

Consumption Commitments and Asset Prices*

Raj Chetty

UC-Berkeley and NBER

Adam Szeidl

Harvard University

February 2004

Abstract

This paper studies portfolio choice and asset prices in a model with two consumption goods, one of which involves a commitment in that its consumption can only be adjusted at a cost. Commitments effectively make investors more risk averse: they invest less in risky assets and smooth total consumption more. Aggregating over a population of such consumers implies dynamics that match those of a representative consumer economy with habit formation. Calibrations show that the model can resolve the equity premium puzzle. We test the key prediction that an exogenous increase in economic commitments (e.g., housing) causes a more conservative portfolio allocation using a novel instrumental variables strategy related to age at marriage. We find that a \$1 increase in housing causes a 50-70 cent reallocation from stocks to bonds for the average investor. Exploiting differences in the variance of home prices across cities, we show that this effect is due to commitments and not greater exposure to housing price risk.

JEL Classification: *G11, G12*

Keywords: *Habit formation, Portfolio choice, Housing*

*E-mail addresses: chetty@econ.berkeley.edu and szeidl@fas.harvard.edu. We would like to thank Gary Chamberlain, Tom Davidoff, Ed Glaeser, Elhanan Helpman, Larry Katz, Miklos Koren, Emmanuel Saez, Jesse Shapiro, Todd Sinai, Jeremy Stein, Moto Yogo, and especially John Campbell and Drew Fudenberg for helpful comments and suggestions. Kory Kroft and Ankur Patel provided capable research assistance. Chetty acknowledges funding from the National Science Foundation and National Bureau of Economic Research. Szeidl thanks the Social Science Research Council and the Institute for Humane Studies for support.

1 Introduction

Much of the modern literature on asset pricing has focused on explaining the equity premium puzzle: Consumption growth covaries only weakly with stock returns – implying that stocks are not very risky – yet the risk premium commanded by stocks is very large.¹ Most existing models of asset pricing assume that agents consume one good, i.e., that there is a composite commodity. This assumption requires that agents can substitute costlessly across different types of consumption within a given period. But in practice, many goods, such as housing and other illiquid durables, involve “commitments” – a transaction cost must be paid to change consumption of these goods. This paper proposes a solution to the equity premium puzzle by showing that such consumption commitments have significant effects on portfolio choice and asset pricing both theoretically and empirically.

The basic intuition underlying our analysis is straightforward. Consider an individual who consumes commitment (e.g., housing) and non-commitment (e.g., food) goods in equal shares.² After making his commitment, suppose the individual faces a shock that necessitates a 10% reduction in expenditure. Since commitments are difficult or costly to adjust in the short run, the individual may rationally choose not to move out of his house, making the shock translate into a 20% drop in food consumption. The presence of commitments magnifies the impact of shocks, effectively making the consumer more risk averse. As a result, agents with commitments may rationally hold less risky portfolios and demand a higher equity risk premium.

We formalize this logic by modelling the portfolio choice decision of an individual who has neoclassical preferences over the two types of consumption goods. A number of assumptions are necessary to make the model analytically tractable. We assume that consumers have simple, constant elasticity of substitution preferences over commitment and non-commitment goods. To model costly adjustment, we use a time-dependent adjustment framework, rather than the more natural but intractable state-dependent setup. We consider two different time-dependent adjustment rules. The first, which is similar

¹See Kocherlakota (1996) and Campbell (2002) for recent surveys of the literature on the equity premium and consumption based asset pricing.

²Tabulations from the Consumer Expenditure Survey (see Table 1) indicate that across a broad range of income categories, approximately 50-60% of expenditure is committed for the average household.

to Taylor-pricing, allows the consumer to reset consumption of the commitment good after a fixed number of periods. The second has a stochastic reset date and resembles the Calvo-pricing rule.³ Our main theoretical results are similar in both models.

We first show that individuals who have more commitments act as if they are more risk averse. They invest less in risky assets and have a smoother path of total consumption. We then consider a population of such individuals and derive aggregate dynamics. The key theoretical result is that aggregate dynamics in a model where individuals have heterogeneity in commitments are identical to the dynamics that arise from a representative consumer model with habit formation utility in the spirit of Constantinides (1990) and Campbell and Cochrane (1999). Moreover, under certain conditions in the stochastic adjustment framework, the aggregate economy coincides exactly with the representative consumer economy of Constantinides, with the exception that we have external rather than internal habit formation. The commitments model therefore provides non-psychological micro-foundations for representative-consumer habit formation models.

The connection between commitments at the microeconomic level and habit formation in the aggregate is quite intuitive. In habit formation models, consumer well-being is determined by surplus consumption over current habit, which is a slow-moving time average of past consumption levels. In the commitments model, part of consumer well-being is determined by consumption over the level of commitments. For the household, these commitments are slow moving in that they are unchanged except at adjustments. When an adjustment takes place, the new commitment level reflects the household's current prosperity, and hence recent levels of consumption. Our aggregation result essentially states that an economy of individuals with fixed costs and infrequent adjustments in consumption look in the aggregate like an individual facing a smooth, slow-moving habit that depends on past consumption levels.

To evaluate the performance of the commitments model in explaining asset pricing puzzles, we perform a calibration exercise similar to Constantinides (1990). We find that our model can match the first and second moments of stock returns and consumption as well as the riskfree rate. The unconstrained coefficient of relative risk aversion, relevant for households when they are free to adjust their commitments, can be as low as $\gamma =$

³Time-dependent adjustment rules have a long history in macroeconomics (see Taylor, 1979, Calvo, 1983 or Blanchard and Fischer, 1989).

2. However, generating the equity premium with such a low γ requires the share of commitments in individual consumption to be about 80%. The model is also capable of matching the equity premium at a share of commitment consumption (50-60%) that matches the actual share of physical commitments if γ is between 4 and 5. Even at these higher levels of risk aversion, the model explains much of the riskfree rate puzzle, though it does not completely resolve it: we are about 0.4 – 0.6 percentage points short of matching mean growth in aggregate consumption. Borrowing constraints and time non-separable utility are likely to further improve the model’s fit.

The central benefit of our micro-foundations is that they yield a natural set of testable predictions. The prediction that we test is the following: Holding all else fixed, an exogenous increase in an agent’s commitments (e.g., size of home mortgage) should make his optimal portfolio allocation shift away from risky assets and toward safe assets. Operationalizing this test is difficult because commitments are chosen endogenously. In particular, agents who are more risk loving or face less background risk are more likely to commit themselves to bigger houses and also invest more in stocks relative to bonds. Not surprisingly, an OLS regression of stockholding on housing commitment, controlling for all observables, shows that house size and stockholding are positively associated.

In view of this endogeneity problem, we employ a novel instrumental variables strategy to create exogenous variation in home size. We exploit the fact the individuals tend to buy homes around the time of marriage, and those who marry later tend to buy bigger homes and therefore have a larger home mortgage outstanding, holding total wealth fixed, at any given age. This effect presumably arises because those who buy later tend to have the resources and credit necessary to buy a larger home.

The central identifying assumption in our empirical analysis is that age at marriage is not directly related to portfolio choice. We provide evidence in support of this orthogonality condition by examining two “control” groups: homeowners who have been married for a long time and renters. Marriage age is unrelated to mortgage size for both of these groups. We find no relationship between age at marriage and stockholding in these groups, confirming that there is no direct association between age at marriage and portfolio choice for these groups. These results support the validity of the instrument; our IV results could only be inconsistent if there is a time-varying relationship between marriage age and risk preferences unique to homeowners. Using age at marriage as an

instrument for housing, we find that a \$1 exogenous increase in mortgage size causes a 50-70 cent shift in portfolio allocation from stocks to bonds. These results are robust to a rich set of controls and specification checks for sample selection and other potential biases.

Having identified what appears to be a strong causal relationship between house size and portfolio choice, we address a competing explanation for these empirical results that has been proposed in the recent literature. Flavin and Yamashita (2002) and Yamashita (2003) observe that homes are a risky asset and show that an exogenous increase in housing will lead to a sharp shift from stocks to bonds under a mean-variance optimization framework. Their argument is that greater risk exposure in the housing market causes less risk taking in other dimensions of the portfolio. To distinguish this theory from ours, we use data on the variance of home prices by city to permit variation in exposure to housing risk while holding home value fixed and vice versa. We find that a \$1 increase in home value continues to cause a large reduction in stockholding even when exposure to home price risk is held fixed. While house price variance does not appear to significantly affect the portfolio allocation of homeowners – perhaps because they have a natural hedge against this risk in their own home – renters do hold less stocks when living in high-risk areas. We conclude that the empirical evidence on housing and portfolio choice strongly supports the commitments hypothesis proposed in this paper.

In addition to the literature on housing risk, our paper also builds on a number of other strands of the asset pricing and consumption literatures. The importance of commitments for asset pricing was documented in a seminal paper by Grossman and Laroque (1990). Their model has a single, durable consumption good, and they do not discuss the relationship between commitments and habit formation. Marshall and Parekh (1999) aggregate and calibrate the Grossman-Laroque model. Maintaining the single (durable) good assumption, they argue that the model can explain approximately half of the equity premium puzzle. In contrast, we are able to fully match the equity premium in our two good framework and test microeconomic implications of the model. More recently, Fratantoni (2001) and Li (2003) analyze two-good models where one of the goods involves adjustment costs. They demonstrate the portfolio choice implications of their models using numerical simulations. These papers do not discuss aggregate dynamics or the empirical connection between housing and portfolio choice documented here. Our

model is also related to Gabaix and Laibson (2001), who study asset pricing with a single good and infrequent adjustments in portfolio choice. In their model, households do not adjust consumption in the short run because they do not observe their risky financial holdings. This generates low consumption volatility over short horizons. In our model, consumers would like to (but cannot) adjust their consumption commitments, which effectively makes them more risk averse.⁴

Risk preferences in the two-good commitments model have also received attention in other contexts. Chetty (2002) studies preferences over wealth in a state-dependent framework and shows that agents with consumption commitments exhibit significantly higher degrees of risk-aversion to moderate-stake wealth fluctuation than they do to large-stake wealth fluctuations using data from labor markets. Olney (1999) finds that large exposure to installment finance, a form of consumption commitments, forced households to cut back on other consumption and was therefore responsible for a significant share of the welfare loss during the Great Depression.

Finally, our results contribute to the growing literature which shows that decomposing aggregate consumption into components sheds light on various asset pricing puzzles. Piazzesi, Schneider and Tuzel (2003) and Yogo (2003) examine the effects of composition risk (fluctuations in housing or durable consumption relative to non-durable consumption), while Lustig and Nieuwerburgh (2003) consider the collateral value of housing. These papers abstract from the commitment feature of consuming these goods, while we abstract from composition risk and collateral value effects. Ait-Sahalia, Parker and Yogo (2003) consider the consumption of luxury and basic goods, and argue that the volatility of luxury goods consumption helps resolving the equity premium puzzle. To the extent that luxury goods are non-commitment, their findings support the hypothesis of this paper.

The remainder of the paper is organized as follows. The next section develops the model and presents the key aggregation results, first with deterministic adjustment and then with stochastic adjustment. Section 3 demonstrates in a calibration that the model can resolve the equity premium puzzle. The fourth section presents empirical evidence for the model and shows how the key risk aversion parameter used in the calibrations

⁴Time dependent adjustment rules are also studied in an asset pricing context by Koren and Szeidl (2003), who use them to model financial illiquidity.

can be imputed from these estimates. The final section concludes.

2 Portfolio Choice and Consumption with Commitments

2.1 A Model of the Household

Consider a consumer with preferences both for consumption commitments, such as housing (x), and non-commitment consumption, such as food (f). The per-period utility function over these two goods is assumed to be a constant elasticity of substitution aggregate with elasticity ε ,

$$u(f, x) = \left(f^{1-\frac{1}{\varepsilon}} + \mu x^{1-\frac{1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}. \quad (1)$$

The consumer maximizes expected lifetime utility given by

$$\max E \int_0^\infty e^{-\rho t} \frac{u(f_t, x_t)^{1-\gamma}}{1-\gamma} dt$$

where ρ is the discount rate. Since utility is assumed to be time-separable, γ measures the relative risk aversion of an individual who is free to adjust both commitment and food consumption, as well as the elasticity of intertemporal substitution. In the following, we further assume that $\gamma = 1/\varepsilon$, that is, the intertemporal and across-good elasticities of substitution are equal. This simplification is necessary to make the model tractable.

Under this assumption, the maximization problem becomes

$$\max E \int_0^\infty e^{-\rho t} \left(\frac{f_t^{1-\gamma}}{1-\gamma} + \mu \frac{x_t^{1-\gamma}}{1-\gamma} \right) dt. \quad (2)$$

Our next modelling choice concerns commitments. We assume that consumption commitments are costly to adjust in the following sense. Every T periods, the consumer is free to adjust her level of commitment consumption x ; however, between adjustment dates x is fixed, and no adjustment is possible.⁵

There are two assets traded in this economy. The price process of the riskless bond is given by

$$\frac{dB_t}{B_t} = r dt$$

⁵In Section 2.3 we consider a stochastic version of this adjustment rule.

where r is the riskfree rate, which is assumed to be constant. The dynamics of the risky asset are given by the standard exponential Brownian motion

$$\frac{dS_t}{S_t} = (r + \pi)dt + \sigma dz_t$$

where z_t is a standard Brownian motion, π is the expected excess return (equity premium), and σ is the standard deviation of asset returns.

Let w_t denote the wealth of the consumer in period t and α_t the share of the risky asset in the consumer's wealth portfolio. Define total consumption as $c = f + x$. We assume that the relative price of commitment and food goods is fixed at one. Then the dynamic budget constraint of the consumer is

$$dw_t = [(r + \alpha_t \pi) w_t - c_t] dt + \alpha_t w_t \sigma dz. \quad (3)$$

Intuitively, the total mean return on the wealth portfolio is $r + \alpha_t \pi$ and a wealth share of α_t in risky assets gives rise to a standard deviation of $\alpha_t \sigma$ in the growth rate of wealth.

The consumer thus maximizes (2) subject to the reset rule for consumption commitments and the budget constraint (3). Let Δ denote the time elapsed since the consumer last adjusted her consumption commitment; clearly $0 \leq \Delta < T$. Note that between reset dates, $\Delta = \Delta(t)$ depends linearly on calendar time, that is, $d\Delta = dt$.

Theorem 1 *The optimal consumption and investment rule between two reset dates is characterized by the policies*

$$f_t = \psi_\Delta \cdot (w_t - \eta_\Delta x_t) \quad (4)$$

$$\alpha_t w_t = \frac{\pi}{\gamma \sigma^2} \cdot (w_t - \eta_\Delta x_t) \quad (5)$$

where $\psi_\Delta = \psi_{\Delta(t)}$ and $\eta_\Delta = \eta_{\Delta(t)}$ are deterministic functions of the time elapsed since the last reset date, given in the Appendix.

At a reset date, the consumer chooses

$$x_t = \chi \cdot w_t \quad (6)$$

where χ depends on the underlying parameters of the model.

The consumption dynamics between two reset dates is given by

$$\frac{dc_t}{c_t} = \left\{ r - \psi_{\Delta(t)} + \frac{\pi^2}{\gamma \sigma^2} + \frac{\psi'_\Delta}{\psi_\Delta} \right\} \frac{f_t}{c_t} dt + \frac{\pi}{\gamma \sigma} \frac{f_t}{c_t} dz. \quad (7)$$

Proof. See the Appendix. ■

To understand this result, note that between reset dates, the model has two state variables, wealth (w) and the current level of consumption commitments (x). The optimal consumption and investment rules turn out to be linear functions of these state variables. More importantly, the expression $w_t - \eta_\Delta x_t$ governs the optimal policy of both consumption and investment. This quantity can be interpreted as net wealth since the value of outstanding future commitments is exactly $\eta_\Delta x_t$. Intuitively, since commitment consumption can only be reset at particular dates, the consumer has to be certain that she has enough funds to finance her outstanding commitments until the next reset date. She allocates an amount corresponding to outstanding commitments in bonds, and uses that money exclusively to finance future commitments. She then decides on how to invest the rest of her wealth in stocks and bonds and how much to spend on food consumption.

According to (5), as η_Δ is always non-negative, an individual with a high level of commitment consumption invests relatively less in risky assets.⁶ Furthermore, the effect of commitments on stockholding is more pronounced when individuals are less risk averse (i.e., when γ is lower). To see why, observe that one dollar less in commitments implies one dollar more net wealth. Clearly, a less risk averse individual invests a higher share of her marginal dollar into stocks; therefore the portfolio of a less risk averse individual is more sensitive to the level of commitments. Less risk averse individuals also hold more stocks on average. In contrast, the optimal portfolio of a more risk averse individual has a lower slope with respect to commitments, and a lower intercept reflecting lower average stockholdings.

We remark that (4) and (5) continue to be the optimal policies even if the current level of commitment was allocated exogenously to the consumer, as long as she is free to choose commitments at the next reset date. This is important because our empirical strategy is to test whether portfolio allocation varies as the model predicts when commitments are assigned exogenously.

We now turn to the implied dynamics of consumption. First, observe that at a reset date, the consumer has a single state variable, her current level of wealth. The optimal level of consumption commitment undertaken will be proportional to wealth because at

⁶Note that this effect is purely a consequence of the level of commitments, as opposed to fluctuations in the relative price of commitments (such as house price fluctuations).

this stage the value function is homogenous of degree $1 - \gamma$ in wealth.

Between two adjustment dates, the instantaneous standard deviation of log consumption growth between two adjustment dates is given by

$$\frac{\pi}{\gamma\sigma} \cdot \frac{f_t}{c_t}. \quad (8)$$

The first term in this product is the standard term that arises in a single good economy (e.g., in the Merton consumption problem). The second term is the ratio of food to total consumption, which is less than one when the agent has commitments. This term has a clear intuitive meaning. Given that a share of total consumption is not adjustable in the short run, the volatility of log consumption growth is smaller in this model than in a model with a single, freely adjustable consumption good. Moreover, the effect is proportional to the share of consumption that is not committed. Imagine that half of consumption is committed: since all shocks affect only the non-commitment part of consumption, along the optimal path that part will have a standard deviation of $\pi/\gamma\sigma$. Consequently, log total consumption growth will have a standard deviation that is only 1/2 times that.

Put differently, commitments effectively increase the coefficient of relative risk aversion, γ , by a factor of c/f . This is apparent in (8), where c/f multiplies γ , reducing consumption volatility. This suggests that the model could help explain the equity premium puzzle. We revisit this point in Section 3.

2.2 Aggregation

One shortcoming of the household model is that consumption jumps discontinuously on adjustment dates. These jumps make it difficult to evaluate the model's implications for aggregate consumption smoothness over long horizons. To address this issue, we now consider a heterogeneous population of individuals and show that in such a population the jumps are averaged out and aggregate consumption becomes smooth.

To explore aggregate dynamics, consider an economy populated by a continuum of agents who have preferences over consumption commitments. Assume that the individual reset dates are uniformly distributed across the population, and normalize the total mass of agents to one. The agents with different adjustment dates may also differ in their wealth levels; however their preference parameters and adjustment horizons are

the same. Our assumptions imply that during any time interval of length T everyone resets exactly once, and during a time interval $[s, t]$, exactly $(t - s)/T$ agents adjust. Let capital variables denote aggregate quantities, so that X_t , F_t and C_t stand for aggregate commitment, food and total consumption, where for example

$$X_t = \int_{\Delta=0}^T x_t(\Delta) d\Delta$$

if we index the individuals by the time Δ elapsed since they last adjusted. We denote the net wealth of an individual Δ at time t by $w_t^{\text{net}}(\Delta)$, and the aggregate of this quantity across the population by W_t^{net} . Finally, we define τ to be the value of t modulo T . This notation allows us to state the following result.

Proposition 1 *For an arbitrary initial net wealth distribution, individual net wealth $w_t^{\text{net}}(\Delta)$ can be written as*

$$w_t^{\text{net}}(\Delta) = g(\Delta, \tau) \cdot W_t^{\text{net}} \quad (9)$$

where $g(\Delta, \tau)$ is a deterministic function that depends on the initial net wealth distribution. Moreover at any point in time $t > T$, aggregate commitment and non-commitment consumption can be expressed as

$$X_t = \frac{1}{T} \int_0^T a(\tau - u) \cdot W_{t-u}^{\text{net}} du \quad (10)$$

and

$$F_t = b(\tau) \cdot W_t^{\text{net}} \quad (11)$$

where $a(\cdot)$ and $b(\cdot)$ are appropriate functions that depend on the initial distribution of net wealth in the population.

The first equation says that individual net wealth is proportional to aggregate net wealth where the factor of proportion depends only on the type of the individual (i.e., Δ) and the value of time modulo T (i.e., τ). Individual and aggregate net wealth are proportional because of power utility, and the factor of proportion naturally depends on the type of the individual. Variable τ affects this factor because the aggregate economy exhibits a T -long cycle, and τ is the measure of how far we are in that cycle. The presence of a T -long cycle is not surprising given that T is every household's adjustment horizon. The remaining results of the proposition follow easily. To show that (10) holds, note

that aggregate consumption commitment is the sum of individual commitment levels. The commitment level of an agent who last adjusted on some past date $t - u$ will be proportional to her wealth level on that date by (6). By equation (9), this individual wealth level on date $t - u$ will be proportional to aggregate net wealth in the population on date $t - u$, with the factor of proportion depending on $\tau - u$, the measure of how far the economy was in the aggregate cycle on date $t - u$. Thus aggregate commitment will be a weighted sum of past levels of aggregate net wealth. The intuition for the final equation is similar.

It is possible to find an initial wealth distribution such that the aggregate T -long cycle washes out. This is the content of the next result.

Proposition 2 *There exists an initial net wealth distribution, the balanced wealth distribution, such that if the economy starts from that distribution at date zero, then for all $t > T$ aggregate commitment consumption can be written as*

$$X_t = \frac{a}{T} \cdot \int_0^T W_{t-u}^{net} du \quad (12)$$

and aggregate food consumption can be written as

$$F_t = b \cdot W_t^{net} \quad (13)$$

with positive constants a and b .

In the special case when the economy is started from the balanced wealth distribution, we have no aggregate cycle, and the coefficients a and b do not depend on τ any more. In the appendix we explicitly construct the balanced net wealth distribution. The intuition behind that construction is the following. We know that with an arbitrary initial distribution, the economy exhibits cycles. However, it is possible to add up a continuum of aggregate cycles, all shifted in time a bit relative to each other. The resulting overall dynamics will then exhibit no cycles. Basically, we aggregate over the aggregate cycles; equivalently, we add up a unit mass of a unit mass of consumers. Of course, the total measure of consumers continues to be one. This procedure results in the balanced wealth distribution of the proposition.

Using the propositions we can prove the following aggregation theorem.

Theorem 2 *The aggregate dynamics of consumption are the optimal policy of a representative consumer with external habit formation utility function*

$$E \int_0^\infty e^{-\rho t} \frac{(C_t - X_t)^{1-\gamma}}{1-\gamma} dt$$

where habit X_t evolves according to the path of committed consumption, given by equation (10) in the general case, or

$$X_t = \frac{a}{T} \cdot \int_0^T W_{t-u}^{net} du. \quad (14)$$

if the economy starts from the balanced net wealth distribution.

Proof. Because the only source of uncertainty in this model is the risky asset, the two traded assets in the economy span the whole asset space. It follows that there exists a unique stochastic discount factor (state price density). Call it $m_{s,t}(\omega)$, then for any trade asset return $R_{s,t}$ between dates s and t , we have

$$1 = E_s m_{s,t} \cdot R_{s,t}.$$

Since consumers are free to adjust their consumption on the food margin, for any individual, the (discounted) ratio of marginal utilities over food consumption between dates s and t has to equal $m_{s,t}$

$$e^{-\rho(t-s)} \frac{f_t^{-\gamma}}{f_s^{-\gamma}} = m_{s,t}$$

hence we have

$$f_t = f_s \cdot \left(e^{\rho(t-s)} m_{s,t} \right)^{-\frac{1}{\gamma}}.$$

Aggregating this equation over the population yields

$$F_t = F_s \cdot \left(e^{\rho(t-s)} m_{s,t} \right)^{-\frac{1}{\gamma}}$$

so that

$$e^{-\rho(t-s)} \frac{(C_t - X_t)^{-\gamma}}{(C_s - X_s)^{-\gamma}} = m_{s,t}$$

which implies

$$(C_s - X_s)^{-\gamma} = e^{-\rho(t-s)} \cdot E_s (C_t - X_t)^{-\gamma} R_{s,t}.$$

But this is the Euler equation for optimality for a representative consumer with habit formation utility, where habit is X . The evolution of habit is governed by a moving

average of past net wealth levels, as shown by the previous proposition. It follows that the aggregate consumption dynamics satisfies the Euler equation for the representative consumer habit formation model. ■

Grossman and Shiller (1982) prove the existence of a representative consumer in an asset pricing context with one consumption good, if individual consumption dynamics follow Ito-processes. Although their result cannot be applied here because individual consumption has jumps and there are two goods, the intuition is similar: individual marginal utilities can be aggregated, and that defines the marginal utility of the representative consumer. The second contribution of the theorem is that when individuals care about commitment consumption, the representative agent will have a habit-formation utility function. More generally, as long as individual marginal utility depends only on surplus consumption over commitments (i.e., $c - x$), in other words, if utility is separable in commitment and non-commitment consumption, the above aggregation argument goes through. Under these circumstances it can be shown that the representative consumer's utility will only depend on aggregate surplus consumption $C - X$. In this sense, the theorem is not a consequence of the exact functional form specification.

The intuitive connection between commitments and habit is straightforward: high commitments, like a high level of habit, make the individual more risk averse. The lumpy adjustment dynamics that are present at the individual level due to commitments are “smoothed out” in the aggregate.⁷ This smoothing effect of aggregation is important because aggregate consumption does not exhibit discontinuous jumps in the data.

While commitments lead to a slow moving habit in the aggregate, the dynamic of habit in this model is somewhat different from existing habit formation specifications, such as Constantinides (1990) or Campbell and Cochrane (1999). In those models, habit is a slow moving time average of past levels of aggregate consumption. Here habit is a slow moving time average of a different aggregate variable, net wealth. But even in our model, habit can be rewritten as an average of past levels of aggregate consumption. This is the content of the next result.

Proposition 3 *In the balanced wealth distribution case, as long as the underlying parameters of the model are such that $a < b$, aggregate commitment consumption can be*

⁷The literature on aggregating agents with state-dependent adjustment rules also finds a similar smoothing effect; see for example Bertola and Caballero (1990).

written as

$$X_t = \kappa(t)W_0^{net} + \int_0^t \zeta(u)C_{t-u}du \quad (15)$$

where

- As u goes to infinity, both $\kappa(u)$ and $\zeta(u)$ go to zero at a geometric rate,
- For $0 \leq u < T$, $\zeta(u) = \frac{a}{Tb} \cdot \exp(-\frac{a}{Tb}u)$,
- The function $\zeta(u)$ is bounded and switches sign at least once on any interval longer than T .

Proof. See the Appendix. ■

There are three linear equations linking the variables C , F , X and W^{net} : an accounting identity, the habit rule (12), and the consumption rule (13). These allow us to express any of the four variables as a linear function of current and lagged values of any other. This is the idea of the proof. The condition $a < b$ in the proposition is used to ensure that the coefficients $\zeta(\cdot)$ go to zero asymptotically.⁸ This condition roughly corresponds to the case where commitment consumption is typically less than food consumption, as shown by equations (12) and (13). Since the parameters a and b are derived from the underlying parameters of the model, we need to demonstrate that the proposition has content, that is, there are underlying parameters for which $a < b$ holds. It is easy to see that when the utility weight of commitment consumption, $\mu = 0$, there is no commitment consumption, therefore $a = 0$. By continuity, for small enough μ the consumer cares relatively little about commitments, hence $a < b$ continues to hold.

However, the condition $a < b$ is restrictive because the data suggests that commitments may constitute a higher share of total consumption than non-commitments. In the next subsection we consider a different version of the model, where the equivalent of Proposition 3 holds without such a restriction. Nevertheless, the result of the proposition is interesting in that it provides a close link to the Constantinides model of habit formation, where habit is a geometrically weighted average of past consumption levels. Here the geometric decay is explicitly present for the near past and asymptotically for the distant past. However, in between the weight function $\zeta(u)$ fluctuates and periodically becomes negative. Thus habit in the commitments model corresponds more closely

⁸We believe that the proposition is true under more general conditions.

to the average past *growth rate* of total consumption: a weighted average of consumption in the past is subtracted from a weighted average of nearer-term past consumption. Because of the geometric decay, consumption levels in the distant past matter little.

2.3 Model and Aggregation with a Stochastic Adjustment Rule

In this section, we continue to assume that consumption commitments are adjusted according to a time dependent rule, but allow that rule to be stochastic. At the cost of some additional assumptions, this new model yields a more useful aggregation result, which makes the connection with habit persistence stronger.

In particular, suppose that during a short interval dt , the consumer can adjust her level of commitment consumption with probability λdt . With remaining probability, she cannot adjust. One interpretation of the stochastic adjustment rule can be random shocks to the household that make moving out of commitments costless. An example of such a shock can be a new job offer with relocation benefits. This specification is similar to the Calvo-pricing rule common in macroeconomics.

The preferences and technology of the consumer are otherwise the same as before, as summarized in (2). We assume that there is a continuum of consumers whose adjustment dates are independently distributed, and that markets are complete. Note that market completeness held by design in the previous model, because the only source of uncertainty was the risky asset. In contrast, here consumers face additional uncertainty in that adjustment dates arrive randomly. Thus the market completeness assumption is potentially quite restrictive. It implies that consumers can insure all the risk that is coming from the uncertainty related to *when exactly* they can adjust their commitment consumption. The risky and the riskless asset introduced in the previous section continue to be traded, and are in perfectly elastic supply. All other insurance is in zero net supply, but because consumers are heterogenous, they will form individual-specific portfolios of those assets. The complete markets assumption is a modelling tool; it is useful because it allows for a simple aggregation of consumers.

Due to power utility and the stationarity of the problem, on an adjustment date, a consumer will find it optimal to choose a level of commitment consumption that is a constant proportion of her current wealth (the only state variable when adjusting). Call this proportion a' .

This yields the following aggregation result.

Proposition 4 *The aggregate consumption dynamic of the economy is the same as the optimal consumption of a representative consumer economy with habit formation utility function*

$$\max E \int_0^\infty e^{-\rho t} \frac{(C_t - X_t)^{1-\gamma}}{1-\gamma} dt$$

where habit X_t evolves according to the path of aggregate commitment consumption

$$X_t = e^{-\lambda t} X_0 + a' \lambda \cdot \int_0^t e^{-\lambda u} W_{t-u} du. \quad (16)$$

Proof. Aggregation follows the same way as in Theorem 2. Because markets are complete, there is a unique stochastic discount factor, which equals the ratio of marginal utilities over food consumption for each individual consumer; and these ratios can be aggregated the same way as earlier.

The only part that needs to be proved is that aggregate consumption indeed evolves according to (16). Note that at time t , the number of people who last adjusted during the short time interval $[s, s + dt]$ is $\lambda dt \cdot e^{-\lambda(t-s)}$. These people were representative of the population at time s , because the adjustment shock is independent of past history. Thus their average wealth at time s was exactly W_s , and their average commitment consumption $a' \cdot W_s$. But this continues to be their commitment consumption at time t too, because they have not adjusted since s . In addition there are $e^{-\lambda t}$ people who have not adjusted since date zero; because they were representative of the whole population at date zero, their contribution to aggregate commitments is $e^{-\lambda t} X_0$. This explains the formula, and the proof is complete. ■

Under some conditions, the aggregate representative consumer model can be solved analytically. Define η' to be the smaller real root of

$$0 = \eta'^2 a' \lambda - \eta'(\lambda + r) + 1$$

if it exists, and let

$$\psi' = \frac{\gamma/(1-\gamma)}{\frac{\rho}{1-\gamma} - \frac{1}{2} \frac{\pi^2}{\gamma \sigma^2} - r + a' \eta' \lambda} \quad (17)$$

and

$$A = r - a' \eta' \lambda + \frac{\pi^2}{\gamma \sigma^2} - \psi'^{-1} + \lambda a' \psi'. \quad (18)$$

Then

Theorem 3 *If η' exists, and*

$$\frac{\rho}{\gamma - 1} + \frac{1}{2} \frac{\pi^2}{\gamma \sigma^2} + r - a' \eta' \lambda > 0 \quad (19)$$

holds, then the optimal consumption and investment policy are

$$C_t = X_t + \psi'^{-1}(W_t - \eta X_t) \quad (20)$$

and

$$\alpha_t = \frac{\pi}{\gamma \sigma^2} \left(1 - \eta' \frac{X_t}{W_t} \right) \quad (21)$$

and the aggregate consumption dynamic can be written as

$$\frac{dC_t}{C_t} = \left\{ A - \frac{X_t}{C_t} [\lambda - a' \eta' \lambda + A] \right\} dt + \frac{\pi}{\gamma \sigma} \frac{C_t - X_t}{C_t} dz. \quad (22)$$

Moreover, if

$$(\gamma - 1) \frac{\pi^2}{\gamma^2 \sigma^2} + \lambda + r - 2\eta' \lambda a' - \psi'^{-1} > 0 \quad (23)$$

holds, then the share of commitment consumption $Z_t = X_t/C_t$ has an invariant distribution with density $p_Z(Z)$ given in the Appendix.

Proof. See the Appendix. ■

The intuition behind the result is similar to that of Theorem 1. Because of habit, the individual holds a certain amount of wealth $\eta' X$ in the riskless asset, to guarantee that she is able to finance habit consumption in all circumstances. Optimal policies then become a function of the state variable $W - \eta' X$, and the linearity of the policy rules is a consequence of power utility. The parameter ψ' , a measure of the marginal propensity to consume, is identical to its equivalent in the single-good Merton consumption problem, except for the additional term $a' \eta' \lambda$. That term decreases the marginal propensity to consume due to the presence of commitments. Once the optimal policy has been identified, we follow Constantinides (1990) in deriving the invariant distribution of the habit share.

Using the results of the theorem, we can calculate the unconditional mean and variance of aggregate consumption growth to be

$$\frac{E(dC_t/C_t)}{dt} = A - [\lambda - a' \eta' \lambda + A] \int_0^1 Z \cdot p_Z(Z) dZ \quad (24)$$

and

$$\frac{\text{var}(dC_t/C_t)}{dt} = \left(\frac{\pi}{\gamma\sigma}\right)^2 \int_0^1 (1-Z)^2 \cdot p_Z(Z) dZ. \quad (25)$$

As in the previous model, under some circumstances the habit dynamics can be rewritten to depend only on past levels of aggregate consumption.

Proposition 5 *If η' exists and (19) holds, and further*

$$a'\eta' - a'\psi' - 1 < 0 \quad (26)$$

then aggregate commitment consumption can be expressed as

$$X_t = e^{-dt} X_0 + D \cdot \int_0^\infty e^{-du} C_{t-u} du \quad (27)$$

where $D, d > 0$ are constants that depend on the underlying parameters of the model.

Proof. See the Appendix. ■

Using the analytic solution of the model, total consumption can be expressed as a function of current and past levels of wealth. The proof inverts that formula and uses the habit rule (16).

This last proposition underlines the strong connection between the commitments model and habit formation. Under some conditions, the preferences and the habit dynamics in our model are *exactly the same* as those in Constantinides (1990). The only difference is that Constantinides has internal habit formation while aggregating the commitments model yields external habit.

A few words on modelling choices: The first adjustment specification is convenient in that the individual consumer problem can be solved analytically, and no additional assumptions are required about the existence of insurance markets. The cost of this model is that in the aggregate, we need to worry about the initial wealth distribution and that the habit adjustment rule is somewhat unusual. The convenience of the second adjustment specification is twofold. First, the aggregate habit rule is simple and easily linked to the literature. Second, the aggregate model can be solved analytically, lending itself to easier calibration. These advantages come at the cost of assuming a very rich set of markets. The fact that both models yield similar suggests that the restrictive assumptions are only necessary for technical reasons. The two models thus complement each other in illustrating the key intuition of the paper.

3 Calibration and the Equity Premium Puzzle

We have shown that in an economy populated by a continuum of agents with consumption commitments, aggregate dynamics coincide with those a representative consumer economy with habit formation. This result has some interesting implications.

First, our model provides micro-foundations for habit persistence. We arrive at habit formation in the aggregate by assuming standard preferences over two goods at the individual level and frictions in adjusting commitment consumption in the short run. Costly adjustment thus suffices to provide micro-foundations for habit persistence while retaining the standard framework of neoclassical economics.

Second, our results suggest that consumption commitments can be viewed as part of habit. Certain goods have physical commitment features, such as adjustment costs. Other goods may have mental or psychological adjustment costs, and yet another part of consumption may be pure habit. Hence, the distinction between habit and consumption commitments is somewhat blurred: total habit could be a sum of physical commitments, mental commitments and pure habit.⁹

These observations suggest that a “blind” calibration of the commitments model may be useful in evaluating its performance. By “blind” calibration we mean that we do not aim to pin down the level of X from data, but rather construct a time series X such that the implied consumption and return dynamics match aggregate data.

The objective of the calibration exercise is to match the first and second moments of consumption, stock returns and bond returns. We take the data on these returns from Constantinides (1990), who relies on the estimates of Mehra and Prescott (1985). Constantinides uses a riskfree rate of $r = .01$ per year and an equity premium of $\pi = .06$, with an annual standard deviation of stock returns $\sigma = .165$. The mean and variance of annual consumption growth are $E(dC/C)/dt = .0183$ and $\text{var}(dC/C)/dt = (.0357)^2$. Campbell (2002) reports consumption moments that are roughly in line with these data for the 1889-1998 sample.¹⁰

⁹Reis (2003) argues that when there are information processing or planning costs to adjusting consumption, the optimal policy is a time dependent rule. Thus mental commitments may correspond to goods with higher planning costs.

¹⁰Looking only at postwar data, Campbell reports a much lower variance of consumption growth of $(.01073)^2$. Part of the difference may be due to the fact that his measure only includes nondurable consumption. Given that the focus of this paper is commitment consumption, we choose statistics for

We calibrate the model with the stochastic adjustment rule because it admits an analytic solution in the aggregate. We need to choose two standard preference parameters, the risk aversion parameter γ and the rate of time preference ρ . Note that γ here measures the relative risk aversion of an unconstrained consumer, who is free to adjust both kinds of consumption. There are also two commitment related parameters to be chosen. Parameter λ measures the horizon of commitment consumption: $1/\lambda$ is the mean waiting time between two adjustments. Parameter a' measures the preference for consumption commitments, in the sense that at every adjustment date the household sets commitments to equal a share a' of current wealth.¹¹

In this model, equity returns are exogenous and aggregate consumption is endogenous. Thus, we fix the parameters r , π and σ at the levels described above, and for each set of values $(\lambda, a', \gamma, \rho)$ calculate the implied mean and variance of consumption growth. We check that η' exists and conditions (19) and (23) are met for all sets of parameters. Since (26) also holds for all parameters we consider, aggregate habit can be written as a time average of past consumption levels.

We first discuss a set of baseline results for the case in which the consumption variances estimated in the data coincide with the instantaneous variance in the model. We then turn to results of a simulation in which we incorporate the fact that consumption data is only observed at a low frequency and therefore does not directly relate to the instantaneous variance implied by the model.

3.1 Baseline Results

Table 2 reports the implied mean and standard deviation of consumption growth for a set of parameter vectors. The mean share of commitment consumption in total consumption is also reported. In the first four columns, the mean and standard deviation of consumption growth are calculated using formulas (24) and (25), with the integration performed numerically. In the first two columns, the unconstrained individual level risk aversion parameter is set at $\gamma = 2$, while we vary λ . Under these circumstances, with appropriate choice of a' and ρ , we are able to match both the .0357 standard deviation

the broader consumption measure. However, we note that a lower consumption volatility can be matched with a sufficiently high commitment share in total consumption.

¹¹Note that a' is a function of the underlying parameters of the model, in particular it is increasing in μ . Since we do not have a closed form expression for a' , we use it directly in the calibration.

and the .0183 mean of log consumption growth. The results are not particularly sensitive to the choice of λ , the expected time between adjustments. The value of $\lambda = .33$ in the first column corresponds to an adjustment on average every three years; whereas a value of $\lambda = 1$ in the second column means that households update on average every year. We are able to achieve almost identical results for intermediate values of λ as well as a λ as high as 2, meaning adjustments on average every six months. If we interpret the average time between adjustments as a measure of adjustment costs, the table suggests that even small adjustment costs can make aggregate consumption fairly smooth in the short run.¹²

Note that in the first two specifications, the rate of time preference, ρ , is chosen to be quite small (around .007). The reason is the risk-free rate puzzle of Weil (1989): in the standard consumption based asset pricing model, high risk aversion implies a low intertemporal elasticity of substitution. Thus, consumption does not grow fast along the optimal path. Although our model relaxes the direct connection between the risk aversion parameter and the intertemporal elasticity of substitution, this is insufficient to completely account for the riskfree rate. In order to make consumption grow faster, we require patient consumers. This motivates our choice of low levels of ρ throughout the calibration.

In order to match the low standard deviation of aggregate consumption with small risk aversion, we need a high share of commitments in total consumption. As the last row in the table shows, when $\gamma = 2$, approximately 80% of total consumption must be committed. According to the summary statistics for expenditure in Table 1, physical commitment goods constitute less than 80% of the typical household's consumption. However, in line with our interpretation of commitment consumption as part of habit, the rest of the 80% may be constituted of goods that involve mental or psychological adjustment costs.¹³

The model is also capable of matching the equity premium with a lower share of commitments in total consumption if we assume that unconstrained risk aversion is

¹²This is in line with the findings of Marshall and Parekh (1999), who report calibration results with small adjustment costs.

¹³The 80% share of habit in total consumption is the roughly same as what Constantinides (1990) finds in his calibration with $\gamma = 2.2$. This is not surprising, since the only difference between his model and the one calibrated here is that ours has external habit formation.

higher. The empirical results presented in Section 4 provide some evidence that the unconstrained individual risk aversion parameter is in the ballpark of $\gamma = 4.17$. For this reason, columns 3 and 4 of Table 2 report calibration results with $\gamma = 4$ and $\gamma = 5$. With these parameters, the equity premium puzzle can be resolved with a consumption share of commitments as low as 60% or 50%. Remarkably, these numbers are similar to the share of physical commitments in total consumption, when commitment goods are defined as housing, transportation excluding gas and maintenance, utilities, health insurance and education (see Table 1). All of these goods involve either transaction costs or contracts that make their adjustment costly and infrequent.

However, even at the low rate of time preference $\rho = .005$, we are not able to match mean consumption growth perfectly: we only get annual growth of 1.2-1.4% as opposed to the true value of about 1.83%. Again, the reason for this is the riskfree rate puzzle. A higher γ implies more consumption smoothing, making it more difficult to match the observed mean consumption growth in the data even with patient consumers. We discuss some alternative explanations of the riskfree rate puzzle below.

3.2 Calibration to Annual Data

One potential problem with the results discussed thus far is that the volatility of consumption growth calculated using (25) may not correspond to the volatility measured in the data. There are two sources of differences. First, (25) expresses the instantaneous variance, which may differ from the variance of quarterly or annual consumption growth unless the consumption process follows a random walk. Since consumption and wealth are co-integrated, we expect the variance of consumption to rise to the variance of wealth over longer horizons in the model (i.e., the random walk condition may fail to hold). In this situation, the consumption volatility reported in columns 1-4 may understate consumption volatility measured over a quarterly or annual horizon. Second, (25) is an unconditional variance, which takes into account the unconditional uncertainty in the surplus consumption ratio $1 - Z$. However, the surplus consumption process is mean reverting. Therefore we expect its variance to fall over longer horizons. This effect would bias the reported consumption volatilities in columns 1-4 upwards.

For these reasons, in the final two specifications of Table 2 we simulate the aggregate consumption dynamics as given by equation (7), with a sampling frequency of one day.

We then aggregate over time to get *annual* consumption, and calculate the implied mean and variance of (annual log) consumption growth. For each of the final two columns, the representative consumer economy is simulated fifty times for a 100 year horizon. The measured consumption means and variances are then averaged across all runs.

The results of columns 3 and 4 largely carry over to our new measure of annual consumption, if we decrease the parameter λ to 0.2. This corresponds to an adjustment of commitments on average every five years. A lower λ is required because less frequent adjustments make aggregate consumption react more slowly to shocks. This mitigates the problem of increasing variance over longer horizons, while allowing for the effect of mean reversion in surplus consumption.

Of course, at even longer horizons, the volatility of consumption growth will rise to match the volatility of wealth, which in particular implies positive autocorrelation in consumption growth at lower frequencies. As emphasized for example in Campbell (2002), there is no evidence for such autocorrelation in the data. We can present two arguments in defense of our results. First, in our model, although the volatility of wealth over longer horizons is high, it is far from being as high as the volatility of the stock market, because consumers hold a large share of safe assets in their portfolio. For example, in the simulation corresponding to column 6 in the table, the (annualized) ten year standard deviation of log wealth growth is about .055. This is about fifty percent higher than the measured standard deviation of consumption, .0357, but still much lower than the standard deviation of the stock market, .165. The reason is that the representative consumer holds more than two-thirds of her wealth in safe assets.

Second, the increase in consumption volatility over longer horizons is a common property of asset pricing models with exogenous i.i.d. returns (including Constantinides, 1990 and Gabaix and Laibson, 2001). One way the literature has dealt with this problem is by introducing mean reversion in returns, for instance in the habit formation model of Campbell and Cochrane (1999). Thus, the increase in consumption variance is not a consequence of our commitments story per se, but rather a consequence of our modelling choice of exogenous, i.i.d. returns. It is likely that the commitments model with mean reversion in asset returns would help resolving asset pricing puzzles while avoiding increasing consumption variance over longer horizons. However, in this paper we restrict attention to the analytically tractable i.i.d. specification.

3.3 The riskfree rate puzzle

We found that the commitments model is able to resolve both the equity premium and riskfree rate puzzles only with a high share of commitment consumption. With a lower commitment share of 50-60%, the model only partially resolves the riskfree rate puzzle.

It is helpful to explore some omissions from our highly parameterized model that could account for the riskfree rate puzzle. First, note that Campbell reports an estimated riskless interest rate of .0202 for the time period 1881-1998, which is approximately double the rate we use in the calibration. In other words, the true riskfree rate is measured with noise in the data. A riskfree rate higher than .01 would obviously increase the performance of our model in matching all moments of consumption, even with a low share of commitments. Second, there may be other frictions in the economy, not modelled in our paper, that generate high consumption growth. For instance, borrowing constraints increase precautionary savings and lead to higher consumption growth (see e.g., Gollier, 2001, pages 269-283). Introducing borrowing constraints into the commitments model may increase implied consumption growth, providing a better fit to the data. The role of borrowing constraints in mitigating the riskfree rate puzzle has also been emphasized in a recent paper by Constantinides, Donaldson and Mehra (2002). Their simulations indicate that the mean bond return roughly doubles when borrowing constraints are relaxed, a large effect. Finally, relaxing the link between the EIS and risk aversion at the individual level could directly allow us to explain the riskfree rate puzzle.

In summary, the calibration exercises indicate that when appropriately parameterized, the commitments model can generate risk premia and consumption patterns that are consistent with aggregate data. Of course, the empirical importance of commitments as an explanation for these puzzles depends on the extent to which they affect households' portfolio allocation decisions in practice. The remainder of the paper addresses this question.

4 Empirical Evidence

The key testable prediction of our model is that an exogenous increase in an agent's level of commitment consumption – e.g., housing – causes him to choose to hold fewer risky assets (stocks) relative to safe assets (bonds). When x is determined exogenously,

(5) continues to hold and implies the following estimating equation for our empirical analysis:

$$\text{portfolio risk} = \alpha + \beta \times \text{commit} + \theta \times \text{wealth} + \text{controls} + \varepsilon. \quad (28)$$

Here portfolio risk denotes a measure of the riskiness of a portfolio (e.g., value of stocks owned or share of stocks in the portfolio), commit measures the agent’s level of commitment consumption (e.g., size of house or mortgage), and wealth is a measure of the lifetime wealth of the agent. In the estimation procedure, we will consider both simple wealth measures and more flexible parametrizations of liquid and illiquid wealth. In addition, we control flexibly for other factors that could be correlated with risk aversion and portfolio choice such as age, education, occupation, etc. The noise term ε captures unobserved individual characteristics that may affect portfolio choice, such as heterogeneity in risk aversion or background income risk, as well as measurement error in wealth. The theory predicts that $\beta < 0$, and the magnitude of β quantifies the importance of commitments for portfolio allocation.

Our empirical representation of commitment is the mortgage debt outstanding on the agent’s house. Note that higher mortgage debt is equivalent to owning a more expensive home when home equity and the interest rate are held fixed, since home equity is defined as property value minus outstanding mortgage. As our benchmark specifications hold home equity fixed and do not exploit variation in interest rates, unless otherwise noted, we will use the terms “house size” and “mortgage” interchangeably.¹⁴

It should be noted that identifying the exact empirical counterpart of the commit variable in (28) is somewhat difficult. In a purely time-dependent framework, the natural measure of commitment is the total expected user cost of the commitment good until the next reset date. If an agent consumes his entire house before moving, home value would be an appropriate measure of commitment. If he were to sell before this point, commitment would rise by less than \$1 for every \$1 increase in property value. Hence, the correct measure of commitment is some proportion of housing value. Rather than attempting to calculate this proportion, we observe that our estimates using variation in mortgage debt – which is equivalent to using variation in property value, as noted above – will understate the effect of the commit variable on portfolio choice when an agent can

¹⁴When interest rates are held fixed, households that have larger mortgages are not poorer, since they also have a more expensive housing asset in their portfolios.

resell his house.

The main difficulty in estimating (28) is that our commitment measure, mortgage debt, is itself endogenously chosen by the household. It may therefore be correlated with both unobserved household characteristics and measurement error in wealth, violating the orthogonality condition necessary for OLS. To deal with this endogeneity problem, we seek exogenous variation in our commitment measure. We generate such variation in mortgage debt using an individual’s age at marriage as an instrument for his house size. Recognizing that age at marriage may be directly related to portfolio choice, we test the hypothesis that marriage age and ε are uncorrelated by using “control” groups for whom our first-stage relation breaks down. In addition, we address the concern that age at marriage is correlated with measurement error in wealth directly by examining its correlation with the consumption of cars, which are less durable than housing. As discussed in detail below, our results support the claim that the variation in house size that we exploit is truly exogenous and its effects on portfolio choice are causal.

An additional complication in using housing as a measure of commitment is that the house is itself an asset in the agent’s portfolio. Consequently, leaving aside commitment effects, owning a larger home can directly affect an agent’s optimal portfolio. We address this issue by showing empirically that houses appear to be treated as riskless assets in portfolio choice. If houses are indeed riskless, a \$1 increase in mortgage causes a \$1 increase in riskless debt coupled with a \$1 increase in riskless assets and leaves the risk properties of the overall portfolio unchanged. Hence, in this case, every dollar shifted from stocks to bonds when an agent owns a larger home reflects a \$1 shift in the overall portfolio towards safer assets.

The discussion of estimating β is organized as follows. The data and key sample selection procedures are described in the next subsection. We then discuss the age at marriage instrument for house size, test its validity by estimating first-stage and reduced-form equations for control groups, and address concerns about measurement error in wealth using data on vehicle consumption. In the third subsection, we present two-stage least squares estimates of (28). The fourth subsection uses data on cross-city differences in home price risk to show that the shift to safer assets when one owns a bigger house is due to the commitment effect rather than greater exposure to risk in the housing market. Finally, we show how the risk aversion parameter (γ) used in the calibration exercises

can be imputed from our empirical estimates of β .

4.1 Data and Sample Selection

The data used in this study are from the 1990-1996 panels of the Survey of Income and Program Participation. The SIPP collects asset and liabilities information from a sample of approximately 30,000 households at least once in each panel. Other data about the demographic and economic characteristics of each household are also collected. The main advantages of the SIPP relative to other commonly used micro datasets on financial characteristics such as the SCF and PSID are its large sample size and the ability to identify a household's geographic location up to the metropolitan statistical area (MSA) level. These geographic identifiers turn out to be useful in distinguishing the commitment effect of owning a bigger house from the effect of greater exposure to price risk in the housing market.

We make two exclusions on the original sample of survey respondents to arrive at the core sample used in our instrumental variables analysis. First, since our model abstracts from transaction costs for participation in asset markets, it has no direct implications for households not participating in the stock market. We therefore focus only on stockholders, who constitute approximately 20% of all households. Second, since our IV strategy relies on changes in the outstanding home mortgage due to variation in age at marriage, we are forced to restrict attention to homeowners who are or were married.¹⁵ Approximately 84% of families who own stocks own homes, and of these, 90% are or were married. This leaves 15,297 households in the core sample.

One potential concern is that selective inclusion into the sample might bias estimation results. To deal with this problem, we repeat all the regressions reported below for the full sample that includes non-stockholders and renters.¹⁶ All the qualitative results we find for the core sample of households are preserved in the full sample. The effect of commitment on stockholding becomes smaller – since non-stockholders cannot respond at all – but remains statistically significant. Importantly, *all of* our validity checks for the marriage age instrument also continue to hold on the full sample. These findings

¹⁵We also look at renters to test the exclusion restriction that age at marriage is not directly related to portfolio choice.

¹⁶We continue to select households who are or were married because our IV strategy forces us to do so.

suggest that the selection bias introduced by our sample exclusions is negligible.

Table 3 gives summary statistics for our core sample of ever-married homeowners who are stock market participants in 1990 dollars. Since stockholders tend to be wealthier than the average family in the US, mean net worth is quite high in this sample (\$248,361). On average, approximately 1/3 of this wealth is held in the form of home equity and another 15% is held in illiquid assets such as cars and other real estate. Of the remaining \$123,795 held in liquid wealth, approximately 55% is in stocks, 15% in savings accounts, 10% in bonds, 15% in IRA assets, and 5% in “other” liquid assets. These other assets can be further broken down into checking accounts (5%), US savings bonds (10%), debt owed to the household (35%), and equity in other financial investments (50%).¹⁷ Importantly, 401k wealth is omitted, as the 90-93 SIPP panels do not collect data for this category. While it is possible that our estimates are biased due to this omission, we find that investments in other retirement accounts (IRAs and Keoghs) do not offset the portfolio-shifting effects documented below. To the extent that investment behavior does not vary across different types of retirement accounts, the lack of data on 401k investments is not highly problematic.

The individuals in this sample have higher incomes, more education, and are older than the average individual in the US, but presumably reflect the characteristics of the stockholding population in whom we are most interested. Note, however, that our data understates the true skewness of the distribution of income and wealth because these variables are topcoded to protect confidentiality.¹⁸ Most importantly for our purposes, our key independent variable, home mortgage, is topcoded for 4% of our sample, while our key independent variable, stockholding, is not topcoded. All results reported below are robust to inclusion or exclusion of the topcoded group.

4.2 Estimation Strategy

We begin our analysis by illustrating the endogeneity problem that plagues the OLS estimation of (28). In this regression, as in most others below, we use the total value of stocks as the dependent variable rather than the portfolio share of stocks. Portfolio-share

¹⁷Since it is difficult to classify equity in “other financial investment” as safe or risky, we show that our results are robust to the classification of this asset.

¹⁸For instance, the 1996 panel topcodes primary home property value at \$550,000 – any individual who owns a home that costs more than \$550,000 has home value coded as \$550,000.

regressions effectively weight high-wealth households less than low wealth households, but high-wealth households are those that matter most from an asset pricing perspective. Indeed, it is shown below that portfolio-share regressions understate the aggregate shift in portfolios caused by commitments since this shift is much larger for high-wealth households. Since the error terms of the levels regressions are heteroskedastic, we always report Huber-White sandwich estimates of standard errors.

Strikingly, even when control carefully for liquid and illiquid wealth using 10 piece linear splines, and include other demographics such as a spline for age, industry and occupation dummies, and controls for education and income, the estimated effect of owning a large home (i.e., having a bigger mortgage) on stockholding is *positive*. The coefficient on mortgage, which is marginally significant, implies that a \$1 increase in home value is correlated with holding 8 cents more in stocks, all else held equal. The corresponding regression of bonds on the same right hand side variables reveals a *negative* relationship between home mortgage size and investment in safe assets.

As noted earlier, the reason that OLS is likely to yield biased estimates of the causal effect of commitments on portfolio allocation is that the choice of house size is endogenous: the error term ε is correlated with mortgage size. Agents who choose to commit a large fraction of their expenditure to a big house are likely to be more risk loving and have less background income risk than others. These agents are also the ones who are likely to hold the most risky portfolios, creating a positive correlation between house size and stockholding. In addition, a more subtle endogeneity problem arises because lifetime wealth is noisily measured: A higher level of commitment could proxy for having higher unobservable safe assets, such as a larger stream of anticipated stable labor income, and may therefore be associated with holding a riskier portfolio of liquid assets.¹⁹

In view of these biases, we need a source of variation in housing size that is *exogenous* in that it is uncorrelated with the error term, ε , in the estimating equation. We create such variation using age at (most recent) marriage as an instrument for housing commitment, conditioning on current age to eliminate life-cycle effects. This results in the following first-stage regression specification:

¹⁹See e.g., Davidoff (2003), who shows that higher covariance between labor income and house prices is associated with smaller house purchases.

$$\text{mortgage debt} = \delta + \xi \times \text{marriage age} + \omega \times \text{wealth} + \text{controls} + \nu \quad (29)$$

The intuition underlying this first-stage relationship is as follows. Marriage is a strong determinant of the home purchase decision: in our sample, 30% of married stockholders bought their current home within five years after getting married. Now, consider two fifty year old married homeowners, Adam and Bob, who are identical except for the year they got married. Adam got married at age 30 and Bob got married at age 35. The data reveals the following strong and highly statistically significant relationship: Holding fixed home equity, Bob has a bigger mortgage than Adam; in other words, Bob owns a bigger house. This is presumably because households are more likely to have the credit and resources necessary to buy a bigger home later in life, and adjustment costs in housing create stickiness in housing consumption over time.

Given this first-stage relationship, the key identifying assumption for our strategy to yield a consistent estimate of the casual effect of mortgage size on portfolio choice is that age at marriage is uncorrelated with unobserved individual characteristics and measurement error in wealth:

$$E[\text{marriage age} \times \varepsilon] = 0. \quad (30)$$

4.3 Does Age at Marriage Directly Affect Portfolio Choice?

A potentially important concern with our empirical analysis is that the orthogonality condition is violated because age at marriage itself is not randomly assigned. Marriage age may directly affect portfolio choice because it is correlated with unobserved heterogeneity in risk attitudes or background income risk. We test the orthogonality condition (30) by examining two “control” groups for whom the first-stage relationship between age at marriage and house size breaks down. The only way marriage age could affect portfolios for these groups is through a direct effect, i.e., if age at marriage is correlated with ε . In both control groups, we find insignificant effects of marriage age on portfolio risk. Under the assumption that the relationship between age at marriage and ε is the same in our control groups and the “treatment” group used in our main analysis, these results confirm the validity of our instrument.

We first define a “treatment” group for whom the first-stage relationship between age at marriage and house size is strong. To do so, we split our core sample into two

by the median number of years since the last marriage (25 years). The treatment group is the set of more recently married individuals, for whom the effect of marriage age on mortgage (house size) is strong and positive. Estimates of this first stage regression is reported in column 1 of Table 4. For this group, when controlling flexibly for wealth, home equity, and other observables such as age, occupation, industry, year, income, etc., marrying one year later is associated with owning a home that is \$481 larger on average. The instrument is statistically powerful: the Huber-White t-statistic is 5.1.²⁰

Importantly, for the treatment group, a reduced form regression of stocks on marriage age reveals a strong and significant negative relationship. This is illustrated in columns 2 and 3 in Table 4. Column 2 is a minimal specification with only age and liquid wealth as controls; this specification should allay concerns that our results are due to controlling for a large set of endogenous variables. Column 3 has the full set of controls. The coefficient of marriage age is virtually identical in the two regressions: we find that getting married one year later implies about \$630-\$640 less in stockholding. The first stage and reduced form regression are the underpinnings of our two-stage estimates in the next section.

We now turn to the control groups to test the orthogonality condition for the instrument. The first control group is homeowners who got married a long time ago (more than the median of 25 years). According to column 4 in Table 4, the first stage regression of mortgage on age at marriage breaks down in this group. This finding is intuitive: People who have been married for a long time are likely to have moved out of the home they bought when getting married, breaking the link between age at marriage and current home size. Hence, marriage age cannot affect portfolio choice through the commitment channel here. To test whether marriage age has a direct effect on portfolio choice, we ran the reduced form regression of stocks on marriage age for this control group. The point estimate, which is statistically insignificant, suggests that in this group, marrying one year later is associated with \$253 *more* invested in stocks (column 5).

Our second control group is married renters. There are 2,037 stockholders in the data who do not currently own homes (but may have in the past) and are or were married. Since these individuals have not purchased homes, there is by construction no first-stage

²⁰The effect is *not* due to variations in interest rates; when we hold the interest rate fixed, the relationship between age at marriage and home mortgage is unchanged. The reason we do not control for interest rates in our primary specifications is that doing so forces us to drop all households who have paid off their mortgage (for whom we have no interest rate data).

relationship between age at marriage and housing commitment for this group. Column 6 reports estimates of the reduced-form relationship between stocks and age at marriage for this group. Marrying a year later has a statistically insignificant *positive* effect of \$250 on stockholding for renters.

The data thus suggests that there is *no* direct effect of marriage age on stockholding in both control groups. This supports the key identification condition (30), provided that age at marriage has the same correlation with the error term ε , i.e., with risk preferences and unobserved wealth, for the treatment and control groups. Renters or homeowners who married a long time ago are not perfect control groups in that their average characteristics differ from those homeowners who married recently. Renters tend to have less liquid wealth and stocks, while long-married homeowners have more liquid wealth and more stocks. The recent-married have on average about \$48,000 in stockholdings, while long married have \$81,000 and renters have \$31,000. But the standard deviation of stockholding in each of these groups is quite large (\$250,000 in the treatment and \$357,000 and \$85,000, respectively, in the control groups) suggesting there is quite a bit of overlap across groups. Moreover, the fact that our treatment group is sandwiched in the middle along these demographics implies that an endogeneity story that is to explain away our findings would have to be fairly complicated. Specifically, it would require a correlation between marriage age and stockholding that changes non-linearly along the above demographics and is unique to homeowners.

Importantly, the reduced form estimates for our treatment and control groups not only have the opposite sign but are also statistically distinguishable. The 95% confidence interval of the reduced-form treatment estimate does not overlap with the corresponding confidence interval in either of the control groups. Moreover, the point estimates of the two treatment groups are positive and virtually identical, suggesting that our two control groups are quite similar along the relevant dimensions. If there is a direct effect of marriage age on stockholdings, it is slightly positive, working against us finding support for the commitments hypothesis.²¹

Theories of age at marriage provide some intuition for why the correlation between marriage age and ε could be positive. The “sociological” theory of marriage suggests that

²¹More precisely, if the correlation between marriage age and the error term ε is positive, the IV regression yields an estimated β coefficient that is biased upward.

individuals who are less affluent and educated tend to choose to marry earlier for cultural reasons. The “economic” theory of marriage, pioneered by Becker (1973), argues that individuals weigh search costs against the benefits of being married in determining when to get married. To the extent that more informed and well educated people have higher opportunity costs of search and take longer to establish observable wealth prospects, they are likely to marry later. Empirical studies of marriage age corroborate the predictions of both theories: the key determinants of marriage age are education, occupation, and mother’s age at marriage (see e.g., Kiernan and Eldridge (1987) and Keeley (1979)).

Both of these theories imply that marrying later is related to being more sophisticated and well educated. More sophisticated individuals typically face less background wealth risk and have greater future labor income, usually thought of as a “safe” asset. For both of these reasons, selection effects could result in late marriers having more stocks relative to bonds.

4.4 Age at Marriage and Measurement Error in Wealth

In our sample, higher liquid wealth is correlated with a higher share of bonds relative to stocks, and higher age at marriage is correlated with higher wealth. These findings suggest that if we do not control properly for wealth, an increase in age at marriage (and commitment) could be spuriously correlated with a shift to more conservative portfolios. This possibility is particularly worrisome because wealth is measured noisily: aside from classical measurement error, we cannot directly observe all components of lifetime wealth, such as future labor income. We are forced to make do with proxies such as education and occupation, leaving open the question of whether we have controlled adequately for wealth.

The control group evidence above should mitigate concerns about the violation of orthogonality through this channel to some extent. However, we can also test the hypothesis that age at marriage is proxying for wealth in our specifications more directly. In particular, if age at marriage is proxying for wealth, it should also predict other less durable aspects of consumption, such as food or vehicle consumption. Although we do not have data on purely nondurable consumption such as food, we can examine the relationship between the total value of a household’s current car stock and age at marriage. As one would expect, there are strong positive correlations between the value of a house-

hold's car stock, property value, and liquid wealth. But cars are much less durable than homes: The median age of cars in our sample is 4.3 years, compared with 15.2 years for houses. Moreover, marriage is a much weaker predictor of when a household buys a car than a first house. Consequently, absent a spurious correlation due to measurement error in wealth, we should expect no relationship between age at marriage and total car value at a later age.

To test this hypothesis, we estimate an equation analogous to the reduced-form specification with rich controls in column 3 of table 4, replacing the dependent variable with the total value of cars currently owned. The coefficient in this regression is -23.9 and the standard error is 15.83. Similar results are obtained when we replicate the more minimal specification in column 2 of Table 4 with car value as the dependent variable. The lack of a relationship between age at marriage and car value suggests that age at marriage is indeed picking up a commitment effect that arises from the durability of housing rather than a component of unobserved wealth.

Based on the analysis of the control groups and vehicle stocks, we proceed to use age at marriage as an instrument for property value under the identification assumption that there is no time-varying relationship between marriage age and portfolio preferences unique to homeowners.

4.5 Instrumental Variables Estimates

Table 5 presents IV estimates of the effect of house size on stock ownership. In these specifications, we use the core sample of ever-married homeowners (both recent and long-married) to avoid unnecessary sample selection. Results are even stronger if we restrict attention to recent-married homeowners, the subsample that drives most of the first stage relationship.

We begin by discussing the estimates of specification (1), which are typical, and then discuss a series of robustness checks. As above, all standard errors are robust to arbitrary heteroskedasticity of error terms. In specification (1), we regress the level of stockholding on outstanding home mortgage, which is instrumented using age at marriage in this and all subsequent specifications. To address concerns that endogenous regressors may be affecting results, this specification has a minimal set of controls: liquid wealth and age. The estimates indicate that at the mean, a \$1 exogenous increase in home value, holding

fixed home equity and other wealth, causes a 95 cent reduction in stockholding.²² This estimate is statistically significant at the 1% level.

Specification (2) shows that this result is robust to the inclusion of a rich set of controls: ten piece linear splines for liquid wealth, home equity, and age; controls for education and income; and year, occupation, and industry dummies. Under this specification, a \$1 increase in home value, holding total wealth fixed, is estimated to cause a 74 cent shift out of stockholding with a standard error of 34 cents.²³ This effect is extremely large. To see this, recall that the standard deviation of home mortgages is approximately 50,000, implying that a one standard deviation increase in mortgage leads to a \$37,000 portfolio reallocation from stocks to less risky assets. Although it is dangerous to extrapolate to the case of zero commitment, which is far out of sample, these results suggest that models of portfolio choice and asset pricing that ignore commitment effects may drastically overstate the level of risk that agents should be willing to take.

Specification (3) reports estimates of the regression that corresponds to (2) with bonds as the dependent variable. It shows that a \$1 increase in home size causes a 52 cent increase in holding bonds. These results confirm that we are identifying an actual shift in portfolio composition and not spurious correlations in levels. The discrepancy between the 74 cent shift out of stocks and 52 shift into bonds is accounted for by an increase in debt owed to the household. Since it is unclear whether this debt is risky or safe, we conclude that the shift in asset composition from risky to safe assets is between 52-74 cents.

To shed further light on the mechanism through which higher commitments lead to safer portfolio allocation, specification (4) of Table 5 estimates the effect of owning a larger home on the probability of participation in the bonds market. We estimate a linear probability model for the bondholder dummy, with the same independent variables as in specification (3). The participation effect is very strong and statistically significant: a

²²As emphasized by Imbens and Angrist (1994), IV only estimates a "local" average treatment effect for the "compliers" who respond to the treatment, in this case the group of individuals who buy larger homes as a result of marrying later. However, the estimates here may be close to average treatment effects in the population since our first-stage relationship remains strong across income groups and various other demographics.

²³Since we are controlling for wealth, this portfolio shift is *not* the result of having more debt. The household's balance sheet is unaffected if the home is a riskless asset because home mortgage debt is secured by the house itself (see subsection 4 below).

\$10,000 increase in mortgage causes a 3.6 percentage point increase in the probability of owning bonds; for comparison, the average probability of owning bonds in this sample is 22%. Hence, a significant fraction of the shift from stocks to bonds when individuals have more commitments appears to be driven by the opening of bond accounts.

These results are also robust to a number of other specification checks that are not reported in the table. First, since we do not have data on the portfolio composition of IRAs, one might worry that households with larger commitments are taking more risk in retirement accounts to offset less risk exposure elsewhere. We establish that this is not driving our results by restricting attention to the subsample who has zero IRA assets (approximately 45% of households) and finding the same estimates. Second, since the risk properties of “other assets” such as other financial equity and debt owed to the household are ambiguous, we drop the households who report such wealth and find that the results continue to hold. Third, results remain similar if we restrict attention to those homeowners who still have a strictly positive amount of mortgage debt to be paid off and control for the interest rate on their loan. Fourth, to allay the concern that our results may be driven by transitory effects when individuals buy houses, we condition on having a tenure of at least five years in the current home and find similar results. Fifth, as noted above, our results are qualitatively unchanged in the large sample that includes non-stockholders. However, it should be noted that estimating a linear probability model for the stock market participation reveals no relationship between mortgage size and stock market participation (potential reasons are discussed below). In summary, the data reveal a strong and robust effect of exogenous changes in the level of commitment on portfolio choice: holding all else fixed, a \$1 increase in commitment causes a shift of approximately 50-70 cents in an agent’s portfolio from risky to bonds.²⁴

To benchmark our results relative to the more common portfolio-share specifications in the existing literature, we also present IV results for regressions run in shares. In these regressions, the dependent variable is the share of stocks in total wealth (stocks/total

²⁴To the extent that housing commitments are positively correlated with other commitments – e.g., furniture and utilities – our estimates overstate the effect of commitments on stockholding. If housing is complementary to these other commitment goods, a \$1 increase in home value corresponds to more than a \$1 increase in commitment. Although nonzero, this effect is unlikely to change the magnitude of the estimates significantly insofar as the elasticities of furniture and utility expenses with respect to house value are small.

wealth) and all the independent wealth-related variables are included as shares of total wealth as well. As in specifications (1)-(3), we control flexibly for home equity, so that the change in mortgage share corresponds to owning a larger house. Since wealth is in the denominator, we drop the 66 households that report negative or extremely low values of total wealth (less than \$1000) in specification (5).²⁵ Though statistically significant, the coefficient on home mortgage is now only -0.07. This is because portfolio-share regressions weight high-wealth households less than low wealth households. Indeed, a levels regression similar to that in column 2 on the restricted sample of households with liquid wealth below the median (\$50,000) shows that the effect of commitments on stockholding is small and statistically insignificant for this group. This reasoning is also confirmed by specification (6), which shows that the shares regression yields estimates similar to those of the levels regression when we restrict the sample to households that have total wealth above \$150,000. Given that people who do not hold stocks tend to be poorer, this evidence is in line with the lack of a significant relationship between mortgage size and stock market participation.

There are a number of reasons why the portfolios of the relatively poor may not be very sensitive to variations in commitments. First, there is a fair amount of evidence that the savings and portfolio decisions of the rich are substantially different from those with moderate wealth. Carroll (2000) and Guiso, Japelli and Terlizzese (1996) find that the share of risky assets in a household's portfolio rises, perhaps dramatically, with wealth. One potential reason emphasized by Carroll is decreasing relative risk aversion, which is also supported by the findings of Ogaki and Zhang (2001) in a risksharing context. With DRRA utility, poorer households will exhibit lower sensitivity of portfolios to variations in commitment levels. This is because commitments essentially reduce effective wealth, and households with higher risk aversion invest a smaller fraction of their marginal dollar of wealth in risky assets, as shown by (5).

Second, poor households are more likely to be borrowing constrained. When close to the constraint, such households have a high marginal propensity to consume out of cash-on-hand, and accordingly have a low marginal propensity to invest. This effect also attenuates the sensitivity of stockholding to changes in commitments.

²⁵Inclusion of these households causes standard errors to explode since they are heavily-weighted outliers in a least-squares regression.

Third, as argued by Gabaix and Laibson (2001), there may be fixed costs to adjusting financial portfolios. These fixed costs matter less for wealthy people with large portfolios. Thus we expect moderate-wealth households to adjust their portfolios less frequently, especially if the required adjustments are also small, as suggested both by the borrowing constraint and DRRA utility stories.

Finally, mismeasurement of wealth may be proportionally much larger for poor than for rich households; note that the sample contains 66 individuals with wealth less than \$1000, and several of them actually have negative wealth. Such measurement error would bias the coefficient of wealth towards zero for low wealth households. Since mortgage debt is correlated with wealth, its coefficient would partially compensate for that effect, and thus get biased upward toward zero. In fact, we do find a much lower coefficient on liquid wealth in the poorer half of the sample, confirming this argument.

It should be noted that this is not the first study to examine the relationship between housing and portfolio choice. Fratantoni (1998), Cocco (2000), Flavin and Yamashita (2002), and Kullmann and Siegel (2002) document a weak relationship between housing and portfolio allocation in the cross section. However, as we have noted, the cross-sectional results are strongly biased by endogeneity. Another important distinction between our results and theirs is that prior work has regressed stock ownership on home value, *without* holding home equity fixed. This makes these estimates difficult to interpret in the context of testing the commitments hypothesis: as mentioned above, we would ideally compare households that have the same level of liquid and illiquid wealth.

More recently, Yamashita (2003) also fits TSLS models using age, age², family size, tenure in home, and home price growth rates as instruments and obtains comparable estimates to his OLS regressions. However, we note that the validity of these instruments is questionable insofar as age may directly affect portfolio choice (e.g., individuals shift to bonds as a guaranteed source of income later in life). Indeed, when using Yamashita's instruments on our sample, we find a *positive* statistically significant coefficient of 0.2 on the home mortgage variable, controlling for home equity.²⁶ Moreover, conditional on the validity of our age at marriage instrument, a Hausman specification test rejects the hypothesis that TSLS estimates using Yamashita's set of instruments are consistent with

²⁶When we replicate Yamashita's specification of regressing stock ownership on property value without controlling for home equity, we find a coefficient of -0.01 on property value, with a confidence interval that contains his estimate of -0.07.

a p value of 0.028.

4.6 Commitments or Housing Price Risk?

The empirical evidence indicates that an exogenous increase in house size causes a significant portfolio reallocation toward safer assets. If homes were riskless investments, the agent's overall portfolio – including the housing asset – would be less risky after an exogenous increase in mortgage debt. This would confirm our hypothesis that larger commitments cause agents to hold less risky portfolios. However, homes are actually volatile assets: Flavin and Yamashita (2002) report a standard deviation of 16% on the growth rate of home prices. Hence, it is not immediately clear that households actually choose a lower overall exposure to risk when they experience an exogenous increase in home size. Put differently, the reallocation from stocks to bonds may occur not because of the commitment effect but simply because greater exposure to risk in the housing market causes risk-averse agents to hold allocate the remainder of their portfolio more conservatively. This raises the question of whether the evidence reported in the previous section is actually due to commitments or the competing theory of risk exposure in the housing market.

The implications of the existing theoretical literature for the magnitude of the housing risk effect on portfolio choice are mixed. Flavin and Yamashita report simulations of optimal portfolio choice in a mean-variance framework and find that a larger house implies a sharp reduction in optimal stockholding, even in a neoclassical model where housing is freely adjustable. One shortcoming of their analysis is that it does not incorporate the natural hedge homeowners have in owning an asset whose value fluctuates with the price of housing, which could greatly reduce the actual riskiness of owning a home. Sinai and Souleles (2003) present evidence suggesting that this hedging motive is quite important: they find that people are more likely to buy homes (rather than rent) in cities where price volatility is high, suggesting that homes may not be as “risky” as suggested by Flavin and Yamashita.

Discerning the relative importance of these two effects requires independent variation in exposure to risk and house size. We obtain such variation by exploiting regional differences in the volatility of house prices. In particular, we use annual home price data from 1975-2002 from the Freddie Mac repeat sales house price index at the metropolitan

statistical area (MSA) level to construct measures of the standard deviation of the real growth rate of house prices, σ_{hg} , in each MSA.²⁷ We have home price data on 91 MSAs, which account for 9,465 of the households in our original sample. The mean value of σ_{hg} in our sample is 5% and the standard deviation of σ_{hg} is 2.5%, creating a considerable amount of exploitable variation in the riskiness of housing.

Before turning to the empirical evidence, it is helpful to formalize our method of distinguishing the housing-price risk theory from the commitments effect. Let H denote an agent's initial house value and g_t denote the growth rate of housing prices in year t . This agent's exposure to risk in the housing market is captured by the standard deviation of the annual change in his housing assets:

$$\sigma_H = \text{stdev}(Hg) = H\sigma_{hg}. \quad (31)$$

The existing literature on housing risk uses variation in H to estimate the effect of σ_H on portfolio choice. Those estimates are biased insofar as H directly affects portfolio choice through the commitment effect modelled in this paper. Our strategy is to instead use cross-city variation in σ_{hg} to create variation in σ_H while holding the value of H fixed.²⁸ This permits us to separate the commitment and housing risk effects and identify their relative importance for portfolio choice.

We begin by demonstrating that our measure of housing market price volatility affects behavior by showing that demand for housing is weaker in cities that have high price volatility. Column 1 of Table 6 reports OLS estimates of the effect of an increase in home price volatility on the value of the property households purchase. This and all subsequent specifications in Table 6 have a rich set of controls: liquid wealth and age splines, industry, occupation, and year dummies, income education, average home price growth in MSA, and other measures of wealth.²⁹ In addition, all standard errors are

²⁷Using alternate measures of housing risk, such as the covariance of house prices with consumption does not affect our results.

²⁸One potential weakness of using city-level measures of price risk is that they do not capture idiosyncratic price risk for individual homeowners, which could potentially have different effects on portfolio choice decisions. However, under the reasonable assumption that the degree of idiosyncratic risk is proportional to the degree of MSA-level risk, finding that MSA level house risk does not significantly impact portfolio choice also implies that idiosyncratic risk is not an important factor in these decisions.

²⁹The key coefficient estimates reported in Table 6 are robust to a variety of different specifications and controls for MSA and household-level characteristics.

robust to arbitrary heteroskedasticity and serial correlation within MSAs. The estimates indicate that a 1 percentage point increase in σ_{hg} is associated with a \$4,112 reduction in home prices on average.³⁰ This estimate is statistically significant with $p < 0.01$. A similar effect is observed for renters: a one percentage point increase in σ_{hg} is associated with a \$25 reduction in the monthly rent an individual pays.

Given that agents consume less housing in areas of high price volatility, the key question for our purposes is the extent to which agents living in high price volatility cities also choose to take less risk in the stock market. Specification (2) replicates (1) with level of stocks as the dependent variable. The effect of home price volatility on stockholding is statistically insignificant. The lower bound of the 95% confidence interval implies that at most, a 1 percentage point increase in σ_{hg} causes a \$900 reduction in stockholding. We can gauge the magnitude of this effect using (31). Since the mean of σ_{hg} is 5%, a 1 percentage point increase in σ_{hg} creates an increase in σ_H that is equivalent to buying a home that is 20% larger. The mean home value is \$150,000, implying that a \$30,000 increase in home value will cause the same change in exposure to housing risk as a 1 percentage point increase in σ_{hg} . It follows that a \$30,000 increase in home value causes at most a \$900 reduction in stockholding due to greater exposure to housing risk. This effect is an order of magnitude smaller than the effect we estimate due to commitments, which would cause roughly a \$15,000 reduction in stockholding.³¹ These results therefore strongly suggest that the estimates of Table 5 identify the causal effect of commitments on portfolio choice rather than variations in housing risk.

It is informative to contrast the effect of housing market price risk on portfolio choice for homeowners with that for renters. Specification (3) replicates (2) for the set of ever-married renters for whom we can identify MSA. Strikingly, a one percentage point increase in σ_{hg} causes a \$2,812 reduction in the value of stocks held by renters. This

³⁰These estimates simply confirm that a rise in price volatility causes an inward shift in the demand for housing. Since the price of housing is endogenous, they do not tell us how much less an individual will spend on a house when the standard deviation of house prices increases, holding the mean price fixed.

³¹Taking the behavioral response of spending less on housing when price volatility is high into account does not change this conclusion. Based on the estimates of specification (1), a 1 percentage point increase in σ_{hg} causes approximately a \$5,000 decrease in home spending, which means that the \$900 reduction would occur after a \$25,000 increase in home value through the risk exposure channel. The commitment effect continues to dwarf this estimate.

effect is statistically significant with a robust t-value of 2.6. This result should allay the concern that MSA price volatility is not a precise measure of an agent's exposure to risk in the housing market; were this the case, we would not see an effect on portfolio choice for renters.

An important identification assumption in the preceding analysis is that the volatility of home prices in an individual's city of residence is not directly related to his risk preferences. This assumption would be violated if, for instance, more risk averse individuals migrate away from high volatility cities while risk lovers go in the opposite direction. This selection effect could bias our estimates of the effect of price volatility on stockholding upward and artificially work against the housing risk theory.

To overcome this form of selection bias, we exploit the strong relationship between an individual's state of birth and current residence. In particular, we instrument for housing risk in an individual's MSA by the housing risk in his state of birth.³² The exclusion restriction for the IV is that state of birth is uncorrelated with risk preferences, conditional on all observables. Specification (4) shows that the point estimate for the effect of housing risk on stockholding in the IV regression for homeowners remains essentially unchanged relative to the OLS specification, although standard errors rise given the weakness of the instrument. Specification (5) shows that even in the IV, high housing risk appears to cause lower stockholding for renters, although this effect is only marginally significant because of the imprecision of the estimates. These results suggest that selection bias is not a first-order concern and that housing risk does in fact have a relatively minor effect on behavior in asset markets.

In our view, it is not surprising that the portfolio of renters is highly sensitive to the degree of housing market risk whereas the portfolio of homeowners is not. Since they own an asset whose price fluctuates contemporaneously with their wealth, homeowners have a natural hedge against price risk insofar as they plan to live in the same house or neighborhood for a long time. In other words, the covariance properties of housing risk with wealth are quite favorable. On the other hand, since renters do not have this type of insurance against price fluctuations, it is intuitive that they choose to bear less risk in financial markets when living in high price-volatility areas. This intuition is not reflected in Flavin and Yamashita's calibrations since they consider a static mean-

³²State level data is used because more detailed data on birthplace is unavailable.

variance optimization framework instead of using the derived utility of consumption.

The final specification of Table 6 reports estimates of an IV regression that attempts to directly contrast the explanatory power of the commitments hypothesis with the competing housing-price risk theory put forth in the recent literature by Flavin and Yamashita (2002), Kullmann and Siegel (2002), and others. The dependent variable in this regression is again the level of stockholding. The key explanatory variables of interest are outstanding home mortgage debt and exposure to home price risk, σ_H . We instrument for mortgage debt and σ_H using age at marriage, σ_{hg} , and the interaction of these two variables. Thus, conditional on the exclusion restrictions for these two instruments, the specification gives separate estimates of the causal effects of changes in house size and housing risk on portfolio choice. The point estimate on the mortgage debt coefficient actually *rises* relative to that reported in Table 5 when the degree of housing risk is held fixed.³³ A \$1 increase in home size is estimated to cause a \$1.13 dollar shift out of stocks. In contrast, the estimate on the σ_H coefficient is small and statistically insignificant. The magnitude of σ_H can be interpreted by using (31) to observe that a \$1 increase in home value causes a .05 increase in σ_H . Based on the point estimate in (6), this yields a $.05 \times 3.4 = 17$ cent reduction in stockholding. In summary, a horse race of the commitments hypothesis versus the housing-price risk hypothesis favors the former as the primary reason that housing size has a causal effect on portfolio choice.

4.7 Imputing Risk Aversion from the Estimates

Our empirical specifications coincide precisely with equation (5) for optimal portfolio choice derived from our model. A structural interpretation of the regression equation allows us to pin down the coefficient of the commitment variable as a function of underlying parameters of the model.

The key step in connecting the empirical estimates to the model's parameters is identifying how the measure of commitments in the model, x , relates to our empirical measure, home mortgage debt. In equation (5), the measure of commitments is ηx . This is the outstanding value of commitments the individual needs to pay before the next adjustment. If we interpret the house one owns as the physical commitment good,

³³Standard errors rise because we have lost more than one third of the sample in focusing on MSAs for which we have price data and because of clustering by MSA.

the outstanding value of commitments roughly corresponds to the outstanding value of mortgage debt, as long as the individual plans to move out of the house only in the distant future. The model therefore implies that the coefficient of home debt in the regressions equals $\pi/\gamma\sigma^2$.

The most conservative point estimate for the commitment coefficient is .527 in Column 3 of Table 5, the specification with bonds. Substituting in the estimates $\pi = .06$ and $\sigma = .165$ for the equity premium and market volatility used in Section 3, the imputed value of $\gamma = 4.17$. As noted in the calibration, values of γ between 4 and 5 generate observed equity premiums and consumption volatility when the share of commitments in total consumption is 50-60%, which is consistent with the expenditure data reported in Table 1.

5 Conclusion

This paper has shown that consumption commitments – goods whose consumption can be adjusted only at a cost – make individuals more risk averse and demand higher equity premiums. The key theoretical result is that aggregating a microeconomic model of commitments yields aggregate dynamics that coincide with those that arise from a representative consumer model with habit formation.

In calibrating this model to aggregate data, we find that it can resolve both the equity premium and riskfree rate puzzles with a high commitment share (80%). With a lower commitment share (50%), the model still explains the equity premium puzzle, but only partially resolves the riskfree rate puzzle.³⁴

The commitment microfoundations generate a wealth of testable implications. Using age at marriage as an instrumental variable, we found that a \$1 exogenous increase in commitment consumption (housing) causes agents to shift 50-70 cents from stocks to bonds, consistent with the model’s key prediction. We reject the hypothesis that this portfolio reallocation occurs primarily because of greater exposure to housing risk, in contrast to recently proposed theories of the relationship between housing and portfolio choice.

³⁴Allowing for a more flexible specification of utility that breaks the link between intertemporal substitution and risk aversion or incorporating borrowing constraints would presumably bring us closer to explaining both puzzles even with a low commitment share.

Other implications for time- and cross-sectional return variation in asset pricing and macroeconomic consumption dynamics remain to be explored.

Appendix

Proof of Theorem 1

Denote the value function of the consumer by $V_t(w_t, x_t)$. For the purposes of this proof only, assume that t stands for time elapsed since the last reset date, as opposed to calendar time. The Bellman equation for the maximization problem between two reset dates is

$$\rho V_t = \max_{f, \alpha} \left\{ \left(\frac{f_t^{1-\gamma}}{1-\gamma} + \mu \frac{x_t^{1-\gamma}}{1-\gamma} \right) + E dV \right\}$$

which yields, using Ito's lemma, to

$$\rho V_t = \max_{f, \alpha} \left\{ \left(\frac{f_t^{1-\gamma}}{1-\gamma} + \mu \frac{x_t^{1-\gamma}}{1-\gamma} \right) + \frac{dV}{dt} + \frac{dV}{dw} [(r + \alpha_t \pi) w_t - c_t] + \frac{1}{2} \frac{dV^2}{d^2 w_t} (\alpha_t \sigma w_t)^2 \right\}. \quad (32)$$

We guess that the value function is of the form

$$V_t(w_t, x_t) = \varphi_t \frac{(w_t - \eta_t x_t)^{1-\gamma}}{1-\gamma} + \mu \frac{x_t^{1-\gamma}}{1-\gamma} \cdot \left[\frac{1 - e^{-\rho(T-t)}}{\rho} \right]$$

where φ_t and η_t are deterministic functions to be determined. Note that the second term is just the utility value of outstanding consumption commitments before the next reset date (discounted by the subjective discount factor ρ).

The first order condition from maximizing (32) yields the consumption rule

$$f_t^{-\gamma} = \frac{dV_t}{dw_t} = \varphi_t (w_t - \eta_t x_t)^{-\gamma}$$

or equivalently

$$f_t = \varphi_t^{-\frac{1}{\gamma}} (w_t - \eta_t x_t) \quad (33)$$

and the investment rule

$$\alpha_t = \frac{\pi}{\gamma \sigma^2} \left(1 - \eta_t \frac{x_t}{w_t} \right). \quad (34)$$

Plugging these back into the Bellman equation and simplifying yields

$$\begin{aligned} \rho \varphi_t \frac{1}{1-\gamma} &= \varphi_t^{\frac{\gamma-1}{\gamma}} \frac{1}{1-\gamma} + \frac{d\varphi_t}{dt} \cdot \frac{1}{1-\gamma} - \varphi_t^{1-\frac{1}{\gamma}} + \frac{1}{2} \varphi_t \frac{\pi^2}{\gamma \sigma^2} \\ &\quad + \varphi_t r (w - \eta_t x)^{-1} \left\{ w - \frac{x}{r} - \frac{d\eta_t}{dt} \frac{x}{r} \right\}. \end{aligned}$$

In order for this equation to hold, we need that

$$w_t - \eta_t x_t = w_t - \frac{x_t}{r} - \frac{d\eta_t}{dt} \frac{x_t}{r_t}$$

is satisfied. Equivalently,

$$\frac{d\eta_t}{dt} = r\eta_t - 1$$

which can be solved and gives

$$\eta_t = \frac{1}{r} - K_1 \cdot e^{rt}$$

where K_1 is a constant of integration. Because we need $\eta_T = 0$, the solution has to be

$$\eta_t = \frac{1}{r} \left\{ 1 - e^{-r(T-t)} \right\}.$$

Note that η_t as defined here is the present discounted value of a cash-flow of 1 every period up to the next reset date. This squares with our intuition that the consumer needs to put away ηx dollars in the riskless asset to be able to finance commitment consumption until the next reset date comes.

Using the formula for η , the Bellman equation implies (after some calculations)

$$\frac{\dot{\varphi}_t}{\varphi_t} = -\gamma \varphi_t^{\frac{-1}{\gamma}} + \rho - (1 - \gamma) \left\{ \frac{\pi^2}{2\gamma\sigma^2} + r \right\}.$$

This is a differential equation. Denote

$$\rho - (1 - \gamma) \left\{ \frac{\pi^2}{2\gamma\sigma^2} + r \right\} = L,$$

then the solution can be written as

$$\varphi_t = \left(\frac{\gamma}{L} + e^{\frac{L}{\gamma}t} K_2 \right)^\gamma$$

where K_2 is a constant. The value of K_2 will be pinned down by the value matching condition that needs to hold at a reset date. Note that

$$\varphi_0 = \left(\frac{\gamma}{L} + K_2 \right)^\gamma.$$

Let us now examine what happens on a reset date. We have

$$V_0(w_0) = \max_x \varphi_0 \frac{(w_0 - \eta_0 x)^{1-\gamma}}{1-\gamma} + \mu \frac{1 - e^{-\rho T}}{\rho} \cdot \frac{x^{1-\gamma}}{1-\gamma}.$$

The first order condition from the maximization gives

$$\varphi_0 \eta_0 (w_0 - \eta_0 x)^{-\gamma} = \mu \frac{1 - e^{-\rho T}}{\rho} \cdot x^{-\gamma}$$

so that

$$x = w_0 \cdot \frac{1}{\left[\mu \frac{1 - e^{-\rho T}}{\rho} \right]^{\frac{-1}{\gamma}} \eta_0^{\frac{1}{\gamma}} \left(\frac{\gamma}{L} + K_2 \right) + \eta_0}. \quad (35)$$

If we denote the constant multiplying w_0 by K_3 , then the value becomes

$$V_0(w_0) = \frac{w_0^{1-\gamma}}{1-\gamma} \left\{ \varphi_0 (1 - \eta_0 K_3)^{1-\gamma} + \mu \frac{1 - e^{-\rho T}}{\rho} \cdot K_3^{1-\gamma} \right\}.$$

If our guess for the value function was correct, then value matching on the reset date needs to be satisfied

$$\varphi_T \frac{w_0^{1-\gamma}}{1-\gamma} = V_0(w_0).$$

This, coupled with the previous equation, implies that

$$\varphi_0(1 - \eta_0 K_3)^{1-\gamma} + \mu \frac{1 - e^{-\rho T}}{\rho} \cdot K_3^{1-\gamma} = \varphi_T$$

and using our formulas for φ_t

$$\left(\frac{\gamma}{L} + K_2\right)^\gamma (1 - \eta_0 K_3)^{1-\gamma} + \mu \frac{1 - e^{-\rho T}}{\rho} \cdot K_3^{1-\gamma} = \left(\frac{\gamma}{L} + e^{\frac{L}{\gamma} T} K_2\right)^\gamma.$$

Note that K_3 depends on K_2 in this equation. It is easy to check that as K_2 goes to infinity, the right hand side dominates. Moreover, if K_2 is chosen such that the right hand side is zero, the left hand side, which can be calculated to equal

$$\begin{aligned} & \left(\frac{1}{\left[\mu \frac{1-e^{-\rho T}}{\rho} \right]^{\frac{1}{\gamma}} \eta_0^{1-\frac{1}{\gamma}} \left(\frac{\gamma}{L} + K_2 \right)^{1/(\gamma-1)} + \left(\frac{\gamma}{L} + K_2 \right)^{\gamma/(\gamma-1)}} \right)^{1-\gamma} + \\ & + \mu \frac{1 - e^{-\rho T}}{\rho} \cdot \left(\frac{1}{\left[\mu \frac{1-e^{-\rho T}}{\rho} \right]^{\frac{-1}{\gamma}} \eta_0^{\frac{1}{\gamma}} \left(\frac{\gamma}{L} + K_2 \right) + \eta_0} \right)^{1-\gamma} \end{aligned}$$

is still positive. It follows that there is at least one value of K_2 for which the equation is satisfied. Thus we have found an optimal policy of the maximization problem.

Now go back to our assumption that t measures the time elapsed since the last reset date. Note that the functions φ and η only depend on the time elapsed since the last reset date, Δ . Defining $\psi_{\Delta(t)} = \varphi_{\Delta(t)}^{-\frac{1}{\gamma}}$ and $\chi = K_3$ formulas (33), (34) and (35) above show that the optimal policy is indeed as claimed in the theorem.

The implied dynamics for the net wealth of a particular cohort between adjustments is

$$d(w_t - \eta_{\Delta(t)} x_t) = [(r + \alpha\pi) w_t - c_t] dt + \alpha\sigma w_t dz - d\eta_{\Delta(t)} x_t$$

and using the ODE for η we get

$$\frac{d(w_t - \eta_{\Delta(t)} x_t)}{w_t - \eta_{\Delta(t)} x_t} = \left\{ r - \varphi_{\Delta(t)}^{-\frac{1}{\gamma}} + \frac{\pi^2}{\gamma\sigma^2} \right\} dt + \frac{\pi}{\gamma\sigma} dz.$$

Denoting $w_t^{\text{net}} = w_t - \eta_{\Delta(t)} x_t$ also yields

$$w_t^{\text{net}} = w_s^{\text{net}} \exp \left\{ \frac{\pi}{\gamma\sigma} (z_t - z_s) + \left(r + \frac{\pi^2}{\gamma\sigma^2} - \frac{1}{2} \frac{\pi^2}{\gamma^2\sigma^2} \right) (t - s) + \int_s^t \varphi_{\Delta(u)}^{-\frac{1}{\gamma}} du \right\} \quad (36)$$

for $s \leq t$ provided there is no adjustment between s and t .

As to consumption, we have

$$\begin{aligned} df_t &= \psi'_{\Delta(t)}(w_t - \eta_{\Delta(t)}x_t)dt + \psi_{\Delta(t)}d(w_t - \eta_{\Delta(t)}x_t) = \\ &= \left\{ \frac{\psi'_{\Delta(t)}}{\psi_{\Delta(t)}} + r - \psi_{\Delta(t)} + \frac{\pi^2}{\gamma\sigma^2} \right\} f_t dt + \frac{\pi}{\gamma\sigma} f_t dz \end{aligned}$$

so that

$$\frac{dc_t}{c_t} = \left\{ \frac{\psi'_{\Delta(t)}}{\psi_{\Delta(t)}} + r - \psi_{\Delta(t)} + \frac{\pi^2}{\gamma\sigma^2} \right\} \frac{f_t}{c_t} dt + \frac{\pi}{\gamma\sigma} \frac{f_t}{c_t} dz.$$

This completes the proof.

Proof of Proposition 1 and Proposition 2

At any point in time t , index each agent by Δ , that is, by the time elapsed since the agent last made a reset.³⁵ We can also index agents by their cohorts. We say that an agent is in cohort q if at time zero, the time elapsed since her last reset is exactly q : if her Δ at zero is q . Clearly, $t - \Delta \equiv q$ modulo T . Moreover, denote the value of t modulo T by τ . When we need to emphasize that Δ depends on calendar time, we write $\Delta(t)$; occasionally we write $\Delta(t, q)$ if we are interested in the Δ of a particular cohort.

In the rest of this proof, when we use individual level variables, we will explicitly refer to the individual in notation, unless not doing so causes no confusion. We can refer to an individual in two ways: by referring to her Δ , or by referring to her cohort q . For instance the food consumption of individual Δ at time t would be $f_t(\Delta)$, but occasionally we may use $f_t(q)$ too.

Start by fixing a cohort q . Note that at an adjustment, net wealth is multiplied by a factor $1 - \eta_0 K_2$. Call this factor ξ , then from (36)

$$w_t^{\text{net}} = w_s^{\text{net}} \cdot \xi^{\lceil \frac{t-\Delta-s}{T} \rceil + 1} \exp \left\{ \frac{\pi}{\gamma\sigma} (z_t - z_s) + \left(r + \frac{\pi^2}{\gamma\sigma^2} - \frac{1}{2} \frac{\pi^2}{\gamma^2\sigma^2} \right) (t - s) + \int_{\Delta-(t-s)}^{\Delta} \varphi_u^{-\frac{1}{\gamma}} du \right\}$$

where we have rewritten the integral of φ to make it depend on calendar time (extending periodically the function so that it is defined on the whole real line). Let

$$Z_t = \exp \left\{ \frac{\pi}{\gamma\sigma} z_t + \left(r + \frac{\pi^2}{\gamma\sigma^2} - \frac{1}{2} \frac{\pi^2}{\gamma^2\sigma^2} \right) t \right\}$$

then

$$w_t^{\text{net}} = w_s^{\text{net}} \cdot \xi^{\lceil \frac{t-\Delta-s}{T} \rceil + 1} \exp \left\{ \int_{\Delta-(t-s)}^{\Delta} \varphi_u^{-\frac{1}{\gamma}} du \right\} \cdot Z_t / Z_s.$$

³⁵Note that typically, the agents indexed by Δ in times t and s will be different.

Define $s^+(q, s)$ to be the first adjustment date after s of the cohort q we are focusing on. Clearly that only depends on q and s . Then the exponential term can be written as

$$\begin{aligned} \exp \left\{ \int_{\Delta-(t-s)}^{\Delta} \varphi_u^{-\frac{1}{\gamma}} du \right\} &= \exp \left\{ \int_{\Delta-(t-s)}^{\Delta-(t-s^+(q,s))} \varphi_u^{-\frac{1}{\gamma}} du \right\} \\ &\times \exp \left\{ \int_{\Delta-(t-s^+(q,s))}^0 \varphi_u^{-\frac{1}{\gamma}} du \right\} \times \exp \left\{ \int_0^{\Delta} \varphi_u^{-\frac{1}{\gamma}} du \right\} \\ &= m(q, s) \times \xi_1^{\left[\frac{t-\Delta-s}{T}\right]+1} \times y(\Delta) \end{aligned}$$

where $m(q, s)$ clearly only depends on q and s , $\xi_1 = \exp \left\{ \int_0^T \varphi_u^{-\frac{1}{\gamma}} du \right\}$ and $y(\Delta)$ only depends on Δ . Define $\xi_2 = \xi \cdot \xi_1$ so we can write

$$w_t^{\text{net}} = w_s^{\text{net}} \cdot m(q, s) \cdot \xi_2^{\left[\frac{t-\Delta-s}{T}\right]+1} y(\Delta) \cdot Z_t/Z_s. \quad (37)$$

We first aggregate the population starting from an arbitrary initial wealth distribution. Choose $s = 0$, and aggregate across all cohorts q to get

$$W_t^{\text{net}} = Z_t \xi_2^{\frac{t}{T}} \int_q w_0^{\text{net}}(q) \cdot m(q, 0) \cdot \xi_2^{\left[\frac{t-\Delta(t,q)}{T}\right]-\frac{t}{T}+1} f(\Delta(t, q)) dq.$$

Note that the integral term on the right hand side only depends on t modulo T , because the integrand is easily seen to be pointwise identical at $t + T$ to its value at t . Recall that τ is the value of t modulo T , then we can write

$$W_t^{\text{net}} = Z_t \cdot \xi_2^{\frac{t}{T}} \cdot b(\tau)$$

with an appropriate function $b(\cdot)$. Individual net wealth of cohort q is then

$$w_t^{\text{net}}(q) = w_0^{\text{net}}(q) \cdot m(q, 0) \cdot \xi_2^{\left[\frac{t-\Delta(t,q)}{T}\right]-\frac{t}{T}+1} \frac{y(\Delta)}{b(\tau)} \cdot W_t^{\text{net}}.$$

Because the cohort of an agent can be deduced from her Δ and calendar time t , the factor multiplying W_t^{net} in this expression only depends on τ and Δ , hence with appropriately chosen $g(\Delta, \tau)$ we have

$$w_t^{\text{net}}(\Delta) = g(\Delta, \tau) \cdot W_t^{\text{net}}.$$

Then aggregate commitment consumption can be written as

$$\begin{aligned} X_t &= \int_{\Delta} x_t(\Delta) d\Delta = \int_{\Delta} \chi w_{t-\Delta}^{\text{net}}(\Delta) d\Delta = \int_{\Delta} \chi g(0, \tau - \Delta) \cdot W_{t-\Delta}^{\text{net}} d\Delta = \\ &= \frac{1}{T} \int_0^T a(\tau - \Delta) \cdot W_{t-\Delta}^{\text{net}} d\Delta \end{aligned}$$

which corresponds is equation (10). Aggregate food consumption is

$$\begin{aligned} F_t &= \int_{\Delta} f_t(\Delta) d\Delta = \int_{\Delta} \psi_{\Delta(t)} w_t^{\text{net}}(\Delta) d\Delta = W_t^{\text{net}} \cdot \int_{\Delta} \psi_{\Delta(t)} g(\Delta, \tau) d\Delta = \\ &= b(\tau) \cdot W_t^{\text{net}} \end{aligned}$$

for an appropriate $b(\cdot)$ function, because the integral only depends on τ .

We now turn to the balanced wealth distribution case. Recall that the net wealth of cohort q at time t is given by (37). We will consider a double continuum of these agents, indexed by their cohort q and s , where s ranges from zero to T . Let the initial wealth levels be

$$w_s^{\text{net}}(q, s) = \frac{Z_s}{m(q, s)} \xi_2^{s/T}$$

then

$$w_t^{\text{net}}(q, s) = \xi_2^{t/T} \cdot Z_t \cdot \xi_2^{\left[\frac{t-\Delta-s}{T}\right]+1-\frac{t-s}{T}} y(\Delta). \quad (38)$$

We need to aggregate wealth over q and s . Clearly, integrating over s is the same as integrating over $t - s$. Let us integrate then first over $t - s$, then over Δ . The only thing that depends on $t - s$ is the power term; and importantly, the integral of that will only depend on Δ , but neither on t nor on s . Integrating that over Δ gives us that aggregate wealth is

$$W_t^{\text{net}} = M \cdot \xi_2^{t/T} \cdot Z_t$$

for some constant M . The wealth of the cohort corresponding to $\Delta = 0$ is given by

$$(w_t - \eta_t x_t)|_{\Delta=0} = \xi_2^{t/T} Z_t \cdot N$$

for some N , because again, the integral of $\xi_2^{\left[\frac{t-s}{T}\right]+1-\frac{t-s}{T}}$ over $t - s$ will not depend on t or s . Hence the wealth of that cohort can be written as

$$(w_t - \eta_t x_t)|_{\Delta=0} = \frac{N}{M} W_t^{\text{net}}.$$

Next consider commitment consumption:

$$x_t = w_{t-\Delta} \cdot \chi$$

and aggregating gives

$$X_t = \chi \frac{N}{M} \int_0^T W_{t-u}^{\text{net}} du.$$

This shows that with an appropriate initial wealth distribution the desired habit rule can be achieved.

Regarding food consumption, from (38) we have that

$$f_t(q, s) = \xi_2^{t/T} \cdot Z_t \cdot \xi_2^{\left[\frac{t-\Delta-s}{T}\right]+1-\frac{t-s}{T}} y(\Delta) \cdot \psi_\Delta$$

and integrating first over $t - s$ then over Δ we get that

$$F_t = b \cdot W_t^{\text{net}}$$

for some constant b .

Proof of Proposition 3

Let us assume that the model was started some time before date zero, so that all variables are defined for $-T < t < 0$ too. We need an $\zeta(u)$ function such that for $t > T$

$$\begin{aligned}
X_t &= \kappa(t)W_0^{\text{net}} + \int_0^t \zeta(u)C_{t-u}du = \kappa(t)W_0^{\text{net}} + \int_0^t \zeta(u) [F_{t-u} + X_{t-u}] du \\
&= \kappa(t)W_0^{\text{net}} + \int_0^t \zeta(u) \cdot bW_{t-u}^{\text{net}} du + \int_0^t \zeta(u) \cdot \frac{a}{T} \int_{t-u-T}^{t-u} W_s^{\text{net}} ds du \\
&= \int_0^T W_{t-u}^{\text{net}} \left[b \cdot \zeta(u) + \frac{a}{T} \int_0^u \zeta(s) ds \right] du + \int_T^t W_{t-u}^{\text{net}} \left[b \cdot \zeta(u) + \frac{a}{T} \int_{u-T}^u \zeta(s) ds \right] du + \\
&\quad + \kappa(t)W_0^{\text{net}} + \int_t^{t+T} W_{t-u}^{\text{net}} \frac{a}{T} \int_{u-t}^u \zeta(s) ds du.
\end{aligned}$$

By equation (12), we need that

$$\frac{a}{T} = b \cdot \zeta(u) + \frac{a}{T} \int_0^u \zeta(s) ds \quad (39)$$

for $0 \leq u < T$,

$$0 = b \cdot \zeta(u) + \frac{a}{T} \int_{u-T}^u \zeta(s) ds \quad (40)$$

for $T \leq u$, and

$$\kappa(t) = \frac{1}{W_0^{\text{net}}} \cdot \frac{a}{T} \int_{-T}^0 W_s^{\text{net}} \int_{t-T-s}^t \zeta(u) ds du. \quad (41)$$

From the first equation, with $u = 0$ we get that $\zeta(0) = a/Tb$. Differentiating that equation with respect to u , we get an ODE

$$\zeta'(u) = -\frac{a}{Tb} \zeta(u)$$

thus for $0 \leq u < T$ we have that

$$\zeta(u) = \frac{a}{Tb} \exp\left(-\frac{a}{Tb}u\right).$$

From equation(40) we have that when $T < u$

$$\zeta(u) = -\frac{a}{Tb} \int_{u-T}^u \zeta(s) ds.$$

This implies that $\zeta(u)$ cannot have the same sign on any interval of length larger than T : if it were e.g., positive on $[u-T, u]$ then $\zeta(u)$ would have to be negative. The equation also implies that as long as $a < b$, $\zeta(u)$ goes to zero geometrically, because

$$|\zeta(u)| \leq \frac{a}{Tb} \cdot \max_{[u-T, u]} |\zeta(v)|.$$

let us turn to the boundedness issue. Differentiating (40) in u we get

$$\zeta'(u) = -\frac{a}{b} (\zeta(u) - \zeta(u-T)).$$

Define $\beta(u) = \exp\left(\frac{a}{Tb}u\right) \zeta(u)$, then the previous ODE implies after some calculations that

$$\beta'(u) = \frac{a}{Tb} e^{\frac{a}{b}} \cdot \beta(u - T).$$

One solution of this ODE is $\beta(u) = \exp\left(\frac{a}{Tb}u\right)$. Using Gronwall's lemma, we can then bound the absolute value of the true solution by $K_4 \cdot \exp\left(\frac{a}{Tb}u\right)$ with some positive constant K_4 . It follows that $|\zeta(u)| < K_4 \cdot \exp\left(\frac{a}{Tb}u\right) \cdot \exp\left(-\frac{a}{Tb}u\right) = K_4$ which shows that $\zeta(u)$ is indeed bounded.

Finally, let us turn to $\kappa(t)$. The formula is given in equation (41), and clearly, as long as $\zeta(\cdot)$ goes to zero geometrically, so will $\kappa(\cdot)$. The reader may have noticed that the formula for $\kappa(\cdot)$ involves levels of net wealth before date zero. However, in fact the shape of $\kappa(\cdot)$ only depends on the cross sectional distribution of net wealth at date zero, because that completely determines the stochastic dynamics of all variables for all $t > 0$. Since $\kappa(\cdot)$ can be expressed as the difference of functions of those variables, $\kappa(\cdot)$ itself is completely determined by them. Furthermore, $\kappa(\cdot)$ is a deterministic function, as shown by the formula (41).

Proof of Theorem 3

Let us first solve the aggregate model

$$\max \int_0^\infty \frac{(C_t - X_t)^{1-\gamma}}{1-\gamma}$$

where

$$X_t = e^{-\lambda t} X_0 + a' \lambda \int_0^\infty e^{-\lambda u} W_{t-u} du.$$

The Bellman equation is

$$\rho V(W_t, X_t) = \max_{C, \alpha} \left\{ \frac{(C_t - X_t)^{1-\gamma}}{1-\gamma} + \frac{dV}{dW_t} [(r + \alpha_t \pi) W_t - C_t] + \frac{dV}{dX_t} \frac{dX_t}{dt} + \frac{1}{2} \frac{dV^2}{d^2 W} (\alpha_t \sigma W)^2 \right\}.$$

Note from the habit rule that

$$\begin{aligned} \frac{dX_t}{dt} &= -\lambda e^{-\lambda t} X_0 + a' \lambda \cdot \frac{d}{dt} \int_{-\infty}^t e^{-\lambda(t-u)} W_u du = \\ &= -\lambda e^{-\lambda t} X_0 + a' \lambda \cdot \left[W_t - \lambda \int_{-\infty}^t e^{-\lambda(t-u)} W_u du \right] = \lambda [a' W_t - X_t]. \end{aligned}$$

Now guess the solution to the problem

$$V(W, X) = \varphi' \frac{(W - \eta' X)^{1-\gamma}}{1-\gamma}$$

then the implied consumption and investment rules (from the first order conditions) are

$$C_t - X_t = \varphi'^{-\frac{1}{\gamma}} (W_t - \eta X_t)$$

and

$$\alpha_t = \frac{\pi}{\gamma \sigma^2} \left(1 - \eta' \frac{X_t}{W_t} \right).$$

Plugging back into the Bellman equation and simplifying yields

$$\begin{aligned} (W_t - \eta' X_t)^{1-\gamma} \left\{ \frac{\rho\varphi'}{1-\gamma} - \frac{\varphi'^{\frac{\gamma-1}{\gamma}}}{1-\gamma} - \varphi' \frac{\pi^2}{\gamma\sigma^2} + \varphi'^{\frac{\gamma-1}{\gamma}} + \frac{1}{2}\gamma\varphi' \frac{\pi^2}{\gamma^2\sigma^2} \right\} = \\ = \varphi' (W_t - \eta' X_t)^{-\gamma} [rW_t - X_t - \eta'\lambda(a'W_t - X_t)]. \end{aligned}$$

This requires that

$$[rW_t - X_t - \eta'\lambda(a'W_t - X_t)] = K_5 \cdot (W_t - \eta' X_t)$$

holds, where K_5 is some constant, so that

$$\eta' = \frac{1 - \eta'\lambda}{r - a'\eta'\lambda}.$$

Equivalently, we have

$$\begin{aligned} 0 &= \eta'^2 a' \lambda - \eta'(\lambda + r) + 1 \\ \eta' &= \frac{(\lambda + r) \pm \sqrt{(\lambda + r)^2 - 4a'\lambda}}{2a'\lambda} \end{aligned}$$

and a solution exists as long as

$$\lambda^2 + (2r - 4a')\lambda + r^2 > 0.$$

When $\lambda = 0$, the correct solution has to be $\eta' = 1/r$; this is the one that comes from the smaller root of the quadratic equation. By continuity, we should always select the smaller root.

When there is a solution, we have

$$K_5 = r - a'\eta'\lambda = \frac{1 - \eta'\lambda}{\eta'}.$$

Going back to the Bellman equation, we have

$$(W_t - \eta' X_t)^{1-\gamma} \left\{ \varphi' \left[\frac{\rho}{1-\gamma} - \frac{1}{2} \frac{\pi^2}{\gamma\sigma^2} \right] - \varphi'^{\frac{\gamma-1}{\gamma}} \frac{\gamma}{1-\gamma} \right\} = \varphi' K_5 (W_t - \eta' X_t)^{1-\gamma}.$$

or equivalently

$$\varphi' \left[\frac{\rho}{1-\gamma} - \frac{1}{2} \frac{\pi^2}{\gamma\sigma^2} - K_5 \right] - \varphi'^{\frac{\gamma-1}{\gamma}} \frac{\gamma}{1-\gamma} = 0$$

so that

$$\varphi'^{\frac{1}{\gamma}} = \frac{\gamma/(1-\gamma)}{\frac{\rho}{1-\gamma} - \frac{1}{2} \frac{\pi^2}{\gamma\sigma^2} - r + a'\eta'\lambda}$$

which has to be a positive number, in order for a solution to exist. This is condition (19).

Just like in the proof of Theorem 1, we can derive the wealth dynamic

$$\frac{d(W_t - \eta' X_t)}{W_t - \eta' X_t} = \left[r - a'\eta'\lambda + \frac{\pi^2}{\gamma\sigma^2} - \varphi'^{-\frac{1}{\gamma}} \right] dt + \frac{\pi}{\gamma\sigma} dz$$

and therefore the consumption dynamic

$$\frac{dC_t}{C_t} = \{A - Z_t [\lambda - a' \eta' \lambda + A]\} dt + \frac{\pi}{\gamma \sigma} (1 - Z_t) dz.$$

From these we can show that

$$dZ_t = (1 - Z_t) \left\{ \varphi'^{\frac{1}{\gamma}} \lambda a' + Z_t \left[\eta' \lambda a' - \lambda - A + \left(\frac{\pi}{\gamma \sigma} \right)^2 (1 - Z_t) \right] \right\} dt - \frac{\pi}{\gamma \sigma} Z_t (1 - Z_t) dz.$$

Define $Y_t = Z_t / (1 - Z_t)$. The dynamics of Y_t can be shown to be

$$dY_t = \left\{ \varphi'^{\frac{1}{\gamma}} \lambda a' + Y_t \left[\varphi'^{\frac{1}{\gamma}} \lambda a' + \eta' \lambda a' - \lambda - A + \left(\frac{\pi}{\gamma \sigma} \right)^2 \right] \right\} dt - \frac{\pi}{\gamma \sigma} Y_t dz.$$

The Kolmogorov forward equation for the distribution $p(Y)$ of this process is the following:

$$\begin{aligned} \frac{1}{2} \frac{d^2}{dY_t^2} \left[\left(\frac{\pi}{\gamma \sigma} Y_t \right)^2 p(Y_t) \right] - \frac{d}{dY_t} \left[\left\{ \varphi'^{\frac{1}{\gamma}} \lambda a' + Y_t \left[\varphi'^{\frac{1}{\gamma}} \lambda a' + \eta \lambda a' - \lambda - A + \left(\frac{\pi}{\gamma \sigma} \right)^2 \right] \right\} p(Y_t) \right] \\ = \frac{dp(Y_t)}{dt}. \end{aligned}$$

Applying the technique developed by Wong (1964), the stationary distribution is the solution to the following Pearson equation

$$\frac{1}{2} \frac{d}{dY_t} \left[\left(\frac{\pi}{\gamma \sigma} Y_t \right)^2 p(Y_t) \right] = \left\{ \varphi'^{\frac{1}{\gamma}} \lambda a' + Y_t \left[\varphi'^{\frac{1}{\gamma}} \lambda a' + \eta \lambda a' - \lambda - A + \left(\frac{\pi}{\gamma \sigma} \right)^2 \right] \right\} p(Y_t)$$

subject to the normalization

$$\int_0^\infty p(Y) dY = 1.$$

Rewriting the ODE gives

$$\frac{1}{2} \left(\frac{\pi}{\gamma \sigma} Y_t \right)^2 p'(Y_t) = \left\{ \varphi'^{\frac{1}{\gamma}} \lambda a' + Y_t \left[\varphi'^{\frac{1}{\gamma}} \lambda a' + \eta \lambda a' - \lambda - A \right] \right\} p(Y_t).$$

The solution is given by

$$p(Y) = K_6 \exp \left\{ - \frac{2 \varphi'^{\frac{1}{\gamma}} \lambda a'}{\left(\frac{\pi}{\gamma \sigma} \right)^2 Y} \right\} \cdot Y^{2 \left[\varphi'^{\frac{1}{\gamma}} \lambda a' + \eta \lambda a' - \lambda - A \right] / \left(\frac{\pi}{\gamma \sigma} \right)^2}$$

where the constant K_6 is

$$K_6^{-1} = \left(\frac{2 \varphi'^{\frac{1}{\gamma}} \lambda a'}{\left(\frac{\pi}{\gamma \sigma} \right)^2} \right)^{1+2 \left[\varphi'^{\frac{1}{\gamma}} \lambda a' + \eta \lambda a' - \lambda - A \right] / \left(\frac{\pi}{\gamma \sigma} \right)^2} \Gamma \left[-2 \left[\varphi'^{\frac{1}{\gamma}} \lambda a' + \eta \lambda a' - \lambda - A \right] / \left(\frac{\pi}{\gamma \sigma} \right)^2 - 1 \right]$$

where $\Gamma(\cdot)$ is the gamma function. Then the density of $Z = Y / (1 + Y)$ is

$$p_Z(Z) = \frac{1}{(1 - Z)^2} p_Y\left(\frac{Z}{1 - Z}\right)$$

which gives

$$p_Z(Z) = K_6 \exp \left\{ 2\varphi'^{\frac{1}{\gamma}} \lambda a' \left(\frac{\pi}{\gamma\sigma} \right)^{-2} \right\} (1-Z)^{-2 \left[\varphi'^{\frac{1}{\gamma}} \lambda a' + \eta \lambda a' - \lambda - A + \left(\frac{\pi}{\gamma\sigma} \right)^2 \right] / \left(\frac{\pi}{\gamma\sigma} \right)^2} \\ \cdot Z^{2 \left[\varphi'^{\frac{1}{\gamma}} \lambda a' + \eta \lambda a' - \lambda - A \right] / \left(\frac{\pi}{\gamma\sigma} \right)^2} \exp \left\{ - \frac{2\varphi'^{\frac{1}{\gamma}} \lambda a'}{\left(\frac{\pi}{\gamma\sigma} \right)^2 Z} \right\}.$$

Proof of Proposition 5

Define

$$\tilde{X}_t = X_0 e^{-dt} + D \int_0^t e^{-du} C_{t-u} du$$

with

$$d = -\lambda \left[a' \eta' - a' \varphi'^{\frac{1}{\gamma}} - 1 \right]$$

and

$$D = a' \lambda \varphi'^{\frac{1}{\gamma}}.$$

Then

$$\frac{d\tilde{X}_t}{dt} = -dX_0 e^{-dt} + D \cdot \frac{d}{dt} \int_0^t e^{-d(t-u)} C_u du = -dX_0 e^{-\omega t} + D \cdot \left[C_t - d \int_0^t e^{-d(t-u)} C_u du \right] = \\ = DC_t - d\tilde{X}_t.$$

Using the consumption rule and the formulas for d and D we can write

$$DC_t - d\tilde{X}_t = D \left(X_t + \varphi'^{-\frac{1}{\gamma}} (W_t - \eta' X_t) \right) - d\tilde{X}_t = D\varphi'^{-\frac{1}{\gamma}} W_t + (D - d - D\varphi'^{-\frac{1}{\gamma}} \eta') X_t + d(X_t - \tilde{X}_t) = \\ = a' \lambda \varphi'^{\frac{1}{\gamma}} \varphi'^{-\frac{1}{\gamma}} W_t + (a' \lambda \varphi'^{\frac{1}{\gamma}} + \lambda \left[a' \eta' - a' \varphi'^{\frac{1}{\gamma}} - 1 \right] - a' \lambda \varphi'^{\frac{1}{\gamma}} \varphi'^{-\frac{1}{\gamma}} \eta') X_t + d(X_t - \tilde{X}_t) = \\ = \lambda [a' W_t - X_t] + d(X_t - \tilde{X}_t).$$

Because $\tilde{X}_0 = X_0$ and $dX_t/dt = \lambda [a' W_t - X_t]$, the above equation shows that $d\tilde{X}_t/dt = dX_t/dt$. Thus the two processes are identical with probability one. Equivalently, we can write X_t in the form stated in the proposition. To ensure that $d > 0$ we require condition (26) to hold.

References

- Ait-Sahalia, Y., Parker, J. A. and Yogo, M. (2003). Luxury goods and the equity premium, *Journal of Finance*. Forthcoming.
- Becker, G. (1973). A theory of marriage: Part i., *Journal of Political Economy* **81**: 813–846.

- Bertola, G. and Caballero, R. (1990). Kinked adjustment costs and aggregate dynamics, *NBER Macroeconomics Annual*.
- Blanchard, O. J. and Fischer, S. (1989). *Lectures on Macroeconomics*, The MIT Press, Cambridge, MA and London, UK.
- Calvo, G. (1983). Staggered prices in a utility maximizing framework, *Journal of Monetary Economics* **12**: 383–398.
- Campbell, J. Y. (2002). Consumption-based asset pricing, Forthcoming in the Handbook of Economics and Finance, North-Holland, Amsterdam.
- Campbell, J. Y. and Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* **107**: 205–251.
- Carroll, C. D. (2000). Portfolios of the rich. Manuscript, John Hopkins University.
- Chetty, R. (2002). Consumption commitments, unemployment durations, and local risk aversion. Working paper, Harvard University.
- Cocco, J. (2000). Portfolio choice in the presence of housing. Working paper, London Business School.
- Constantinides, G. (1990). Habit formation: a resolution of the equity premium puzzle, *Journal of Political Economy* **98**: 519–543.
- Constantinides, G. M., Donaldson, J. B. and Mehra, R. (2002). Junior can't borrow, *Quarterly Journal of Economics* **117**: 269–196.
- Davidoff, T. (2003). Labor income, housing prices and homeownership. Working paper, Haas School of Business, UC-Berkeley.
- Flavin, M. and Yamashita, T. (2002). Owner-occupied housing and the composition of the household portfolio, *American Economic Review* **92**: 345–362.
- Fratantoni, M. C. (1998). Homeownership and investment in risky assets, *Journal of Urban Economics* **44**: 27–42.

- Fratantoni, M. C. (2001). Homeownership, committed expenditure risk, and the stock-holding puzzle, *Oxford Economic Papers* **53**: 241–259.
- Gabaix, X. and Laibson, D. (2001). The 6D bias and the equity premium puzzle, *NBER Macroeconomics Annual* .
- Gollier, C. (2001). *The Economics of Risk and Time*, The MIT Press, Cambridge, MA and London, England.
- Grossman, S. J. and Laroque, G. (1990). Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods, *Econometrica* **58**(1): 25–51.
- Grossman, S. J. and Shiller, R. J. (1982). Consumption correlatedness and risk measurement in economies with non-traded assets and heterogeneous information, *Journal of Financial Economics* **10**: 195–210.
- Guiso, L., Japelli, T. and Terlizzese, D. (1996). Income risk, borrowing constraints and portfolio choice, *American Economic Review* **86**: 158–172.
- Imbens, G. W. and Angrist, J. D. (1994). Identification and estimation of local average treatment effects, *Econometrica* **62**: 467–475.
- Keeley, M. C. (1979). An analysis of the age pattern of first marriage, *International Economic Review* **20**: 527–544.
- Kiernan, K. E. and Eldridge, S. M. (1987). Age at marriage: Inter and intra cohort variation, *The British Journal of Sociology* **38**: 44–65.
- Kocherlakota, N. (1996). The equity premium: It’s still a puzzle, *Journal of Economic Literature* **34**: 42–71.
- Koren, M. and Szeidl, A. (2003). Portfolio choice with illiquid assets. CEPR Working Paper 3795.
- Kullmann, C. and Siegel, S. (2002). Real estate and its role in asset allocation. CUER Working paper, University of British Columbia.
- Li, R. (2003). Household portfolio choice with precommitment. Working paper, University of Toronto.

- Lustig, H. and Nieuwerburgh, S. V. (2003). Housing collateral, consumption insurance and risk premia. Working paper, University of Chicago and NYU Stern.
- Marshall, D. A. and Parekh, N. G. (1999). Can costs of consumption adjustment explain asset pricing puzzles?, *Journal of Finance* pp. 623–654.
- Mehra, R. and Prescott, E. (1985). The equity premium: A puzzle, *Journal of Monetary Economics* **15**: 145–161.
- Ogaki, M. and Zhang, Q. (2001). Decreasing risk aversion and tests of risk sharing, *Econometrica* **69**: 515–526.
- Olney, M. (1999). Avoiding default: The role of credit in the consumption collapse of 1930, *Quarterly Journal of Economics* pp. 319–335.
- Piazzesi, M., Schneider, M. and Tuzel, S. (2003). Housing, consumption and asset pricing. Working paper, University of Chicago.
- Reis, R. (2003). Inattentive consumers. Working paper, Harvard University.
- Sinai, T. and Souleles, N. S. (2003). Owner-occupied housing as a hedge against risk. Working Paper, The Wharton School, University of Pennsylvania.
- Taylor, J. (1979). Staggered wage setting in a macro model, *American Economic Review* **69**: 108–113.
- Weil, P. (1989). The equity premium puzzle and the risk-free rate puzzle, *Journal of Monetary Economics* **24**: 401–421.
- Wong, E. (1964). The construction of a class of stationary markoff processes, *Proceedings of Symposia in Applied Mathematics* **16**. Providence, R.I.:American Mathematical Society.
- Yamashita, T. (2003). Owner-occupied housing and investment in stocks, *Journal of Urban Economics* **53**: 220–237.
- Yogo, M. (2003). A consumption-based explanation of expected stock returns. Working paper, Harvard University.

TABLE 1
MEAN HOUSEHOLD EXPENDITURE SHARES

Category	Income Group			Overall Mean
	\$30-40K	\$50-70K	>\$70K	
Shelter	0.20	0.19	0.21	0.20
Transport (excluding gas and maint)	0.16	0.17	0.16	0.16
Utilities, fuels, and public services	0.07	0.07	0.06	0.07
Health care	0.06	0.05	0.04	0.06
Education	0.01	0.02	0.02	0.02
Food	0.16	0.15	0.14	0.15
Apparel	0.05	0.05	0.06	0.15
Household supplies and furniture	0.06	0.08	0.10	0.06
Entertainment	0.05	0.06	0.06	0.05
Miscellaneous	0.02	0.02	0.02	0.02
Cash contributions	0.03	0.03	0.05	0.04
Mean annual expenditure	\$15,369	\$32,609	\$64,134	\$35,930
Mean take-home pay	\$6,858	\$29,720	\$90,748	\$41,531

^aSource: BLS tabulations from Consumer Expenditure Survey, 2000

^bTake-home pay defined as gross income net of taxes and mandatory insurance/pension contributions

TABLE 2
CALIBRATION OF THE COMMITMENTS MODEL
TO EQUITY PREMIUM AND CONSUMPTION DYNAMICS

Parameters ^a	Specification					
	(1)	(2)	(3)	(4)	(5)	(6)
λ (adjustment horizon)	0.33	1	.33	.33	.2	.2
α' (pref for commit)	.0315	.033	.0149	.011	.0155	.01
γ (risk aversion)	2	2	4	5	4	5
ρ (time preference)	.007	.0072	.005	.005	.005	.005
Mean Cons. Growth (true value=.0183)	.0182	.0182	.0146	.0127	.0135	.0122
S.D. Cons. Growth (true value=.0357)	.0358	.0359	.036	.036	.0356	.0351
Mean Commit. Share	.8063	.8036	.6054	.506	.6148	.4504

^aOther parameter values are the riskfree rate $r = .01$, the equity premium $\pi = 0.06$ and the annual standard deviation of stock returns $\sigma = 0.165$. All parameters are measured per year.

TABLE 3
DESCRIPTIVE STATISTICS FOR EVER-MARRIED STOCKHOLDING HOMEOWNERS^a

Variable ^b	Mean	Median	Standard Deviation
Household net worth ^c	248,361.40	163,058.20	396,541.90
Home equity	83,510.68	65,808.16	68,895.05
Home mortgage outstanding	40,828.15	24,206.94	48,586.63
Liquid wealth ^d	123,795.20	52,896.43	351,091.80
Stocks	64,837.94	12,864.17	310,008.40
Savings accounts, CDs, and money market	19,082.17	6,432.09	29,405.59
Municipal and corporate bonds	12,215.43	0.00	53,260.76
IRA wealth	17,099.09	1,356.75	34,740.86
Other assets ^e	10,560.54	499.81	86,807.28
Other real estate	24,427.70	0.00	69,579.90
Business equity	10,525.37	0.00	55,687.11
Vehicle equity	10,022.04	8,440.80	8,548.38
Unsecured debt	3,919.61	416.51	25,177.12
Annual income	55,226.93	46,562.19	43,435.50
Years of education	14.13	14.00	2.75
Age	53.59	52.00	14.91
Age at marriage	27.07	24.00	9.18
1 = currently married	0.78	1.00	0.42
Children at home	0.41	0.00	0.84

^aSource: Survey of Income and Program Participation 1990-96 Asset Topical Modules. *N*=15,297

^bAll variables are in 1990 dollars

^cNet worth is total household wealth minus unsecured debt

^dLiquid wealth is total wealth minus home equity, other real estate, business equity, and vehicle equity

^eOther assets comprise money owed to respondent, savings bonds, checking accounts, and equity in other investments

TABLE 4
FIRST STAGE AND REDUCED-FORM REGRESSIONS FOR AGE-AT-MARRIAGE INSTRUMENT

	(1)	(2)	(3)	(4)	(5)	(6)
	Recent-Married (<25 years ago)			Long-Married		Renters
	First-stage	Reduced-form	Reduced-form	First-stage	Reduced-form	Reduced-form
Dependent var:	mortgage	stocks	stocks	mortgage	stocks	stocks
age at marriage	481.451 (93.925)	-645.817 (222.638)	-632.167 (275.741)	86.769 (52.733)	253.221 (203.641)	250.935 (231.245)
liquid wealth		0.758 (0.097)				
age		-539.242 (512.677)				
other real estate	0.032 (0.011)		-0.050 (0.059)	0.027 (0.007)	-0.074 (0.023)	-0.095 (0.062)
business equity	0.007 (0.013)		-0.010 (0.039)	-0.005 (0.011)	0.012 (0.023)	-0.034 (0.028)
vehicle equity	-0.048 (0.073)		-0.279 (0.269)	-0.098 (0.053)	-0.274 (0.153)	0.253 (0.276)
unsecured debt	0.104 (0.044)		-0.180 (0.217)	0.017 (0.044)	0.019 (0.017)	0.095 (0.061)
annual income	0.262 (0.021)		-0.092 (0.060)	0.213 (0.021)	-0.290 (0.049)	-0.157 (0.075)
years of schooling	3,149.289 (262.629)		-944.527 (482.985)	1,152.793 (145.136)	-1,516.937 (370.781)	742.053 (559.834)
Sample Size	7565	7565	7565	7369	7369	2037

NOTE-Huber-White standard errors are reported in parentheses. Columns 1-3 include stockholders who own homes and got married less than 25 years ago. Columns 4 and 5 include homeownership stockholders who were married more than 25 years ago. Column 6 includes married renters who own stocks. Columns 1, 3-6 include 10-piece linear splines for liquid wealth, home equity, and age. These columns also include year, occupation and industry dummies.

TABLE 5
HOUSING COMMITMENT AND PORTFOLIO ALLOCATION
INSTRUMENTAL VARIABLES ESTIMATES

	(1)	(2)	(3)	(4)	(5)	(6)
	Base case	With Controls	With Controls	Bond Participation	Portfolio Shares	Shares Wealthy only
Dependent variable:	stocks	stocks	bonds	bondholder*	stk share	stk share
mortgage debt	-0.950 (0.249)	-0.738 (0.344)	0.527 (0.224)	0.036 (0.012)		
mortgage debt share					-0.071 (0.035)	-0.446 (0.196)
liquid wealth	0.839 (0.043)					
age	-2464.354 (456.606)					
other real estate		-0.064 (0.025)	-0.001 (0.011)	0.000 (0.000)		
business equity		0.003 (0.025)	-0.013 (0.010)	-0.001 (0.001)		
vehicle equity		-0.356 (0.161)	0.173 (0.073)	0.015 (0.005)		
unsecured debt		-0.009 (0.053)	-0.047 (0.032)	-0.001 (0.003)		
years of schooling		430.888 (795.225)	-249.5 (499.741)	0.010 (0.002)	-0.000 (0.001)	0.003 (0.001)
annual income		0.009 (0.100)	-0.051 (0.058)	-0.005 (0.003)	-0.000 (0.000)	-0.000 (0.000)
other real estate share					-0.459 (0.146)	-0.273 (0.230)
business equity share					-0.459 (0.146)	-0.281 (0.227)
vehicle equity share					-0.346 (0.166)	-0.248 (0.285)
unsecured debt share					0.046 (0.023)	0.143 (0.058)
Sample size	15297	15297	15297	15297	15206	8280

NOTE-For scaling purposes, coefficients in column 4 are for a \$10,000 increase in monetary variables. Huber-White standard errors are reported in parentheses. Columns 2-6 include industry, occupation and year dummies. Columns 2-4 have a 10-piece linear spline for liquid wealth, age, and home equity. Columns 5 and 6 have splines for age and shares of liquid wealth and home equity. Shares are defined as level divided by total wealth. Column 5 excludes observations with total wealth below \$1000; column 6 those with wealth below \$150,000. All columns use age at marriage to instrument for mortgage debt.

TABLE 6
COMMITMENTS VS. HOUSING RISK

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	OLS		Birthplace IV		IV
	House size	Homeown	Renters	Homeown	Renters	Horse race
Dependent variable:	propty value	stocks	stocks	stocks	stocks	stocks
σ_{hg} = MSA stdev house price	-4,112.332	312.063	-2,812.232	689.550	-8,520.640	
growth	(1,608.475)	(598.978)	(1,078.113)	(1,779.665)	(5,278.069)	
mortgage debt						-1.125
						(0.568)
σ_H = MSA mortgage risk						-0.034
						(0.019)
MSA avg house price growth	35,777.740	-2,841.184	339.894	-7,651.679	2,475.511	13,032.296
	(9,477.428)	(2,994.141)	(5,201.115)	(15,592.905)	(31,261.779)	(9,705.645)
MSA avg house price growth ²	-1,679.179	920.679	823.117	2,191.792	1,331.480	1,225.211
	(1,887.182)	(779.864)	(1,427.407)	(3,939.442)	(7,838.833)	(1,656.425)
other real estate	0.102	-0.080	-0.096	-0.072	-0.109	-0.041
	(0.012)	(0.029)	(0.143)	(0.032)	(0.057)	(0.038)
business equity	0.061	0.018	-0.021	0.001	-0.033	0.056
	(0.019)	(0.019)	(0.035)	(0.014)	(0.053)	(0.033)
vehicle equity	0.796	-0.307	0.316	-0.370	0.285	-0.242
	(0.113)	(0.207)	(0.398)	(0.297)	(0.484)	(0.211)
unsecured debt	0.007	-0.051	0.149	-0.171	0.211	-0.021
	(0.038)	(0.025)	(0.082)	(0.202)	(0.119)	(0.075)
years of schooling	4,413.739	-1,163.990	1,009.525	-1,315.577	1,511.899	1,735.308
	(212.063)	(445.766)	(959.259)	(515.083)	(1,116.449)	(1,629.492)
annual income	0.370	-0.180	-0.228	-0.163	-0.244	0.141
	(0.026)	(0.060)	(0.112)	(0.051)	(0.116)	(0.158)
Sample size	9465	9560	1322	8440	1154	9560

NOTE-Huber-White standard errors, reported in parentheses, are clustered by MSA in 1-3,6 and by birth state in 4-5. All columns include splines for age and liquid wealth and industry, occupation, and year dummies. Column 6 includes a spline for home equity. Columns 1-3 are OLS. In columns 4-5, MSA vars are instrumented using corresponding state of birth variables as described in text. In column 6, mortgage debt and mortgage risk are instrumented using σ_{hg} , age at marriage, and σ_{hg} *age at mar, while controlling for home equity risk. Mortgage risk is product of mortgage and σ_{hg} , home equity risk is product of home equity and σ_{hg} .