saving.

Laibson argues that a considerable amount of introspective and experimental evidence supports his formulation of

time-inconsistent preferences. However, switching from a standard modeling strategy to one with time-inconsistent preferences can dramatically change a model's implications for economic policy.² So to make such a switch, we need to have more than introspective and experimental evidence that preferences are time-inconsistent; we need supportive evidence from actual choices that affect actual outcomes.

In this study, I ask, Can we see that sort of evidence in asset market data? My answer is that we cannot see this evidence in the prices of retradable assets. Rather, we need to look at the prices and holdings of what I call *commitment assets*.

I determine this by building and analyzing the implications of a deterministic, three-period general equilibrium model. In this model economy, agents can, in the initial period, trade a one-period (*short-term*) bond, a two-period (*long-term*) bond, and a commitment asset which, as the name implies, is an investment that must be held for its

*The author thanks V. V. Chari, Erzo Luttmer, Barbara McCutcheon, and Art Rolnick for their comments. He thanks Yan Bai for catching a subset (improper, the author hopes) of his mistakes.

¹The notion of time-inconsistent preferences over consumption profiles was first formalized by Strotz (1955–1956).

²For example, Laibson (1997) demonstrates that with time-inconsistent preferences, it may be optimal for the government to restrict financial market innovation. And Jovanovic and Stolyarov (2000) show that societies in which agents have time-inconsistent preferences may find it optimal to reject technological progress.

lifetime. The long-term bond can be retraded in the second period; the commitment asset cannot be. Also, agents cannot borrow in the second period against the future proceeds of the commitment asset.³

When the agents in this model economy have time-inconsistent preferences, they have three utility discount factors. They have two in period 1, one to discount the utility of consumption between periods 1 and 2 and one to discount that utility between periods 2 and 3. And they have another in period 2, also to discount utility between periods 2 and 3. Preferences are time-inconsistent if and only if the utility discount factor between periods 2 and 3 is different in period 2 than in period 1. I follow Laibson by restricting attention to the case in which the value of this discount factor in period 1 is greater than or equal to its value in period 2; over time, that is, the discount factor may decline.

In this model economy, I prove three results. First I examine the informational content of the prices of the short- and long-term bonds. Intuition might suggest that these prices are enough to tell us whether or not preferences are time-inconsistent. Again, the economy has three bond prices and three discount factors. Since we have three observable variables and three unknown parameters, it might seem plausible that we should be able to figure out whether preferences are time-consistent or time-inconsistent.

As my first result, I prove that this line of reasoning is wrong: in any equilibrium, bond prices are consistent with the discount factor between periods 2 and 3 being the same in period 1 as in period 2. The mistake in the intuitive reasoning is that it ignores dynamic arbitrage. Regardless of the form of preferences, the period 1 relative price between the two bonds must equal the period 2 price of the long-term bond; otherwise, agents can make arbitrage profits. The two bond prices are thus not independent sources of information about the two discount factors. To try to learn about time-inconsistent preferences, we must turn instead to price data on the commitment asset.

My second result is that we can, indeed, tell from the price of the commitment asset whether or not agents' preferences are time-inconsistent. In particular, I prove that the price of the commitment asset is higher than the present value of its future payoffs if and only if preferences are time-inconsistent. In effect, time-inconsistent people value commitment, and this value shows up in the price of the commitment asset.

My final result concerns agents' holdings of the commitment asset and the long-term bond. I prove that if preferences are time-inconsistent, then all agents' asset hold-

ings are *exclusive*: in period 3, agents receive all income from either the commitment asset or the long-term bond; they don't receive income from both types of assets. If some agents held both the commitment asset and the long-term bond at the end of period 2, then, on the margin, the commitment asset would provide no commitment. The agents could always reduce their consumption in period 3 by lowering their holdings of the bond. Because in this case, the commitment asset and the long-term bond would be marginally equivalent, they would have the same price. But this contradicts the second result. Therefore, in period 3, if preferences are time-inconsistent, the holdings of the commitment asset and the long-term bond must be exclusive.

How do these implications of time-inconsistent preferences compare to empirical observations from the U.S. asset market? To answer that question, I examine evidence about two good examples of commitment assets: education and individual retirement accounts (IRAs). Contrary to the implications of time-inconsistent preferences, neither asset seems to have an unusually low after-tax return. And virtually all agents with education or IRAs also have highly liquid bank accounts or highly collateralizable housing. This contradicts the exclusive holdings result. I conclude that there is little evidence from data on these two commitment assets against the traditional assumption that preferences are time-consistent.⁴

A Model With Time-Inconsistent Preferences

I start by developing a general equilibrium model of asset pricing in which people have time-inconsistent preferences.

³My concept of a commitment asset differs from Laibson's (1997) of an illiquid asset. He models illiquidity by assuming that an agent must commit to selling the asset one period before the sale takes place.

Another major difference between my analysis and Laibson's is that he analyzes a dynamic decision problem in which illiquid and liquid assets have the same rate of return in every period. I consider a general equilibrium in which the price of the commitment asset and the price of retradable bonds adjust so as to clear markets.

⁴My work here is quite different from that of Barro (1999) and Luttmer and Mariotti (2000). These researchers do not allow for commitment assets, and they restrict attention to stationary discounting (so that the utility discount factor between two periods depends only on the amount of time between the two periods). These researchers ask whether data on aggregate quantities and asset returns can be used to reject the traditional assumption of time-consistent preferences. Despite the absence of commitment assets, they find that the answer to their question is generically yes.

The key to their result is their assumption of stationary discounting. If discounting is stationary and preferences are time-consistent, then agents in any year have the same utility discount factor between the years 2010 and 2011 as between the years 2001 and 2002. Time-consistency requires only that the utility discount factor between 2010 and 2011 be the same in 2001 as in 2010. Because it is more restrictive, the joint hypothesis of stationary discounting and time-consistency is easier to refute than the single hypothesis of time-consistency.

The model has a unit measure of people, indexed by $j \in [0,1]$, who all live for three periods.⁵ The world is deterministic and has a single perishable consumption good. Each agent is endowed with y_1 units of consumption in period 1.

Agents are each also endowed with three assets. I refer to the first two assets as *bonds*. Each agent is endowed with \bar{b}_2 units of a short-term bond that pays off one unit of consumption in period 2; each agent is also endowed with \bar{b}_3 units of a long-term bond that pays off one unit of consumption in period 3 but can be retraded in period 2. The last asset is a *commitment asset* which cannot be retraded. It pays off one unit of consumption in period 3, and each agent is endowed with \bar{b}_3^{com} units of it. These three assets are the entire endowment of the J agents; hence, the per capita endowment in periods 2 and 3 is given by $y_2 = \bar{b}_2$ and $y_3 = \bar{b}_3 + \bar{b}_3^{com}$.

All agents have identical preferences over future consumption streams. However, these preferences may change over time. In particular, the agents' preferences over consumption streams (c_1, c_2, c_3) in period 1 are representable by the utility function

$$u(c_1) + \beta_{12}[u(c_2) + \beta_{23}u(c_3)]$$

where the β 's here represent discount factors in period 1, β_{12} between periods 1 and 2 and β_{23} between periods 2 and 3. The agents' preferences over consumption streams (c_2, c_3) in period 2 are representable by this utility function:

$$u(c_2) + \beta'_{23}u(c_3)$$

for $\beta'_{23} \leq \beta_{23}$, where β'_{23} is the discount factor in period 2 between periods 2 and 3.

If $\beta'_{23} = \beta_{23}$, then the agents' preferences are *time-consistent*; the discount factor between periods 2 and 3 is the same in period 2 as in period 1. Otherwise, preferences are *time-inconsistent*. To understand this assertion, suppose that $\beta'_{23} < \beta_{23}$, and suppose that one individual—say, Paul—is endowed with one unit of consumption in each period. Paul is asked, would you be willing to give up ϵ units of consumption in period 3 in exchange for ϵ/R units of consumption in period 2, given that $\beta'_{23}R < 1 < \beta_{23}R$? In period 1, Paul would respond negatively to this question; the rate of return R is higher than the discount rate between periods 2 and 3, so he would not want a loan on these terms in period 2. In period 2, though, he would respond affirmatively. Just like Jan in the introduction,

Paul would like to sign a contract in period 1 preventing himself from doing what he thinks he will otherwise do in period 2, here, accept the period 2 loan.

The restriction that $\beta'_{23} \leq \beta_{23}$ guarantees that in period 1, agents want more period 2 saving than they do when they actually get to period 2. If $\beta_{12} = \beta_{23} = \beta'_{23}$, then preferences are said to exhibit *exponential discounting*. I assume that $u', -u'' > 0$ and that $u'(0) = \infty$.

Agents engage in trade in both periods 1 and 2. In period 1, after receiving their endowments, the agents can trade the three assets and consumption. They cannot short sell the assets (so that holdings are restricted to being non-negative). In period 2, after receiving their payoffs from their holdings of the short-term bond, the agents can trade consumption and the long-term bond that pays off in period 3. Again, they cannot short sell the traded asset or trade the commitment asset in period 2. Thus, holding a units of the commitment asset commits the agent to consuming no less than a units of consumption in period 3.

In both periods, consumption is the numeraire. I use the following notation for prices: q_{st} is the price in period s of a bond that pays off in period t , and p_{com} is the period 1 price of the commitment asset.

Decision Problems

Given this trading protocol, I can now write agent j 's decision problems in periods 2 and 1.

□ Period 2

Agent j enters period 2 with b_{12}^j units of the bond that pays off in period 2, b_{13}^j units of the bond that pays off in period 3, and a_{com}^j units of the commitment asset. Then, given a price q_{23} , the agent has to choose b_{23}^j units of period 3 bonds and c_2^j units of consumption.

Given q_{23} , b_{12}^j , b_{13}^j , and a_{com}^j , then, let $q_{23}b_{13}^j + b_{12}^j = W_{liq}^j$, or *liquid wealth*, and let $q_{23}a_{com}^j = W_{com}^j$, or *committed wealth*. Now I can write agent j 's period 2 decision problem (DP2) as

$$\max_{c_2^j, c_3^j, b_{23}^j} [u(c_2^j) + \beta'_{23}u(c_3^j)]$$

subject to

$$c_2^j + q_{23}b_{23}^j = W_{liq}^j$$

$$c_3^j = b_{23}^j + (W_{com}^j/q_{23})$$

$$c_2^j, b_{23}^j \geq 0.$$

⁵Kocherlakota (1996) and Krusell and Smith (2000) show that the implications of time-inconsistent preferences may be quite different when the horizon is infinite.

I define $c_2^*(W_{liq}^j, W_{com}^j; q_{23})$ and $c_3^*(W_{liq}^j, W_{com}^j; q_{23})$ to be the solution of DP2.

□ *Period 1*

In period 1, agent j chooses consumption c_1^j and asset holdings $(a_{com}^j, b_{12}^j, b_{13}^j)$ taking as given the current prices $(p_{com}, q_{12}, q_{13})$ and the future price q_{23} . The key part of this decision problem is the fact that the agent realizes that the agent's own preferences may change next period. Hence, the agent cannot directly plan for period 2 and period 3 consumption. Instead, consumption in these periods is essentially under the control of the agent's future self, a person in the same physical body but with different preferences. Agent j must choose period 1 asset holdings taking into account this future self's response to those asset holdings next period (that is, taking as given the response functions c_2^* and c_3^*).⁶

This logic produces the following period 1 decision problem (DP1):

$$\max_{(c_1^j, b_{12}^j, W^j)} \{ u(c_1^j) + \beta_{12} u(c_2^*(W_{liq}^j, W_{com}^j; q_{23})) + \beta_{12} \beta_{23} u(c_3^*(W_{liq}^j, W_{com}^j; q_{23})) \}$$

subject to

$$c_1^j + p_{com} a_{com}^j + q_{12} b_{12}^j + q_{13} b_{13}^j = y_1 + p_{com} \bar{b}_3^{com} + q_{13} \bar{b}_3 + q_{12} \bar{b}_2$$

$$W_{liq}^j = q_{23} b_{13}^j + b_{12}^j$$

$$W_{com}^j = q_{23} a_{com}^j$$

$$c_1^j, b_{12}^j, b_{13}^j, a_{com}^j \geq 0.$$

Here, the objective includes how consumption (c_2, c_3) responds to the agent's choices of liquid and committed wealth.

Equilibrium

I define an *equilibrium* in the natural way, as a specification of consumption $(c_1^j, c_2^j, c_3^j)_{j \in [0,1]}$, asset holdings $(b_{12}^j, b_{13}^j, b_{23}^j, a_{com}^j)_{j \in [0,1]}$, and prices $(q_{12}, q_{13}, q_{23}, p_{com})$ that satisfies three criteria. First, (c_2^j, c_3^j) solves DP2 given q_{23} , $W_{liq}^j = q_{23} b_{13}^j + b_{12}^j$, and $W_{com}^j = q_{23} a_{com}^j$. Second, $(c_1^j, b_{12}^j, b_{13}^j, a_{com}^j)$ solves DP1 given q_{12} , q_{13} , q_{23} , and p_{com} . And third, markets clear, so that

$$\int c_1^j dj = y_1$$

$$\int c_2^j dj = \bar{b}_2$$

$$\int c_3^j dj = \bar{b}_3^{com} + \bar{b}_3$$

$$\int b_{12}^j dj = \bar{b}_2$$

$$\int b_{23}^j dj = \int b_{13}^j dj = \bar{b}_3$$

$$\int a_{com}^j dj = \bar{b}_3^{com}.$$

This definition may seem overly elaborate, because it allows for asymmetric equilibria. After all, this is a simple economy in which agents are identical in tastes and endowments. The equilibrium would seem to naturally have identical allocations across agents. However, we will see that any equilibrium is asymmetric if $\beta'_{23} < \beta_{23}$.

If preferences are time-consistent, however, so that $\beta'_{23} = \beta_{23}$, then we can price assets using the marginal rates of substitution of a representative agent.

PROPOSITION 1. *Suppose that $\beta'_{23} = \beta_{23}$. Then autarky is an equilibrium, and in any equilibrium,*

$$c_t^j = y_t$$

for all t and

$$q_{12} = \beta_{12} u'(y_2) / u'(y_1)$$

$$q_{13} = \beta_{12} \beta_{23} u'(y_3) / u'(y_1)$$

$$q_{23} = \beta_{23} u'(y_3) / u'(y_2)$$

$$p_{com} = q_{13}.$$

Proof. Let $\beta'_{23} = \beta_{23}$. Then solving DP1 is equivalent to solving

$$(c_1, c_2, c_3) \in \arg \max_{(c_1, c_2, c_3, a_{com})} [u(c_1) + \beta_{12} u(c_2) + \beta_{12} \beta_{23} u(c_3)]$$

subject to

$$c_1 + q_{12} c_2 + (c_3 - a_{com}) q_{13} + p_{com} a_{com} \leq y_1 + q_{12} \bar{b}_2 + q_{13} \bar{b}_3 + p_{com} \bar{b}_3^{com}$$

$$c_3 \geq a_{com}$$

$$c_1, c_2, c_3, a_{com} \geq 0.$$

⁶Formally, as Pollak (1968) does, I treat the outcome of the decision problem as being the subgame perfect equilibrium outcome of a game between the agent in period 1 and the agent in period 2.

If $p_{com} < q_{13}$, it is optimal to set $c_3 = a_{com}$. If $p_{com} > q_{13}$, it is optimal to set $a_{com} = 0$. Thus, DP1 is equivalent to

$$(c_1, c_2, c_3) \in \arg \max_{(c_1, c_2, c_3)} [u(c_1) + \beta_{12}u(c_2) + \beta_{12}\beta_{23}u(c_3)]$$

subject to

$$c_1 + q_{12}c_2 + \min(p_{com}, q_{13})c_3 \leq y_1 + q_{12}\bar{b}_2 + q_{13}\bar{b}_3 + p_{com}\bar{b}_3^{com}$$

$$c_1, c_2, c_3 \geq 0.$$

This problem has a convex constraint set and a strictly concave objective. All agents have the same solution, which means that in equilibrium they must consume autarky. The rest of the proposition follows from the fact that all assets are in positive supply. Q.E.D.

Here, the bonds are priced in the usual way, using the marginal rates of substitution of the representative agent. Note that when preferences are time-consistent, the commitment asset is priced in the same way as the long-term bond; there is no value to commitment.

Implications of the Model

Now I use the general equilibrium model just developed to attempt to understand whether data on asset prices and holdings can tell us if agents have time-inconsistent preferences. Throughout, I assume that we know that the above general equilibrium model is true, we know the aggregate quantities $(y_t)_{t=1}^3$, and we know the function u' . I ask whether with this knowledge we can determine if $\beta'_{23} = \beta_{23}$, based on asset prices and holdings. My answer is both no and yes.

Bonds

First I argue that characteristics of bond prices cannot reveal whether preferences are time-inconsistent.

Note that in the model, agents have two ways to generate a unit of consumption in period 3 using the available bonds. One is to buy a short-term bond in period 1 and roll it over in period 2. The other is to simply buy a long-term bond in period 1. The key result here is that in any equilibrium, the costs of these strategies are the same:

$$q_{12}q_{23} = q_{13}.$$

The proof of this is a typical kind of arbitrage argument. However, we must go through the argument carefully be-

cause of the time-consistency issue. Suppose that $q_{12}q_{23} > q_{13}$. Because of market-clearing, some agent j has bond holdings $b_{12}^j > 0$. I claim that there is an element of the constraint set of DP1 that improves this agent's objective: lower the agent's short-term bond holdings b_{12}^j by ε and raise the agent's long-term bond holdings b_{13}^j by ε/q_{23} . This leaves W_{liq}^j unchanged, but frees up resources in period 1 because $\varepsilon q_{12} > \varepsilon q_{13}/q_{23}$.

Similarly, suppose that $q_{12}q_{23} < q_{13}$. Some agent j exists who has bond holdings $b_{13}^j > 0$ (again, by market-clearing). To improve this agent's objective, lower the agent's holdings of long-term bonds b_{13}^j by ε and raise the short-term bond holdings b_{12}^j by εq_{23} . Again, this change in holdings generates extra resources in period 1.

Hence, in this economy, the usual arbitrage arguments apply to bond pricing. But this immediately means that the economy never has evidence in bond prices against the hypothesis that preferences are time-consistent.

Specifically, suppose that bond prices are given by (q_{12}, q_{13}, q_{23}) . We see these prices and know u' and $(y_t)_{t=1}^3$. Now define the discount factors

$$\beta_{12} = q_{12}^{-1}u'(y_1)/u'(y_2)$$

$$\beta_{23} = \beta'_{23} = q_{23}^{-1}u'(y_2)/u'(y_3).$$

Then from Proposition 1, we know that if agents have these discount factors, equilibrium bond prices are given by (q_{12}, q_{13}, q_{23}) . Bond prices never contradict the hypothesis of time-consistency. (Obviously, if we did not know u' or $(y_t)_{t=1}^3$, rejecting the hypothesis of time-consistency would be harder.)

Why don't bond prices reveal whether preferences are time-consistent or time-inconsistent? In period 1, the relative price between period 2 and period 3 consumption is given by q_{13}/q_{12} . Some might think that this relative price contains information about β_{23} , that is, the agents' period 1 willingness to substitute between period 2 and period 3 consumption. But this thinking ignores dynamic arbitrage. If agents anticipate that the period 2 relative price between periods 2 and 3 consumption is q_{23} , then the period 1 relative price must also be q_{23} . This means that q_{23} and q_{13}/q_{12} have the same information about β_{23} .

The result here can be generalized: in any economy, the prices of one-period (short-term) assets and costlessly retraded (long-term) assets do not contradict the hypothesis of time-consistency. The key is that even if preferences are time-inconsistent, prices must not admit any dynamic ar-

bitrage opportunities. From the work of Hansen and Richard (1987), we know that this implies that there is a stochastic pricing kernel representation for asset prices. This pricing kernel can immediately be translated into a stochastic discount factor for agents. (Indeed, these concepts are identical if we allow agents to have linear utility.)⁷

The Commitment Asset

Now I argue that we can learn whether preferences are time-inconsistent from some characteristics of commitment assets. I prove two results about this type of asset.

□ High Price.

My first result here is that if preferences are time-inconsistent, then the commitment asset's price is higher than the present value of its future payoffs, which is to say, in equilibrium, higher than the price of the long-term, tradable bond.

The proof works as follows. Suppose that the long-term bond and the commitment asset have the same return in period 1. Then the agent in period 1 can translate period 1 consumption into period 3 consumption at the same rate using either bonds or the commitment asset. Given the time-consistency problem, it is then optimal for any agent to commit to a particular level of c_3 by holding just enough W_{liq} to fund the agent's period 2 consumption, but no more. This cannot be an equilibrium, though, because no agent is holding any of the long-term bond at the end of period 2.

PROPOSITION 2. *If $\beta'_{23} < \beta_{23}$, then in any equilibrium, $p_{com} > q_{13}$.*

Proof. Suppose not, and let $(c^*, b_{12}^*, b_{13}^*, b_{23}^*, a^*)$, $(q_{12}, q_{13}, q_{23}, p_{com})$ be an equilibrium in which $p_{com} \leq q_{13}$. From market-clearing, we know that for some j , $b_{23}^{j*} > 0$. Then $\beta_{23} u'(c_3^{j*}) > \beta'_{23} u'(c_3^{j*}) = u'(c_2^{j*}) q_{23}$. Let (c_2^j, c_3^j) be the solution to this problem:

$$\max_{c_2^j, c_3^j} [u(c_2^j) + \beta_{23} u(c_3^j)]$$

subject to

$$c_2^j + q_{23} c_3^j \leq c_2^{j*} + q_{23} c_3^{j*}$$

$$c_2^j, c_3^j \geq 0.$$

Since $\beta_{23} u'(c_3^{j*}) > u'(c_2^{j*}) q_{23}$, we know that $u(c_2^j) + \beta_{23} u(c_3^j) > u(c_2^{j*}) + \beta_{23} u(c_3^{j*})$. We also know that $c_3^{j*} < c_3^j$, which implies in turn that $a_{com}^{j*} < c_3^j$.

Now I claim that this plan

$$c_1^{j'} = c_1^{j*}$$

$$b_{12}^{j'} = c_2^{j'}$$

$$b_{13}^{j'} = 0$$

$$a_{com}^{j'} = c_3^{j'}$$

lies in the constraint set of DP1, given equilibrium prices $(p_{com}, q_{12}, q_{13}, q_{23})$. This is demonstrated by the following chain of logic:

$$\begin{aligned} & c_1^{j'} + q_{12} b_{12}^{j'} + p_{com} a_{com}^{j'} \\ &= c_1^{j'} + q_{12} c_2^{j'} + p_{com} c_3^{j'} \\ &= c_1^{j'} + q_{12} c_2^{j'} + q_{13} c_3^{j'} + (p_{com} - q_{13}) c_3^{j'} \\ &= c_1^{j*} + q_{12} c_2^{j*} + q_{13} c_3^{j*} + (p_{com} - q_{13}) c_3^{j'} \\ &= c_1^{j*} + q_{12} b_{12}^{j*} + q_{13} a_{com}^{j*} + q_{13} b_{13}^{j*} \\ &\quad + (p_{com} - q_{13}) c_3^{j'} \\ &= c_1^{j*} + q_{12} b_{12}^{j*} + q_{13} b_{13}^{j*} + p_{com} a_{com}^{j*} \\ &\quad + (p_{com} - q_{13}) (c_3^{j'} - a_{com}^{j*}) \\ &< c_1^{j*} + q_{12} (b_{12}^{j*} + q_{23} b_{13}^{j*}) + p_{com} a_{com}^{j*}. \end{aligned}$$

The one inequality comes from the assumption that $p_{com} \leq q_{13}$ and the result that $c_3^{j'} > a_{com}^{j*}$.

We also know that

$$\beta'_{23} u'(c_3^{j'}) q_{23}^{-1} < u'(c_2^{j'}).$$

This implies that $c_2^*(b_{12}^{j'}, c_3^{j'} / q_{23}) = c_2^{j'}$ and $c_3^*(b_{12}^{j'}, c_3^{j'} / q_{23}) = c_3^{j'}$. We have a contradiction: the plan $(c_1^{j'}, b_{12}^{j'}, b_{13}^{j'}, a_{com}^{j'})$ lies in the constraint set of DP1 and delivers a higher value of the objective in DP1 than $(c_1^{j*}, b_{12}^{j*}, b_{13}^{j*}, a_{com}^{j*})$.

Q.E.D.

This result is intuitive. All of the agents' preferences are time-inconsistent. The agents want to commit to a certain

⁷Contrast this result with the logic of Barro (1999) and Luttmer and Mariotti (2000). These researchers impose a stationarity restriction on discounting: the discount factor between periods t and s is required to be a function of $t - s$. In my model, this restriction is equivalent to assuming that $\beta_{23} = \beta_{12}$. Time-consistency ($\beta_{23} = \beta_{23}$) then implies that

$$q_{12}^{-1} u'(y_1) / u'(y_2) = q_{23}^{-1} u'(y_2) / u'(y_3).$$

This restriction is falsifiable given data on $\{u'(y_t)\}_{t=1}^3$, q_{12} , and q_{23} . Moreover, the restriction is not generally valid if preferences are time-inconsistent. (See the Appendix.) Hence, if discounting is restricted to satisfy stationarity, bond pricing data can be used to reject time-consistency in favor of time-inconsistency.

amount of consumption in period 3. The commitment asset is a better way to do so than the long-term bond; hence, in equilibrium, the price of the commitment asset should be higher than the price of the long-term bond.

□ *Exclusive Holdings*

My second result here is that asset holdings must be exclusive in equilibrium: If $\beta'_{23} < \beta_{23}$, then for all agents, either $b_{23}^j = 0$ or $a_{com}^j = 0$. The proof is simple. If $b_{23}^j > 0$ and $a_{com}^j > 0$ for some agent j , then the long-term bond and the commitment asset are, on the margin, equivalent. Their rates of return must then be the same, which contradicts Proposition 2.

PROPOSITION 3. *Suppose that $\beta'_{23} < \beta_{23}$. Then in any equilibrium, $b_{23}^j a_{com}^j = 0$ for all j .*

Proof. Consider an arbitrary equilibrium $(c^*, b_{12}^*, b_{13}^*, b_{23}^*, a^*)$, $(q_{12}, q_{13}, q_{23}, p_{com})$. We know from Proposition 2 that $p_{com} > q_{13}$. Let $b_{23}^{j*} > 0$ and $a_{com}^{j*} > 0$ for some j . Then set $a_{com}^{j'} = a_{com}^{j*} - \varepsilon$, $b_{12}^{j'} = b_{12}^{j*} + \varepsilon q_{23}$, and $c_1^{j'} = c_1^{j*} + \varepsilon p_{com} - \varepsilon q_{13}$, with ε sufficiently small that all the primed variables are positive. Clearly, this new plan satisfies the budget constraint in DP1 and delivers more consumption in period 1. I claim that

$$c_t^*(b_{12}^{j'} + q_{23}b_{13}^{j*}, q_{23}a_{com}^{j'}) = c_t^{j*}$$

for $t > 1$, so that this new plan also provides more utility to the agent in period 1.

To prove my claim, we need to solve DP2:

$$\max_{c_2^j, b_{23}^j, c_3^j} [u(c_2^j) + \beta'_{23}u(c_3^j)]$$

subject to

$$c_2^j + q_{23}b_{23}^j = q_{23}b_{13}^{j*} + b_{12}^{j*} + \varepsilon q_{23}$$

$$c_3^j = b_{23}^j + a_{com}^{j*} - \varepsilon$$

$$c_2^j, b_{23}^j \geq 0.$$

It is straightforward to see that if an element of the constraint set satisfies

$$\beta'_{23}u'(c_3^j) = u'(c_2^j)q_{23}$$

then that element solves DP2. But because $b_{23}^{j*} > 0$, we know that (c_2^{j*}, c_3^{j*}) satisfies this first-order condition. By setting $b_{23}^j = b_{23}^{j*} + \varepsilon$, we can see that (c_2^{j*}, c_3^{j*}) lies in the constraint set, too. It follows that $(c_1^{j'}, b_{12}^{j'}, b_{13}^{j*}, a_{com}^{j'})$ lies in

the constraint set of DP1 and improves the period 1 utility of agent j .
Q.E.D.

An immediate consequence of Proposition 3 is that there is no symmetric equilibrium if $\beta'_{23} < \beta_{23}$. In a symmetric equilibrium, all agents must hold their endowments, including having $a_{com}^j = \bar{b}_3^{com}$ and $b_{23}^j = \bar{b}_3$. But this is impossible.

Propositions 2 and 3 suggest a natural question: Is there any equilibrium at all if $\beta'_{23} < \beta_{23}$? I do not know the answer to this question for general specifications of u . In the Appendix, though, I construct an equilibrium for log utility. In that equilibrium, the commitment asset pays a low return relative to the long-term, retradable bond. Between periods 2 and 3, a fraction θ of the agents choose to hold the commitment asset and a fraction $1 - \theta$ of the agents choose to hold the long-term bond. The first strategy pays a low return, but allows the agent in period 1 to lock in a high level of saving from period 2 to period 3. The second strategy does not deliver the lock-in effect, but gives the agent a high return. The fraction θ adjusts until the agents are indifferent between the two strategies.

My guess is that Propositions 2 and 3 are not as general as Proposition 1. I can imagine a situation that would contradict them. For example, suppose that agents face shocks to their income that are not directly insurable. In this kind of economy, agents smooth their consumption by selling assets when their income is low. Because commitment assets can't be sold but retradable assets can, commitment assets might be worth less than retradable assets. Similarly, in such an economy, agents might hold both types of assets in equilibrium.⁸

At the same time, Propositions 2 and 3 should be useful benchmarks as long as time-inconsistency is actually quantitatively significant. More specifically, I make the following conjecture about economies in which agents have uninsurable income shocks: If β_{23}/β'_{23} is sufficiently large, then commitment assets have a higher price than do retradable assets, and agents' portfolios consist nearly entirely of commitment assets or nearly entirely of retradable assets.

Evidence

So, to uncover evidence in favor of time-inconsistency, we need to look at actual prices and holdings of commitment assets. If preferences are time-inconsistent, the price of a

⁸For a numerical analysis of a calibrated decision problem, see the work of Angelos et al. (2001). They do not, however, solve for the full general equilibrium.

commitment asset is higher than that of comparable retradable assets and agents do not simultaneously hold both the commitment asset and retradable assets. To what extent are these implications supported by available empirical evidence? As best I can tell, not at all.

Overpriced Commitment Assets?

Commitment assets have two defining features. First, they are hard to sell. Second, they are hard to use as collateral: borrowing against their future payoffs is costly for some reason. Identifying actual commitment assets is not easy. No thickly traded asset is a commitment asset, and even thinly traded assets can generally be used as collateral for loans. Still, I have come up with what I think are two good examples of commitment assets.

One is education. Educated people pay a cost (primarily, forgone wages) to acquire increased earning power. Because of legal prohibitions against indentured servitude, people cannot directly capitalize this increased earning power. Moreover, for the same reason, borrowing against it is often difficult.

Because education has both features of a commitment asset, we know from Proposition 2 that education's return should be low. A considerable amount of effort has been made to estimate the return to education. Card (1999, Table 4) provides a partial survey of these estimates.⁹ In his survey, the (instrumental variable) estimates from the United States range widely, from 6 percent to 15 percent. In contrast, real stock returns in the United States over the past century have averaged around 8 percent. I have not found good measures of education's risk as an asset. Nonetheless, it is hard to conclude from this kind of analysis that education returns are strikingly low.

My other example of commitment assets is perhaps more natural: IRAs. These are part of a U.S. government program designed to encourage saving. People can invest up to a fixed amount per year in IRAs. Once people turn age 59 1/2, they can withdraw funds from IRAs without cost. However, they face penalties for early withdrawal and for borrowing against the contents of the accounts.

IRAs closely resemble commitment assets: they are designed to allow agents to pre-commit to a lower floor for consumption when they are retired. From Proposition 2, then, if preferences are time-inconsistent, we would expect these assets to pay a low return. Of course, they don't: IRAs pay a return equivalent to that in the marketplace—or higher, since IRAs are tax-free.

This would appear to be strong evidence against time-inconsistency. We do have to be cautious, though. Unlike

the agents in the general equilibrium model, governments are not wealth-maximizing entities. It is quite possible (indeed, more than likely) that some large portion of the IRAs' high return represents a government subsidy to savers. To evaluate time-inconsistency, then, the appropriate question to ask is, would the IRA program attract a large amount of participation if people had to pay an especially high (instead of especially low) tax on their returns? There is, of course, no way to know for sure, but with fairly high confidence, I think the answer is no.

Exclusive Holdings?

Proposition 3 demonstrates that in equilibrium, if preferences are time-inconsistent, agents do not receive income in period 3 from both retradable assets and commitment assets. This characteristic of time-inconsistency is clearly not supported by the evidence. Virtually all owners of investment assets—including those who are educated or who own IRAs—also have a fairly large liquid or collateralizable reserve, in the form of either bank accounts or equity in their homes. There is thus little or no evidence consistent with the exclusiveness implication of time-inconsistency.

But what if we changed the model economy so that some agents have time-inconsistent preferences and others have time-consistent preferences? Then if any agent simultaneously holds commitment assets and retradable assets, the same logic as in Proposition 3 implies that commitment assets will not be overpriced. And the reasoning in the proof of Proposition 2 then implies that the agents with time-inconsistent preferences will buy only commitment assets. It follows that if most people have time-inconsistent preferences, most people should own only commitment assets. We do not see this kind of asset holding behavior.

Conclusions

In this study, I have looked for evidence of time-inconsistent preferences in asset market data. I have shown that be-

⁹Card (1999) literally reports estimates of the percentage increase in wages from an additional year of schooling. We can interpret this as a return in the financial sense if most schooling costs are forgone wages and the increase in wages persists for a long time. Under those assumptions, if the estimated coefficient is γ , then people give up a wage w per unit of time and perpetually gain γw per unit of time. The internal rate of return to this investment is γ .

Other than the usual selection issues (which the econometricians do their best to control), there are various biases in this analysis that are easy to sign. If people work while going to school, then the return to education is higher than γ . If the increase in earnings does not last long, then the return to education is lower than γ . Finally, if education involves significant costs other than time, then the return is lower than γ .

cause of dynamic arbitrage, there can be no evidence of time-inconsistency in the prices of one-period assets or in the prices of long-term, retradable assets. However, if people have time-inconsistent preferences, commitment assets are systematically overpriced, and people do not hold both them and retradable assets in equilibrium.

The question of whether or not preferences are actually time-inconsistent may appear to be arcane and technical. It is not. A major question—perhaps the major question—in socioeconomic policy concerns the extent to which governments should regulate market interactions and individual choices. Under traditional economic models of individual behavior, people act in their own interest. Any kind of choice must improve their welfare; any kind of trade must be mutually beneficial. Under this assumption about human behavior, the main rationale for government intervention is to cure externalities.

This changes if preferences are time-inconsistent. Now the government has a new role to play: it can improve current welfare by limiting future choices. Jan would welcome the government limiting her to four beers. Paul and other consumers in the model economy would like the government to stop sales of the retradable bond that pays off in period 3. With time-inconsistent preferences, it is beneficial for the government to restrict individual choices in order to guard people against the excesses of their future selves.

Thus, the policy recommendations based on models with time-inconsistent preferences may be quite radical. Economists need to have strong evidence supporting time-inconsistency before making such recommendations. The data analysis here demonstrates that this kind of strong evidence is not currently available.

Appendix An Equilibrium When Utility Is Logarithmic

In this appendix, I construct an equilibrium for the case in which $u(c) = \ln(c)$.

It is simple to show that in this case,

$$\begin{aligned} c_2^*(W_{liq}, W_{com}) &= \min[(W_{liq} + W_{com})/(1 + \beta'_{23}), W_{liq}] \\ c_3^*(W_{liq}, W_{com}) &= \max[(W_{liq} + W_{com})q_{23}^{-1}\beta'_{23}/(1 + \beta'_{23}), W_{com}]. \end{aligned}$$

The function c_2^* is not concave; it is this nonconcavity that generates the asymmetry in the equilibrium. I proceed by first guessing the equilibrium and then verifying my guess.

Guessing . . .

Define θ to be the solution to the equation

$$\begin{aligned} \ln(1 - \theta) - \ln(\theta) \\ = \ln(\bar{b}_3^{com}) - \ln(\bar{b}_3) - \beta_{23}^{-1}\ln(1 + \beta_{23}) + \beta_{23}^{-1}\ln(1 + \beta'_{23}). \end{aligned}$$

Given θ , define

$$\begin{aligned} q_{12} &= \beta_{12}y_1\bar{b}_2^{-1}[1 + \theta\beta_{23} + (1 - \theta)\beta'_{23}]/(1 + \beta'_{23}) \\ q_{13} &= y_1(\bar{b}_3/\theta)^{-1}\beta_{12}\beta'_{23}(1 + \beta_{23})(1 + \beta'_{23})^{-1} \\ q_{23} &= q_{13}/q_{12} \\ p_{com} &= (1 - \theta)y_1\beta_{12}\beta_{23}\bar{b}_3^{com}. \end{aligned}$$

For $0 \leq j \leq \theta$, define

$$\begin{aligned} c_1^j &= y_1 \\ c_2^j &= y_1\beta_{12}q_{12}^{-1}(1 + \beta_{23})/(1 + \beta'_{23}) \\ c_3^j &= \bar{b}_3/\theta \\ b_{12}^j &= c_2^j \\ b_{13}^j &= b_{23}^j = \bar{b}_3/\theta \\ a_{com}^j &= 0. \end{aligned}$$

For $\theta < j \leq 1$, define

$$\begin{aligned} c_1^j &= y_1 \\ c_2^j &= y_1\beta_{12}q_{12}^{-1} \end{aligned}$$

$$\begin{aligned} c_3^j &= \bar{b}_3^{com}/(1 - \theta) \\ b_{12}^j &= c_2^j \\ b_{13}^j &= b_{23}^j = 0 \\ a_{com}^j &= \bar{b}_3^{com}/(1 - \theta). \end{aligned}$$

. . . And Verifying the Equilibrium

Now I will verify that this is actually an equilibrium.

To do that, first I need to show that $p_{com} > q_{13}$. Note that

$$\begin{aligned} \ln(p_{com}) - \ln(q_{13}) \\ = \ln(1 - \theta) - \ln(\theta) - \ln(\bar{b}_3^{com}) + \ln(\bar{b}_3) + \ln(1 + \beta'_{23}) \\ - \ln(1 + \beta_{23}) + \ln(\beta_{23}) - \ln(\beta'_{23}) \\ = \beta_{23}^{-1}[\ln(1 + \beta'_{23}) - \ln(1 + \beta_{23})] + \ln(1 + \beta'_{23}) \\ - \ln(1 + \beta_{23}) + \ln(\beta_{23}) - \ln(\beta'_{23}). \end{aligned}$$

Differentiating this expression with respect to β'_{23} yields

$$\begin{aligned} (\beta_{23}^{-1} + 1)/(1 + \beta'_{23}) - 1/\beta'_{23} \\ = (\beta'_{23}/\beta_{23} - 1)/[\beta'_{23} + (\beta'_{23})^2] \\ < 0 \end{aligned}$$

if $\beta'_{23} < \beta_{23}$. Hence, $p_{com} > q_{13}$ if $\beta'_{23} < \beta_{23}$. What this means is that if $j \leq \theta$, agents are following a low-return strategy in order to commit themselves to a high level of period 3 consumption. If $j > \theta$, agents are following a high-return strategy that does not generate as much period 3 consumption.

I can now show that these prices and quantities constitute an equilibrium. It is straightforward but tedious to show that markets clear and that

$$c_t^j = c_t^*(q_{23}b_{13}^j + b_{12}^j, q_{23}a_{com}^j)$$

for $t = 2, 3$ and for all j .

What is more difficult to demonstrate is that $(c_1^j, a_{com}^j, b_{12}^j, b_{13}^j)$ solves DP1 for all j . Because $p_{com} > q_{13}$, I know from the proof of Proposition 3 that there cannot be a solution to DP1 in which $a_{com}^j > 0$ and $b_{23}^j > 0$. This means that any solution to DP1 must lie in the portion of the constraint set in which $a_{com}^j = 0$ or in the portion of the constraint set in which $W_{com} \geq \beta'_{23}W_{liq}$.

I first show that the candidate equilibrium allocation for $j \leq \theta$ solves a version of DP1 in which a_{com}^j is fixed at 0. Let DP1a be the problem

$$\max_{c_1, b_{12}, b_{13}, c_2, c_3} [\ln(c_1) + \beta_{12} \ln(c_2) + \beta_{12} \beta_{23} \ln(c_3)]$$

subject to

$$c_1 + q_{12}b_{12} + q_{13}b_{13} = y_1 + p_{com} \bar{b}_3^{com} + q_{12} \bar{b}_2 + q_{13} \bar{b}_3$$

$$c_2 = (b_{12} + q_{23}b_{13}) / (1 + \beta'_{23})$$

$$c_3 = q_{23}^{-1} \beta'_{23} (b_{12} + q_{23}b_{13}) / (1 + \beta'_{23})$$

$$c_1, b_{12}, b_{13} \geq 0.$$

I can substitute out (b_{12}, b_{13}) so that this problem has a strictly concave objective over (c_1, c_2, c_3) and a convex constraint set in (c_1, c_2, c_3) :

$$\max_{c_1, c_2, c_3} [\ln(c_1) + \beta_{12} \ln(c_2) + \beta_{12} \beta_{23} \ln(c_3)]$$

subject to

$$c_1 + q_{12}(c_2 + q_{23}c_3) = y_1 + p_{com} \bar{b}_3^{com} + q_{12} \bar{b}_2 + q_{13} \bar{b}_3$$

$$c_3 = q_{23}^{-1} \beta'_{23} c_2$$

$$c_1, c_2, c_3 \geq 0.$$

The first-order conditions are necessary and sufficient, and the candidate equilibrium allocation for $j \leq \theta$ satisfies these first-order conditions.

I next show that the candidate equilibrium allocation for $j > \theta$ solves a version of DP1 in which $W_{com} \geq \beta'_{23} W_{liq}$. Let DP1b be the problem

$$\max_{c_1, b_{12}, b_{13}, a_{com}, c_2, c_3} [\ln(c_1) + \beta_{12} \ln(c_2) + \beta_{12} \beta_{23} \ln(c_3)]$$

subject to

$$c_1 + q_{12}b_{12} + q_{13}b_{13} + p_{com} a_{com} = y_1 + p_{com} \bar{b}_3^{com} + q_{12} \bar{b}_2 + q_{13} \bar{b}_3$$

$$c_2 = b_{12} + q_{23}b_{13}$$

$$c_3 = a_{com}$$

$$q_{23} a_{com} \geq (b_{12} + q_{23}b_{13}) \beta'_{23}$$

$$c_1, b_{12}, b_{13}, a_{com} \geq 0.$$

We can rewrite this as

$$\max_{c_1, c_2, c_3} [\ln(c_1) + \beta_{12} \ln(c_2) + \beta_{12} \beta_{23} \ln(c_3)]$$

subject to

$$c_1 + q_{12}c_2 + p_{com}c_3 = y_1 + p_{com} \bar{b}_3^{com} + q_{12} \bar{b}_2 + q_{13} \bar{b}_3$$

$$q_{23}c_3 = \beta'_{23}c_2$$

$$c_1, c_2, c_3 \geq 0.$$

This problem has a strictly concave objective and a convex constraint set. The consumption allocation for $j > \theta$ is the unique solution, because it satisfies the first-order necessary conditions.

I now claim that the candidate equilibrium allocations for all θ solve DP1. But this claim follows trivially from the fact that both consumption allocations generate the same period 1 utility. Thus, I have verified that the candidate equilibrium allocations satisfy period 1 optimality, period 2 optimality, and market-clearing. Q.E.D.

Note that the equilibrium allocations actually provide less period 1 utility to the agents than do their original endowment streams. This may seem strange, but given the equilibrium prices, an agent cannot credibly commit to not trading. In particular, an agent who does not trade in period 1 will want to sell some of the long-term bond in period 2.

Note also that if $\bar{b}_3^{com} = 0$, there is no trade in equilibrium. Hence, it would be Pareto-improving if all of the agents' period 3 endowments came in the form of a long-term, retradable asset.

References

- Angeletos, George-Marios; Laibson, David; Repetto, Andrea; Tobacman, Jeremy; and Weinberg, Stephen. 2001. The hyperbolic buffer stock model: Calibration, simulation, and empirical evaluation. Manuscript. Harvard University.
- Barro, Robert J. 1999. Ramsey meets Laibson in the neoclassical growth model. *Quarterly Journal of Economics* 114 (November): 1125–52.
- Card, David. 1999. The causal effect of education on earnings. In *Handbook of Labor Economics*, ed. Orley C. Ashenfelter and David Card, Vol. 3A, Chap. 30, pp. 1803–63. Amsterdam: North-Holland.
- Hansen, Lars Peter, and Richard, Scott F. 1987. The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica* 55 (May): 587–613.
- Jovanovic, Boyan, and Stolyarov, Dmitriy. 2000. Ignorance is bliss. Manuscript. New York University.
- Kocherlakota, Narayana R. 1996. Reconsideration-proofness: A refinement for infinite horizon time inconsistency. *Games and Economic Behavior* 15 (July): 33–54.
- Krusell, Per, and Smith, Anthony A., Jr. 2000. Consumption and savings decisions with quasi-geometric discounting. Manuscript. University of Rochester.
- Laibson, David. 1997. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112 (May): 443–77.
- Luttmer, Erzo G. J., and Mariotti, Thomas. 2000. Subjective discounting in an exchange economy. Manuscript. University of Minnesota.
- Pollak, R. A. 1968. Consistent planning. *Review of Economic Studies* 35 (April): 201–8.
- Strotz, R. H. 1955–1956. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 23 (Issue 3): 165–80.