

Asset Pricing Implications of Precommitted Consumption

Kenichi Fukushima*

June 15, 2005

When certain consumption needs exhibit strong complementarities with household characteristics that are not easily adjustable over time, those consumption components that are required to fulfill such needs are subject to higher adjustment frictions than those that are not. This paper develops a consumption based asset pricing model that incorporates such frictions and shows that the model can partially explain a number of asset pricing facts that have proved puzzling from the perspective of standard consumption based models.

JEL classification: E21, E32, E44, G12.

Keywords: consumption precommitments, asset pricing.

*University of Tokyo. E-mail: kenichi@fukushimas.net. I am grateful to Toni Braun and Fumio Hayashi for their encouragement and support throughout the project. I would also like to thank Deborah Lucas for her comments and criticism, Kenji Wada for his suggestions, and John Cochrane, Lars Hansen, Hiro Ishise, Masao Ogaki, Yuichiro Waki, and workshop participants at the University of Tokyo for useful conversations. Remaining errors are mine.

1 Introduction

Casual observations suggest that certain consumption needs exhibit strong complementarities with household characteristics—such as family size, location of residence, or occupation of the wage-earner—that are not amenable to smooth adjustments over time for a variety of reasons. The number of children is closely related to the amount of food, clothing, housing, and medical services that must be consumed to maintain decent standards of living, but cannot be continuously updated because it takes time for children to grow up. The transportation services required to make a living will necessarily depend on where one lives, but moving homes too often can be unaffordable because of trading frictions in the housing market. And being a professional athlete may require more health maintenance than being an academic, but switching from one to the other can be too costly because of the associated losses in occupation-specific skills. A straightforward implication of these observations is that certain consumption components are harder to adjust over time than others.

The purpose of this paper is to show how the adjustment frictions that arise in this way can help explain a number of time series properties of aggregate consumption and asset prices that have proved puzzling from the perspective of standard consumption based asset pricing models. The paper develops a two-good consumption based model that incorporates this type of friction and compares it to U.S. data on aggregate stock price indexes and short-term interest rates as well as several alternative measures of aggregate and disaggregated consumption. The model maps a stream of near-i.i.d. aggregate consumption growth rates with low volatility into a process representing stock returns and interest rates. It is shown that the stochastic properties of these variables are reasonably in line with their data counterparts. While only partial, the model explains to a quantitatively significant extent the following facts: the high equity premium; the high volatility of stock prices; the low level and volatility of risk-free interest rates; the high persistence of price-dividend ratios and their ability to forecast future excess returns; and the fact that expected returns, stock market volatility, and the market price of risk seem to move countercyclically over the business cycle. The model does not require high relative risk aversion or negative time discount rates to do this.

The analysis in this paper relates to the literature as follows. First, it builds on the basic framework of Lucas (1978), and extends it by subjecting households to certain types of consumption adjustment frictions. A similar line of inquiry was pursued by Grossman and Laroque (1990), Lynch (1996), Marshall and Parekh (1999), and Gabaix and Laibson (2002). While close in spirit, this paper departs from this earlier literature in that the model determines asset prices endogenously and that it features two types of consumption goods. The former feature makes it easier to deal sensibly with the rich temporal dependence of certain asset return characteristics that have been documented in the empirical literature, and the latter is essential for deriving the connections with habit persistence and luxury goods that are discussed below.

The second strand in the literature to which the analysis relates is that of habit persistence, due to Constantinides (1990) and Campbell and Cochrane (1999) among others. In the model, the consumption component that is subject to adjustment frictions behaves as a time-varying subsistence level that adjusts

slowly in response to past consumption shocks. This helps the model explain the behavior of asset prices through a mechanism that resembles habit persistence. The idea that consumption adjustment frictions could generate habit-like effects was anticipated by a number of authors and was recently confirmed by Chetty and Szeidl (2004). The relationship between the two mechanisms as derived in this paper is in fact almost an exact replication of what Chetty and Szeidl arrived at, and not surprisingly, since the model is one that embeds a discrete-time analogue of the household decision problem they studied in a general equilibrium framework.

Finally, the model relates to the luxury-goods consumption based model of Ait-Sahalia, Parker, and Yogo (2004) in an important way. Both models share the property that certain consumption components—luxury goods consumption in their model and the consumption component that is “smoothly adjustable” in the present model—provide the relevant measure of household marginal utilities, and that the stochastic properties of those components are sufficiently different from those of aggregate consumption to justify the behavior of asset prices. The mechanisms behind this common conclusion, however, are different for the two models; in particular, Ait-Sahalia, Parker, and Yogo resort to limited stock market participation and non-homothetic preferences, neither of which are used in the present study. Also, while Ait-Sahalia, Parker, and Yogo provide direct evidence that luxury goods consumption does in fact behave differently from aggregate consumption, the economic forces behind this difference are not well understood. This paper complements their study by providing a possible explanation for this “puzzle” as well.

The paper is organized as follows. Section 2 sets up the model, and Section 3 characterizes its equilibrium. Section 4 compares to model to several related models in the literature. Section 5 provides a quantitative assessment of the model. The final section concludes.

2 The Model

2.1 Description of the Economy

The economy is populated by a unit measure of households indexed by $i \in (0, 1)$. Their consumption in each period $t = 0, 1, 2, \dots$ consists of two components $C_{1t}(i)$ and $C_{2t}(i)$, which will be referred to as *discretionary consumption* and *precommitted consumption*, respectively. Preferences are identical:

$$E_0 \left[\sum_{t=0}^{\infty} \beta^t U(C_{1t}(i), C_{2t}(i)) \right]. \quad (1)$$

Here, $\beta \geq 0$ is a subjective time discount factor, $E_0[\cdot]$ is an expectations operator conditional on information available at time 0, and the utility kernel $U(\cdot, \cdot)$ is taken to have Houthakker’s (1960) addilog form

$$U(C_1, C_2) = \sum_{\ell=1}^2 \psi_{\ell} \frac{C_{\ell}^{1-\gamma} - 1}{1-\gamma} \quad (\psi_1 \equiv \psi, \psi_2 = 1). \quad (2)$$

Households receive Y_t units of endowments in each period, which can be converted into consumption goods using a linear “shopping technology”; let $P_{\ell t}$

for $\ell = 1, 2$ be the “relative prices” of the two goods.¹ Households also have access to a complete set of state contingent claims paying off in endowments. Each household faces a sequence of budget constraints

$$P_{1t}C_{1t}(i) + P_{2t}C_{2t}(i) + E_t[M_{t+1}A_{t+1}(i)] \leq Y_t + A_t(i) \quad (3)$$

where M_{t+1} is the stochastic discount factor between periods t and $t + 1$ and $A_{t+1}(i)$ is household i 's asset holdings between periods t and $t + 1$. A borrowing limit

$$W_{t+1}(i) \equiv A_{t+1}(i) + E_{t+1} \left[\sum_{j=1}^{\infty} M_{t+1 \rightarrow t+j} Y_{t+j} \right] \geq 0, \quad (4)$$

imposed for each t over all states of nature in $t + 1$, prohibits households from running Ponzi schemes. Here, $M_{t+1 \rightarrow t+j} \equiv \prod_{k=1}^j M_{t+k}$ is the stochastic discount factor over a horizon of j periods, and $W_{t+1}(i)$ has the interpretation of present-value wealth in period $t + 1$.

In addition to this, household decisions are also subject to frictions of the type mentioned in the Introduction, which will be formulated as follows. First, a vector $X_t(i)$ records household i 's various family characteristics such as family size, area of residence, or occupation of the wage-earner. The consumption component $C_{2t}(i)$, on the other hand, represents a composite of consumption service flows that exhibit complementarities with $X_t(i)$, such as certain types of food, housing, medical services, transportation services, or health care. Deviations from a relationship of the form

$$C_{2t}(i) = h[X_t(i)] \quad (5)$$

are then subject to certain costs, such as those associated with failing to feed one's children, commute to work, or keep one's health conditions suitable for one's occupational endeavors.² We capture this by assuming that these costs are sufficiently large so that households have no incentive to deviate from (5) at any time; households thus take the relation as a *constraint* in their decision problems. And finally, the choice of $X_t(i)$ is subject to adjustment frictions to capture the fact that it is costly, and in some cases impossible, to instantaneously adjust the number of children, move homes, or switch occupations. For the sake of tractability, we choose to apply a Calvo (1983) type adjustment mechanism to capture these costs:³ $X_t(i)$ is predetermined as of date t , and in any period, only a random fraction θ of the whole population—chosen independently of any other event or variable—is given the chance to choose $X_{t+1}(i)$ differently from $X_t(i)$; those who are not given this chance must set $X_{t+1}(i) = X_t(i)$. From the

¹In the terminology of general equilibrium theory, the P_ℓ 's are “technology parameters” rather than “prices,” since the model will be closed as an endowment economy.

²The interpretation of the nature of these “costs” is, because they are not explicitly modelled, flexible. They may be thought of as coming in pecuniary terms (e.g., an income loss occurs if one doesn't commute to work) or utility terms (e.g., it is painful to see one's children starve).

³The use of a Calvo type adjustment rule distinguishes the present model from most of its predecessors: Grossman and Laroque (1990) and Marshall and Parekh (1999) both feature explicit adjustment costs, and Lynch (1996) and Gabaix and Laibson (2002) allow households to adjust their consumption at fixed intervals. While these alternative formulations have their virtues, so far no model with such frictions has been able to successfully determine asset prices endogenously.

assumptions made above then follows the simplification that households will behave as if their choice of $C_{2t}(i)$ is subject to the same adjustment frictions as $X_t(i)$ is, and that once this is taken into account, $X_t(i)$ can be substituted out from their decision problems. The constraint on the choice of $C_{2t}(i)$ that arises thus will be referred to as the *precommittedness* of $C_{2t}(i)$. For later reference, let $I_t(i)$ be an “indicator” such that $I_t(i) = 1$ if household i is given a chance to choose of $C_{2,t+1}(i)$ differently from $C_{2t}(i)$ and $I_t(i) = 0$ if not.

In summary, then, households in this economy take as given a sequence $\{Y_t, I_t(i), P_{1t}, P_{2t}, M_{t+1}\}_{t=0}^{\infty}$ and initial conditions $A_0(i), C_{2,0}(i)$, and choose a plan $\{C_{1t}(i), C_{2,t+1}(i), A_{t+1}(i)\}_{t=0}^{\infty}$ to maximize lifetime utility (1) subject to the budget constraints ((3) and (4)) and $C_{2,t+1}(i) = C_{2,t}(i)$ if $I_t(i) = 0$. The household decision problem, simplified in this way, is essentially a discrete-time analogue of one studied by Chetty and Szeidl (2004).

We close the model as an endowment economy by an aggregate resource constraint

$$P_{1t}C_{1t} + P_{2t}C_{2t} = Y_t \quad (6)$$

where $C_{\ell t} \equiv \int_0^1 C_{\ell t}(i) di$ for $\ell = 1, 2$. The model takes an initial condition and $\{Y_t, P_{1t}, P_{2t}\}_{t=0}^{\infty}$ as given, and determines the consumption allocation and asset prices endogenously. In principle, closing the model in this way does not entail a conceptual loss of generality: the allocations and asset prices in this economy will coincide with those in a two-good production economy if the statistical model of the exogenous sequence $\{Y_t, P_{1t}, P_{2t}\}_{t=0}^{\infty}$ is the same as the equilibrium sequence of aggregate consumption and good prices in the production economy.

2.2 Definition of Equilibrium

Given an initial condition $\{A_0(i), C_{2,0}(i), I_0(i)\}_{i \in (0,1)}$ and exogenous process $\{Y_t, P_{1t}, P_{2t}, \{I_{t+1}(i)\}_{i \in (0,1)}\}_{t=0}^{\infty}$, an equilibrium for this economy is a stochastic process $\{\{C_{1t}(i), C_{2t}(i), A_{t+1}(i)\}_{i \in (0,1)}, R_{t+1}\}_{t=0}^{\infty}$ such that: (i) for each i , the plan $\{C_{1t}(i), C_{2,t+1}(i), A_{t+1}(i)\}_{t=0}^{\infty}$ solves the corresponding household’s problem; and (ii) for each $t \geq 0$ there holds (6) and $\int_0^1 A_{t+1}(i) di = 0$ almost surely.

2.3 Specification of Exogenous Variables

In order to allow for a numerical solution and evaluation of the model, we specify the initial condition $\{A_0(i), C_{2,0}(i), I_0(i)\}_{i \in (0,1)}$ and the distribution of the exogenous variables $\{Y_t, P_{1t}, P_{2t}\}_{t=0}^{\infty}$ as follows. First, for simplicity, the “relative prices” of the consumption goods are fixed to unity for all t : $P_{1t} = P_{2t} = 1$.⁴ Although this assumption is somewhat restrictive, it provides a convenient context for demonstrating that the model generates all of its interesting implications from fluctuations in aggregate consumption endowments without resorting to additional sources of risk.

Second, the law of motion for aggregate and individual endowments is taken to be

$$\Delta y_t = g + \varepsilon_t, \quad (7)$$

where $y_t \equiv \log Y_t$; lower cases will represent variables in logs hereafter. Here, $\{\varepsilon_t\}_{t=0}^{\infty}$ follows a finite-support distribution that approximates $\varepsilon_t \sim \text{i.i.d. } N(0, \sigma^2)$

⁴This would arise in a production economy if each of the two consumption goods are produced using the same technology.

in a certain sense; the exact specification will be given in Section 3.1.⁵ This specification appears to be the simplest way to capture the salient features of aggregate consumption data, and is used frequently in the consumption based asset pricing literature.

Finally, to simplify the exposition to some extent, we take individual initial conditions so that all households are treated symmetrically at time 0. It is not possible to assume that households are perfectly ex-ante identical, however, since a fraction θ of them must start with $I_0(i) = 1$ whereas others begin with $I_0(i) = 0$. We therefore make the following assumptions. First, we set $C_{2,0}(i) = C_{2,0}$ for all i , which is to say that all households start with the same amount of precommitted consumption. Second, to “counteract” the initial heterogeneity in $I_0(i)$, we take the initial wealth distribution $\{A_0(i)\}_{i \in (0,1)}$ so that household decisions at $t = 0$ satisfy $C_{1,0}(i) = C_{1,0}$ for all i .

3 Characterization of Equilibrium

As in Lucas (1978), the equilibrium is characterized in two steps. In the first step, the consumption allocation is determined without explicit calculation of asset prices. The second step, then, takes this allocation as given and reads asset prices off the implied marginal utilities. The former step requires some elaboration in this setting, since the model endogenously determines the intratemporal allocation of consumption expenditures among the two components in a non-trivial way.

3.1 The Aggregate Consumption Allocation

We begin our characterization of the equilibrium consumption allocation by examining the household’s first order conditions (derivations are given in the Appendix).

The condition for $A_{t+1}(i)$, which applies to all households, gives the following representation of the stochastic discount factor:

$$M_{t+1} = \beta \left(\frac{C_{1,t+1}(i)}{C_{1t}(i)} \right)^{-\gamma}. \quad (8)$$

This is the standard condition that equates the prices of state contingent claims with intertemporal marginal rates of substitution. To understand this point, it is convenient to define household i ’s “total consumption” by $C_t(i) \equiv C_{1t}(i) + C_{2t}(i)$ and note that its “atemporal indirect utility” $\bar{U}(C_t(i))$ (the maximized value of the period utility U when total consumption is $C_t(i)$ and the value of precommitted consumption is $C_{2t}(i)$) is

$$\bar{U}(C_t(i)) \equiv U(C_t(i) - C_{2t}(i), C_{2t}(i)). \quad (9)$$

Expression (8) is then equivalent to $M_{t+1} = \beta \bar{U}'(C_{t+1}(i)) / \bar{U}'(C_t(i))$, a condition familiar from the standard, single-good/power-utility, consumption based

⁵Specifically, $\{\varepsilon_t\}_{t=0}^{\infty}$ will be constructed using the residuals from a finite-state Markov chain approximation of a first-order autoregressive process. This specification, although somewhat roundabout, is adopted to prevent situations in which extremely low realizations of ε cause $C_{1t}(i) < 0$, in which case utility is not defined.

model. The economics behind (8) is exactly the same as that underlying the standard model.

The first order condition for the optimal adjustment of $C_{2,t+1}(i)$ applies to households for which $I_t(i) = 1$ (i.e., those who are given the chance to choose their $C_{2,t+1}(i)$'s differently from $C_{2t}(i)$):

$$C_{2,t+1}(i) = \left\{ (1 - \rho) \text{E}_t \sum_{j=1}^{\infty} \rho^{j-1} \psi C_{1,t+j}(i)^{-\gamma} \right\}^{-1/\gamma} \quad (10)$$

where $\rho \equiv (1 - \theta)\beta$. An interpretation of this condition is that households make their consumption precommitments in a forward-looking manner so as to minimize expected distortions in the intratemporal allocation of consumption expenditures. This condition in conjunction with the adjustment rule that $C_{2,t+1}(i) = C_{2t}(i)$ if $I_t(i) = 0$ determines the time path of precommitted consumption $\{C_{2t}(i)\}_{t=0}^{\infty}$.

It now follows from (8) and the completeness of financial markets together with the assumption $C_{1,0}(i) = C_{1,0}$ for all i that $C_{1t}(i) = C_{1t}$ for all i and t in equilibrium. From this and (10), all households who readjust their $C_{2,t+1}(i)$'s choose the same value; call this value $C_{2,t+1}^*$. Aggregating across households then gives the law of motion for aggregate precommitted consumption:

$$C_{2,t+1} = \theta C_{2,t+1}^* + (1 - \theta)C_{2t}. \quad (11)$$

As demonstrated in the Appendix, substituting the definition of $C_{2,t+1}^*$ in this equation and dividing both sides by Y_t gives an equilibrium characterization of the aggregate expenditure share of precommitted consumption $S_t \equiv C_{2t}/Y_t$. A recursive law of motion $S_{t+1} = f(S_t, \varepsilon_{t+1})$ satisfying that restriction is then found numerically. The procedure, which is detailed in the Appendix, involves a log-linearization of the equilibrium condition to get an AR(1) representation for $\hat{s}_t \equiv \log(S_t) - \log(\bar{S})$

$$\hat{s}_{t+1} = \alpha \hat{s}_t - \varepsilon_{t+1} \quad (12)$$

and a finite-state Markov chain approximation of this law of motion using Tauchen's (1986) methodology. For each t , then, we use the residuals of that chain to determine $\varepsilon_{t+1} = -\hat{s}_{t+1} + \alpha \hat{s}_t$. The law of motion $S_{t+1} = f(S_t, \varepsilon_{t+1})$ in conjunction with (7) then determines the aggregate consumption allocation $\{C_{1t}, C_{2t}\}_{t=0}^{\infty}$.

3.2 Asset Prices

A favorable property of this model is that asset prices can be determined directly from the aggregate allocation $\{C_{1t}, C_{2t}\}_{t=0}^{\infty}$ without finding the consumption profiles of individual households. To see this point, substitute $C_{1t}(i) = C_{1t}$ in (8) and use the definition of S_t to get

$$M_{t+1} = \beta \left(\frac{Y_{t+1}}{Y_t} \right)^{-\gamma} \left(\frac{1 - S_{t+1}}{1 - S_t} \right)^{-\gamma}. \quad (13)$$

The aggregate allocation $\{C_{1t}, C_{2t}\}_{t=0}^{\infty}$ thus determines the stochastic discount factor M_{t+1} via (13), which in turn pins down the gross rate of return R_{t+1} on

any asset through the “basic pricing equation”

$$1 = E_t(M_{t+1}R_{t+1}). \quad (14)$$

In the following, we use (14) and (13) in conjunction with the law of motion $S_{t+1} = f(S_t, \varepsilon_{t+1})$ to solve for asset prices as functions of S_t . These functions in conjunction with the law of motion for S_t fully characterize the data generating process of asset prices.

3.2.1 The Risk Free Asset

The risk free rate R_{t+1}^f is given explicitly by

$$R_{t+1}^f(S_t) = \frac{1}{E_t(M_{t+1})}.$$

This easily follows from (14) and the fact that R_{t+1}^f belongs to the time t information set. The explicit formulas used in the computation are presented in the Appendix.

3.2.2 The Consumption Claim

Following a common practice in the consumption based asset pricing literature, we take a claim to the aggregate consumption endowment stream $\{Y_t\}_{t=0}^{\infty}$ to stand for stocks in the economy. Let Q_t be its price in period t . The one-period return on this asset is $R_{t+1}^s \equiv (Y_{t+1} + Q_{t+1})/Q_t$, so from (14), its price/dividend ratio (or, equivalently, its price/consumption ratio) Q_t/Y_t is characterized as the solution to the functional equation

$$\frac{Q_t}{Y_t}(S_t) = E_t \left[M_{t+1} \left(\frac{Y_{t+1}}{Y_t} \right) \left[1 + \frac{Q_{t+1}}{Y_{t+1}}(S_{t+1}) \right] \right].$$

With the state space of S_t discretized, this reduces to a system of linear equations, which is easy to solve. Stock returns and related statistics were then computed from this function; see the Appendix for details.

4 Comparison to Alternative Models

As is well-known, the asset pricing implications of *any* model is completely summarized by its stochastic discount factor (see Cochrane (2005) for a detailed exposition). Here, we exploit this fact to compare the present model to three classes of models: the standard consumption based model, the habit persistence models of Constantinides (1990) and Campbell and Cochrane (1999), and the luxury goods consumption based model of Ait-Sahalia et al. (2004).

4.1 Difference From the Standard Model

The standard consumption based asset pricing model with a single consumption good and power utility prices assets through a stochastic discount factor of the form

$$M_{t+1}^S = \beta \left(\frac{Y_{t+1}}{Y_t} \right)^{-\gamma} \quad (15)$$

where Y_t is aggregate consumption. There is now significant evidence against this specification (see, e.g., Campbell (1999) for a survey), and therefore any model must imply a stochastic discount factor different from M^S if it is to successfully account for the behavior of asset prices and aggregate consumption as observed in the data.

The stochastic discount factor for the present model M , given by (14), differs from M^S in that

$$M_{t+1} = M_{t+1}^S \left(\frac{1 - S_{t+1}}{1 - S_t} \right)^{-\gamma}. \quad (16)$$

The factor $((1 - S_{t+1})/(1 - S_t))^{-\gamma}$ multiplying M^S on the right hand side is therefore responsible for any difference from, or improvement over (if any), the standard model. This factor can be thought of summarizing the contributions of consumption precommitments: since C_{2t} can adjust only sluggishly, the expenditure share $S_t = C_{2t}/Y_t$ fluctuates over time, which in turn affects the asset pricing implications of the model through (16).

4.2 Similarity to Habit Persistence

In the habit persistence models of Constantinides (1990) and Campbell and Cochrane (1999), assets are priced through a stochastic discount factor of the form

$$M_{t+1}^H = \beta \left(\frac{Y_{t+1} - Z_{t+1}}{Y_t - Z_t} \right)^{-\gamma}, \quad (17)$$

where Y_t is aggregate consumption and Z_t is a time-varying subsistence level, also referred to as a “habit stock,” which adjusts slowly in response to past consumption. Constantinides (1990) and Campbell and Cochrane (1999) show that this specification does substantially better than the standard model with (15) in explaining the joint time evolution of aggregate consumption and asset prices.

The stochastic discount factor for the present model, on the other hand, can be written as:

$$M_{t+1} = \beta \left(\frac{Y_{t+1} - C_{2,t+1}}{Y_t - C_{2t}} \right)^{-\gamma}. \quad (18)$$

Comparing this with (17) reveals an obvious similarity that was discovered in the earlier work of Chetty and Szeidl (2004): aggregate precommitted consumption C_{2t} behaves as a “habit stock” that moves slowly over time—though in a forward-looking rather than a backward-looking manner—affecting asset prices through a functional form identical to (17). The mechanics of the model that drive its implications with regard to asset prices are therefore very similar to those of the habit persistence models of Constantinides (1990) and Campbell and Cochrane (1999).

4.3 Relation to Luxury Goods

The luxury-goods consumption based model of Ait-Sahalia et al. (2004) considers an economy in which households have non-homothetic preferences over two consumption goods—basic goods and luxury goods—and stock market participation is limited. They then argue that the stochastic discount factor that is

relevant for pricing stocks takes the form:

$$M_{t+1}^L = \beta \left(\frac{L_{t+1} + b}{L_t + b} \right)^{-\gamma} \left(\frac{P_{Lt}}{P_{L,t+1}} \right), \quad (19)$$

where L_t is luxury goods consumption, P_{Lt} is the relative price of luxury goods, and b is a small constant parameterizing the utility function (a negative “subsistence level”). Using several measures of luxury goods consumption, they obtain substantial evidence in favor of this specification by estimating the model’s Euler equations and expected-beta representations.

The stochastic discount factor for the present model, on the other hand, can be written as

$$M_{t+1} = \beta \left(\frac{C_{1,t+1}}{C_{1t}} \right)^{-\gamma}. \quad (20)$$

Under the assumption of constant relative prices $P_{Lt}/P_{L,t+1} = 1$ (as assumed in this paper) and $b \approx 0$ (as assumed in Ait-Sahalia et al. (2004)), this is nearly identical to (19) when $C_{1t} = L_t$. Identifying C_1 with L in this way is justified when the measures of luxury goods used by Ait-Sahalia et al. (2004)—including luxury retail sales (Bulgari, Gucci, Hermès, LVMH, Saks, Tiffany, and Waterford Wedgwood), imported luxury automobiles (BMW, Mercedes, Jaguar, and Porsche), and charitable contributions by wealthy households—can be thought of as “purely discretionary” (i.e., do not exhibit complementarities with family size etc.), which seems reasonably appropriate. Under this interpretation, Ait-Sahalia, Parker, and Yogo’s evaluation of (19) can also be interpreted as an evaluation of (20), and their empirical findings can be taken as direct evidence in favor of the present model’s specification. Several facts related to this observation will be exploited in the quantitative analysis that follows.

5 Quantitative Analysis

This section solves the model numerically and evaluates its ability to explain the time series properties of aggregate consumption and asset prices as observed in a century-long U.S. data set covering years 1890-1995 used by Campbell (1999). This data is available on John Campbell’s website. The main attraction of this data set is its length, which is important given the high persistence of many asset pricing statistics. The time period for the model is one year.

5.1 Calibration

The model has six free parameters that need to be calibrated in order to solve the model numerically: g and σ are the mean and standard deviation of the log consumption endowment growth rate, γ determines the curvature of the utility kernel, β is the subjective discount factor, θ is the reciprocal of the adjustment frequency of the family status variable (and hence of precommitted consumption), and ψ governs the average expenditure share of precommitted consumption. In the following, we find it convenient to re-parameterize ψ by \bar{S} , using a one-to-one mapping between the two implied by the model. The parameter choices are summarized in Table 1.

Table 1: Parameter Choices

Parameter	Variable	Value
Mean consumption growth (%)	g	1.80
Standard deviation of consumption growth (%)	σ	3.20
Utility curvature	γ	2.50
Subjective discount factor	β	0.99
Adjustment frequency parameter	θ	1/15
Steady state expenditure of C_2	\bar{S}	0.70

Note: All values are annual.

The first four parameters are chosen as follows. We take g and σ to match the mean and standard deviation of aggregate consumption log growth rates Δy as observed in the data. We also set $\gamma = 2.50$ to keep the model’s implied relative risk aversion coefficient around 3 on average (see Section 5.3.2 for details). And for the subjective discount factor, we simply set $\beta = 0.98$.

In order to calibrate θ , we exploit the implication of the theory that a household’s family status $X_t(i)$ is adjusted every $1/\theta$ years on average, and take $1/\theta = 15$ as a benchmark. This choice is mainly motivated by the consideration that a primary component of $X_t(i)$ should be family size, which is arguably very persistent: raising children, for example, can take around 20 years in total. This suggests $1/\theta = 20$. An adjustment to $1/\theta = 15$ is made in view of the fact that there are other elements in $X_t(i)$, such as residential or occupational status, that are presumably easier to adjust—for instance, Kambourov and Manovskii (2004) report that the average fraction of workers changing occupations or industries within a year has ranged between 10 to 20 percent for the U.S. during the period 1968-1997.

Finally, drawing on an interpretation discussed above in Section 4.3, we choose \bar{S} so that the volatility of discretionary consumption growth Δc_1 implied by the model roughly matches that of luxury goods consumption growth as measured by Ait-Sahalia et al. (2004). Under the benchmark choice $\bar{S} = 0.70$, the former is about 11 percent, while the latter ranges between 10 and 20 percent depending on the specific measure. The motivation for adopting this calibration strategy derives from the following two considerations. The first is that the model appears to be reasonably well-identified along this dimension in that a high choice of \bar{S} consistently implies a high volatility of Δc_1 . The intuition for this is simple: if a large fraction of Y is occupied by a slowly-moving component C_2 , the residual $C_1 = Y - C_2$ has no choice but to vary a lot. The second consideration is that since the model’s stochastic discount factor has the form (20) with a modest curvature parameter $\gamma = 2.50$, the model’s quantitative performance as an asset pricing model depends crucially on the volatility of Δc_1 . Choosing \bar{S} so that this volatility has direct empirical support from the data therefore disciplines the analysis to a considerable extent. A minor drawback of this strategy is that it is “indirect” in the sense that it chooses an expenditure share parameter without using any data on the expenditure shares of actual consumption categories. A preliminary attempt in addressing this point is made in Section 5.4 below, where we draw on data evidence from the U.S. National Income Accounts (NIA) on how aggregate consumption breaks up into various

Table 2: Means and Standard Deviations

Statistic	Data	Model
$E(r^s)$	6.59	5.11
$\sigma(r^s)$	18.5	20.4
$E(r^f)$	2.00	1.41
$\sigma(r^f)$	8.84	5.07
$E(r^s - r^f)$	4.59	3.71
$\sigma(r^s - r^f)$	18.4	19.5
$E(r^s - r^f)/\sigma(r^s - r^f)$	0.25	0.19
$\exp(E(p - d))$	21.9	31.2
$\sigma(p - d)$	0.28	0.31

Notes: E and σ represent means and standard deviations, respectively. r^s is the log real return on stocks, r^f is the log real return on risk-free assets, and $p - d$ is the log price/dividend ratio. All returns are in annual percentages.

components—such as food or housing—to obtain an estimate of \bar{S} .

5.2 Asset Pricing Implications

We now simulate the model using the parameters discussed in the previous section and compare the results to Campbell’s (1999) data. In the following exercise, the risk free rate R^f is taken to stand for real returns on 3-month T-bills (1931-1995), Treasury Certificates (1920-1930), and Prime Commercial Paper (1890-1920), and the price of the consumption claim Q is interpreted as the Standard and Poors 500 stock price index. Model statistics are sample averages over 100,000 simulations, each of length 106 (equal to the number of observations in the data sample).

5.2.1 Means and Standard Deviations

Table 2 reports the means and standard deviations of several variables of interest from the data together with their model counterparts. Judged from the first four rows in Table 2, the model captures a salient characteristic of the data that stock returns are significantly higher in level and more volatile than the risk-free interest rate. The high average return on stocks is in part reflected in the equity premium $E(r^s - r^f)$, where the model accounts for about 80 percent of its historical value. This, combined with an excess return volatility slightly higher than in the data, makes the model generate a log Sharpe ratio $E(r^s - r^f)/\sigma(r^s - r^f)$ somewhat lower than in the data. The model’s risk-free rate, on the other hand, is low and smooth, and in fact slightly errs on the side of interest rate smoothness—a point which may be contrasted with a general tendency in the literature for otherwise successful models to generate risk-free rates that are too volatile, and sometimes even more volatile than stock returns. Since the model generates these features with $\beta < 1$, it partially resolves the risk-free rate puzzle of Weil (1989). Another interesting characteristic of the data, also captured by the model, is the high volatility of stock prices as measured by the standard deviation of the log price/dividend ratio. How this finding

Table 3: Autocorrelations

Variable	Lag (Years)				
	1	2	3	4	5
$p - d$					
Data	0.80	0.61	0.57	0.50	0.39
Model	0.80	0.63	0.50	0.39	0.30
$r^s - r^f$					
Data	0.02	-0.21	0.10	-0.02	-0.16
Model	-0.03	-0.03	-0.02	-0.02	-0.02
$\sum_{k=1}^j \rho(r_t^e, r_{t-i}^e)^*$					
Data	0.02	-0.19	-0.10	-0.11	-0.27
Model	-0.03	-0.06	-0.09	-0.11	-0.13
$ r^s $					
Data	0.06	-0.11	0.08	0.05	-0.11
Model	0.02	0.02	0.01	0.00	0.00

Note: $r^e \equiv r^s - r^f$ is the log excess real return on stocks. The statistic * is the partial sum of its autocorrelation up to lag j .

relates to the volatility test literature will be discussed in Section 5.2.3 below. The average level of the price/dividend ratio implied by the model, however, is somewhat higher than in the data, possibly reflecting the fact that the model does not provide a full account of the average level of stock returns.

5.2.2 Autocorrelations and Cross-correlations

Table 3 summarizes the autocorrelations of several variables from the data and the model. The autocorrelations of the log price/dividend ratio indicate that the model captures its high persistence that characterizes the data. The negative autocorrelations of excess stock returns show mean-reversion as documented by Fama and French (1988b) and Poterba and Summers (1988). Perhaps this point

Table 4: Cross-Correlations

Variable	Lag (Years)				
	1	2	3	4	5
$p_t - d_t, r_{t+j}^e$					
Data	-0.24	-0.22	-0.08	-0.19	-0.09
Model	-0.10	-0.08	-0.07	-0.06	-0.06
$r_t^e, r_{t+j}^e $					
Data	-0.19	0.05	0.17	-0.05	-0.01
Model	-0.08	-0.07	-0.06	-0.05	-0.04
$p_t - d_t, r_{t+j}^e $					
Data	-0.16	0.02	0.04	-0.17	-0.10
Model	-0.17	-0.13	-0.11	-0.09	-0.07

Table 5: Predictability of Excess Returns

$$r_{t \rightarrow t+k}^e = a + b(p_t - d_t)$$

Horizon k	Data		Model	
	$10 \times b$	R^2	$10 \times b$	R^2
1	-1.7	0.06	-0.7	0.02
2	-3.4	0.12	-1.4	0.03
3	-4.2	0.13	-2.0	0.05
4	-5.3	0.16	-2.5	0.06
5	-6.8	0.22	-2.9	0.07
6	-7.6	0.25	-3.3	0.07
7	-8.5	0.28	-3.7	0.08

Notes: The dependent variable in the regression above is the k -year log excess return on stocks over the risk free security, and the regressors are a constant and the log price/dividend ratio.

is more vivid in the partial sums of these autocorrelations, where the model roughly matches the quantitative magnitudes and pattern observed in the data. Also, although somewhat small in quantitative magnitude, the model generates a positive autocorrelation in absolute stock returns, indicating volatility persistence that has been emphasized in the autoregressive conditional heteroskedasticity (ARCH) literature (see Bollerslev, Chou, and Kroner (1992, Section 3.6) for a survey).

Table 4 reports the cross-correlation structure of several variables. Here, the negative correlation between the log price/dividend ratio and future excess stock returns shows that the model captures a well-documented pattern of return predictability that unusually low stock prices correspond with unusually high returns in subsequent periods. This point will be examined from another point of view in Section 5.2.3 below. The cross-correlation between excess stock returns or price/dividend ratios with absolute excess returns in future periods, both of which are negative, indicate that model also fits another familiar pattern that large declines in stock prices are associated with high return volatility in future periods—a “leverage effect” documented by Black (1976) and many others.

5.2.3 Return Predictability and Stock Price Volatility

Table 5 reports results from regressions of long-horizon excess stock returns over future periods on current log price/dividend ratios. The data exhibits a familiar pattern documented by Campbell and Shiller (1988) and Fama and French (1988a): the coefficients are negative so that low stock prices predict high excess returns over future periods, and the R^2 statistics rise sharply with the horizon. The model replicates this pattern to a reasonable extent, although the magnitude of the both the coefficients and the R^2 statistics are smaller than in the data.

Closely related to the issue of return predictability is the “excess” volatility of stock prices. A striking finding from the volatility test literature starting with LeRoy and Porter (1981) and Shiller (1981) is that stock prices are too

Table 6: Variance Decompositions

Source	Returns (%)	Dividends (%)
Data	0.96	-0.16
Model	1.04	-0.06

Note: The table reports the sample analogues of the first and second terms in (21), both truncated at 15 lags and multiplied by 100.

volatile to be explained by varying expectations of future dividend growth and interest rates. Instead, as pointed out by Cochrane (1992), nearly all variation in the price/dividend ratio appears to be coming from changes in expected returns. Cochrane’s analysis is based on a log-linearization of the identity $1 = (R_{t+1}^s)^{-1}R_{t+1}^s$ around $P_t/D_t = P/D$, which yields, in the absence of bubbles:

$$\text{Var}(p_t - d_t) \approx \sum_{j=1}^{\infty} \xi^j \text{Cov}(p_t - d_t, \Delta d_{t+j}) - \sum_{j=1}^{\infty} \xi^j \text{Cov}(p_t - d_t, r_{t+j}^s) \quad (21)$$

where $\xi \equiv (P/D)/[1 + (P/D)]$. This formula states that, ignoring the possible contributions of “market irrationality,” the volatility of the log price/dividend ratio is fully explained by the predictable variation of future dividend growth (the first term) and that of future returns (the second term). Table 6 reports the estimated contributions of each of these two components. The results show that the model generates a pattern similar to that in the data, with almost all variation in the log price/dividend ratio coming from predictable changes in expected returns.

5.2.4 Asset Returns Over the Business Cycle

Figures 1 and 2 summarize the time dependence of several conditional moments from the model by showing their values as functions of the single state variable S_t . In interpreting these figures, it is useful to note the following. First, the stationary distribution of S is nearly normal, peaks at the steady state value $\bar{S} = 0.70$, and assigns virtually zero probability to S outside of the region (.55, .85). The nature of the time evolution of S_t can be inferred from this in conjunction with the fact that it exhibits slow mean reversion (which follows from (12) and $\alpha \approx 0.87$). Second, S has the interpretation of a “recession status variable” that is high during recessions and low during booms: since C_2 moves sluggishly over time, a recession that causes Y to drop sharply makes $S = C_2/Y$ go up. Combining these two observations with the two figures, one can get an idea of how the conditional moments they depict vary over the business cycle.

Figure 1 depicts two curves, representing the conditional expectation (lower line) and standard deviation of excess stock returns (upper line) respectively, both conditional on the current state being S_t . Consistent with the empirical findings of Fama and French (1989) and Schwert (1989) among others, the curves are upward-sloping, indicating that the two variables they depict vary countercyclically and persistently over the business cycle. The two curves do not move one-for-one, however, and the ratio of the former to the latter rises with S . This is reflected in Figure 2, where the two curves depict the two variables

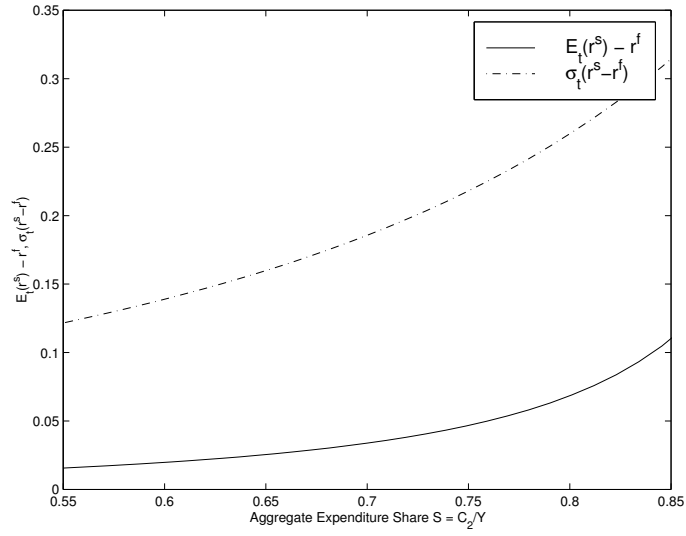


Figure 1: Conditional expectation and standard deviation of excess stock returns as functions of S_t .




Figure 2: The conditional Sharpe ratio and the slope of the mean-variance frontier (or the market price of risk) as functions of S_t .

Table 7: Correlations with the Stochastic Discount Factor

Correlation of m with:		
Δy	r^s	Δc_1
-0.973	-0.997	-1.000

Note: $m \equiv \log M$ is the log stochastic discount factor.

that appear in the Hansen and Jagannathan (1991) inequality:

$$\frac{E_t(R_{t+1}^s - R_{t+1}^f)}{\sigma_t(R_{t+1}^s - R_{t+1}^f)} \leq \frac{\sigma_t(M_{t+1})}{E_t(M_{t+1})}. \quad (22)$$

The lower line in the figure is the left hand side of this equation—the conditional Sharpe ratio—and the upper line is the right hand side—the slope of the mean-variance frontier, or the “market price of risk”—both as functions of S_t . Here, again, the pattern is countercyclical, consistent with evidence documented by Chou, Engle, and Kane (1992).

5.2.5 Correlations Between the Stochastic Discount Factor and Its Proxies

Table 7 reports the correlation between the log stochastic discount factor m and three variables that major alternative models take as proxies of m : the Capital Asset Pricing Model (CAPM) of Sharpe (1964) and Lintner (1965) prices assets using returns on the “market portfolio,” which may be interpreted as $m \propto r^s$ in the present model; the standard consumption based model with power utility assumes $m \propto \Delta y$; and the luxury-goods consumption based model of Ait-Sahalia et al. (2004) can be interpreted as $m \propto \Delta c_1$.

Consistent with evidence from Mankiw and Shapiro (1986) among others, the model implies that the “market return” r^s gives a better proxy of m than does aggregate consumption growth Δy , so that the CAPM will do a better job at pricing assets than does the standard consumption based model. The implication that Δc_1 is even better correlated with m than the other two alternatives is also consistent with Ait-Sahalia, Parker, and Yogo’s finding that their luxury-goods consumption based model does a better job at pricing the cross-section of stock returns than the CAPM or the standard consumption based model.

5.3 Microeconomic Implications

This section presents several implications of the model with regard to consumption and risk aversion. These results serve two purposes: first, they allow for a further assessment of the quality of the model, and second, they help build some intuition on how the model works.

5.3.1 Volatility of Discretionary Consumption

As noted earlier in Section 5.1, one of the most important factors underlying the model’s ability to account for the behavior of asset prices is the volatility of discretionary consumption growth Δc_1 . This assertion follows from Hansen and Jagannathan’s (1991) observation that the standard consumption based model




Figure 3: Conditional standard deviations of Δy and Δc_1 as function of S .

with power utility is unable to account for the observed Sharpe ratio because the low volatility of aggregate consumption growth Δy together with (15) makes the stochastic discount factor insufficiently volatile to satisfy the inequality (22). The present model, whose stochastic discount factor can be written as (20), avoids this problem by making Δc_1 volatile. Its standard deviation, as discussed in Section 5.1, was chosen to be approximately 11 percent based on Ait-Sahalia, Parker, and Yogo’s (2004) finding that the standard deviation of luxury goods consumption growth has been about that magnitude. This confirms the model’s connection to Ait-Sahalia, Parker, and Yogo’s luxury goods model: both models “price” assets using a consumption component that is very volatile, and that makes it possible for them to reconcile a high equity premium with smooth aggregate consumption growth and a low curvature parameter γ .

An issue related to these is the conditional, as opposed to the unconditional, volatility of discretionary consumption growth Δc_1 . As discussed in Cochrane (2005, p. 463), empirical evidence on the countercyclical variation of the market price of risk taken together with the smoothness of risk-free interest rates suggest—through the inequality (22)—that the stochastic discount factor must be conditionally heteroskedastic. The standard consumption based model again fails to fit this pattern because data on aggregate consumption growth Δy does not show strong evidence of conditional heteroskedasticity. The present model gets around this problem by making discretionary consumption growth conditionally heteroskedastic as well. This point is demonstrated in Figure 3, which plots the conditional standard deviations of aggregate consumption growth Δy and discretionary consumption growth Δc_1 as functions of the recession state variable S . As is clear from the figure, discretionary consumption is always conditionally more volatile than aggregate consumption, and especially so during recessions when S is high, exhibiting strong conditional heteroskedasticity. The intuition for this is fairly straightforward: since aggregate consumption Y




Figure 4: Local curvature of the utility function, relative risk aversion, and the curvature parameter of the “average” household as functions of S .

moves around by “about the same amount” at any time—it is nearly conditionally homoskedastic—the larger the fraction occupied by the sluggishly-moving C_2 , the more volatile becomes the residual $C_1 = Y - C_2$. From the perspective of the model, therefore, the time variation of the market price of risk can be understood as a consequence of a time-varying “quantity of risk,” as opposed to time-varying risk aversion, once Δc_1 is viewed as the fundamental risk factor driving asset returns.

5.3.2 Risk Aversion and Utility Curvature

A bulk of the literature on consumption based asset pricing has consisted of attempts to resolve the equity premium puzzle of Mehra and Prescott (1985). The nature of the puzzle is that it is difficult to produce equilibrium asset pricing models that are consistent with the level of the equity premium that is observed in the data without assuming extremely high risk aversion on the part of households. Risk aversion is conventionally measured by the Arrow-Pratt measure of relative risk aversion to wealth bets, which in the present context becomes

$$\text{RRA}_t(i) \equiv -\frac{W_t(i)V_{WW}}{V_W}, \quad (23)$$

where V is the value function for household i . The value function, whose precise definition is given in the Appendix, depends on variables that characterize the state of the individual household— $W_t(i)$, $C_{2t}(i)$, and $I_t(i)$ —as well as those about the aggregate economy— Y_t and C_{2t} . Conventional wisdom suggests that this value should not be much larger than 2 or 3. The middle line in Figure 4 displays the $\text{RRA}_t(i)$ ’s for time- t “average” households—defined as i ’s for which $C_{2t}(i)/C_t(i) = S_t$, $A_t(i) = 0$, and $I_t(i) = 0$ —as a function of the aggregate state

S_t . The curve begins from slightly less than 3 and rises gradually to 4, taking approximately 3.2 around the steady state $\bar{S} = 0.70$. This, together with the high equity premium implied by the model, shows that the model partially resolves the equity premium puzzle.

The top line in Figure 4, on the other hand, plots the local curvature of the household utility function, or the Arrow-Pratt measure of aversion towards fluctuations in total consumption $C_t(i) = C_{1t}(i) + C_{2t}(i)$, defined as

$$\eta_t(i) \equiv -\frac{C_t(i)\bar{U}''(C_t(i))}{\bar{U}'(C_t(i))} \quad (24)$$

where \bar{U} is the atemporal indirect utility function defined in (9). The values reported in the figure are again those of time- t “average” households (in the same sense as above) when the aggregate recession state is S_t . As is clear from the figure, this quantity is significantly higher than both γ and the $\text{RRA}_t(i)$ ’s, and exhibits a strong countercyclical variation over the business cycle. Households in this economy are therefore highly averse to consumption risk, and especially so during severe recessions, even though their attitude toward wealth bets— gambles or lotteries in the usual sense—is only moderate. The model, viewed in this way, interprets the high equity premium as a consequence of this attitude: households *are* risk averse to consumption risk; it is just that that aspect of risk aversion is not reflected in their behavior toward ordinary bets. The time-variation in $\eta_t(i)$ also suggests that the cyclical behavior of the market price of risk can be interpreted as a consequence of time-varying risk aversion, as opposed to one of a time-varying quantity of risk, so long as Δy is viewed as the fundamental risk factor driving asset returns. Viewed in this way, the time-series behavior of the utility curvature confirms the model’s analogy with habit persistence: in both cases, households become more risk averse to consumption risk (i.e., have higher curvature $\eta_t(i)$) during recessions, which allows the model explain the cyclical variation of various asset return characteristics.⁶

5.3.3 Cross-sectional Heterogeneity

Another interesting feature of the model that deserves some attention is that household consumption profiles are ex-post heterogeneous, even though their treatment is ex-ante “symmetrical”—in the sense described earlier in Section 2.3—and financial markets are complete. This point is most clearly understood by observing that the consumption profile of any individual household can be backed out from the aggregate consumption allocation by first constructing the sequence $\{C_{2,t+1}^*\}_{t=0}^\infty$ from $\{C_{2t}\}_{t=0}^\infty$ using (11), and then setting for all

⁶The two setups are not observationally equivalent, however. In particular, it appears that no habit persistence model so far has been able to reconcile smooth i.i.d. consumption growth with a high equity premium, low interest rate volatility, and low relative risk aversion: versions of Constantinides’s (1990) specification generate too much interest rate volatility when consumption growth is i.i.d. and Campbell and Cochrane’s (1999) implies high relative risk aversion. The present model thus, in a sense, achieves a “cream skimming” of the two. This may be thought of as following from fact that in the model, what corresponds to the “habit” (precommitted consumption $C_{2t}(i)$) evolves “internally” (in accordance with the circumstances of individual households) as in Constantinides’s specification but responds non-linearly to past consumption as in Campbell and Cochrane’s. The stochastic discount factor retains a simple form despite these features because the “habit” evolves in a forward-looking, rather than a backward-looking, manner.

$t \geq 0$: $C_{1t}(i) = C_{1t}$, $C_{2,0}(i) = C_{2,0}$, $C_{2,t+1}(i) = C_{2,t+1}^*$ if $I_t(i) = 1$, and $C_{2,t+1}(i) = C_{2t}(i)$ if $I_t(i) = 0$. The ex-post heterogeneity of $\{C_{1t}(i), C_{2t}(i)\}_{t=0}^{\infty}$ then follows from the fact that realizations of $\{I_t(i)\}_{t=0}^{\infty}$ differ across each i .

The source of this unusual—and perhaps seemingly odd—phenomenon is that, roughly speaking, the intratemporal marginal rates of substitutions U_1/U_2 between the two consumption components are not equated across households in equilibrium:⁷ situations can arise in which two households i_1 and i_2 have time t consumption profiles such that $C_{1t}(i_1) = C_{1t}(i_2)$ and $C_{2t}(i_1) > C_{2t}(i_2)$, yet no further trade takes place. The key to understanding the economics behind this result is to recognize that the inequality $C_{2t}(i_1) > C_{2t}(i_2)$ is an empirical manifestation of a deeper picture expressed by $C_{2t}(i_1) = h[X_t(i_1)] > h[X_t(i_2)] = C_{2t}(i_2)$. The first and last equalities say, for example, that each of the two households live in houses that are large enough to accommodate their families, and the middle inequality says that household i_1 has a bigger family than does household i_2 . In order to realize the apparent gain from trade in the situation above, then, household i_1 would have to give away one of its family members to household i_2 . Such “gains” can remain unrealized in equilibrium because households in the model are not allowed to make such transfers smoothly enough in each period.⁸

An interesting corollary of the cross-sectional heterogeneity in consumption that arises in this way is that, since the relative risk aversion coefficient $\text{RRA}_t(i)$ and the local curvature of the utility function $\eta_t(i)$ both depend on $C_{2t}(i)$, the two quantities are endogenously heterogeneous as well. The intuition for this is simple: since these quantities rise as $C_{2t}(i)$ gets closer to $C_t(i)$, households who happen to have a high value of $C_{2t}(i)$ at any point in time become more averse to both wealth risk and consumption risk than the rest of the population. The existence of this cross-sectional heterogeneity appears to be consistent with a number of empirical studies (see, e.g., Barsky, Juster, Kimball, and Shapiro (1997) for experimental evidence), and, together with the time-variation of these quantities over the business cycle as discussed earlier, seems also to suggest a possible line of explanation as to why most studies have had substantial difficulty in obtaining sharp point estimates of these “parameters.” Also, of some potential theoretical interest is that the model provides an example in which utility curvature heterogeneity can be reconciled with a complete-markets, infinite horizon framework that features growth without predicting that a subset of all households will eventually own all wealth and consume all resources in the long-run.

5.4 Inferring \bar{S} From Disaggregated Consumption Data

In the benchmark calibration, the choice $\bar{S} = 0.70$ was based exclusively on information about the volatility of Δc_1 . Here, we make a preliminary attempt

⁷The discussion here is heuristic in that, strictly speaking, it makes sense only if the costs associated with deviations from $C_{2t}(i) = h[X_t(i)]$ are interpreted as coming in pecuniary terms. If those costs come in utility terms, the utility function (of the “underlying model” in which those costs are explicitly modelled) can be discontinuous around $C_{2t}(i) = h[X_t(i)]$ and the intratemporal marginal rate of substitution (again, of the “underlying model”) may not be well-defined.

⁸However, one can think of a portion of the infrequent adjustments in $X_t(i)$ as involving such transfers, so the model is not inconsistent with the fact that such transfers can and do take place in various forms from time to time.

Table 8: Properties of Various Consumption Components

Consumption Category	$E(C^n/Y)$	$\sigma(\Delta c^n)/\sigma(\Delta y)$
Food + Medical care + Recreation	0.44	0.98
Clothing, accessories, and jewelry	0.05	1.77
Personal care	0.02	2.13
Housing	0.17	0.84
Household operation	0.09	1.60
Personal business	0.08	1.38
Transportation	0.08	4.08
Education and research	0.03	1.77
Religious and welfare activities	0.02	1.93

Notes: Data is from the National Income Accounts, annual 1929-2003. $E(C^n/Y)$ is the average expenditure share of the consumption component, and $\sigma(\Delta c^n)/\sigma(\Delta y)$ is the standard deviation of its log growth rates relative to total consumption.

in inferring this quantity using a more “direct” method that makes explicit use of expenditure share data.

In this attempt, we compare the model to NIA data on how aggregate consumption breaks up into various components such as food, clothing, housing, medical care, transportation, recreation, education, and religious and welfare activities. In relating the model variables to these consumption components, we imagine that the time series of each of these components are generated by the model, appended with the relationship

$$C_t^n = \lambda_1^n C_{1t} + \lambda_2^n C_{2t} \quad (25)$$

where C_t^n is the n -th consumption category ($n = 1, \dots, N$) and λ_ℓ^n 's ($\ell = 1, 2$) are constants such that $\sum_{n=1}^N \lambda_\ell^n = 1$ for $\ell = 1, 2$. The time series implied by this relationship may be interpreted as an equilibrium in a $2N$ -good economy in which each C_t^n consists of a “discretionary component” C_{1t}^n and “precommitted component” C_{2t}^n (so that $C_t^n = C_{1t}^n + C_{2t}^n$) and the period utility function is given by

$$\tilde{U}(C_1^1, \dots, C_1^N, C_2^1, \dots, C_2^N) = \sum_{\ell=1}^2 \tilde{\psi}_\ell \frac{[\prod_{n=1}^N (C_\ell^n)^{\lambda_\ell^n}]^{1-\gamma} - 1}{1-\gamma} \quad (26)$$

where $\tilde{\psi}_\ell = \psi_\ell / [\prod_{n=1}^N \lambda_\ell^n]^{1-\gamma}$. In this setting, $C_{\ell t}^n(i) = \lambda_\ell^n C_{\ell t}(i)$ follows from second stage of a two-stage budgeting procedure because the intratemporal aggregator is Cobb-Douglas. Summing this over $\ell = 1, 2$ and integrating over $i \in (0, 1)$ gives (25).

The data facts we use to pin down the parameters—the average expenditure shares and standard deviations of the log growth rates of these disaggregated consumption components relative to that of total consumption—are reported in Table 8. Here, food, recreation, and medical care are lumped into one category to eliminate the strong time trends in the expenditure shares of these categories—which is negative for food and positive for recreation and medical

care—so as to make the mean expenditure share a meaningful statistic.⁹ We limit our focus to nondurables and services consumption at this stage.

Using this data, then, we choose \bar{S} and λ_ℓ^n 's jointly to minimize the sum of squared percentage prediction errors for the statistics reported in Table 8, subject to the requirements that $\bar{S} \in (0, 1)$, $\lambda_\ell^n \in (0, 1)$, and $\sum_{n=1}^N \lambda_\ell^n = 1$ for $\ell = 1, 2$. Notice here that there are $2N-1$ “free” parameters and $2N-1$ independent data facts to be matched. This procedure gives a point estimate of $\bar{S} = .66$. Given this, we adjust this value to correct for the contributions of durable goods consumption—such as furnitures and automobiles—which intuition suggests are important components of precommitted consumption. Since the expenditure share of durable goods has averaged roughly 13 percent during the sampling period, Slesnick’s (1992) finding that the service flows derived from consumer durables goods are fairly close in magnitude to the expenditures on those durable goods suggests that the adjustment be approximately $(.66 + .13)/(1.00 + .13) \approx .70$. The procedure therefore gives some support for the benchmark choice $\bar{S} = 0.70$ from an alternative point of view.

6 Conclusion

This paper developed and evaluated a consumption based asset pricing model in which household consumption choices are subject to a certain type of adjustment friction. The model was solved numerically using parameters values that were selected based on evidence on aggregate and disaggregated consumption as well as several other observations. The time series properties of the model’s endogenous variables were shown to be reasonably consistent with the salient features of asset prices and consumption, explaining substantive portions of many of the associated statistics. The side-effects arising from the frictions incorporated in the model also seem to be minor: the crucial mechanism—the volatility of a certain consumption component—has some direct support from the data on luxury goods consumption presented by Ait-Sahalia et al. (2004), and the model does not require implausibly high relative risk aversion or negative time discounting on the part of households.

As discussed in the paper, the model achieved many of these results through a mechanism that exhibits a striking similarity to habit persistence. In a sense, this finding is reassuring in that it provides an explanation as to why certain macroeconomic time series seem to behave “as if” consumers are forming habits, even though microeconomic evidence in favor of such preference specifications is rather slim (e.g., Dynan (2000)). At a deeper level, however, the connection between the two is potentially disruptive: Since what corresponds to the “habit stock” in the model is endogenously determined by the households’ optimal decision rules, its time evolution will most likely be affected in non-trivial ways when there are important shifts in government policy regimes. To the extent that the habit-like effects captured by and incorporated in many macroeconomic models are in fact manifestations of a mechanism similar to one described in this paper, then, it is conceivable that policy analyses carried out in such frameworks can potentially fall prey to a very subtle version of the Lucas (1976) critique.

⁹A theoretical justification for this procedure may be possible if one is willing to accept the position that the “discretionary/precommitted components” of these categories are perfect substitutes.

Given the wide spread use of habit persistence in a variety of contexts, whether this issue can be safely ignored is certainly of much importance.

Another interesting connection to the literature, also heavily exploited in the analysis, was the near-equivalence between the model's stochastic discount factor and that of Ait-Sahalia, Parker, and Yogo's (2004) luxury goods model. Both models theoretically predict that, measurement issues aside, the growth rates of particular types of consumption goods should serve as excellent risk factors for pricing assets. Further investigations on the extent to which this proposition is "true" enough to be useful are well deserved.

Perhaps the most controversial ingredient of the analysis is the adoption of a Calvo (1983) type adjustment rule. Several authors have noted, primarily in the context of sticky-price monetary models, that this formulation can overstate the associated adjustment frictions in quantitatively important ways. Assessing the robustness of the model along this dimension, which would require a model that features a more "realistic" formulation of the adjustment frictions considered here, would also be an interesting line of further inquiry.

References

- Ait-Sahalia, Yacine, Jonathan A. Parker, and Motohiro Yogo (2004) "Luxury Goods and the Equity Premium", *Journal of Finance*, 59(6): 2959–3004.
- Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro (1997) "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study", *Quarterly Journal of Economics*, 112(2): 537–579.
- Black, Fischer (1976) "Studies of Stock Price Volatility Changes", in *Proceedings of the Business and Economic Statistics Section*, pages 177–181, American Statistical Association, Washington: American Statistical Association.
- Bollerslev, Tim, Ray Y. Chou, and Kenneth F. Kroner (1992) "ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence", *Journal of Econometrics*, 52(1-2): 5–59.
- Calvo, Guillermo A. (1983) "Staggered Prices in a Utility-Maximizing Framework", *Journal of Monetary Economics*, 12(3): 383–398.
- Campbell, John Y. (1999) "Asset Prices, Consumption, and the Business Cycle", in John B. Taylor and Michael Woodford (eds.), *Handbook of Macroeconomics*, volume 1C, Amsterdam: North-Holland.
- Campbell, John Y. and John H. Cochrane (1999) "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior", *Journal of Political Economy*, 107(2): 205–251.
- Campbell, John Y. and Robert J. Shiller (1988) "Stock Prices, Earnings, and Expected Dividends", *Journal of Finance*, 43(3): 661–676.
- Chetty, Raj and Adam Szeidl (2004) "Consumption Commitments: Neoclassical Foundations for Habit Formation", Working paper, UC Berkeley.

- Chou, Ray Y., Robert F. Engle, and Alex Kane (1992) “Measuring Risk Aversion from Excess Returns on a Stock Index”, *Journal of Econometrics*, 52(1-2): 201–24.
- Cochrane, John H. (1992) “Explaining the Variance of Price-Dividend Ratios”, *Review of Financial Studies*, 5(2): 243–280.
- (2005) *Asset Pricing*, Princeton: Princeton University Press, revised edition.
- Constantinides, George M. (1990) “Habit Formation: A Resolution of the Equity Premium Puzzle”, *Journal of Political Economy*, 98(3): 519–543.
- Dynan, Karen E. (2000) “Habit Formation in Consumer Preferences: Evidence from Panel Data”, *American Economic Review*, 90(3): 391–406.
- Fama, Eugene F. and Kenneth R. French (1988a) “Dividend Yields and Expected Stock Returns”, *Journal of Financial Economics*, 22(1): 3–25.
- (1988b) “Permanent and Temporary Components of Stock Prices”, *Journal of Political Economy*, 96(2): 246–273.
- (1989) “Business Conditions and Expected Returns on Stocks and Bonds”, *Journal of Financial Economics*, 25(1): 23–49.
- Gabaix, Xavier and David Laibson (2002) “The 6D Bias and the Equity-Premium Puzzle”, in Ben S. Bernanke and Kenneth Rogoff (eds.), *NBER Macroeconomics Annual*, Cambridge, MA: MIT Press.
- Grossman, Sanford J. and Guy Laroque (1990) “Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods”, *Econometrica*, 58(1): 25–51.
- Hansen, Lars P. and Ravi Jagannathan (1991) “Implications of Security Market Data for Models of Dynamic Economies”, *Journal of Political Economy*, 99(2): 225–262.
- Houthakker, Hendrik S. (1960) “Additive Preferences”, *Econometrica*, 28(2): 244–257.
- Kambourov, Gueorgui and Iourii Manovskii (2004) “Rising Occupational and Industry Mobility in the United States: 1968-1997”, Working paper, University of Pennsylvania.
- LeRoy, Stephen F. and Richard D. Porter (1981) “The Present-Value Relation: Tests Based on Implied Variance Bounds”, *Econometrica*, 49(3): 555–574.
- Lintner, John (1965) “Security Prices, Risk, and Maximal Gains From Diversification”, *Journal of Finance*, 20(4): 587–615.
- Lucas, Robert E., Jr. (1976) “Econometric Policy Evaluation: A Critique”, in Karl Brunner and Allan H. Meltzer (eds.), *The Phillips Curve and Labor Markets*, volume 1 of *Carnegie-Rochester Conference Series on Public Policy*, pages 19–46, Amsterdam: North-Holland.

- (1978) “Asset Prices in an Exchange Economy”, *Econometrica*, 46(6): 1429–1445.
- Lynch, Anthony W. (1996) “Decision Frequency and Synchronization across Agents: Implications for Aggregate Consumption and Equity Return”, *Journal of Finance*, 51(4): 1479–1497.
- Mankiw, N. Gregory and Matthew D. Shapiro (1986) “Risk and Return: Consumption Beta versus Market Beta”, *Review of Economics and Statistics*, 68(3): 452–59.
- Marshall, David A. and Nayan G. Parekh (1999) “Can Costs of Consumption Adjustment Explain Asset Pricing Puzzles?”, *Journal of Finance*, 54(2): 623–654.
- Mehra, Rajnish and Edward C. Prescott (1985) “The Equity Premium: A Puzzle”, *Journal of Monetary Economics*, 15(2): 145–161.
- Poterba, James M. and Lawrence H. Summers (1988) “Mean Reversion in Stock Prices: Evidence and Implications”, *Journal of Financial Economics*, 22(1): 27–59.
- Schwert, G. William (1989) “Why Does Stock Market Volatility Change over Time?”, *Journal of Finance*, 44(5): 1115–53.
- Sharpe, William F. (1964) “Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk”, *Journal of Finance*, 19(3): 425–442.
- Shiller, Robert J. (1981) “Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?”, *American Economic Review*, 71(3): 421–436.
- Slesnick, Daniel T. (1992) “Aggregate Consumption and Saving in the Postwar United States”, *Review of Economics and Statistics*, 74(4): 585–97.
- Tauchen, George (1986) “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions”, *Economics Letters*, 20(2): 177–181.
- Weil, Philippe (1989) “The Equity Premium Puzzle and the Risk-Free Rate Puzzle”, *Journal of Monetary Economics*, 24(3): 401–421.

Appendix

A Derivations and Model Solution

A.1 Deriving the Household First Order Conditions

To obtain the first order conditions for household optimization, note that the household’s decision problem can be formulated as a standard dynamic programming problem in which the “state” at time t consists of $A_t(i)$, $C_{2t}(i)$, $I_t(i)$, Y_t , and C_{2t} , and the “controls” are $C_{1t}(i)$, $A_{t+1}(i)$, and $C_{2,t+1}(i)$. Denote the value function by $V(A_t(i), C_{2t}(i), I_t(i); Y_t, C_{2t})$.

The Bellman equation can be written as

$$\begin{aligned} & V(A(i), C_2(i), I(i); Y, C_2) \\ &= \max_{C_1(i), A'(i), C_2'(i)} \{U(C_1(i), C_2(i)) + \beta E[V(A'(i), C_2'(i), I'(i); Y', C_2')]\} \end{aligned} \quad (27)$$

subject to the constraints

$$C_1(i) + C_2(i) + E[M' A'(i)] \leq Y + A(i) \quad (28)$$

$$C_2'(i) = C_2(i) \text{ if } I(i) = 0 \quad (29)$$

the equilibrium transition laws

$$Y' = e^{\theta + \varepsilon'} Y \quad (30)$$

$$C_2' = Y' f(C_2/Y, \varepsilon') \quad (31)$$

and the equilibrium “pricing function”

$$M' = \beta \left(\frac{Y'}{Y} \right)^{-\gamma} \left(\frac{1 - C_2'/Y'}{1 - C_2/Y} \right)^{-\gamma} \quad (32)$$

where “primes” denote next-period values and the expectation on the right hand side is taken over ε' and $I'(i)$.

Under suitable regularity conditions, the value function V satisfies the Bellman equation and the household’s optimal policy is characterized as one that maximizes the right hand side of the Bellman equation. Assuming that those conditions are satisfied, we can break up the household’s problem into a sequence of “time- t problems”: for each t , choose $C_{1t}(i)$, $A_{t+1}(i)$, and $C_{2,t+1}(i)$ to maximize

$$E_t \left[\sum_{j=0}^{\infty} \beta^j U(C_{1,t+j}(i), C_{2,t+j}(i)) \right] \quad (33)$$

subject to

$$C_{1t}(i) + C_{2t}(i) + E_t[M_{t+1} A_{t+1}(i)] \leq Y_t + A_t(i) \quad (34)$$

$$C_{2,t+1}(i) = C_{2t}(i) \text{ if } I_t(i) = 0 \quad (35)$$

taking the transition laws, equilibrium stochastic discount factor, and the “initial conditions” $W_t(i)$, $C_{2t}(i)$, and $I_t(i)$ as given. This characterization is obtained by expressing $V(A_{t+1}(i), C_{2,t+1}(i), I_{t+1}(i); Y_{t+1}, C_{2,t+1})$ as the utility function evaluated at the optimal decision rule and then substituting it in the right hand side of the Bellman equation.

Clearly, the budget constraint (34) must hold with equality at the optima, so we can substitute it into the objective function (33) to get

$$E_t \left[\sum_{j=0}^{\infty} \beta^j U(Y_{t+j} + A_{t+j}(i) - C_{2,t+j}(i) - E_{t+j}[M_{t+j+1} A_{t+j+1}(i)], C_{2,t+j}(i)) \right]. \quad (36)$$

The time- t problem then reduces to choosing $A_{t+1}(i)$ and $C_{2,t+1}(i)$ to maximize (36) subject to (34) and (35), taking $W_t(i)$, $C_{2t}(i)$, and $I_t(i)$ as given.

A.1.1 Deriving Equation (8)

To derive (8), collect the terms in (36) that involve $A_{t+1}(i)$ to get

$$\begin{aligned} & \mathbb{E}_t[U(Y_t + A_t(i) - C_{2t}(i) - \mathbb{E}_t[M_{t+1}A_{t+1}(i)], C_{2t}(i)) \\ & \quad + \beta U(Y_{t+1} + A_{t+1}(i) - C_{2,t+1}(i) - \mathbb{E}_{t+1}[M_{t+2}A_{t+2}(i)], C_{2,t+1}(i))]. \end{aligned}$$

The first order condition with respect to $A_{t+1}(i)$ is then

$$\begin{aligned} 0 &= M_{t+1}U_1(Y_t + A_t(i) - C_{2t}(i) - \mathbb{E}_t[M_{t+1}A_{t+1}(i)], C_{2t}(i)) \\ & \quad - \beta U_1(Y_{t+1} + A_{t+1}(i) - C_{2,t+1}(i) - \mathbb{E}_{t+1}[M_{t+2}A_{t+2}(i)], C_{2,t+1}(i)). \end{aligned}$$

Substituting (34) and using $U_1(C_1, C_2) = \psi C_1^{-\gamma}$ gives (8).

A.1.2 Deriving Equation (10)

To derive (10), collect the terms in (36) that involve $C_{2,t+1}(i)$ to get

$$\frac{1}{1-\theta} \mathbb{E}_t \left[\sum_{j=1}^{\infty} \rho^j U(Y_{t+j} + A_{t+j}(i) - C_{2,t+1}(i) - \mathbb{E}_t[M_{t+j+1}A_{t+j+1}(i)], C_{2,t+1}(i)) \right]$$

where $\rho \equiv (1-\theta)\beta$. Differentiating this expression with respect to $C_{2,t+1}(i)$ and setting it to zero gives

$$\begin{aligned} 0 &= \mathbb{E}_t \left[\sum_{j=1}^{\infty} \rho^j \{ U_1(Y_{t+j} + A_{t+j}(i) - C_{2,t+1}(i) - \mathbb{E}_t[M_{t+j+1}A_{t+j+1}(i)], C_{2,t+1}(i)) \right. \\ & \quad \left. - U_2(Y_{t+j} + A_{t+j}(i) - C_{2,t+1}(i) - \mathbb{E}_t[M_{t+j+1}A_{t+j+1}(i)], C_{2,t+1}(i)) \} \right] \end{aligned}$$

Substitute back (34), noting that $C_{2,t+j}(i) = C_{2,t+1}(i)$ for all $j \geq 0$ in the expression since the summation and expectation are over those states of nature in which $I_{t+1}(i) = \dots = I_{t+j}(i) = 0$, and use $u_1(C_1, C_2) = \psi C_1^{-\gamma}$ and $u_2(C_1, C_2) = C_2^{-\gamma}$ to get

$$\begin{aligned} 0 &= \mathbb{E}_t \left[\sum_{j=1}^{\infty} \rho^j \{ \psi C_{1,t+j}(i)^{-\gamma} - C_{2,t+1}(i)^{-\gamma} \} \right] \\ &= \mathbb{E}_t \left[\sum_{j=1}^{\infty} \rho^j \psi C_{1,t+j}(i)^{-\gamma} \right] - \frac{\rho}{1-\rho} C_{2,t+1}(i)^{-\gamma} \end{aligned}$$

Rearranging gives (10).

A.2 Characterization of S_t

The condition characterizing the law of motion of S_t is:

$$\begin{aligned} S_{t+1} \left(\frac{Y_{t+1}}{Y_t} \right) - (1 - \theta)S_t \\ = \theta \left\{ \frac{1 - \rho}{\rho} \mathbb{E}_t \sum_{j=1}^{\infty} \rho^j \psi (1 - S_{t+j})^{-\gamma} \left(\frac{Y_{t+j}}{Y_t} \right)^{-\gamma} \right\}^{-1/\gamma} \end{aligned} \quad (37)$$

To derive this, substitute the definition of $C_{2,t+1}^*$ in (11) to get

$$C_{2,t+1} = \theta \left\{ (1 - \rho) \mathbb{E}_t \sum_{j=1}^{\infty} \rho^{j-1} \psi C_{1,t+j}^{-\gamma} \right\}^{-1/\gamma} + (1 - \theta)C_{2t}.$$

From the resource constraints $C_{1,t+j} + C_{2,t+j} = Y_{t+j}$ we have

$$C_{2,t+1} = \theta \left\{ (1 - \rho) \mathbb{E}_t \sum_{j=1}^{\infty} \rho^{j-1} \psi (Y_{t+j} - C_{2,t+j})^{-\gamma} \right\}^{-1/\gamma} + (1 - \theta)C_{2t}.$$

Now, divide both sides by Y_t to get

$$\frac{C_{2,t+1}}{Y_t} = \theta \left\{ (1 - \rho) \mathbb{E}_t \sum_{j=1}^{\infty} \rho^{j-1} \psi \left(\frac{Y_{t+j}}{Y_t} - \frac{C_{2,t+j}}{Y_t} \right)^{-\gamma} \right\}^{-1/\gamma} + (1 - \theta) \frac{C_{2t}}{Y_t}.$$

Rewriting,

$$\begin{aligned} \frac{C_{2,t+1}}{Y_{t+1}} \frac{Y_{t+1}}{Y_t} - (1 - \theta) \frac{C_{2t}}{Y_t} \\ = \theta \left\{ \frac{1 - \rho}{\rho} \mathbb{E}_t \sum_{j=1}^{\infty} \rho^j \psi \left(1 - \frac{C_{2,t+j}}{Y_{t+j}} \right)^{-\gamma} \left(\frac{Y_{t+j}}{Y_t} \right)^{-\gamma} \right\}^{-1/\gamma}. \end{aligned}$$

Using the definition $S_t \equiv C_{2t}/Y_t$ then gives the desired equation (37).

A.3 Steady State and Log-Linearization of (37)

To log-linearize (37), first note that

$$\frac{Y_{t+j}}{Y_t} = e^{gj + \varepsilon_{t+1} + \dots + \varepsilon_{t+j}}.$$

Substitute this into (37) to get

$$\begin{aligned} S_{t+1} e^{g + \varepsilon_{t+1}} - (1 - \theta)S_t \\ = \theta \left\{ \frac{1 - \rho}{\rho} \mathbb{E}_t \sum_{j=1}^{\infty} \tilde{\rho}^j \psi (1 - S_{t+j})^{-\gamma} e^{-\gamma(\varepsilon_{t+1} + \dots + \varepsilon_{t+j})} \right\}^{-1/\gamma} \end{aligned} \quad (38)$$

where $\tilde{\rho} \equiv \rho e^{-\gamma g}$. The steady state value \bar{S} is therefore the solution to

$$\bar{S}e^g - (1 - \theta)\bar{S} = \theta \left\{ \frac{1 - \rho}{\rho} \sum_{j=1}^{\infty} \tilde{\rho}^j \psi (1 - \bar{S})^{-\gamma} \right\}^{-1/\gamma}.$$

This equation is linear in \bar{S} and can thus be solved explicitly as

$$\bar{S} = \frac{\Xi}{e^g - (1 - \theta) + \Xi}$$

where

$$\Xi \equiv \theta \left(\psi \frac{1 - \rho}{\rho} \frac{\tilde{\rho}}{1 - \tilde{\rho}} \right)^{-1/\gamma}.$$

To log-linearize (38) around this \bar{S} , rewrite it as

$$\begin{aligned} & \left\{ \frac{1}{\theta} (\bar{S} e^{\hat{s}_{t+1} + g + \varepsilon_{t+1}} - (1 - \theta) \bar{S} e^{\hat{s}_t}) \right\}^{-\gamma} \\ &= \frac{1 - \rho}{\rho} \psi \mathbb{E}_t \left[\sum_{j=1}^{\infty} \tilde{\rho}^j (1 - \bar{S} e^{\hat{s}_{t+j}})^{-\gamma} e^{-\gamma(\varepsilon_{t+1} + \dots + \varepsilon_{t+j})} \right]. \end{aligned}$$

where $\hat{s}_{t+j} \equiv \log(S_{t+j}) - \log(\bar{S})$. Linearizing this expression around $\hat{s}_{t+j} = 0$, $\varepsilon_{t+j} = 0$ all j and rearranging gives

$$e^g \hat{s}_{t+1} + e^g \varepsilon_{t+1} - (1 - \theta) \hat{s}_t = -\Xi \frac{1 - \tilde{\rho}}{\tilde{\rho}} \mathbb{E}_t \sum_{j=1}^{\infty} \tilde{\rho}^j \hat{s}_{t+j}.$$

To solve this by the method of undetermined coefficients, substitute $\hat{s}_{t+1} = \alpha \hat{s}_t + \zeta \varepsilon_{t+1}$ in the expression (where α and ζ are the undetermined coefficients) to get

$$e^g \alpha \hat{s}_t + e^g \zeta \varepsilon_{t+1} + e^g \varepsilon_{t+1} - (1 - \theta) \hat{s}_t = -\Xi \frac{1 - \tilde{\rho}}{\tilde{\rho}} \frac{\tilde{\rho} \alpha}{1 - \tilde{\rho} \alpha} \hat{s}_t.$$

Comparing coefficients on ε_{t+1} , we have

$$\zeta = -1.$$

Comparing coefficients on \hat{s}_t gives the quadratic equation

$$\phi_0 \alpha^2 + \phi_1 \alpha + \phi_2 = 0$$

where

$$\begin{aligned} \phi_0 &\equiv -e^g \tilde{\rho}, \\ \phi_1 &\equiv e^g + (1 - \theta) \tilde{\rho} + \Xi (1 - \tilde{\rho}), \\ \phi_2 &\equiv -(1 - \theta). \end{aligned}$$

Choosing the $\alpha \in (0, 1)$ that solves this quadratic equation, we obtain the AR(1) law of motion (12). We then apply Tauchen's (1986) methodology to approximate this AR(1) by a 15-state Markov chain, choosing suitably the support of \hat{s}_t so that $S_t = \bar{S} e^{\hat{s}_t} \in (0, 1)$ for any realization of \hat{s}_t . In the following discussion, we will denote the transition and stationary probabilities as $\pi(\hat{s}' | \hat{s}) \equiv \Pr(\hat{s}_{t+1} = \hat{s}' | \hat{s}_{t+1} = \hat{s})$ and $\pi(\hat{s}) \equiv \Pr(\hat{s}_t = \hat{s})$, respectively.

A.4 Asset Prices

To price assets, rewrite the stochastic discount factor as

$$\begin{aligned} M_{t+1} &= \beta \left(\frac{C_{1,t+1}}{C_{1t}} \right)^{-\gamma} \\ &= \beta \left(\frac{Y_{t+1}}{Y_t} \right)^{-\gamma} \left(\frac{1 - S_{t+1}}{1 - S_t} \right)^{-\gamma} \\ &= \beta e^{-\gamma(g - \hat{s}_{t+1} + \alpha \hat{s}_t)} \left(\frac{1 - \bar{S} e^{\hat{s}_{t+1}}}{1 - \bar{S} e^{\hat{s}_t}} \right)^{-\gamma}. \end{aligned}$$

Using this, we can price assets as functions of \hat{s}_t (and hence of S_t) from the pricing equation

$$1 = E_t(M_{t+1} R_{t+1}).$$

A.4.1 The Risk Free Rate

The risk free rate R_{t+1}^f is obtained from the pricing equation, noting that R_{t+1}^f belongs to the time- t information set, so that

$$1 = E_t(M_{t+1} R_{t+1}^f) = E_t(M_{t+1}) R_{t+1}^f$$

or

$$R_{t+1}^f = \frac{1}{E_t(M_{t+1})}.$$

Taking logs and applying the formula for M_{t+1} , we have

$$\begin{aligned} r_{t+1}^f &= -\log E_t(M_{t+1}) \\ &= -\log \sum_{\hat{s}_{t+1}} \beta e^{-\gamma(g - \hat{s}_{t+1} + \alpha \hat{s}_t)} \left(\frac{1 - \bar{S} e^{\hat{s}_{t+1}}}{1 - \bar{S} e^{\hat{s}_t}} \right)^{-\gamma} \pi(\hat{s}_{t+1} | \hat{s}_t). \end{aligned}$$

This gives an expression for r_{t+1}^f as a function of \hat{s}_t : $r_{t+1}^f = r^f(\hat{s}_t)$.

A.4.2 The Price of the Consumption Claim

To obtain the price of the consumption claim Q_t , apply the pricing equation to the expression for returns on this claim

$$R_{t+1}^s = \frac{Y_{t+1} + Q_{t+1}}{Q_t}$$

to get

$$\frac{Q_t}{Y_t} = E_t \left[M_{t+1} \left(\frac{Y_{t+1}}{Y_t} \right) \left(1 + \frac{Q_{t+1}}{Y_{t+1}} \right) \right]$$

or

$$\frac{Q_t}{Y_t}(\hat{s}_t) = \sum_{\hat{s}_{t+1}} \beta e^{(1-\gamma)(g - \hat{s}_{t+1} + \alpha \hat{s}_t)} \left(\frac{1 - \bar{S} e^{\hat{s}_{t+1}}}{1 - \bar{S} e^{\hat{s}_t}} \right)^{-\gamma} \left(1 + \frac{Q_{t+1}}{Y_{t+1}}(\hat{s}_{t+1}) \right) \pi(\hat{s}_{t+1} | \hat{s}_t).$$

This is a system of linear equations and can be solved easily for Q_t/Y_t as a function of \hat{s}_t . The log price/dividend ratio is then obtained by simply taking logs: $\log(Q_t/Y_t) = pd(\hat{s}_t)$. Log returns on the consumption claim are

$$\begin{aligned} r_{t+1}^s &= \log \left[\frac{Y_{t+1} + Q_{t+1}}{Q_t} \right] \\ &= \log \left[\left(\frac{Y_{t+1}}{Y_t} \right) \left(\frac{1 + Q_{t+1}/Y_{t+1}}{Q_t/Y_t} \right) \right] \\ &= g - \hat{s}_{t+1} + \alpha \hat{s}_t + \log \left[1 + \frac{Q_{t+1}}{Y_{t+1}}(\hat{s}_{t+1}) \right] - \log \left[\frac{Q_t}{Y_t}(\hat{s}_t) \right] \end{aligned}$$

which gives the expression $r_{t+1}^s = r^s(\hat{s}_t, \hat{s}_{t+1})$.

B Computing Relative Risk Aversion

The relative risk aversion coefficient is computed by a method that involves simulations of individual consumption streams. The simulation method is described in the first subsection, and the next subsection explains how to compute the relative risk aversion coefficient using it.

B.1 Simulating Individual Consumption Streams

As mentioned in the main text, an individual household's consumption stream $\{C_{1t}(i), C_{2t}(i)\}_{t=0}^T$ can be easily backed out from the aggregate consumption stream $\{C_{1t}, C_{2t}\}_{t=0}^T$, once a realization of $\{I_t(i)\}_{t=0}^{T-1}$ is given. The simulation algorithm is follows:

1. Specify the initial values Y_0 , $C_{2,0}$, and $C_{2,0}(i)$.
2. Set $S_0 = C_{2,0}/Y_0$ and $\hat{s}_0 = \log(S_0) - \log(\bar{S})$.
3. Starting from \hat{s}_0 , simulate a Markov chain $\{\hat{s}_t\}_{t=0}^T$ and construct $\{S_t\}_{t=0}^T$ by $S_t = \bar{S}e^{\hat{s}_t}$.
4. Also draw an i.i.d. sequence $\{I_t(i)\}_{t=0}^{T-1}$ that takes $I_t(i) = 1$ with probability θ and $I_t(i) = 0$ with probability $1 - \theta$.
5. Construct $\{Y_t\}_{t=0}^T$ recursively by $Y_{t+1} = e^{g - \hat{s}_{t+1} - \alpha \hat{s}_t} Y_t$. Take $C_{1t} = (1 - S_t)Y_t$, $C_{2t} = S_t Y_t$, and $C_{2,t+1}^* = (C_{2,t+1} - (1 - \theta)C_{2t})/\theta$ for each t .
6. For each t , set $C_{1t}(i) = C_{1t}$ and take

$$C_{2,t+1}(i) = \begin{cases} C_{2,t+1}^* & \text{if } I_t(i) = 1 \\ C_{2t}(i) & \text{if } I_t(i) = 0 \end{cases}$$

B.2 Calculating the Relative Risk Aversion Coefficient

In the discussion above, the value function V was defined over financial wealth $A_t(i)$, rather than present value wealth $W_t(i)$. To compute relative risk aversion,

it is convenient to redefine the value function V over W (with an abuse of notation) using the definition

$$W_t(i) \equiv A_t(i) + E_t \sum_{j=0}^{\infty} M_{t \rightarrow t+j} \quad (39)$$

as

$$\begin{aligned} & V(W_t(i), C_{2t}(i), I_t(i); Y_t, C_{2t}) \\ &= V \left(A_t(i) - E_t \sum_{j=0}^{\infty} M_{t \rightarrow t+j} Y_{t+j}, C_{2t}(i), I_t(i); Y_t, C_{2t} \right) \end{aligned} \quad (40)$$

where the V on the right hand side is the value function as defined previously. Clearly, the two functions always return the same value. The newly defined V has the interpretation of

$$\begin{aligned} & V(W_0(i), C_{2,0}(i), I_0(i); Y_0, C_{2,0}) \\ &= \max_{\{C_{1t}(i), C_{2,t+1}(i)\}_{t=0}^{\infty}} E_0 \left[\sum_{t=0}^{\infty} \beta^j u(C_{1t}(i), C_{2t}(i)) \right] \end{aligned} \quad (41)$$

subject to the present-value budget constraint

$$E_0 \sum_{t=0}^{\infty} M_{0 \rightarrow t} [C_{1t}(i) + C_{2t}(i)] \leq W_0(i) \quad (42)$$

and $C_{2,t+1}(i) = C_{2t}(i)$ if $I_t(i) = 0$, since this optimization problem is equivalent to the one that households solve in the model.

Using this newly defined V , a household's relative risk aversion coefficient is defined as a function of its individual state and the aggregate state of the economy as

$$\text{RRA}(W_t(i), C_{2t}(i), I_t(i); Y_t, C_{2t}) \equiv -\frac{W_t(i)V_{WW}}{V_W} = -\frac{\log V_W}{\log W_t(i)}. \quad (43)$$

To compute this value, first use the envelope condition $V_W = U_1$ to get

$$\begin{aligned} \text{RRA}(W_0(i), C_{2,0}(i), I_0(i); X_0) &= -\frac{\partial \log V_W}{\partial \log W_0(i)} \\ &= -\frac{\partial \log U_1}{\partial \log C_{1,0}(i)} \times \frac{\partial \log C_{1,0}(i)}{\partial \log W_0(i)} \\ &= \gamma \times \frac{\partial \log C_{1,0}(i)}{\partial \log W_0(i)}. \end{aligned}$$

where the final line uses $-C_1 u_{11}/u_1 = \gamma$. In order to find the elasticity $\partial \log C_{1,0}(i)/\partial \log W_0(i)$, which is the percentage increase in $C_{1,0}(i)$ that the household would choose in response to a windfall that increases its present value wealth $W_0(i)$ by one percent (with no change at the aggregate level), we adopt the following strategy. First, suppose a household receives a windfall that increases its present value wealth by k percent, and that this induces the household

to increase its $C_{1,0}(i)$ by one percent. Once we find the number k , the elasticity can be estimated as its reciprocal, so that $\partial \log C_{1,0}(i)/\partial \log W_0(i) \approx 1/k$. To find k , let $\{C_{1t}(i), C_{2t}(i)\}_{t=0}^{\infty}$ and $\{\tilde{C}_{1t}(i), \tilde{C}_{2t}(i)\}_{t=0}^{\infty}$ be the household's consumption before and after the windfall, respectively. Since (42) holds with equality, the value k is

$$k = \log \left[\frac{\mathbb{E}_0 \sum_{t=0}^{\infty} M_{0 \rightarrow t} [\tilde{C}_{1t}(i) + \tilde{C}_{2t}(i)]}{\mathbb{E}_0 \sum_{t=0}^{\infty} M_{0 \rightarrow t} [C_{1t}(i) + C_{2t}(i)]} \right] \times 100.$$

But from the previous subsection, we know how to simulate $\{C_{1t}(i), C_{2t}(i)\}_{t=0}^T$ from a given initial state $(W_0(i), C_{2,0}(i), I_0(i); Y_0, C_{2,0})$ under particular realizations of $\{\hat{s}_t\}_{t=0}^T$ and $\{I_t(i)\}_{t=0}^{T-1}$. Moreover, a draw from $\{\tilde{C}_{1t}(i), \tilde{C}_{2t}(i)\}_{t=0}^T$ (under the same initial state $(W_0(i), C_{2,0}(i), I_0(i); X_0)$ and realizations of $\{\hat{s}_t\}_{t=0}^T$ and $\{I_t(i)\}_{t=0}^{T-1}$) can be backed out from this, using the fact that for all $t \geq 0$

$$\tilde{C}_{1t}(i) = e^{0.01} \times C_{1t}(i) \quad (44)$$

and

$$\tilde{C}_{2,0}(i) = C_{2,0}(i) \quad (45)$$

$$\tilde{C}_{2,t+1}(i) = \begin{cases} e^{0.01} \times C_{2,t+1}^* & \text{if } I_t(i) = 1 \\ \tilde{C}_{2t}(i) & \text{if } I_t(i) = 0 \end{cases} \quad (46)$$

Here, (44) follows from $\tilde{C}_{1,0}(i) = e^{0.01} \times C_{1,0}(i)$ (by construction) and the pricing equation (8) together with the fact that asset prices do not change in response to this windfall; (45) holds because $\tilde{C}_{2,0}(i)$ is predetermined as of date 0 and household i is an ‘‘average’’ one; and (46) follows from the adjustment rule of precommitted consumption and (44) plugged into (10). Thus, we can take a large number T and a large number of simulated realizations of $\{C_{1t}, C_{2t}\}_{t=0}^T$, $\{C_{1t}(i), C_{2t}(i)\}_{t=0}^T$, and $\{\tilde{C}_{1t}(i), \tilde{C}_{2t}(i)\}_{t=0}^T$ and use them to numerically estimate the value k by

$$\hat{k} = \log \left[\frac{\hat{\mathbb{E}}_0 \sum_{t=0}^T M_{0 \rightarrow t} [\tilde{C}_{1t}(i) + \tilde{C}_{2t}(i)]}{\hat{\mathbb{E}}_0 \sum_{t=0}^T M_{0 \rightarrow t} [C_{1t}(i) + C_{2t}(i)]} \right] \times 100$$

where $\hat{\mathbb{E}}_0$ represents Monte Carlo integration. Relative risk aversion is then estimated by $\text{RRA} \approx \gamma/\hat{k}$.

The numbers used in Figure 4 were constructed using $T = 200$ and 100,000 Monte Carlo replications, starting from $W_0(i) = \mathbb{E}_0 \sum_{t=0}^{\infty} M_{0 \rightarrow t} Y_t$, $I_0(i) = 0$, $Y_0 = 1.0$,¹⁰ and $C_{2,0}(i) = C_{2,0} = S_t Y_0$ for each S_t .

¹⁰The scale of Y_0 is irrelevant: starting from $W_0(i) = \mathbb{E}_0 \sum_{t=0}^{\infty} M_{0 \rightarrow t} Y_t$, $C_{2,0}(i) = C_{2,0} = \bar{S} Y_0$, $I_0(i) = 0$, $Y_0 = Y$ gives the same result for any $Y > 0$.