Precautionary Saving Unfettered

James Feigenbaum*
Department of Economics
University of Pittsburgh
jfeigen@pitt.edu
http://www.pitt.edu/~jfeigen/

February 12, 2009

Abstract

Precautionary saving due to uninsurable income risk has engendered much interest because it can explain why consumption roughly tracks income over the lifecycle. However, recent findings suggest precautionary saving has negligible macroeconomic effects. In this paper, we show that a steady-state, general-equilibrium lifecycle model can be parameterized to account for the hump-shaped lifecycle profile of mean consumption. As a by-product, we can then determine the contribution of precautionary saving to the capital stock for these parameters. This turns out to be sensitive to the type of borrowing frictions present. If we impose an exogenous borrowing constraint, steady-state macroeconomic variables are insensitive to simultaneous changes in the discount factor and risk aversion that preserve the equilibrium interest rate. Effectively, the borrowing constraint overwhelms any effects of uncertainty. As aggregate precautionary saving is commonly defined, its dependence in this model on risk aversion has more to do with the effect of intertemporal substitution. Meanwhile, an endogenous "borrowing constraint" imposed via assumptions about the income process fails to capture the properties of lifecycle consumption. Finally, if no borrowing constraint is imposed, precautionary saving will, indeed, have significant effects for risk aversion close to 4, and this model can better account for the lifecycle consumption profile than the corresponding model with borrowing constraints. Thus, while it is generally accepted that borrowing constraints and uninsurable risk are complementary frictions, in fact they are not.

JEL Classification: E21

Keywords: consumption hump, aggregate precautionary saving, borrowing constraints, risk aversion

*I would like to thank Frank Caliendo, Christopher Carroll, Juan Carlos Cordoba, Dave DeJong, Geng Li, Josep Pijoan-Mas, and Annie Fang Yang for useful discussions and comments. I would also like to thank participants at seminars at the 2006 Midwest Macro Meeting and Computation in Economics and Finance Conference, Carnegie-Mellon University, and the Federal Reserve Bank of Atlanta for their input. In addition, Jonathan Parker provided his empirical and theoretical consumption and income profiles.

One of the biggest issues in consumption and saving today is an apparent dichotomy between the behavior of consumers at the microeconomic level and the impact of their decisions at the macroeconomic level. According to household data, individual consumers face large, idiosyncratic variation in their income. If we assume they cannot use markets to insure against the possibility of future fluctuations, then consumers can only protect themselves by saving more. Aggregating this behavior over the whole economy, Carroll and Samwick (1998) and Gourinchas and Parker (2001) found that this precautionary saving could account for 45-65% of aggregate wealth, in which case idiosyncratic risk should have an enormous impact on the macroeconomy that is ignored by most models. However, these estimates were obtained in partial equilibrium with an exogenous interest rate. In general equilibrium, increasing the demand for saving will push down interest rates, reducing the demand for nonprecautionary saving until the market reequilibrates. Taking this effect into account, Huggett (1993) and Aiyagari (1994) found that precautionary saving could only account for about 1-2% of aggregage wealth, suggesting it is quite reasonable to ignore the effect of uninsurable risk on the macroeconomy. However, they considered infinite-horizon models, which have no lifecycle implications, even though it was lifecycle behavior that first motivated interest in uninsurable risk.

In its most basic formulation, the Lifecycle/Permanent-Income Hypothesis (LCPIH) predicts that the path of consumption over the lifecycle should be independent of the path of income over the lifecycle. However, the preponderance of empirical evidence suggests that consumption actually does depend on current income because both consumption and income follow hump-shaped paths over the lifecycle (Carroll and Summers (1991), Thurow (1969)). Rather than throw out the LCPIH, Nagatani (1972) argued that allowing for precautionary saving might account for the hump-shaped lifecycle consumption profile. As Leland (1968) and Sandmo (1970) demonstrated, under common assumptions about the structure of preferences, consumers will respond to uncertainty about future income by saving more in the present. This decreases current consumption and, on average, tends to increase future consumption. Whereas in a frictionless model the rate of consumption growth is strictly proportional to the difference between the interest rate and the consumer's discount rate (Yaari (1964)), precautionary saving augments the consumption growth rate by a factor that increases with the variance of future income (Feigenbaum (2008b), Skinner (1988)). Thus if the variance of future income starts out large enough and decreases over the lifecycle and if the consumption growth rate in the absence of risk is negative, the consumption growth rate will decrease over the lifecycle, starting positive and ending negative, yielding a hump-shaped consumption profile.

Gourinchas and Parker (2002) have demonstrated that there is enough idiosyncratic uncertainty about income for the precautionary saving mechanism to account for the consumption hump in partial equilibrium. The question raised by the present paper is whether this result continues to hold in general equilibrium. Feigenbaum (2008a) has shown that the consumption hump can be trivially explained in a partial-equilibrium model with time-varying mortality risk where both the intrinsic preferential discount rate and the interest rate are free parameters. Indeed, it is easily seen that a consumption hump similar to what is found in the data can be generated in partial equilibrium if we weaken any of the three key assumptions of the standard model: complete markets, additive separability of preferences over consumption, and geometric discounting. Here markets are incomplete. Following the examples of Bullard and Feigenbaum (2007), who did a similar exercise with nonseparable preferences, and Hansen and Imrohoroglu (2008), who did the same with nongeometric discounting, in this paper we see what restrictions must be imposed on the parameters of a model with uninsurable income risk to match the behavior of the lifecycle profile of mean consumption in a steady-state general equilibrium. As a byproduct, we are then able to see what the model under these parameters has to say about how much precautionary saving contributes to the aggregate capital stock.

Here we consider a finite-horizon, general-equilibrium model, which allows us to depart from the work of Aiyagari (1994) and Huggett (1993) in two imporant respects. First, we can address the importance of precautionary saving in a model also calibrated to match the lifecycle properties of consumption. Second, we can also relax the requirement that the model include an active borrowing limit. This is necessary in the infinite horizon because the equilibrium interest rate must be less than the discount rate, which implies in the long run that all agents will run their assets down to zero if allowed to do so (Huggett and Ospina (2001)). But in a finite-horizon model, lifecycle considerations will endogenously prevent this counterfactual behavior since the consumer will need assets to fund retirement consumption.

We study the model under three different assumptions about borrowing constraints. The most common type of borrowing constraint, exemplified by Carroll (1997) and Gourinchas and Parker (2002), exploits the endogenous borrowing limit that arises when there is a positive probability of receiving near-zero income in any remaining periods (Aiyagari (1994)). Alternatively, we consider an exogenously imposed no-borrowing constraint as in Deaton (1991). Last, we also consider the case that Leland and Sandmo studied, where borrowing is subject to no exogenous constraints and the lowest possible income realization is large enough that the endogenous limit plays no role.

It turns out the relative importance of precautionary saving is critically sensitive to the assumptions made about the borrowing frictions in the model. Of the three regimes we consider, we find the one that best replicates the profile of mean consumption over the lifecycle is the one that allows borrowing. However, this is only true in general equilibrium. In the neighborhood of Gourinchas and Parker's (2002) baseline values in the larger parameter space of the partial-equilibrium model, we locally recover their result that borrowing frictions are needed to best account for the consumption hump. For the parameters that best account for the consumption hump in general equilibrium, precautionary saving does have significant effects at the macro level, consistent with the findings of Carroll and Samwick (1998), and Gourinchas and Parker (2002).

For both versions of the model without borrowing, the model predicts a mean consumption profile that exhibits a hump similiar to what is found in empirical data for low levels of risk aversion. However, the location of the peak is too early, and the amplitude is too large. With the endogenous borrowing limit, the predicted consumption profile deviates substantially from the empirical profile for high values of risk aversion because the small probability of never earning income again engenders far too much precautionary saving. In contrast, the results for the exogenous borrowing constraint are remarkably insensitive to the risk aversion in general equilibrium. In a perfect-foresight model with or without exogenous borrowing constraints, the risk-aversion parameter and the discount rate cannot be separately identified in a steady-state general equilibrium. This property approximately carries over to the model with uninsurable risk and exogenous borrowing constraints. Holding the technology and equilibrium interest rate fixed, the borrowing constraint forces the allocation of consumption over the lifecycle to follow a path that is virtually the same for any combination of preference parameters consistent with this interest rate. Since these combinations of preference parameters do not confer the same equilibrium interest rate in a frictionless model, aggregate precautionary saving measured as the difference between the capital stocks in these two models is also not identified. Consequently, the debate about how much of aggregate saving is precautionary in a model with an exogenous no-borrowing constraints is artificial.

But this criticism does not apply to a model without an active borrowing limit. Unencumbered by borrowing constraints, the Leland-Sandmo mechanism is sensitive to risk aversion. Holding the equilibrium interest rate fixed, the degree to which the average consumption profile deviates from the monotonic path of the frictionless LCPIH increases for more risk averse households. With this additional degree of freedom, a better fit to the data is possible, and the shape of the profile can be used to pin down a value of 3.75 for the risk aversion parameter. In fact, the model can nearly replicate the lifecycle consumption profile measured by Gourinchas and Parker (2002). Furthermore, with this calibration, precautionary saving accounts for two thirds of the capital stock, consistent with Gourinchas and Parker's (2001) estimate of the magnitude of aggregate precautionary saving, but in general equilibrium. Thus we do see a large deviation between the macroeconomic predictions of the precautionary saving model and a model with the same parameters but without frictions.

Note that precautionary saving is by no means the only mechanism that can account for the consumption hump. Bullard and Feigenbaum (2007) have shown in a calibrated model that Heckman (1974) and Becker and Ghez's (1975) suggestion that substitution between leisure and consumption could account for the hump is also workable. Likewise, mortality risk (Feigenbaum (2008a), Hansen and Imrohoroglu (2008)) has been considered as an explanation for the hump. Within the context of incomplete-markets explanations, Fernandez-Villaverde and Krueger (2005) have also shown how the interaction between durable and nondurable consumption may be important in explaining the hump in a model where durable goods can serve as collateral for loans, and uninsurable risk plays only a secondary role. Moving beyond simple modifications of the standard model, variation in household size (Attanasio et al (1999)), substitution of home production for market consumption (Aguiar and Hurst (2003, 2007), and time-inconsistent preferences (Caliendo and Aadland (2007), Laibson (1997)) have also been proposed as explanations for the hump.

All of these mechanisms account for the hump by interacting the consumption-saving decision with other aspects of the economy, which introduces testable predictions outside the consumption-saving milieu. This paper is part of a continuing program that aims to see which mechanism, or combination thereof, can best account for the consumption hump while also meeting these external predictions. In the case of precautionary saving, the size of the hump is regulated by the amount of income uncertainty that households face. We find that a general-equilibrium precautionary saving model is able to jointly accommodate both Gourinchas and Parker's (2002) estimates of income uncertainty and the lifecycle consumption profile. Moreover, the model without borrowing frictions is competitive with all of these alternative explanations in its ability to match their consumption data. The models with borrowing frictions are not.

The paper is organized as follows. The model is laid out in Section 1, and the calibration strategy is explained in Section 2. Section 3 describes the results for the three frictional models and shows that the model without an active borrowing limit matches the data best. In Section 4, the performance of the model with a hump generated by Leland-Sandmo precautionary saving alone is compared to other mechanisms in the literature. The robustness of the model's results is discussed in Section 5. We conclude in Section 6.

1 The Model

The model is in the style of Huggett (1996). There is a continuum of agents with rational expectations who live for T periods in an overlapping-generations economy where each generation has unit measure. The economy is in a stationary equilibrium with no aggregate uncertainty, so aggregate quantities are time-independent. Quantities pertaining to a specific individual can vary with age. For such quantities, age is indexed by a subscript that runs from 0 to $T-1$.

Agents maximize

$$E_0\left[\sum_{t=0}^{T-1}\beta^t u(\widetilde{c}_t)\right],$$

where c_t is consumption at age t and preferences have a constant relative risk aversion (CRRA) specification:[1]

$$u(c) = \frac{c^{1-\gamma} - 1}{1-\gamma} \quad \gamma > 0. \tag{1}$$

Agents have one unit of labor supplied inelastically. During the working life, an agent at age t will have a stochastic productivity endowment e_t, drawn independently of all other agents both outside and within his cohort. The endowment e_t can be multiplicatively decomposed as

$$e_t = a_t p_t z_t, \tag{2}$$

where a_t is a deterministic, age-dependent component, p_t is a permanent endowment shock such that

$$p_t = q_t p_{t-1}, \tag{3}$$

and z_t is a temporary endowment shock. The innovation q_t to the permanent shock is a unit-mean i.i.d. process, and p_{-1} is fixed at 1. Likewise, the temporary shocks z_t are also a unit-mean i.i.d. process.

Total wage income at age t will be $y_t = we_t$, where w is the real wage per unit of labor productivity. An agent at age t can invest in one-period assets (either bonds or capital which are perfect substitutes[2]) that pay the risk-free gross return R (and net return $r = R - 1$).[3] Let b_{t+1} denote the assets purchased by an agent at age t, which will then pay Rb_{t+1} at $t+1$. Thus the budget constraint is

$$c_t + b_{t+1} = w_t e_t + Rb_t. \tag{4}$$

We also impose an exogenous borrowing limit

$$b_{t+1} \geq -B,$$

where B is set either to 0 or ∞.

We can then write the optimization problem of an agent at age t as a recursive set of Bellman equations. Following Deaton (1991), we define cash on hand as the sum of current income and the value of financial wealth:

$$x_t = we_t + Rb_t.$$

The terminal value function is

$$v_{T-1}(x_{T-1}, p_{T-1}) = u(x_{T-1}) \tag{5}$$

while the value function for $0 \leq t < T-1$ is

$$v_t(x_t, p_t) = \max_{c_t, b_{t+1}} u(c_t) + \beta E\left[v_{t+1}(wa_{t+1}\widetilde{q}_{t+1}p_t\widetilde{z}_{t+1} + Rb_{t+1}, \widetilde{q}_{t+1}p_t)\right] \tag{6}$$

[1] Inside expectation and other moment operators, tildes denote variables that are stochastic with respect to the relevant information set.

[2] The only distinction between bonds and capital is that an agent can hold negative amounts of bonds.

[3] Thus, markets are incomplete since there is only the one, risk-free asset. Claims contingent on realizations of productivity endowments do not exist in the market.

subject to

$$\begin{aligned} c_t + b_{t+1} &= x_t \\ c_t &\geq 0 \\ b_{t+1} &\geq -B, \end{aligned}$$

where $b_0 = 0$.

The aggregate demand for assets is the sum over all individual asset demands. Since all agents are a priori identical, the measure of agents at age t who have realized the endowment history $e^t = (e_0, \ldots, e_t)$ is equal to the unconditional probability of achieving that history. Thus the sum of asset demands equals the unconditional expectation of the sum of asset demands over a lifespan. Bonds are in zero net supply while the capital stock is K, so in equilibrium we must have

$$K = E\left[\sum_{t=1}^{T-1} b_t(\tilde{e}^t)\right]. \tag{7}$$

The gross production function is Cobb-Douglas:

$$F(K, L) = K^\alpha L^{1-\alpha} \tag{8}$$

for $\alpha \in (0, 1)$, where the effective labor supply L is the sum of the labor endowments for all agents in the economy. This can also be written as an unconditional expectation

$$L = E\left[\sum_{t=0}^{T-1} \tilde{e}_t\right]. \tag{9}$$

Capital depreciates at the rate δ. We assume firms behave competitively, so the return on capital and the wage must satisfy the profit-maximization conditions

$$F_K(K, L) + 1 - \delta = R \tag{10}$$

and

$$F_L(K, L) = w. \tag{11}$$

Thus, an equilibrium is a set of consumption demands $\{c_t(x_t, p_t)\}_{t=0}^{T-1}$, asset demands $\{b_{t+1}(x_t, p_t)\}_{t=0}^{T-2}$, a capital stock K, a wage w, and a gross interest rate R that optimize the consumer's problem (5)-(6) and satisfy Eqs. (7)-(11).

For a given set of parameters α, β, γ, and δ, and for a given endowment process, we will consider four versions of this model that vary with respect to the information structure and the borrowing limit B. The simplest model, which provides a baseline for comparison with the standard LCPIH, is a frictionless model (FM). In this model, the borrowing constraint does not bind so $B = \infty$, and all information about the endowment process is revealed at $t = 0$ so there is no uncertainty. Consideration of the frictionless model is necessary because, in addition to any saving motivated by frictions, the consumer will have to save during his working life to finance consumption during retirement, and we wish to distinguish these two sources of saving.

In the other three models, information about the endowment shocks q_t and z_t is not revealed until t, and this idiosyncratic uncertainty about income is uninsurable. These frictional models differ in terms of how and whether borrowing is disallowed. In the Leland-Sandmo Model (LSM), $B = \infty$, so there is no exogenous borrowing limit. There will still be an endogenous borrowing limit determined by the minimum present value of future income, but the income process will be calibrated so this is large enough that the endogenous limit is irrelevant for most agents. In the Exogenous Borrowing-Constraint Model (XBM), $B = 0$, and borrowing is shut down institutionally. Finally, in the Endogenous Borrowing-Limit Model (NBM), $B = \infty$, but an additional temporary shock state with a small probability is added where the consumer receives negligible income.[4] Since the minimum possible present value of future income is essentially zero, the consumer will endogenously choose not to borrow in the NBM.

Following Aiyagari (1994), we define aggregate precautionary saving as saving that would not persist if we remove the frictions from a model. For a given frictional model, if K is the equilibrium capital stock as defined by (7), aggregate precautionary saving will be $K - K_F$, where K_F is the equilibrium capital stock for the corresponding frictionless model with the same values of the exogenous parameters α, β, γ, δ, and $\{a_t\}_{t=0}^{T-1}$. The aggregate precautionary saving rate is then defined as

$$s_P = \frac{K - K_F}{K}. \tag{12}$$

Thus s_p is the fraction of the aggregate capital stock in a frictional model that disappears if we remove borrowing and uninsurability frictions.

Xu (1995) used the FM, the LSM, and the XBM to decompose the effects of borrowing and risk frictions into the effect of precautionary saving alone and the additional effect of the borrowing constraint. Here, that is not our purpose because the question is which of the three frictional models best accounts for the data. Thus, we will separately calibrate the XBM, NBM, and LSM, and see how well each can simultaneously account for the lifecycle profile of mean consumption and macroeconomic data for the US economy. We then use the FM for each calibration to compute how much uninsurable risk–and the borrowing constraint if present–contribute to aggregate saving.

2 Calibration

[4] This extreme state is interpreted as a zero-income state, but income cannot be exactly zero in this state because the model would not be defined for agents who get this shock at $t = 0$.

To quantify the model, we must specify its parameters. We set a period to a year and let agents live for $T = 65$ years from age 25 to 89.[5] To obtain the age-productivity factor, I fit Gourinchas and Parker's estimate of the average after-tax income profile to the quartic polynomial

$$a_t = 1 + 0.018095t + 0.000817t^2 - 5.1 \times 10^{-5}t^3 + 5.36 \times 10^{-7}t^4, \qquad (13)$$

where t runs from 0 to $T_W - 1$. Setting the retirement age at 65 (so $T_W = 40$), we have $a_t = 0$ for $t \geq T_W$.[6] Assuming log-normal distributions for both the temporary and permanent income shocks, Gourinchas and Parker (2002) use PSID data to estimate

$$\ln q_t \sim N\left(-\frac{1}{2}\sigma_p^2, \sigma_p^2\right)$$

and

$$\ln z_t \sim N\left(-\frac{1}{2}\sigma_z^2, \sigma_z^2\right),$$

where $\sigma_p^2 = 0.0212$ and $\sigma_z^2 = 0.0440$.

For this exercise we are *only* interested in mean consumption for each age group and not the whole distribution of consumption or wealth since only the mean has an effect on macroeconomic quantities such as the equilibrium interest rate. Feigenbaum (2008b) has shown it is sufficient to use an income process that matches the variance to compute the mean consumption profile since higher order moments will have a negliglible effect. Thus we discretize the log-normal distribution with a two-state distribution where $q_t \in \{Q_1, Q_2\}$ with probability 1/2, where $Q_1 = 1 - \sigma_p$ and $Q_2 = 1 + \sigma_p$. Likewise, $z_t \in \{Z_1, Z_2\}$ with probability 1/2, where $Z_1 = 1 - \sigma_z$ and $Z_2 = 1 + \sigma_z$.[7] This choice

[5] Gourinchas and Parker (2002) do not explicitly study what happens after age 65, after which policy functions are linear, but they calibrated these policy functions as though agents live till age 88.

[6] In principle Gourinchas and Parker (2002) allow for agents to receive income in their retirement years, either from a pension or Social Security. However, in their baseline model the replacement rate for these postretirement payments is estimated to be only a tenth of a percent of permanent income at retirement. Thus, I disregard Social Security. While this simplification is not entirely innocuous, introducing Social Security can only affect the lifecycle consumption profile before retirement through its impact on the interest rate, which ought to rise because young agents have less need to save. However, in the baseline calibrations of the XBM and LSM, the majority of aggregate saving occurs for precautionary reasons and not to finance retirement, so this should dampen the effect of Social Security on the interest rate in these models. The absence of Social Security would be a major concern if the properties of the consumption profile were sensitive to the interest rate (as they are when mortality risk accounts for the hump (Feigenbaum (2008a))), but that is not the case here as we see in Section 5.

[7] With these distributions, the kurtosis of log income is 1. To our knowledge, no study has actually measured the kurtosis of the unpredictable component of income, although most papers, including Gourinchas and Parker (2002), assume log-income has a normal distribution with a kurtosis of 3. In Section 5, we consider what happens with a finer income distribution. Then we can set the kurtosis of both the temporary and permanent shocks to be 3, and we find that any change is negligible. In the following, we do report the fraction of agents with zero or negative assets, which is one quantile of the wealth distribution. However, this particular quantile is quite robust to how we model the shocks.

also means that, for the LSM, the lower bound on income is high enough that the endogenous borrowing limit has no effect. For the NBM, we modify the temporary shock distribution so that with probability κ the agent will receive $z_t = \exp(-1000)$. The probabilities of receiving $1 \pm \sigma_z$ are then adjusted accordingly to $(1-\kappa)/2$. Following Gourinchas and Parker (2002), I calibrate $\kappa = 0.00302$.

This leaves the model with four remaining scalar parameters: the share of capital α, the discount factor β, the risk aversion coefficient γ, and the depreciation rate δ. In the steady state, the Cobb-Douglas production function implies

$$\frac{C}{Y} = 1 - \delta \frac{K}{Y} \tag{14}$$

and

$$r + \delta = \alpha \frac{Y}{K}, \tag{15}$$

so a choice of K/Y, C/Y, and α determine δ and r. We calibrate the share of capital to $\alpha = 0.3375$, the consumption rate to $C/Y = 0.75$, and the capital to output ratio to $K/Y = 2.5$.[8] This fixes $\delta = 0.1$ and $r = 3.5\%$.

This leaves β and γ. In general equilibrium, K/Y will be jointly determined by these parameters, given α and δ. Thus there is a continuous locus of points in (β, γ)-space that confer $K/Y = 2.5$. In the ensuing section we will see how the predictions of each of the three frictional models vary in equilibrium as we move along this curve. Note that if we wish to account for the lifecycle consumption profile then we have many more targets than parameters since we have data on mean consumption for forty age groups. We will designate a separate baseline for each of the three frictional models, choosing for the baseline calibration of (β, γ) the point on the curve that minimizes the mean squared deviation between the lifecycle consumption profile of the model and the lifecycle profile of mean consumption for each age group, adjusted for household size, as reported by Gourinchas and Parker (2002). To standardize our terminology, we will parameterize the curve in terms of γ, and then β will be determined by the function $\beta(\gamma)$ such that $K/Y = 2.5$.

The computational procedure used to solve these models is described in the Appendix.

3 Baseline Results

First we will consider the behavior of each frictional model as a function of the risk aversion parameter. Then we will compare the three models, after

[8] The C/Y target was taken from Rios-Rull (1996). While he targets K/Y to about 3, since the present model is not able to capture the behavior of the wealthiest consumers (Carroll (2000)), who own a significant fraction of the wealth, we reduce the K/Y target to remove them from consideration. This also has the effect of reducing the interest rate. In Section 5, we consider what happens if we calibrate the LSM with $K/Y = 3$ and see that the consumption hump is robust.

which we will discuss the calibration of each model that best fits to Gourinchas and Parker's (2002) lifecycle consumption data. Finally, we will demonstrate how we would miss these results if we considered the three models in partial equilibrium, focusing on the neighborhood of the parameter space studied by Gourinchas and Parker (2002).

3.1 Exogenous Borrowing-Constraint Model (XBM)

First we consider the Exogenous Borrowing-Constraint Model (XBM) in which income shocks are not revealed until they arrive and no borrowing is allowed by financial institutions. Such a borrowing constraint can account for a consumption hump by interfering with consumption smoothing. If $\beta R < 1$, the optimal consumption profile, absent any frictions, will be monotonically decreasing and start out with a high value of consumption, greater than the consumer's income. With a hump-shaped income profile, this consumption allocation will not be feasible if the consumer cannot borrow. Consumption will be constrained to equal income, which is increasing, early in life and then will follow the decreasing, optimal path later in life when the optimal path falls below the income profile.

Since uninsurable risk appears in this model, the Leland-Sandmo precautionary saving mechanism could conceivably play an important role in shaping the consumption profile too, but, in fact, precautionary saving only has a second-order effect on the profile.[9] This is shown in Fig. 1. Along with Gourinchas and Parker's (2002) empirical consumption data, Fig. 1 plots lifecycle consumption profiles in the XBM for choices of the risk aversion varying from Gourinchas and Parker's (2002) estimate of $\gamma = 0.5$ to a value of $\gamma = 5$. The latter is near the upper limit of what most researchers would consider plausible. If precautionary saving is important, its effects ought to increase with the risk aversion γ, but Fig. 1 shows hardly any change as γ is varied. Between ages 25 and 45, the spread between the consumption profiles is roughly the same as the magnitude of jumps in the empirical consumption profile from one period to the next, which could be viewed as a measure of the noise in the data.

As a practical matter, the discount factor and risk aversion cannot be separately identified using only the first-order moment data of the income and consumption distributions that we consider here.[10] This approximate nonidentification result is another indication of the inconsequentiality of precautionary saving in the presence of a borrowing constraint, for the nonidentification result would hold exactly in the absence of uninsurable risk. In the frictionless model, the rate of consumption growth is simply

$$\frac{c_{t+1}^{FM}}{c_t^{FM}} = (\beta R)^{1/\gamma}. \tag{16}$$

[9] Note that *aggregate* precautionary saving as defined by (12) does not differentiate between saving caused by the borrowing friction and saving caused by the Leland-Sandmo mechanism.

[10] As is discussed in 3.2, this approximate nonidentification result does not hold for the NBM, which was used by Gourinchas and Parker (2002) to estimate the model.

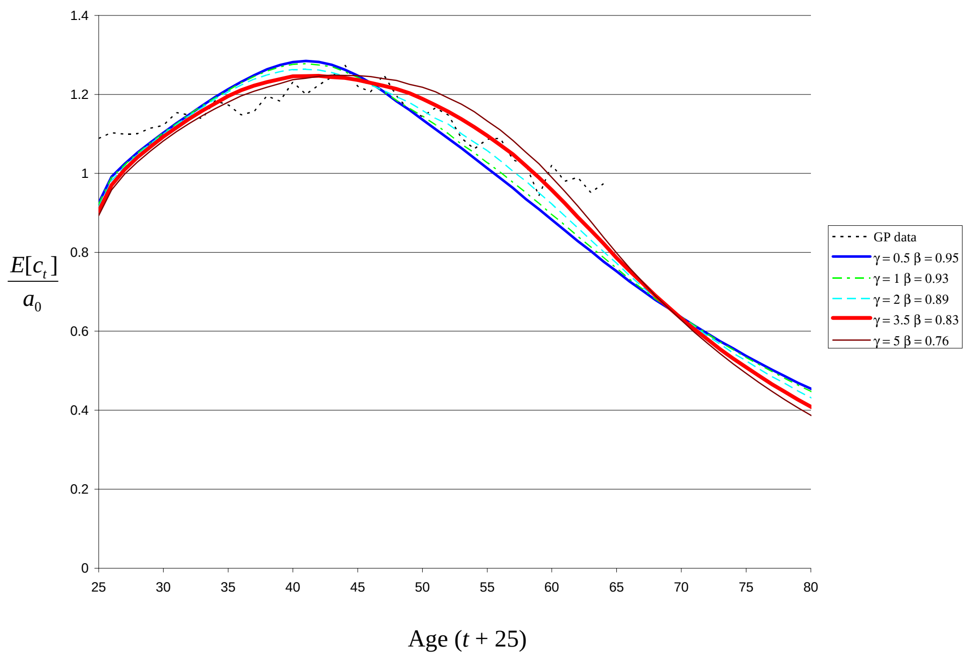

Figure 1: The lifecycle consumption profile for the exogenous borrowing-constraint model (XBM) for several combinations (β, γ) that give an equilibrium $K/Y = 2.5$ along with the Gourinchas and Parker consumption data.

For a given interest rate R, combinations of β and γ that fall on level curves of $(\beta R)^{1/\gamma}$ will all give rise to the same consumption allocation. Since R is also unchanged, if the model is in equilibrium for one choice of β and γ, it will be in an observationally indistinguishable equilibrium for any other combination of β and γ that falls on the same level curve.[11] Thus, if we parameterize the space in terms of $(\alpha, \gamma, \delta, R)$, the predictions of the frictionless model will be completely independent of γ. Adding a borrowing constraint to the model does not modify this result since the growth rate where the borrowing constraint does not bind is still $(\beta R)^{1/\gamma}$, and consumption behavior where the borrowing constraint does bind is independent of preferences.

A variable that has received significant attention in the literature is the aggregate precautionary saving rate, i.e. the fraction of the equilibrium capital stock in the XBM that cannot be accounted for in the corresponding frictionless model, holding all parameters unrelated to the borrowing and risk frictions the same. This ratio is not observable since it involves the behavior of different, hypothetical models. However, it does depend on β and γ since the level curves of $(\beta R)^{1/\gamma}$ for the FM and XBM do not coincide.

The aggregate precautionary saving rate, s_P, as defined by (12) is shown as a function of γ in Fig. 2. For a low value of γ on the order of Gourinchas and Parker's (2002) baseline calibration of 0.5, this fraction is only 10%. This is consistent with the low values obtained by Aiyagari (1994) and Huggett (1993), who also focused on calibrations for γ approximately equal to 1.[12] However, as we increase γ, the fraction increases. For γ on the order of 3, we find that roughly 2/3 of the capital stock is due to frictions, which is consistent with Gourinchas and Parker's (2001) estimate of this fraction. But if γ is effectively unidentified, s_P is also unidentified, in which case this disagreement is artificial.

While Fig. 1 shows that the XBM produces a hump-shaped profile similar to what is seen in the data, there are notable discrepancies. The upslope of the profile is much steeper than in the data. The model predicts that, on average, young consumers should consume less than their income while they actually consume slightly more than their income. It also predicts that the peak of the consumption hump should be around age 41 while the data peak around age 45. Finally, in later years, consumption falls off slightly more rapidly in the model than in the data.[13]

[11] Note that the nonidentifiability of β and γ in the FM is a consequence of the fact that we are only considering steady states without aggregate risk. If we allowed for aggregate shocks, such as to technology, that allow the interest rate R to vary over time, it would be straightforward to estimate β and R by regressing the log of consumption growth with respect to ln R according to (16). In actual practice, however, this exercise has not yielded tight estimates of γ (Hall (1988) reports the coefficient on $\ln R$, corresponding to γ^{-1}, could actually be zero), so we do not consider such aggregate shocks. Indeed, there is considerable controversy in the literature whether it is even possible to estimate β and γ with such a regression if households face uninsurable risk and borrowing frictions (Attanasio and Low (2004), Carroll (2001), Ludvigson and Paxson (2001)).

[12] They obtained fractions closer to 1% in an infinite-horizon model. The larger fraction found here can partly be accounted for by our more persistent income process.

[13] In partial equilibrium, Gourinchas and Parker (2002) are able to exactly match the downslope of the consumption hump, but they have an extra degree of freedom since they can adjust

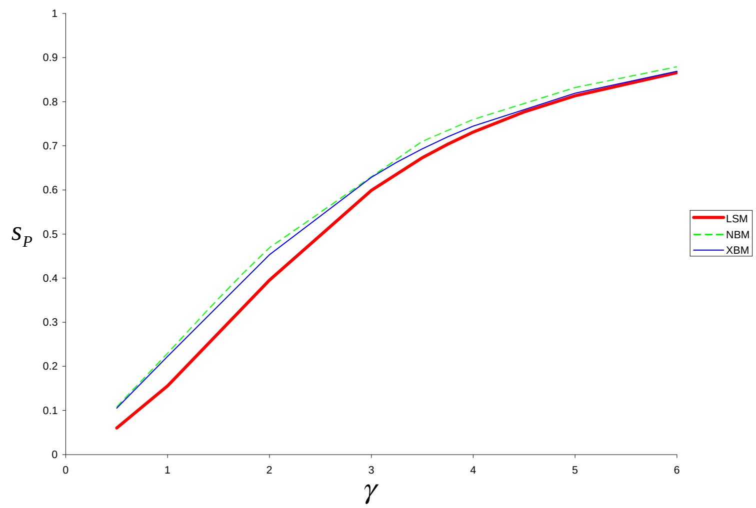

Figure 2: Aggregate precautionary saving rate s_P in each of the frictional models as a function of risk aversion γ.

As we will see, these discrepancies are a special feature of the XBM that can be mitigated or even eliminated if we make different assumptions about what frictions are present. Note that while there is certainly microeconomic evidence that some consumers face constraints on borrowing (Gross and Souleles (2002)), it is also certainly true that some consumers do borrow. The common modeling choice that consumers cannot borrow at all is made for its simplicity and not for its realism.

3.2 Endogenous Borrowing-Limit Model (NBM)

The tradeoff between simplicity and realism is even more pronounced in the Endogenous Borrowing-Limit Model (NBM). In this variation of the model, there is no institutional barrier to borrowing. Instead, in every period consumers face a small probability κ that they will receive no income. Since the possibility exists that a consumer will never receive any income again, and since the model assumes full commitment, consumers will endogenously choose not to borrow.

The advantage of this approach is purely computational. Since the value function in every period obeys an Inada condition in the limit as cash on hand $x \rightarrow 0$, the consumer will always choose to save a positive amount to stay away from $x = 0$ next period. Thus the borrowing limit in the NBM

the interest rate and the discount rate independently.

never actually "binds", unlike the exogenous borrowing constraint of the XBM. Consequently, the NBM consumption function has no kinks like those the XBM consumption function exhibits at wealth values where the borrowing constraint starts to bind either presently or with some probability in the future. Numerical approximation of the consumption functions in the NBM is, therefore, less problematic than in the XBM.

The disadvantage of the endogenous borrowing limit is that it is even less realistic, for what causes people not to borrow is the near inconceivable possibility that they will never get any income ever again. It also depends crucially on the assumption of full commitment, which does not arise in a world of bankruptcy laws.

In the limit as $\kappa \to 0$, the NBM converges to the XBM.[14] For $\kappa > 0$, the two models are virtually indistinguishable for small values of the risk aversion γ. However, as γ increases, they begin to diverge. This can be seen in Fig. 3, which shows lifecycle consumption profiles for different values of γ in the NBM. Gourinchas and Parker (2002) estimated $\gamma = 0.5$ for the NBM using the Method of Simulated Moments. For this low value of risk aversion, the lifecycle consumption profile is essentially the same as the corresponding profile in Fig. 1. However, as γ increases above 1, the consumption of young agents decreases drastically as they have to build up a larger buffer stock of saving to insure themselves against the possibility of hitting the borrowing limit. Thus the approximate nonidentification result for γ that we found in the XBM does not apply to the NBM, although the shape of the consumption hump near its peak is again largely independent of γ. Nevertheless, the fact that young consumers do not save almost all their initial income constrains estimates of γ for the NBM to low values as Gourinchas and Parker (2002) found.[15]

3.3 Leland-Sandmo Model (LSM)

In the Leland-Sandmo Model (LSM), there is no exogenous borrowing constraint, and the lowest income shock is high enough that the endogenous borrowing limit is ineffectual. Thus, precautionary saving induced by the uninsurable income risk is the only mechanism that can cause deviations from perfect consumption smoothing. Since the consumption function is smooth, Taylor's Theorem can be applied to the Euler equation to obtain an approximate analytic expression for the consumption function. To second-order in the moments of the income process, the expected rate of consumption growth is then (Feigenbaum (2008b))

$$E_t\left[\frac{\widetilde{c}_{t+1}}{c_t}\right] \approx (\beta R)^{1/\gamma}\left[1 + \frac{\gamma+1}{2}\frac{V_t[\widetilde{w}_{t+1}]}{E_t[\widetilde{w}_{t+1}]^2}\right], \tag{17}$$

[14] Note that the model behaves discontinuously at $\kappa = 0$. In that case, the NBM is identical to the LSM.

[15] Gourinchas and Parker (2002) do not report such small values of initial consumption. This is presumably because they give agents a random draw of initial assets to match the wealth distribution of age-25 agents so agents do not have to build up their buffer stock from zero. They can do this in partial equilibrium without accounting for where these assets originate, but we cannot do this in general equilibrium.

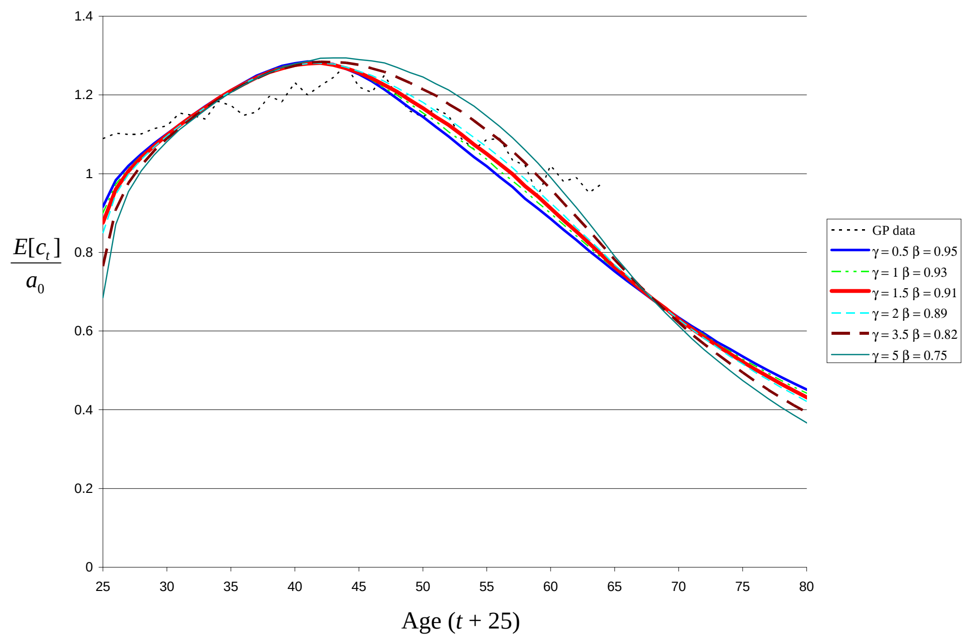

Figure 3: The lifecycle consumption profile for the endogenous borrowing-limit model (NBM) for several combinations (β, γ) that give an equilibrium $K/Y = 2.5$ along with the Gourinchas and Parker consumption data.

where

$$w_t = E_t \left[\sum_{s=t}^{T} \frac{\widetilde{y}_s}{R^{s-t}} \right] + Rb_t$$

is total wealth including the expected present value of the consumer's income stream. In the absence of consumption uncertainty, the growth rate of consumption would be $(\beta R)^{1/\gamma}$ as in the frictionless model. If consumption next period is uncertain, the consumer will save more now, decreasing his current consumption c_t and, on average, raising his consumption next period. Since to zeroth-order (i.e. in the absence of corrections for uncertainty) consumption is proportional to wealth, the growth rate is augmented by a factor proportional to the variance of the ratio of realized consumption to expected consumption (both next period). A consumer who is more risk averse will save more in response to consumption risk, so the variance correction term is also proportional to $\gamma + 1$.[16]

If $(\beta R)^{1/\gamma} < 1$, the lifecycle consumption profile will decrease late in life after all uncertainty is resolved. But if γ is big enough, the variance term in (17) can push the mean rate of consumption growth above 1 early in life. If, moreover, the variance term decreases with age, the rate of consumption growth will be monotonically decreasing, and the lifecycle consumption profile will be concave and, therefore, hump-shaped.

As is shown in Fig. 4, with our calibration of the model precautionary saving can, indeed, account for a hump-shaped consumption profile when $\gamma \gtrsim 2$. Again we consider here only combinations of β and γ that give an equilibrium $K/Y = 2.5$. In contrast to the XBM, however, β and γ are separately identified in the LSM because the hump gets more pronounced as we increase γ and, conversely, goes away if we decrease γ. Thus we can set γ to obtain the best fit between the predicted consumption profile and Gourinchas and Parker's (2002) empirical consumption profile.

3.4 Comparison Across Precautionary Saving Models

Figs. 1, 3, and 4 provide graphical support for the following characterization of what happens in calibrated general equilibria of the three models. In the XBM, the lifecycle consumption profile is extremely robust to changes in the risk aversion γ. In the NBM, the profiles are not robust: as the distinction between the endogenous borrowing limit and an exogenous no-borrowing constraint becomes more severe, the NBM profiles move further away from the data. Meanwhile, in the simplest of the three models, the LSM, the consumption profile is not only sensitive to γ but it can also be made to fit to the empirical profile better than the XBM consumption profile.

The next three graphs provide a more quantitative basis for this characterization by showing how objective properties of the consumption profiles vary with γ. For small γ, the XBM and NBM are, as the theory requires, nearly

[16] The model does not respond continuously to changes in γ at $\gamma = 0$, so the variance factor does not vanish in the limit as $\gamma \to 0$.

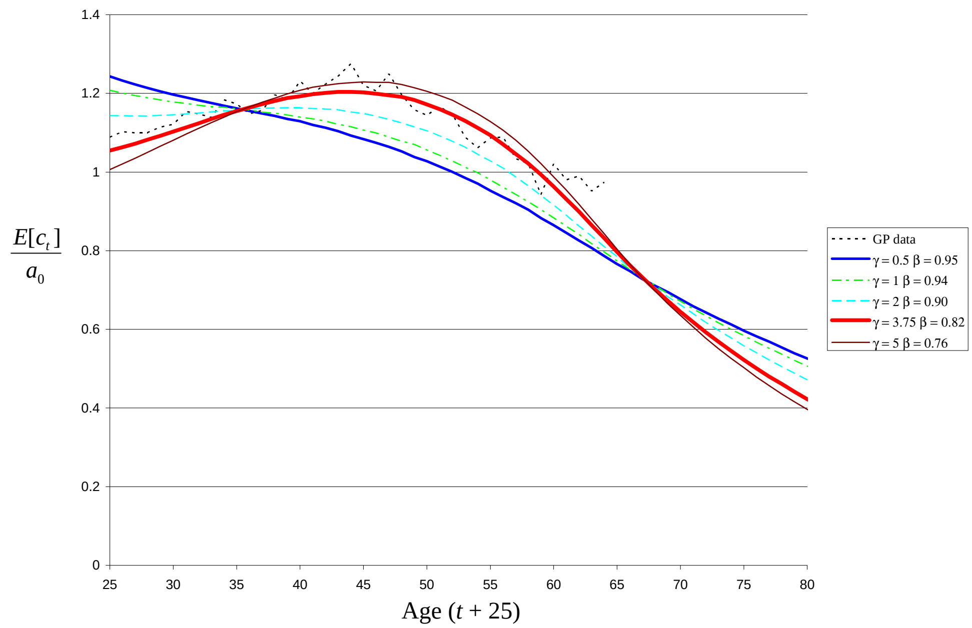

Figure 4: The lifecycle consumption profile for the Leland-Sandmo model (LSM) for several combinations (β, γ) that give an equilibrium $K/Y = 2.5$ along with the Gourinchas and Parker consumption data.

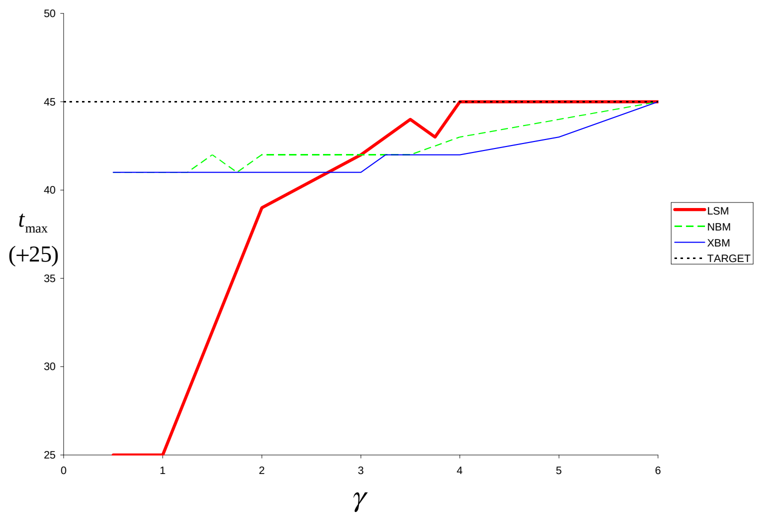

Figure 5: Peak age of the lifecycle consumption profile as a function of risk aversion γ for each frictional model.

identical, whereas the LSM is very different. For large γ, the LSM profile approaches the XBM profile as the endogenous borrowing limit in the LSM becomes important, and so both models are governed by two frictions. Meanwhile the NBM completely diverges from the data in this limit.

Fig. 5 shows how the age of the peak of the consumption hump depends on γ for each of the three models. In the XBM and NBM, the peak age always stays between ages 41 and 45 as γ varies between 0.5 and 6. In contrast, the peak age for the LSM covers the larger range from 25, for small γ where there is no hump, to 45, for large γ when the XBM and LSM start to converge. Note that most studies of the hump in consumption data find a peak age between ages 40 and 50, so all three variations of the model can match this target.

A common measure of the size of the consumption hump, propounded by Bullard and Feigenbaum (2007) and Hansen and Imrohoroglu (2008), is the ratio of mean consumption at the peak of the hump to mean consumption at the starting age of the model. The dependence of this ratio on γ for each model is shown in Fig. 6. The ratio exhibited by Gourinchas and Parker's (2002) empirical consumption profile is 1.15. For the XBM, this ratio stays confined to a band between 1.4 and 1.5. The XBM and NBM have similar ratios for small γ, but the ratio for the NBM increases drastically for γ larger than 1. Only the LSM is able to match the target of 1.15 as the ratio varies smoothly between 1, for small γ where there is no hump, and the XBM ratio, for very large γ.

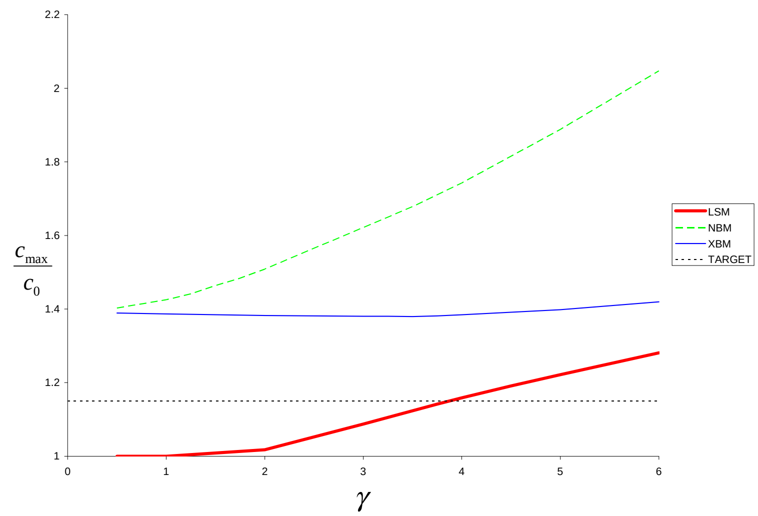

Figure 6: Ratio of peak to initial consumption for the lifecycle consumption profile of each frictional model as a function of risk aversion γ.

To measure how well the three models compare to Gourinchas and Parker's (2002) data along every dimension, we can also compare the root-mean-squared deviation between the predictions of each model and the data. This is shown in Fig. 7. As a sense of how much noise is in the data, the root-mean-squared deviation of Gourinchas and Parker's data from a polynomial fit is 0.024. Consistent with the above findings, the LSM achieves the lowest mean squared deviation while the NBM fares the worst. Interestingly, both the LSM and the XBM minimize the mean squared deviation at values of γ between 3 and 4. Thus, if we suppose that we can measurably distinguish between different choices of β and γ in the XBM, we would obtain roughly the same estimates as we would for the LSM.

One argument for why borrowing constraints should be included in precautionary saving models is because, otherwise, the model will predict that consumers should counterfactually borrow far too much. The source of this intuition can be understood by looking at what happens in the LSM for $\gamma \sim 1$. Fig. 8 shows the fraction of the whole population of agents who have zero or negative financial holdings (i.e. $b_t \leq 0$) for each model as a function of γ. At $\gamma = 1$, roughly a third of agents engage in borrowing whereas, according to the Survey of Consumer Finances, only about 15-20% of the population has zero or negative assets (Wolff (2001)).

However, for larger values of γ, as precautionary saving becomes more significant, consumers will endogenously choose to borrow less. The borrowing

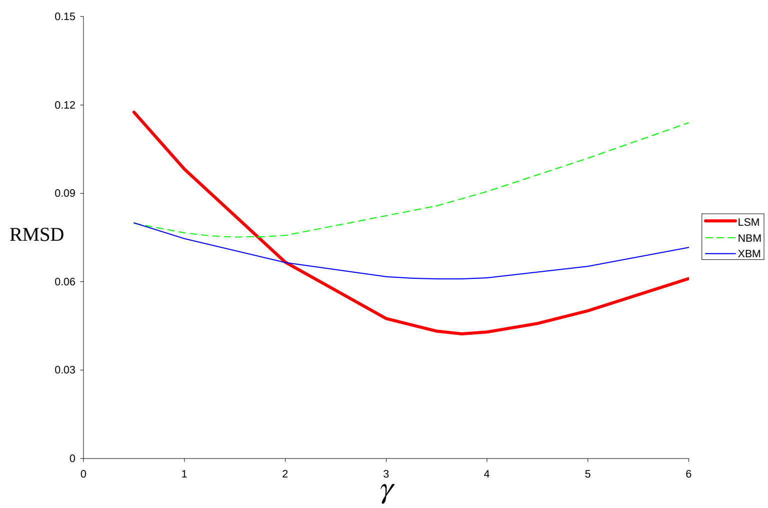

Figure 7: Root mean-squared deviation of best consumption profile for each frictional model from Gourinchas and Parker's lifecycle consumption data.

fraction falls inside the band of 15-20% for $\gamma \sim 3$, which is also roughly the point where the LSM minimizes the mean squared deviation for Gourinchas and Parker's (2002) data. Thus the criticism that consumers will borrow too much if we do not include borrowing constraints is not valid. On the contrary, Fig. 8 shows that consumers borrow too little when we do, either endogenously or exogenously, exclude borrowing.

3.5 Best Model Calibration

To obtain baseline calibrations, we set γ to the value that minimizes the root mean squared deviation between the predicted consumption profile and the empirical consumption profile for each of the three frictional models. The parameters and key observables for each baseline model are given in Table 1 while the consumption profile of each baseline model is shown in Fig. 9. Clearly the LSM does the best of the three models at matching the data.

What happens if we compare across the three frictional models for the same set of scalar parameters, α, β, γ, and δ? Holding these parameters fixed at the baseline LSM values, the resulting macro observables for the three frictional models are given in Table 2 and the mean consumption profiles are shown in Fig. 10. Comparing Tables 1 and 2, we find that for the XBM the equilibrium value of r decreases and K/Y increases as we go from the XBM baseline parameters to the LSM baseline parameters. However, the other variables and the lifecycle

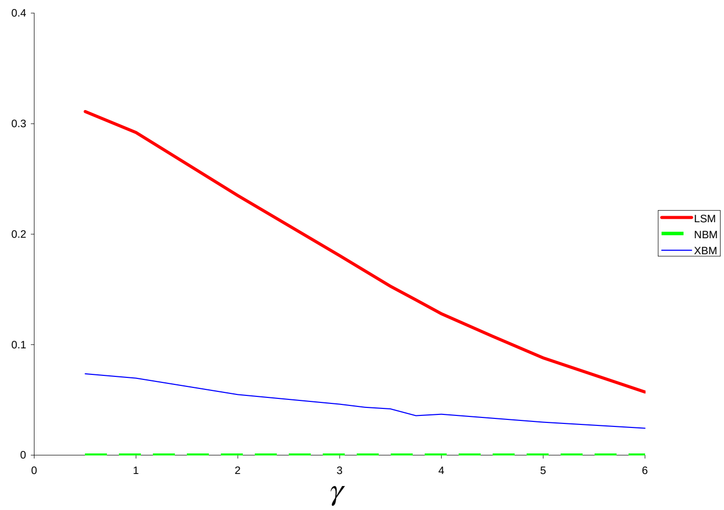

Figure 8: Fraction of agents who borrow in each of the frictional models as a function of risk aversion γ.

Variable	Target	XBM	NBM	LSM
γ	-	3.5	1.5	3.75
β	-	0.830	0.912	0.825
r	3.5%	3.50%	3.50%	3.51%
K/Y	2.5	2.501	2.500	2.498
C/Y	0.75	0.750	0.750	0.750
$t_{\max}(+25)$	45	42	42	43
$c_{\max}/c(0)$	1.15	1.379	1.464	1.141
Borr. Frac.	0.15	0.042	5×10^{-5}	0.140
s_P	-	0.693	0.354	0.704
RMSD	0	0.061	0.075	0.042

Table 1: Parameters and key observables for the baseline calibrations of each frictional model along with the target values.

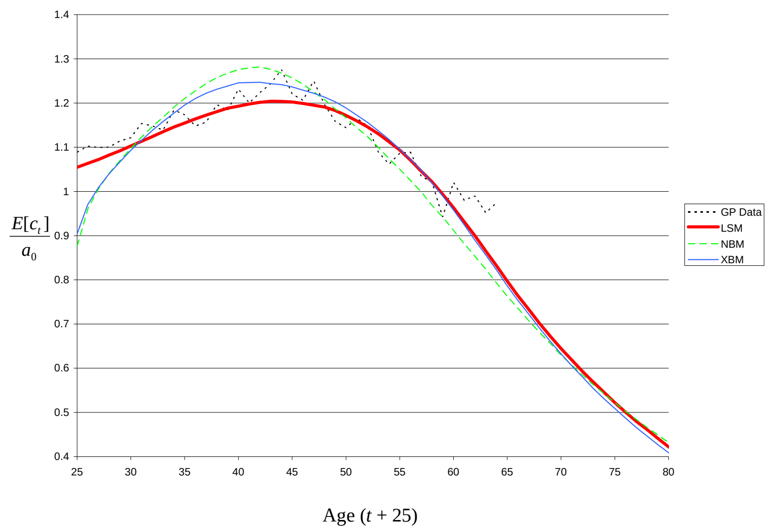

Figure 9: Lifecycle consumption profile for each frictional model with γ calibrated to minimize the mean-squared deviation between the profile and Gourinchas and Parker's (2002) lifecycle consumption data.

Variable	Target	XBM	NBM	LSM
r	3.5%	2.80%	2.19%	3.51%
K/Y	2.5	2.638	2.768	2.498
C/Y	0.75	0.736	0.723	0.750
$t_{\max}(+25)$	45	41	41	43
$c_{\max}/c(0)$	1.15	1.360	1.657	1.141
Borr. Frac.	0.15	0.0376	5×10^{-5}	0.140
s_P	-	0.727	0.746	0.704
RMSD	0	0.064	0.089	0.042

Table 2: Key observables for each frictional model calibrated to the baseline parameters of the LSM, along with target values.

consumption profile are more or less the same, consistent with our finding above that the XBM is insensitive to changes in risk aversion. When we do the same experiment for the NBM, the interest rate decreases and K/Y increases, both more dramatically, since we are increasing risk aversion all the way from 1.5 to 3.5. Initial consumption also decreases drastically.

The XBM, for its baseline parameters, and all three models, with the LSM baseline parameters, give aggregate precautionary saving rates of roughly two thirds, which is in line with the estimate of Gourinchas and Parker (2001). For the XBM one could, nevertheless, discount these estimates because γ and, therefore, s_P is essentially unidentified. But for the LSM, γ is well identified, so if the Leland-Sandmo model is correct the fraction of saving due to precautionary motives must, in fact, be quite large.

Note that Figs. 9 and 10 both show that all three frictional models capture the hump fairly well around its peak. Where the LSM profile matches the data better is at very young ages, when Gourinchas and Parker (2002) find empirically that consumption exceeds income.[17] The NBM and XBM cannot account for this behavior because they both disallow borrowing. The LSM can because it does not.

One aspect of Table 1 that may be troubling is the low discount factor of 0.82 for the baseline LSM and 0.83 for the baseline XBM. Macroeconomists often have a strong prior against discount rates as large as 17%. However, this prior has its roots in frictionless real business cycle (RBC) models for which a calibration of $\beta \sim 0.96$ is necessary to hit standard target values for K/Y and C/Y. It does not come from studies that directly elicit information about consumers' discount rates, which are, of course, unobservable.

Indeed, Warner and Pleeter (2001) found in an experiment that 90% of subjects had a net marginal rate of time preference $D = \rho + \gamma g > 17\%$, where ρ is the discount rate and g is the growth rate of consumption.[18] What we find here

[17] Fernandez-Villaverde and Krueger (2005) also report that initial consumption is higher than initial income.

[18] A risk aversion coeffient of 3.5 is also in line with experimental data. The 1996 Panel

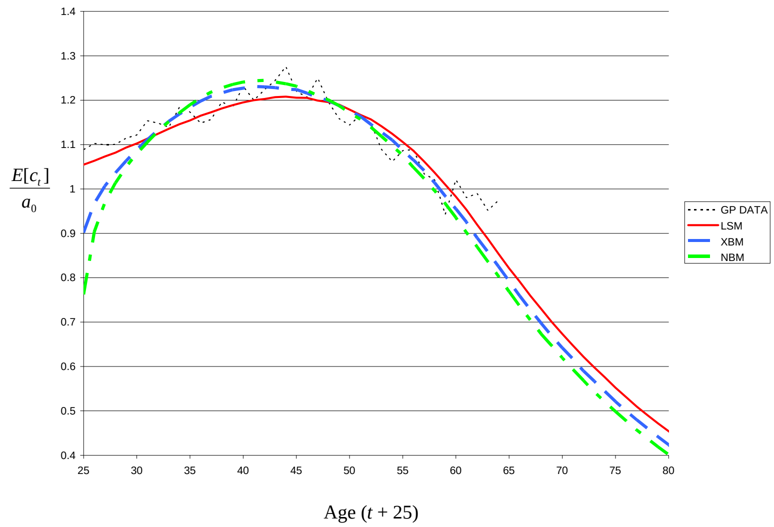

Figure 10: Lifecycle consumption profile for each frictional model with parameters for baseline LSM and Gourinchas and Parker's (2002) lifecycle consumption data.

is that, in macroeconomic OLG models with uninsurable idiosyncratic risk and relative risk aversion $\gamma \gtrsim 3$, there is significant precautionary saving that is not included in frictionless RBC models. Indeed, for both the LSM and the XBM, two thirds of the capital stock can be attributed to aggregate precautionary saving. Because this saving would have a substantial macroeconomic impact, we should not expect that we could match target values for macroeconomic observables with the same preference calibration that would achieve those values in a standard RBC model.

Keep in mind also that only $(\beta R)^{1/\gamma}$ is identifiable in the frictionless model, not β and γ separately. In the frictionless model, γ is best interpreted as the inverse elasticity of intertemporal substitution since there is no uncertainty and, thus, no risk to be averse toward. In the XBM, we need a discount factor of $\beta = 0.93$ to obtain an equilibrium interest rate of $r = 3.5\%$ with $\gamma = 1$. If 3.5% was the equilibrium interest rate in the FM, we would have an observationally equivalent FM equilibrium with $\beta = 0.78$ and $\gamma = 3.5$. Likewise, for the LSM, we need a discount factor of $\beta = 0.94$ to obtain an equilibrium interest rate of $r = 3.5\%$ with $\gamma = 1$. If $r = 3.5\%$ was the equilibrium interest rate in the FM for these parameters, we would also have an FM equilibrium with $\beta = 0.80$ and

Study of Income Dynamics included a set of questions intended to elicit information about risk aversion. Roughly 60% of the subjects between ages 25 and 55 gave answers consistent with a risk aversion greater than 2 and about half gave answers consistent with $\gamma > 3.75$.

$\gamma = 3.5$. Thus the reason why we need β less than 0.9 in the baseline equilibria for the XBM and the LSM is because this is necessary to maintain the same equilibrium interest rate with a lower elasticity of intertemporal substitution. Accounting for the frictions actually raises the required discount factor.

3.6 The Frictional Models in Partial Equilibrium

If we consider only the behavior of households in these models, the general-equilibrium version of each model is nested within the corresponding partial-equilibrium model since we obtain the former by imposing the additional constraint that markets must clear. In the general-equilibrium model, the scalar parameters are $(\alpha, \beta, \gamma, \delta)$, though α and δ only matter within the household's problem to the extent that R is a function of all four of these parameters. In the partial-equilibrium model, the parameters are $(\alpha, \beta, \gamma, \delta, R)$, but α and δ are nuisance parameters that play no role. Thus the preceding results could all be obtained in partial equilibrium as well as in general equilibrium.

Nevertheless, the choices of β and γ that work best in general equilibrium for the LSM and XBM models differ substantially from what macroeconomists would ordinarily consider. Here we consider how the three frictional models would fare in partial equilibrium if we restrict $\beta = 0.96$ and $R = 1.035$ as Gourinchas and Parker (2002) do and then let γ vary. The lifecycle consumption profiles for representative values of γ between 0.5 and 5.0 are shown for the NBM in Fig. 11. Here we see that Gourinchas and Parker's (2002) value of $\gamma = 0.5$ does indeed give a consumption hump most like what is in the data. As we increase γ, the hump shifts later in age and gets much bigger, producing increasingly counterfactual profiles. If we do the same experiment for the XBM, as is shown in Fig. 12, we get quite similar results.

The consumption profiles for the LSM are shown in 13. For large γ, they are essentially the same as in the models without borrowing. For $\gamma = 0.5$, there is only a very slight consumption hump as consumers do borrow when they are young so they can better smooth their consumption. As a consequence, consumption does not peak as high as it does in the XBM or the NBM, so the LSM has a marginally larger root mean squared deviation of 0.09 as compared to 0.08 for the models without borrowing. Thus in partial equilibrium, in the vicinity of Gourinchas and Parker's (2002) baseline values, the XBM and NBM better account for the consumption hump than the LSM.

4 Comparison to Other Mechanisms

Now that we have seen that the Leland-Sandmo Model can fit the data better than the other two frictional models, it is natural to ask how it compares to other models that can account for the consumption hump. We limit our attention to general-equilibrium models because the consumption profile can be

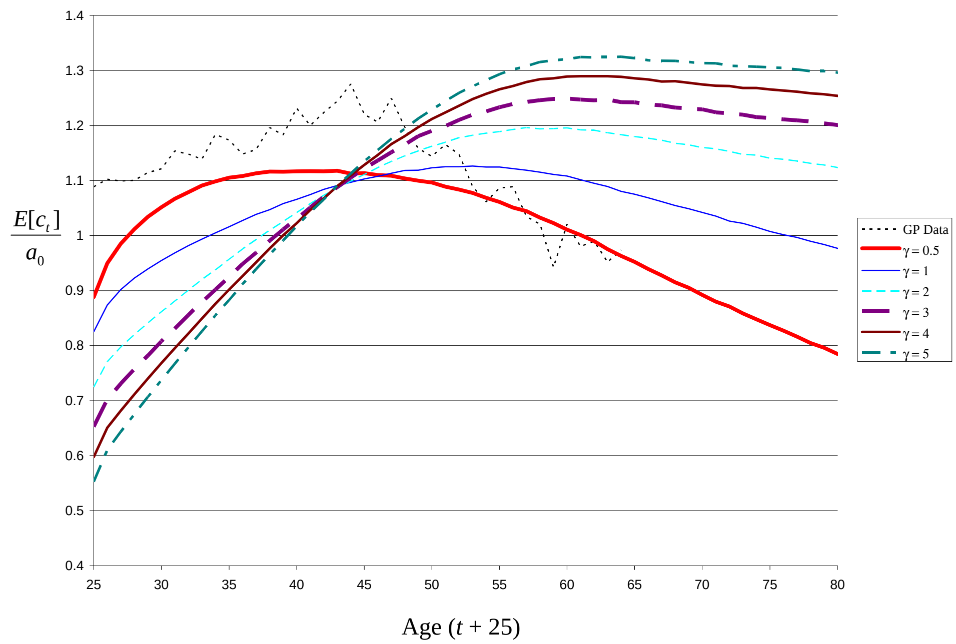

Figure 11: The lifecycle consumption profile for the endogenous borrowing-limit model (NBM) in partial equilibrium for representative values of γ with $\beta = 0.96$ and an interest rate of 3.5% per year along with the Gourinchas and Parker consumption data.

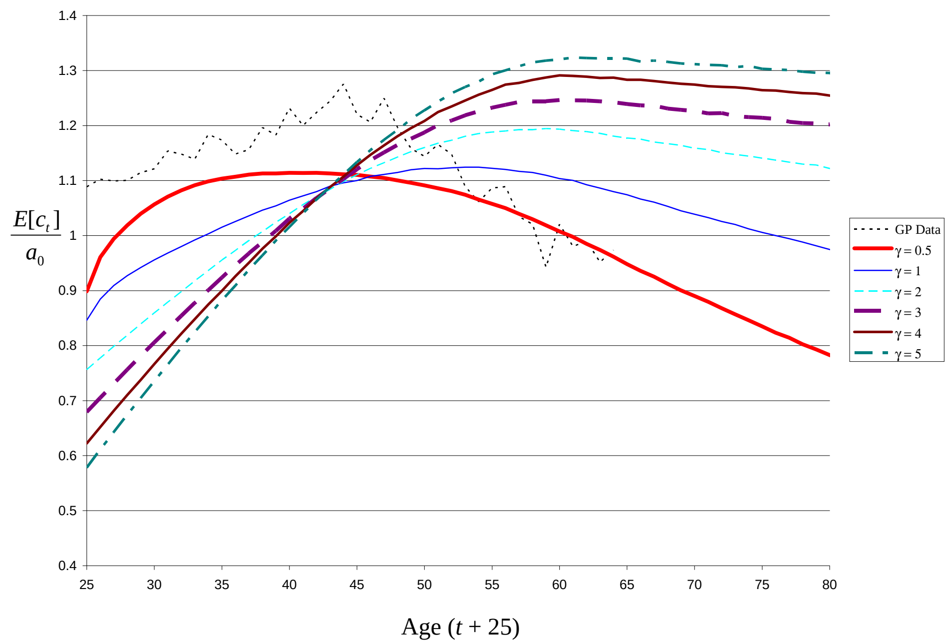

Figure 12: The lifecycle consumption profile for the exogenous borrowing-constraint model (XBM) in partial equilibrium for representative values of γ with $\beta = 0.96$ and an interest rate of 3.5% per year along with the Gourinchas and Parker consumption data.

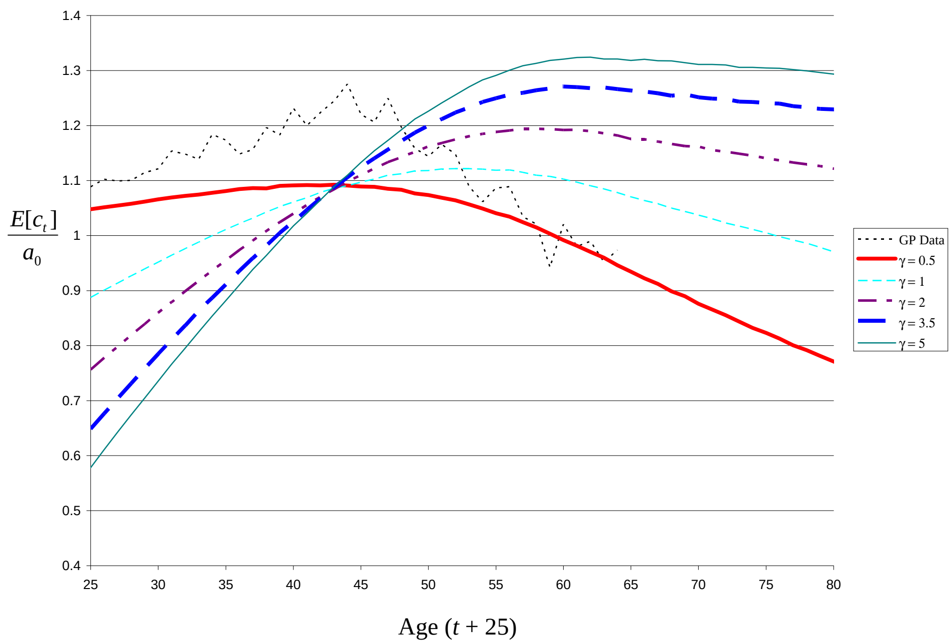

Figure 13: The lifecycle consumption profile for the Leland-Sandmo model (LSM) in partial equilibrium for representative values of γ with $\beta = 0.96$ and an interest rate of 3.5% per year along with the Gourinchas and Parker consumption data.

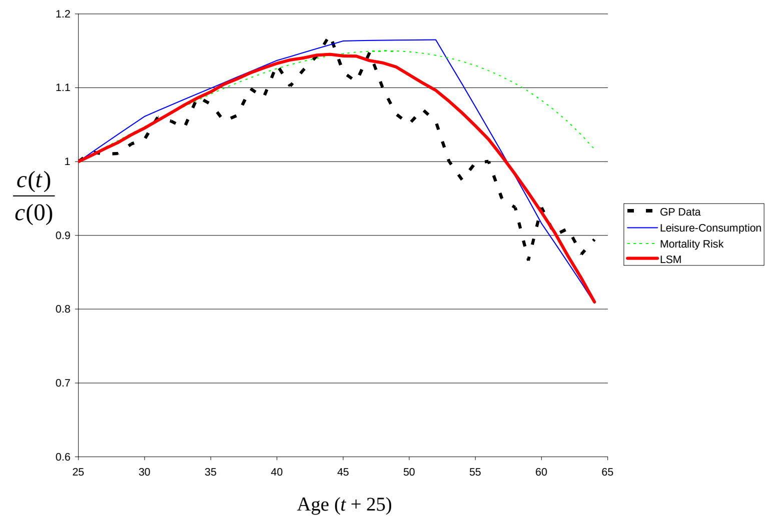

Figure 14: Comparison of lifecycle consumption profiles, normalized to 1 at age 25, for models with mortality risk and leisure-consumption substitution to the baseline LSM model and Gourinchas and Parker's (2002) empirical consumption data.

replicated almost trivially if both the discount factor and the interest rate are free parameters.

Fig. 14 plots the lifecycle consumption profile of the baseline Leland-Sandmo Model alongside the profile of a model with leisure-consumption substitution (Bullard and Feigenbaum (2007)), a model with time-varying mortality risk (Feigenbaum (2008a)), and Gourinchas and Parker's (2002) empirical consumption profile. Note that the primitives of the leisure-consumption model are different from the other two models since it is labor productivity as a function of age that is exogenous rather than labor income. Thus, instead, of comparing the four profiles in absolute terms, we compare them relative to consumption at age 25.

The Leland-Sandmo model obviously fits the data better in Fig. 14. Note, however, that the Bullard-Feigenbaum model faces more constraints since labor supply is also endogenous in that model, so the baseline model was calibrated to fit more observables. As will be shown for the Leland-Sandmo Model, the properties of the consumption hump in both these models are quite robust to small changes in parameters other than the elasticity of intertemporal substitution. This is not the case for the mortality-risk model, which can only fit the data to the extent that it does in Fig. 14 for the specific calibration of the baseline model.

Fernandez-Villaverde and Krueger (2005) also consider a general-equilibrium model of the consumption hump with incomplete markets. Whereas in the present paper we consider what happens if we eliminate unrealistic no-borrowing constraints, they consider the effects of more realistic constraints. There is no unsecured borrowing, but durable goods can be used as collateral for loans. We do not include their consumption profile in Fig. 14 because they calibrate their model to different consumption data (Fernandex-Villaverde and Krueger (2007), for which the consumption profile has its peak at age 52 with a peak to initial ratio of 1.26. In their benchmark model, they find the lifecycle profile for nondurable consumption peaks between ages 40 and 45 with a peak to initial consumption ratio of 1.4. Like Bullard and Feigenbaum (2007), they make a more constrained calibration since the lifecycle profile for durable consumption goods is also endogenous in their model.

5 Robustness Checks

In Section 3.5, we set $(\alpha, \beta, \gamma, \delta)$ to fit the macro variables α, K/Y, and C/Y and to best fit the lifecycle profile of mean consumption to Gourinchas and Parker's (2002) empirical profile. We have seen how these variables and the consumption profile vary with γ, but how do they vary with the other parameters? Are the properties of the consumption hump specific to our baseline calibration?

Table 3 shows they are not. As we vary α and δ around their LSM baseline values, only the interest rate makes huge adjustments. Likewise, if we redo the calibration with $K/Y = 3$, our results are quite similar to the baseline model. A consumption hump with a peak in the 40s and a peak to initial consumption ratio between 1.07 and 1.25 is seen for all of these perturbations of the scalar parameters from their baseline.

The only parameter that our results are greatly sensitive to is the discount factor β, but this is not surprising in the light of Section 3.6. There we saw that with a value of the discount factor that would generally be considered more reasonable the LSM cannot account for the consumption hump. Here we see that we would get a consumption hump in general equilibrium, though the root mean squared deviation is quite large, but the macroeconomic variables are amiss. With $\gamma = 3.75$, there will be so much precautionary saving that equilibrium interest rates will be negative. Yet even with such low interest rates, agents will not borrow because of their fear of uncertainty. Instead, they oversave, and the equilibrium is dynamically inefficient with huge amounts of capital and little consumption.

Another concern is how sensitive our results are to the simplifying assumption of a 2-state process for both the temporary and permanent income shocks, which forces both shock processes to have a kurtosis of 1. In Fig. 15, we compare the LSM lifecycle consumption profiles for income processes where

Model	r	K/Y	C/Y	Borr. Frac.	$t_{\max}(+25)$	$c_{\max}/c(0)$	RMSD
LSM Baseline	3.51%	2.498	0.750	0.140	43	1.141	0.042
$\alpha = 0.3$	1.76%	2.550	0.745	0.170	40	1.071	0.080
$\alpha = 0.36$	4.57%	2.470	0.753	0.122	45	1.188	0.045
$\alpha = 0.3, \delta = 0.07$	4.20%	2.679	0.812	0.129	45	1.170	0.041
$\beta = 0.9$	-3.11%	4.896	0.510	0.0	39	1.082	0.378
$\beta = 0.96$	-4.85%	6.547	0.345	0.0	47	1.279	0.664
$\delta = 0.07$	5.80%	2.636	0.815	0.103	47	1.245	0.072
$\delta = 0.12$	1.94%	2.422	0.709	0.167	41	1.077	0.075
Kurtosis 3	3.29%	2.540	0.746	0.129	43	1.160	0.046
$K/Y = 3$	2.91%	3.002	0.750	0.084	44	1.166	0.052

Table 3: Sensitivity analysis for baseline LSM model. First column specifies which parameters are changed from their baseline values.

the two shock processes have the 2-state representation and where they have a 5-state representation with a kurtosis of 3 (as would arise with the standard assumption of a log-normal distribution for the shocks). Relevant macroeconomic variables are also listed in the bottom row of Table 3. There is only negligible change between the two models, which is consistent with Feigenbaum's (2008b) result that only the variance of the shocks has a significant effect on the mean of consumption for each age group. Thus our assumption of a 2-state process is reasonable.

6 Conclusions

Although the majority of current work on consumption over the lifecycle involves precautionary saving to protect against uninsurable, idiosyncratic risk, there is considerable disagreement about how much influence this saving actually has on the macroeconomy. Here we consider a general-equilibrium lifecycle model and calibrate the model to match the properties of mean consumption over the lifecycle. Surprisingly, we find that the effect of risk aversion on observable macro variables like the interest rate, the capital to output ratio, and the properties of the lifecycle profile of mean consumption crucially depend on the assumptions we make about what borrowing constraints are in the model. If we impose an exogenous no-borrowing constraint, all of these observables are essentially independent of risk aversion. Therefore it is reasonable to suppose that risk aversion is close to 1, as most economic research assumes, in which case precautionary saving would be unimportant at the aggregate level. However, if we include only minimal borrowing frictions in the model, the properties of the mean consumption profile that can be generated by precautionary saving alone do depend on risk aversion. A risk aversion of 3.75 is then best able to match

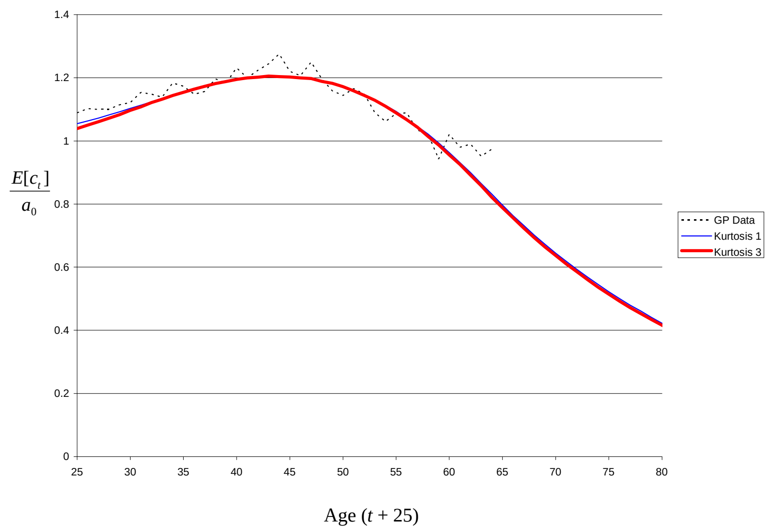

Figure 15: Lifecycle consumption profile for the baseline LSM with income shocks calibrated to have a kurtosis of 1 (2 states) or 3 (5 states) and Gourinchas and Parker's (2002) lifecycle consumption data.

the properties of lifecycle consumption, and this calibration also yields that two thirds of the capital stock is due to precautionary saving.

We should emphasize that the findings of this paper do not repudiate borrowing constraints. The types of borrowing constraints considered in this paper, though common in the literature, are overly strong, designed for simplicity instead of realism. What we find here is that, while such stringent constraints can account for the qualitative properties of the lifecycle consumption profile, a model with no such constraints can better accomodate the quantitative properties of the profile. This does not mean we should abandon borrowing constraints but rather that we should focus our attention on matching the real-world features of constraints faced by borrowers. For example, we could introduce collateralized borrowing as Fernandez-Villaverde and Krueger (2005) have done, or we could allow for a wedge between the interest rate paid by borrowers and the interest rate received by lenders.

In other work, including Bullard and Feigenbaum (2007) and Hansen and Imrohoroglu (2008), researchers have considered the ability of alternative mechanisms, such as leisure-consumption substitution and time-varying mortality risk, to account for the hump-shaped profile of mean consumption over the lifecycle. What is not clear is how different mechanisms that can account for the hump in isolation will interact in combination. Although one might expect these mechanisms to be complementary, preliminary evidence suggests the reality is more complicated. For example, borrowing constraints and uninsurable risk are almost always studied together, but as we have seen here these two mechanisms do not enhance each other. On the contrary, a tight borrowing constraint can suppress the ability of precautionary saving to explain the consumption hump. Likewise, Feigenbaum (2008a) has shown that mortality risk can suppress the ability of a tight borrowing constraint to explain the hump. Future work needs to carefully consider what happens when these different mechanisms, most of which obviously do have a role to play in the real world, are brought together, both in terms of their predictions regarding the lifecycle consumption profile and their predictions about other variables that these mechanisms entangle with consumption.

A Solution Method

The theory regarding the behavior of the consumption function in the frictional models presented here is a straightforward generalization of Carroll (2004). For $t \geq T_W - 1$, there is no future income, so the consumption and value functions have the well-known solution

$$v_t(x_t, p_t) = \left(\frac{1 - \phi^{-(T-t)}}{1 - \phi^{-1}} \right)^{\gamma} \frac{x_t^{1-\gamma}}{1-\gamma} \tag{18}$$

and

$$c_t(x_t, p_t) = \frac{1 - \phi^{-1}}{1 - \phi^{-(T-t)}} x_t, \tag{19}$$

where

$$\phi = (\beta R^{1-\gamma})^{-1/\gamma}. \tag{20}$$

Thus the consumer smooths his remaining wealth x_t over the remaining $T - t$ periods.

To begin with let us consider what happens if there is no exogenous borrowing constraint, i.e. in the LSM and NBM. The Bellman equation for these models at $t < T_W - 1$ can be written for $x_t \geq -wa_t p_t M_t$

$$\begin{aligned} v_t(x_t, p_t) \quad = \quad & \max_{0 \leq c_t \leq x_t + wa_t p_t M_t} \{u(c_t) \\ & + \beta E \left[v_{t+1}(wa_{t+1}\widetilde{q}_{t+1}\widetilde{z}_{t+1} p_t + R(x_t - c_t), \widetilde{q}_{t+1} p_t) \right] \}, \end{aligned} \tag{21}$$

where

$$M_t = \sum_{s=t+1}^{T_W - 1} \frac{a_s}{a_t} Z_1 \left(\frac{Q_1}{R} \right)^{s-t}.$$

The lower limit $-wa_t p_t M_t$ on x_t arises because it is the Aiyagari (1994) endogenous limit on borrowing $-b_{t+1}$. From a strictly mathematical standpoint, it is not necessary to explicitly include the constraint $c_t \leq x_t + wa_t p_t M_t$ in (21) since the consumer would never choose to violate this limit, but the upper limit needs to be identified to solve the model numerically.

The permanent income p_t factors out of the righthand side of (21), so we can reduce the problem to one dimension. Define

$$C_t = \frac{c_t}{wa_t p_t}$$

and

$$X_t = \frac{x_t}{wa_t p_t}.$$

Then we can define $V_t(X_t)$ such that

$$v_t(x_t, p_t) = (wa_t p_t)^{1-\gamma} V_t \left(\frac{x_t}{wa_t p_t} \right),$$

and the Bellman equation for V_t at $X_t \geq -M_t$ is

$$\begin{aligned} V_t(X_t) \quad = \quad & \max_{0 \leq C_t \leq X_t + M_t} \{u(C_t) \\ & + \beta E \left[\left(\frac{a_{t+1}}{a_t} \widetilde{q}_{t+1} \right)^{1-\gamma} V_{t+1} \left(\widetilde{z}_{t+1} + \frac{a_t}{a_{t+1}\widetilde{q}_{t+1}} R(X_t - C_t) \right) \right] \right\}. \end{aligned} \tag{22}$$

The first-order condition for this problem is the Euler equation

$$C_t(X_t)^{-\gamma} = \beta R E \left[\left(\frac{a_{t+1}}{a_t} \widetilde{q}_{t+1} \right)^{-\gamma} C_{t+1} \left(\widetilde{z}_{t+1} + \frac{a_t}{a_{t+1}\widetilde{q}_{t+1}} R(X_t - C_t(X_t)) \right)^{-\gamma} \right]. \tag{23}$$

In the limit of large X_t, the consumption function asymptotes to the linear function that would arise in the absence of uncertainty:

$$C_t(X_t) = \frac{1-\phi^{-1}}{1-\phi^{-(T-t)}}\left(X_t + \sum_{s=t+1}^{T_W-1}\frac{a_s}{a_t}\left(\frac{1}{R}\right)^{s-t}\right) + O\left(\frac{1}{X_t}\right), \qquad (24)$$

which is proportional to cash on hand plus the expected present value of future income. In the limit as $X_t \to -M_t$, the behavior of the consumption function depends on the probability of receiving the lowest income shock. Let $\rho^* = \rho_1$ if $Z_1 > 0$ (in which case the lowest income shock only arises if both the permanent and temporary shocks equal their lowest possible values) and $\rho^* = 1$ if $Z_1 = 0$ (in which case the lowest income shock arises independent of the permanent shock). Then the consumption function approximates to

$$C_t(X_t) = K_t(X_t + M_t) + O\left((X_t + M_t)^2\right), \qquad (25)$$

where

$$K_t^{-1} = 1 + \phi^{-1}(\pi_1\rho^*)^{1/\gamma}K_{t+1}^{-1} \qquad (26)$$

and

$$K_{T_W-1} = \frac{1-\phi^{-1}}{1-\phi^{-(T-T_W+1)}}.$$

Schumaker shape-preserving splines (Judd (1999)) were used to approximate the consumption function for $t < T_W - 1$. The splines were computed using a grid of 1000 nodes between $-M_t$ and $X_{\max}$, where $X_{\max}$ was chosen large enough so the approximation (24) would be valid. The grid points were spaced geometrically so each interval of the spline was 1.01 times wider than the interval to the left. The spline requires us to compute both the value and the slope of the consumption function at each grid point. At $X_t = -M_t$, these were computed using (25). At $X_t = X_{\max}$, they were computed using (24). In between, given the Schumaker spline representing $C_{t+1}(X_{t+1})$, we solved (23) for $C_t(X_t)$ at each gridpoint. Then the slope is given by

$$\begin{aligned} C_t'(X_t)^{-1} &= 1 + \beta^{-1}R^{-2}C_t^{-\gamma-1}(X_t) \qquad (27) \\ &\quad \times E\left[\left(\frac{a_{t+1}}{a_t}\widetilde{q}_{t+1}\right)^{-\gamma-1}\frac{C_{t+1}'\left(X_{t+1}(X_t,\widetilde{q}_{t+1},\widetilde{z}_{t+1})\right)}{C_{t+1}\left(X_{t+1}(X_t,\widetilde{q}_{t+1},\widetilde{z}_{t+1})\right)^{\gamma+1}}\right]^{-1}, \end{aligned}$$

where

$$X_{t+1}(X_t, q_{t+1}, z_{t+1}) = z_{t+1} + \frac{a_t}{a_{t+1}q_{t+1}}R(X_t - C_t(X_t)).$$

For the XBM, we use the same approach except now the consumption and value functions are only defined for $X_t \geq 0$. The borrowing constraint will bind for $0 \leq X_t \leq \underline{X}_t$, where

$$\underline{X}_t = (\beta R)^{-1/\gamma}E\left[\left(\frac{a_{t+1}\widetilde{q}_{t+1}}{a_t}C_{t+1}(\widetilde{z}_{t+1})\right)^{-\gamma}\right]^{-1/\gamma},$$

in which case $C_t(X_t) = X_t$. For $X_t > \underline{X}_t$, the consumption function will still be determined by the Euler equation (23).

The equilibrium interest rate was obtained by solving Eq. (10) for R. The capital stock $K(R)$ was obtained from Eq. (7) by solving the consumption function for each t and then simulating the model with a million simulations per cohort. The large number of simulations was necessary to get reliable lifecycle consumption profiles since the persistence of the income process leads to a large cross-sectional variance in c_t for each t. Computing the Euler error

$$\beta RE\left[\left(\frac{a_{t+1}}{a_t}\widetilde{q}_{t+1}\frac{C_{t+1}\left(C_{t+1}\left(X_{t+1}(X_t,\widetilde{q}_{t+1},\widetilde{z}_{t+1})\right)\right)}{C_t(X_t)}\right)^{-\gamma}\right]-1$$

for each point visited in a simulation, the maximum Euler error obtained is on the order of 10^{-7} for the LSM and NBM, and on the order of 10^{-4} for the XBM. The consumption function is more difficult to compute precisely in the XBM because of the kinks in the consumption function that arise where the borrowing constraint starts to bind.

References

[1] Aguiar, Mark and Erik Hurst, (2003), "Consumption vs Expenditure," Working Paper.

[2] Aguiar, Mark and Erik Hurst, (2007), "Life-Cycle Prices and Production," *American Economic Review* 97: 1533-1559.

[3] Aiyagari, S. Rao, (1994), "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics* 109: 659-684.

[4] Attanasio, Orazio P., James Banks, Costas Meghir, and Gugielmo Weber, (1999), "Humps and Bumps in Lifetime Consumption," *Journal of Business & Economic Statistics* 17: 22-35.

[5] Attanasio, Orazio P. and Hamish Low, (2004), "Estimating Euler Equations," *Review of Economic Dynamics* 7: 406-435.

[6] Becker, Gary S. and Gilbert R. Ghez, (1975), *The Allocation of Time and Goods Over the Life Cycle* (New York: NBER and Columbia Press).

[7] Bullard, James and James Feigenbaum, (2007), "A Leisurely Reading of the Lifecycle Consumption Data," forthcoming in *Journal of Monetary Economics.*

[8] Caliendo, Frank and David Aadland, (2007), "Short-Term Planning and the Life-Cycle Consumption Puzzle," *Journal of Economic Dynamics and Control* 31: 1392-1415.

[9] Carroll, Christopher D., (1997), "Buffer-Stock Saving and the Life-Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics* 112: 1-55.

[10] Carroll, Christopher D., (2000), "Why Do the Rich Save So Much?" in *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*, ed. by Joel B. Slemrod (Harvard University Press: Cambridge, MA).

[11] Carroll, Christopher D., (2001), "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation," *Advances in Macroeconomics* 1: Article 6.

[12] Carroll, Christopher D., (2004), "Theoretical Foundations of Buffer Stock Saving," Working Paper.

[13] Carroll, Christopher D. and Andrew A. Samwick, (1998), "How Important is Precautionary Saving?" *Review of Economics and Statistics* 80(3): 410-419.

[14] Carroll, Christopher D. and Lawrence H. Summers, (1991), "Consumption Growth Parallels Income Growth: Some New Evidence," in *National Saving and Economic Performance* eds. B. Douglas and J. Shoven (Chicago: University of Chicago Press).

[15] Constantinides, George M. and Darrell Duffie, (1996), "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy* 104: 219-240.

[16] Deaton, Angus, (1991), "Saving and Liquidity Constraints," *Econometrica* 59: 1221-1248.

[17] Feigenbaum, James, (2008a), "Can Mortality Risk Explain the Consumption Hump?" *Journal of Macroeconomics* 30: 844-872.

[18] Feigenbaum, James, (2008b), "Information Shocks and Precautionary Saving," *Journal of Economic Dynamics and Control* 32: 3917-3938.

[19] Fernández-Villaverde, Jesús and Dirk Krueger, (2005), "Consumption and Saving over the Life Cycle: How Important are Consumer Durables?" Working Paper.

[20] Fernández-Villaverde, Jesús and Dirk Krueger, (2007), "Consumption over the Life Cycle: Some Facts from Consumer Expenditure Survey Data," *Review of Economics and Statistics* 89: 552-565.

[21] Gourinchas, Pierre-Olivier and Jonathan A. Parker, (2001), "The Empirical Importance of Precautionary Saving," *American Economic Review* 91: 406-412.

[22] Gourinchas, Pierre-Olivier and Jonathan A. Parker, (2002), "Consumption over the Life Cycle," *Econometrica* 70: 47-89.

[23] Gross, David B. and Nicholas S. Souleles, (2002), "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data," *Quarterly Journal of Economics* 67: 149-187.

[24] Hall, Robert E., (1988), "Intertemporal Substitution in Consumption," *Journal of Political Economy* 96: 339-357.

[25] Hansen, Gary D. and Selahattin İmrohoroğlu, (2008), "Consumption over the Life Cycle: The Role of Annuities," *Review of Economic Dynamics* 11: 566-583.

[26] Heckman, James, (1974), "Life Cycle Consumption and Labor Supply: An Explanation of the Relationship Between Income and Consumption Over the Life Cycle," *American Economic Review* 64: 188-194.

[27] Huggett, Mark, (1993), "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control* 17: 953-969.

[28] Huggett, Mark, (1996), "Wealth Distribution in Life-Cycle Economies," *Journal of Monetary Economics* 38: 469-494.

[29] Huggett, Mark and Sandra Ospina, (2001), "Aggregate Precautionary Savings: When is the Third Derivative Irrelevant?" *Journal of Monetary Economics* 48: 373-396.

[30] Judd, Kenneth L., (1999), *Numerical Methods in Economics* (MIT Press: Cambridge, MA).

[31] Laibson, David, (1997), "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics* 112: 443-477.

[32] Leland, Hayne E., (1968), "Saving and Uncertainty: The Precautionary Demand for Saving," *Quarterly Journal of Economics* 82: 465-473.

[33] Ludvigson, Sydney and Christina H. Paxson, (2001), "Approximation Bias in Linearized Euler Equations," *Review of Economics and Statistics* 83: 242-256.

[34] Nagatani, Keizo, (1972), "Life Cycle Saving: Theory and Fact," *American Economic Review* 62: 344-353.

[35] Rios-Rull, Jose-Victor, (1996), "Life-Cycle Economies and Aggregate Fluctuations," *Review of Economic Studies* 63: 465-489.

[36] Sandmo, A., (1970), "The Effect of Uncertainty on Saving Decisions," *Review of Economic Studies* 37: 353-360.

[37] Skinner, Jonathan, (1988), "Risky Income, Life Cycle Consumption, and Precautionary Savings," *Journal of Monetary Economics* 22: 237-255.

[38] Thurow, Lester C., (1969), "The Optimum Lifetime Distribution of Consumption Expenditures," *American Economic Review* 59: 324-330.

[39] Warner, John T. and Saul Pleeter, (2001), "The Personal Discount Rate: Evidence from Military Downsizing Programs," *American Economic Review* 91: 33-53.

[40] Wolff, E. N., (2001), "Recent Trends in Wealth Ownership, 1983-1998," in *Assets for the Poor: The Benefits of Spreading Asset Ownership*, eds. by Thomas M. Shapiro and Edward N. Wolff (Russell Sage Press: New York, NY).

[41] Xu, Xiaonian, (1995), "Precautionary Saving Under Liquidity Constraints: A Decomposition," *International Economic Review* 36: 675-690.

[42] Yaari, Menahem E., (1964), "On the Consumer's Lifetime Allocation Process," *International Economic Review* 5: 304-317.