

Consumption Commitments: Neoclassical Foundations for Habit Formation*

Raj Chetty Adam Szeidl
UC-Berkeley and NBER UC-Berkeley

November 2004

Abstract

This paper studies consumption and portfolio choice in a model where agents have neoclassical preferences over two consumption goods, one of which involves a commitment in that its consumption can only be adjusted infrequently. Aggregating over a population of such agents implies dynamics identical to those of a representative consumer economy with habit formation utility. In particular, aggregate consumption is a slow-moving average of past consumption levels, and risk aversion is amplified because the marginal utility of wealth is determined by excess consumption over the prior commitment level. We test the model's prediction that commitments amplify risk aversion by using home tenure (years spent in current house) as a proxy for commitment: Recent home purchasers are unlikely to move in the near future, and are therefore more constrained by their housing commitment. We use a set of control groups to establish that the timing of marital shocks such as marriage and divorce can be used to create exogenous variation in home tenure conditional on age and wealth. Using these marital shocks as instruments, we find that the average investor reallocates \$1,500 from safe assets to stocks per year in a house. Hence, recent home purchasers have highly amplified risk aversion, suggesting that real commitments are a quantitatively powerful source of habit-like behavior.

JEL Classification: *D81, E21, G11, G12*

Keywords: *Portfolio choice, Housing, Risk Aversion, Microfoundations*

*E-mail addresses: chetty@econ.berkeley.edu and szeidl@econ.berkeley.edu. We would like to thank Gary Chamberlain, Tom Davidoff, Ed Glaeser, Elhanan Helpman, Larry Katz, Miklos Koren, Hanno Lustig, Jonah Rockoff, Emmanuel Saez, Jesse Shapiro, Todd Sinai, Jeremy Stein, Moto Yogo, numerous seminar participants, 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

Representative consumer habit formation models have attracted much interest in many areas of economics. In asset pricing, Sundaresan (1989), Constantinides (1990) and Campbell and Cochrane (1999) use habit formation to understand the mean equity premium and the joint dynamics of equity returns and aggregate consumption. In monetary economics, a growing literature makes use of habit preferences to match the hump-shaped response of consumption to interest-rate shocks, including Fuhrer (2000), Christiano, Eichenbaum and Evans (2004) and McCallum and Nelson (1999). In other applications, Boldrin, Christiano and Fisher (2001) demonstrate that habit preferences can help match a variety of business cycle facts, including the excess sensitivity of consumption. Carroll, Overland and Weil (2000) show that habit sheds light on the joint behavior of savings and growth. Ljungqvist and Uhlig (2000) show that optimal fiscal policies look very different with habit-forming consumers.

Two important features of habit formation preferences distinguish these new theories from previous neoclassical models. First, habit affects risk preferences. A higher level of habit implies that marginal utility is more sensitive to consumption shocks, making the representative consumer more risk averse. Second, habit is a slowly adjusting state variable. Shocks can therefore have lasting effects on future consumption.

Despite the widespread application and success of habit formation, the exact source of “habit” has been elusive. The psychological notion of habit, though intuitive in certain cases, is difficult to pin down in data. Hence, direct empirical evidence for the presence of habit formation is limited (see section 2 for a review of this literature).

This paper identifies a new, non-psychological source of “habit” both theoretically and empirically. We show that habit-like behavior can arise from the rigidities inherent in the consumption of certain goods, such as housing, cars, and contracted services such as college tuition bills. All of these goods involve a “commitment,” since their consumption often cannot be changed in the short run by even a small amount. The main theoretical result of the paper is that aggregating an economy of heterogeneous households with standard, non-habit preferences over such commitment goods yields aggregate dynamics that coincide precisely with those of existing representative-consumer habit models. In this sense, consumption commitments provide neoclassical micro-foundations for habit formation.¹

The intuition underlying the connection between commitments and habit can be seen with an example. Consider an individual who consumes commitment (housing) and non-commitment

¹The qualitative similarity between adjustment costs and habit models has been observed by Dybvig (1995), Flavin (2001), Fratantoni (2001), Li (2003), and Flavin and Nakagawa (2003).

(food or entertainment), goods in equal shares. After making his commitment, suppose the individual faces a shock that necessitates a 10% reduction in expenditure. If the agent cannot move out of his house immediately, he must cut discretionary consumption by 20%, raising his marginal utility of wealth more sharply in the short run. Hence, commitments, like habit, magnify the impact of shocks on marginal utility and effectively make the consumer more risk averse. In the long run, the individual may sell his house and move into a smaller home. Commitment consumption thus acts as a state variable that adjusts with a lag, making shocks have lasting impacts on consumption. These examples illustrate that the commitment model shares the two central features of the habit models identified above.

We formalize the connection between the two theories by first modelling the portfolio choice and consumption decision of an individual who has neoclassical preferences over the two types of consumption goods. To make the model analytically tractable, we assume that consumers have separable, constant elasticity of substitution preferences. We model commitments by a time-dependent adjustment framework intended to capture the fact that households typically move at certain junctures of their life-cycles, such as marriage, child birth, retirement, and job relocations.² We consider two types of time-dependent adjustment rules to model these moves. The first, which is similar to Taylor pricing, allows the consumer to reset consumption of the commitment good after a fixed number of periods. The second, which generalizes the idea of Calvo pricing, permits a stochastic, partially forecastable reset date, but requires a stronger complete markets assumption.³ The main theoretical results are similar in both models. The main aggregation theorem extends to more general environments with arbitrary separable utility and an adjustment rule that is both time and state dependent, but an analytic expression for habit dynamics is not available in the more general case.

We first show that individuals who have more commitments act as if they are more risk averse by investing less in risky assets. We then consider a population of such individuals and derive aggregate dynamics. These dynamics are identical to those that arise from a general equilibrium representative consumer model with habit formation utility. 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. These aggregation results stem from the fact that in habit formation models, consumer well-being is determined by surplus consumption over current habit,

²Merrill (1984), Venti and Wise (1989) and Venti and Wise (2000) document that such shocks are the most important determinants of mobility.

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

which is a slow-moving time average of past consumption levels. In our model, commitments are slow-moving at the microeconomic level in that they are changed only at adjustments. When an adjustment takes place, the new commitment level reflects current prosperity, and therefore recent consumption levels. Consequently, the aggregate commitment stock is a slow-moving function of past aggregate consumption levels. Summing the individual Euler equations yields an aggregate Euler equation where utility is determined by surplus consumption over the level of commitments, exactly as in the habit model.

The central benefit of these microfoundations is that they yield a set of predictions that can be tested with micro data on consumption and portfolio choices. The most substantive predictions relate to the two key features of habit-formation preferences: (1) Commitments should make agents adjust consumption – especially of more committed goods – slowly in response to shocks and (2) Commitments should drive up risk aversion. Many studies (described in section 4) have already shown that consumption of goods such as housing adjusts slowly at the microeconomic level (i.e., that people do not move often). We therefore focus on testing the second hypothesis here.

In our empirical analysis, we focus on housing as the commitment good and examine its effects on portfolio allocations using data from the Survey of Income and Program Participation (SIPP), a nationally representative survey of households. The model implies that the number of years a family has spent in its current home (home tenure) proxies for commitment. This is because those who recently bought a home are unlikely to move in the near future, both in theory and in practice, and are therefore more committed to the stream of financial obligations associated with the house. Since they are more constrained in their ability to adjust their consumption bundle in response to financial shocks, recent home buyers should seek lower risk exposure. Following this intuition, we derive an estimating equation relating stockholding to home tenure and test whether home tenure has a positive causal effect on portfolio risk.

The primary difficulty in testing this prediction is that home tenure is correlated with other characteristics that have a direct effect on portfolio choice. For example, more educated people tend to be more mobile but may also have different risk attitudes, e.g., because they have less background labor income risk. We employ an instrumental variables strategy to overcome this problem. We identify a set of move-inducing shocks that might not have a direct effect on portfolio choice, conditional on appropriate controls. The shocks we examine as candidate instruments are all related to the timing of changes in marital status, holding fixed current wealth and age: age at first marriage, age at first divorce or spouse’s death, and an indicator variable for remarriage. Consistent with prior studies of the determinants of mobility, the data reveal that these marital

shocks frequently induce moves, and are therefore powerful predictors of home tenure.

Of course, one worries that these marital shocks are directly related to portfolio choice. We test whether this orthogonality condition holds by examining several “control” groups where our theory predicts that the first stage relationship between the timing of marital shocks and home tenure should break down: 1) Long-married homeowners, for whom the relevant marital shock happened many years ago, and has no effect on tenure in the current home; 2) Married renters, whose current home (apartment) tenure has little to do with marital shocks because renting is not as much of a long-term commitment as owning; and 3) Unmarried divorced homeowners, for whom the date of first marriage termination does not predict current home tenure, because moves are typically induced only by remarriage. The lack of a first stage relationship in these groups, both in theory and in the data, means that any association between a marital shock and portfolio choice would imply that the exclusion restriction is violated. Fortunately, reduced form regressions of stockholding on each marital shock instrument reveal no such association for any of the candidate instruments in any of the control groups, strongly supporting the exogeneity of the instrument set.

Using the marital shocks as instruments, we find that a one year exogenous increase in home tenure causes a highly statistically significant \$1,500 reallocation from safe assets (such as bonds and savings accounts) to stocks for the average stockholding household, holding fixed total wealth. The estimates are similar irrespective of which of the three marital shocks are used, indicating that the instrument set easily passes standard overidentification criteria. The estimates are also highly robust to the inclusion of various controls and specification checks, including a regression on the full sample of all 55,000 married households in the data. The estimates imply that the average stock market participant reallocates 13% (\$15,000) of wealth from safe to risky assets over a span of ten years in a given home, holding all else fixed. The large magnitude of this effect indicates that recent home buyers – for whom the constraints of a long-term housing commitment loom largest – have highly amplified risk aversion. Hence, real economic commitments appear to be a powerful source of habit-like behavior in practice.

The remainder of the paper is organized as follows. The next section reviews related literature. Section 3 develops the model and presents the aggregation results, first with deterministic and then with stochastic adjustment. Section 4 describes our estimation strategy to test the model. Section 5 presents the empirical evidence, beginning with a graphical overview of the main results and then discussing regression estimates. Section 6 addresses other potential biases in the empirical tests. Section 7 concludes.

2 Related Literature

This paper builds on and contributes to three different literatures: (1) work on durable goods consumption and housing in macroeconomics and asset pricing, (2) recent work on commitments and risk preferences, and (3) the literature on habit formation.

The importance of durable goods for asset pricing was shown in a seminal paper by Grossman and Laroque (1990). They showed that transaction costs in a model with a single durable good can reduce the level of consumption volatility implied by the model. Flavin (2001) analyzed asset pricing in a generalization of the Grossman-Laroque model with two goods, one durable and one non-durable, and observed that this model has properties qualitatively similar to habit formation. For instance, the two-good model generates sluggish adjustment and a coefficient of relative risk aversion that is not directly related to the elasticity of intertemporal substitution. Fratantoni (2001) and Li (2003) analyze portfolio choice in a two-good model with adjustment costs and also note its similarity to habit. In related work, Dybvig (1995) examines ratcheting consumption demand in a model with extreme habit persistence, and remarks informally that such preferences could be motivated by pre-commitment in consumption.

Our approach differs from and complements these papers in two ways. First, we present an aggregation theorem showing the *formal equivalence* between habit formation in total aggregate consumption and commitments at the household level. Hence, commitments are not merely similar to habit, but could constitute the microeconomic foundations of aggregate habit. Second, we directly test for the presence of these microfoundations by estimating the causal links between commitments, risk preferences, and habit-like behavior using household data.

Several other recent studies have also emphasized the importance of durable goods in understanding asset pricing. Flavin and Nakagawa (2003) extend Flavin (2001) and show that their durable goods model outperforms household-level pure habit and neoclassical models in matching the dynamics of food consumption and asset returns. The implication of commitment models that consumption adjusts with a lag is also supported by Parker and Julliard (2004) and Jagannathan and Wang (2004), who show that consumption growth over multiple quarters is more successful in pricing excess returns. Lustig and Nieuwerburgh (2003), Piazzesi, Schneider and Tuzel (2003), and Yogo (2003) explore other channels through which durables affect asset pricing such as composition risk and collateral concerns. All of these ideas are relevant independently of the commitment feature of durables emphasized here.

Risk preferences in the two-good commitments model have also received attention in other recent papers. Chetty (2003) 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. Shore and Sinai (2004) extend this moderate-stake risk aversion idea to show that households subject to exogenously higher income risk optimally undertake larger housing commitments. Postlewaite, Samuelson and Silverman (2004) show that commitments generate incentives to bunch uninsured risks together, potentially explaining real wage rigidities. Olney (1999) gives historical evidence of the macroeconomic importance of commitments by arguing that large exposure to installment finance 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 paper contributes to the large literature on habit formation discussed in the introduction. As we noted there, empirical evidence on habit formation is mixed. Ferson and Constantinides (1991) find large and statistically significant amounts of habit formation in aggregate (NIPA) monthly, quarterly as well as annual consumption. However, Eichenbaum, Hansen and Singleton (1988) and Heaton (1993) use different methodologies and find little evidence for habit formation in aggregate monthly consumption data. Shea (1994) points out that aggregate food consumption appears to exhibit much less evidence of “excess-smoothness” than other, more broadly defined categories of consumption. Consistent with this finding, at the micro level, Dynan (2000) and Flavin and Nakagawa (2003) find no evidence for habit formation in food consumption in the Panel Study of Income Dynamics, and Lupton (2003) and Ravina (2004) find stronger evidence of habit in more broadly defined measures of consumption.

Although the results remain controversial, taken together these studies suggest that there is stronger evidence of habit in NIPA nondurables and services than in the more narrowly-defined food category. Importantly, the definition of nondurables and services in the aggregate NIPA data is very broad: In 2004, nondurables and services accounted for more than 85% of total consumption, and excluded only vehicle and furniture purchases. As a result, this commonly used consumption measure *includes* several commitment and durable goods such as housing services. The finding that food consumption behaves differently from more broadly defined measures of consumption that include commitment goods is consistent with the implication of our theory, which predicts less habit persistence in immediately-adjustable consumption categories relative to other types of consumption.⁴ Testing more systematically how the degree of adjustability of consumption goods correlates with their habit-like dynamics in aggregate data is a useful direction for future work.

⁴The simple model of neoclassical preferences in this paper predicts no habit in non-commitment goods. In a more general analysis that permits habit preferences, the model would generate some habit-like behavior in non-commitment goods along with stronger habit-like behavior in commitment goods.

3 Consumption and Portfolio Choice with Commitments

3.1 A Model of the Household

Most existing models of consumption assume that agents consume one good, implicitly requiring that all the components of a consumption bundle can be adjusted within a given period. In practice, however, it is difficult to adjust the consumption of many goods in the short run. Leading examples include housing and other durables such as cars and furniture. The infrequency of adjustments in these goods is evident in our data, where the median home tenure among homeowners is 11 years and the median car is six years old. Commitments also extend beyond durable goods. For instance, many services require explicit contracts and impose penalties for early termination – examples include college tuition bills, health insurance, health clubs, cellular phones, and cable television.⁵ Commitments are quantitatively important: Table 1 shows that households across all wealth levels in the US allocate roughly half of net-of-tax income to the infrequently-adjusted goods described above.⁶

To analyze the effect of commitments on consumption and portfolio behavior, consider a consumer with preferences over a commitment good such as housing (x) and non-commitment (discretionary) consumption, such as food or entertainment (f).

The consumer maximizes expected lifetime utility given by

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

where ρ is the discount rate and μ measures the relative preference for commitments. Because utility is time-separable, γ measures the elasticity of intertemporal substitution as well as relative risk aversion for an individual who is free to adjust on both the housing and food margins. Because utility is separable across the two goods, the intertemporal and across-good elasticities of substitution are also equal. These simple preferences make the model analytically tractable, for reasons discussed below.

We model commitments by assuming that 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. This time-dependent adjustment rule is a stylized means of capturing the fact that adjustments in housing and other commitments are often triggered by events such as

⁵Late payments do not eliminate consumption commitments. In the event of an income shock, the bill must still be paid. The ability to make late payments only acts as an additional credit channel.

⁶At least the first five categories in Table 1 could involve commitments. Certain expenditures, such as health care and education, may be committed or discretionary. Since the exact fraction of committed expenditures is not crucial to the subsequent analysis, this definition is left open to interpretation.

marriage, childbirth, graduation, job changes, or retirement. In this section, we assume deterministic adjustment, i.e., that the dates on which these shocks occur are perfectly forecastable.⁷ In Section 3.3, we show that similar results are obtained when the dates of shocks can only be partially forecasted.

Ideally, we would also permit households to change their commitment consumption in the interim periods $(0, T)$ by paying a finite transaction cost. Unfortunately, allowing for such adjustments makes the model analytically intractable. However, several empirical studies of durables and housing consumption (e.g., Attanasio, 2000, Eberly, 1994, Martin, 2003) find very wide (s, S) bands over wealth.⁸ In addition, empirical studies of mobility indicate that changes in family composition or employment are much stronger determinants of home purchases than fluctuations in wealth (e.g., Merrill, 1984, Venti and Wise, 1989, Venti and Wise, 2000). These results suggest that the infinite cost approximation may not be unreasonable.

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

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

where r is the riskfree rate, which is assumed to be constant. The dynamics of the risky asset are given by an 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 wealth held in the risky asset. Define total consumption as $c = f + x$. Normalizing the relative price of commitment and food goods at one, the dynamic budget constraint is

$$dw_t = [(r + \alpha_t \pi) w_t - c_t] dt + \alpha_t w_t \sigma dz_t. \tag{2}$$

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 household's problem is to maximize (1) subject to the reset rule for consumption commitments and the budget constraint (2). 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, $d\Delta = dt$.

⁷Gabaix and Laibson (2001) and Reis (2003) use similar time-dependent rules to model limited attention, while Koren and Szeidl (2003) formalize financial illiquidity in a related fashion.

⁸For example, (S, s) bands estimated by Attanasio imply that households allow the share of car consumption in total consumption to fall to 20% of the desired level before adjusting.

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

$$f_t = \psi_\Delta \cdot (w_t - \Lambda(\Delta)x_t) \quad (3)$$

$$\alpha_t w_t = \frac{\pi}{\gamma\sigma^2} \cdot (w_t - \Lambda(\Delta)x_t) \quad (4)$$

where

$$\Lambda(\Delta) = \frac{1}{r} \left\{ 1 - e^{-r(T-\Delta)} \right\}$$

is a decreasing function of Δ and the expression for ψ_Δ is given in the Appendix.

At a reset date, the consumer chooses $x_t = \kappa \cdot w_t$ where κ depends on the underlying parameters of the model.

The dynamic of food consumption on all dates is characterized by

$$\frac{df_t}{dt} = \left\{ \frac{\pi^2}{2\gamma^2\sigma^2} + \frac{1}{\gamma} \left(\frac{\pi^2}{2\sigma^2} + r - \rho \right) \right\} dt + \frac{\pi}{\gamma\sigma} dz \quad (5)$$

Proof. See the Appendix. ■

To understand this result, 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 is proportional to wealth because at an adjustment, the value function is homogenous of degree $1 - \gamma$ in wealth.

Between reset dates, the model has three state variables, wealth (w), the current level of consumption commitments (x), and time elapsed since the last adjustment (Δ). Conditional on Δ , the optimal consumption and investment rules are linear functions of wealth and commitments. More importantly, the expression $w_t - \Lambda(\Delta)x_t$, which we term “net wealth,” governs the optimal policy of both consumption and investment. 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 therefore allocates an amount corresponding to outstanding commitments ($\Lambda(\Delta)x_t$) to the safe asset, and uses that money exclusively to finance future commitments. After setting this money aside, she then decides how to invest the rest of her wealth (net wealth) in stocks and bonds and how much food to eat. This behavior is very sensible: It reflects the common practice of first paying bills such as a home mortgage payment out of a monthly paycheck, and then deciding how to allocate what is left on discretionary consumption or savings.⁹

⁹A riskless labor income stream can trivially be introduced by changing the definition of wealth to include the present value of labor income; Section 3.4 describes how a risky labor income stream can also be permitted.

The key comparative static that follows from the theorem is that when the value of outstanding commitments $\Lambda(\Delta)x$ is higher, the consumer invests a lower share of her wealth in risky assets, holding all else fixed. The value of $\Lambda(\Delta)x$ can be high for two reasons. The first is when the adjustment date is far away – that is, when Δ is low, implying that $\Lambda(\Delta)$ is high. When a consumer moved in to a new house recently, she expects to live in her current home for a long time, generating a higher total outstanding commitment payment. Second, $\Lambda(\Delta)x$ is high when the commitment level x is high, i.e., when the agent owns a big house.¹⁰

An important implication of formula (5) is that the trajectory of food consumption is fully determined by the rate of return of the stock market. This observation is central to the subsequent aggregation results, since it guarantees that the trajectory of each household’s food consumption does not depend on how far a consumer is from her adjustment date (i.e., on Δ). To understand why it is true, observe that in this economy markets are complete, since the stock market is the only source of risk. As a result, the dynamics of returns pin down the value of the intertemporal marginal rate of substitution (IMRS) for the household. Since commitments are fixed in the short run, the IMRS in turn pins down the growth rate of food consumption. It follows that the trajectory of food consumption is continuous on all dates. This is in contrast with the time path of commitment as well as total consumption, both of which adjust discontinuously every T periods.

3.2 Aggregation

Now consider an economy populated by a continuum of agents with individual reset dates that vary across the population. Normalize the total mass of agents to one. Agents with different adjustment dates may also differ in their wealth levels; however their preference parameters and adjustment horizons are the same.¹¹ Individuals can be organized into cohorts measured by their Δ at date zero, which we label q .¹² The range of cohorts is therefore $0 \leq q < T$. Note that households stay in the same cohort throughout their life: there is no entry and exit.

The aggregate dynamics in the economy depend on the relative size of each cohort, quantified by the amount of food it consumes. Let the share of aggregate food consumption consumed by cohorts between zero and q be $Y(q)$. Then $Y(\cdot)$ is one measure of the relative size of cohorts. For example, if $Y(\cdot)$ is the distribution function of a uniform random variable over the $[0, T]$

¹⁰Note that this effect is purely a consequence of commitment consumption, as opposed to fluctuations in the relative price of commitments (such as house price fluctuations).

¹¹We show that the results are robust to introducing other dimensions of heterogeneity in Section 3.4.

¹²Since the economy starts on date zero, there is no previous adjustment date for any household, but their Δ can be inferred from their next adjustment date.

interval, food consumption is equal across cohorts. We assume that $Y(\cdot)$ is absolutely continuous, in which case the density $y(q)$ corresponds to the share of food consumption “eaten” by cohort q . Because markets are complete, all cohorts have proportional trajectories of food consumption, and the distribution $Y(\cdot)$ is constant over time. This follows from Theorem 1, which implies that food consumption growth is perfectly correlated across all households and therefore all cohorts. $Y(\cdot)$ can therefore be interpreted as the initial distribution of relative food consumption in the economy.

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

$$X_t = \int_{q=0}^T x_t(q) dq.$$

Denote the net wealth of a cohort q at time t by $w_t^{\text{net}}(q)$, and the aggregate of this quantity across the population by W_t^{net} . Finally, define τ to be the value of t modulo T . This notation allows us to characterize aggregate dynamics as follows.

Proposition 1 *For an arbitrary initial distribution of food consumption across cohorts, $Y(\cdot)$, at any point in time $t > T$ aggregate commitment and non-commitment consumption can be expressed as*

$$X_t = \int_0^T a(\tau - u) \cdot W_{t-u}^{\text{net}} du \tag{6}$$

and

$$F_t = b(\tau) \cdot W_t^{\text{net}} \tag{7}$$

where $a(\cdot)$ and $b(\cdot)$ are functions that depend on $Y(\cdot)$.

If the initial distribution $Y(\cdot)$ is uniform, both $a(\cdot)$ and $b(\cdot)$ are constants.

The content of this proposition is that we can express aggregate commitment consumption as a time average of past aggregate net wealth levels. Because such a time average is continuous, the proposition implies that aggregate commitments follow a continuous path. Intuitively, the lumpy adjustment dynamics at the household level created by commitments are “smoothed out” in the aggregate.¹³

To understand the intuition for this result, first focus on food consumption. Aggregate food consumption is proportional to aggregate net wealth, because the same is true for each cohort by equation (3). In general, the factor of proportionality b will depend on τ , because the share

¹³The literature on aggregating agents with state-dependent adjustment rules also finds a similar smoothing effect; see for example Bertola and Caballero (1990).

of food consumption corresponding to any given Δ in (3) exhibits a T -long cycle. When $Y(\cdot)$ is uniform, the share of food consumption is constant across cohorts, and b is a constant.

Now consider commitment consumption. People in a cohort that adjusts on date t choose commitment consumption that is proportional to their level of net wealth according to the policy rule of Theorem 1. Likewise, cohorts that adjusted in the recent past chose commitment consumption that was proportional to their net wealth on their adjustment dates. Hence, the sum of commitments across all cohorts in the economy is a time average of past levels of the net wealth of each of these cohorts. This time average can be re-written as a time average of past levels of aggregate net wealth because each cohort's net wealth is proportional to aggregate net wealth on its adjustment date. When $Y(\cdot)$ is uniform, the factor of proportionality $a(\cdot)$ is constant because all the cohorts are identical.

Using the proposition we can prove the following aggregation theorem, which is the main theoretical result of the paper.

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 (6) in the general case, or

$$X_t = a \cdot \int_0^T W_{t-u}^{net} du. \quad (8)$$

if the economy starts from the uniform distribution.

Proof.

There exists a unique stochastic discount factor (state price density) in this economy because markets are complete. Call the SDF $m_{s,t}(\omega)$. Then for any traded asset return vector $R_{s,t}$ between dates s and t , we have

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

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

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

Hence, for every household

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

Aggregating this equation over the population yields

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

which implies

$$(C_s - X_s)^{-\gamma} = e^{-\rho(t-s)} \cdot \mathbb{E}_s [(C_t - X_t)^{-\gamma} R_{s,t}].$$

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 aggregate consumption dynamics satisfy the Euler equation for the representative consumer habit formation model. All optimality conditions hold if we assume that households take the dynamics of commitment consumption as exogenously given; hence habit is external in the aggregate model. ■

This theorem shows that consumption commitments provide microfoundations for habit formation. It states that an economy of households who make infrequent adjustments in commitment consumption look in the aggregate like a single individual facing a smooth, slow-moving habit that depends on past consumption levels. To understand the proof of this result, observe first that all agents in the economy have identical linear Euler equations for food consumption, given complete markets. These equations can therefore be summed, giving an analogous equation for the dynamics of aggregate food consumption.¹⁴ Since aggregate food consumption, F , is the excess of total consumption over commitment consumption, $C - X$, the aggregated Euler equation has the same form as the Euler equation in a representative-consumer habit formation model, where marginal utility is determined by precisely $C - X$. The final step in establishing the equivalence to habit formation is to show that X is a slow-moving average of past net wealth levels. This was established above in Proposition 1.

The proof of Theorem 2 indicates that the idea that commitments provide microfoundations for habit does not hinge on either the particular CES preferences we have assumed or on purely time-dependent adjustment. As long as individual marginal utility depends only on surplus consumption over commitments (i.e., $c - x$), we can aggregate these individual marginal utilities to get a marginal utility for the representative consumer that depends on $C - X$. Hence, a representative consumer with habit preferences defined over $C - X$ exists as long as utility is separable in commitment and non-commitment consumption. The only additional benefit of constant elasticity of substitution preferences and purely time dependent (instead of state-dependent) adjustment is that they allow us to express habit analytically as a time average of past wealth levels as in (6).

¹⁴The logic of the proof is similar to that in Grossman and Shiller (1982), who prove the existence of a representative consumer in an asset pricing context with one consumption good, if individual consumption dynamics follow Ito-processes with no jumps.

Note that the aggregate model can be interpreted as an economy in general equilibrium, in the same spirit as Constantinides (1990) or Cox, Ingersoll and Ross (1985). In the equilibrium interpretation, the stock S and the bond B correspond to two exogenous, constant returns to scale production technologies, both of which use the food good as input and produce the food good as output. The portfolio decisions of the households can be interpreted as investments into these technologies, and total wealth is equal to the capital stock. In addition, there is an irreversible production technology that transforms one unit of the food good into one unit of the commitment good.¹⁵ Finally, due to technological constraints, commitment consumption can only be adjusted on reset dates. With this specification of technology and our earlier description of preferences, the aggregate economy is in general equilibrium.

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 function of past levels of aggregate consumption. Here habit is a slow moving function of a different aggregate variable, net wealth. We next establish that the model also admits a representation where habit is a function of aggregate consumption.

Proposition 2 *When the initial distribution of relative food consumption $Y(\cdot)$ is uniform, aggregate commitment consumption can be written as*

$$X_t = o(t)X_0 + \int_0^t \zeta(u)C_{t-u}du \quad (9)$$

where

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

Proof. See the Appendix. ■

The proof of the result is simple. When $Y(\cdot)$ is uniform, Proposition 1 shows that aggregate food and commitment consumption are linear functions of net wealth. Hence, there are three linear equations linking the variables C , F , X and W^{net} : an accounting identity, the habit rule (6), and the consumption rule (7). These equations allow us to express any of the four variables as a linear function of current and lagged values of the other variables.¹⁶

¹⁵By scaling units, the model can accommodate any linear production function transforming food goods into commitment goods.

¹⁶Hence, commitment consumption can also be written as a time average of past levels of food consumption.

This proposition shows that the utility of the representative agent for the commitments model is a special case of the habit specification advocated in Chen and Ludvigson (2004):

$$X_t = f(C_t, C_{t-1}, \dots, C_{t-L}).$$

The commitments model generates aggregate dynamics that are particularly similar to the Constantinides' formulation of habit, 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. In between, the weight function $\zeta(u)$ fluctuates and periodically becomes negative. Thus habit in the commitments model corresponds more closely 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.

Note that the technical condition $aT < b$ in the proposition is used to ensure that the coefficients $\zeta(\cdot)$ go to zero asymptotically. This restriction roughly requires that commitment consumption is typically less than food consumption, as shown by equations (6) and (7). In the next subsection we consider a different version of the model, where we obtain habit as an exponential time average as in Constantinides (1990) without such restrictions.

3.3 Model and Aggregation with a Stochastic Adjustment Rule

We now introduce the possibility that the reset dates for consumption commitments are stochastic. At the cost of some additional assumptions, this new model yields a more general aggregation result, making the connection with habit persistence stronger.

To model uncertainty in moving dates, suppose that during a short interval dt , a consumer who has been committed for a time period of Δ can adjust her level of commitment consumption with probability $\lambda(\Delta)dt$. With remaining probability, she cannot adjust. When λ does not vary with Δ , this specification is similar to the Calvo-pricing rule common in macroeconomics. Here we allow $\lambda(\Delta)$ to be any weakly increasing function. This is equivalent to assuming that the probability of facing a move-inducing shock is weakly rising with home tenure. This assumption captures the idea that recent home purchasers are relatively unlikely to move; further intuition and evidence for this assumption is given in the context of our estimation strategy below. For technical reasons, we also assume that at a finite date T , the function $\lambda(\Delta)$ reaches a maximum $\lambda(T) = \lambda_H$.

Let $G(s)$ denote the unconditional distribution function of the arrival time for the next ad-

justment date, s . The function $\lambda(\Delta)$ determines $G(s)$ according to the differential equation

$$\lambda(\Delta) = \frac{G'(\Delta)}{1 - G(\Delta)}. \quad (10)$$

Indeed, if arrival dates are drawn from $G(\cdot)$, the right hand side is the likelihood that the next adjustment date arrives immediately conditional on Δ periods having elapsed.

We assume that there is a continuum of consumers whose adjustment dates are independently distributed. Aside from the change in the process that determines arrival dates, the preferences and technology in the economy are the same as in the previous model, as summarized in (1).

As before, we assume that financial markets are complete to facilitate a simple aggregation of consumers. While market completeness held by design in the deterministic model, it is a non-trivial assumption here. This is because consumers now face additional uncertainty in the random arrival of adjustment dates. The market completeness assumption that we impose requires that consumers can insure all financial risk associated with when they can adjust their commitment consumption. In the optimum, consumers will intuitively seek to insure two types of risk. First, they face “payment risk,” which is the randomness in the total payments associated with living in a particular house, given that the total duration spent in the residence is a random variable. With complete markets, this risk can be insured, and households will do so by pooling their risk with others in the cohort of agents who buy a new house at the same time that they do. With the optimal contract, conditional on the value of x , each consumer’s commitment payment for his current house is known with certainty. Second, agents could face fluctuations in food consumption that are induced by uncertainty in their moving dates. For instance, when a household moves to a bigger home, it might need to cut back on food consumption. The household insures this risk by holding a contract that pays off when the moving date comes, avoiding the cut in food consumption. Real-world examples of these two types of risk-sharing arrangements include (1) young adults living with their parents until they get married, effectively being insured against uncertainty in when they find the right spouse and buy a new house, and (2) family support and gifts after shocks that induce new commitments, such as wedding or childbirth.

Although consumers are able to insure against these forms of risk, note that they cannot eliminate one important source of risk arising from the technological constraints of the economy. The adjustment of commitment consumption remains possible only on exogenous, uncertain dates. Because utility over commitment consumption is concave, the household dislikes such risk, but it has no way to insure against it given the lack of a market for housing that permits continuous, costless adjustment of this good.

This model yields a set of results very similar to that of the deterministic adjustment model, which are summarized in the following theorem.

Theorem 3 *The following are true.*

- *The optimal level of stocks in the household portfolio is*

$$\alpha_t w_t = \frac{\pi}{\gamma \sigma^2} \cdot (w_t - \Lambda(\Delta) x_t). \quad (11)$$

where $\Lambda(\Delta) = \int_0^\infty e^{-rs} \frac{1-G(s+\Delta)}{1-G(\Delta)} ds$ is decreasing in Δ .

- *The aggregate consumption dynamic of the economy is the same as the optimal consumption of a representative consumer economy with habit formation utility function where habit X_t evolves according to the path of aggregate commitment consumption.*
- *If the economy is started from the ergodic distribution of relative food consumption (which exists), then*

$$X_t = \int_0^t \zeta(s) \cdot C_{t-s} ds + X_0 \cdot o(t) \quad (12)$$

with appropriate deterministic weight functions $\zeta(\cdot)$ and $o(\cdot)$.

- *In the special case when the adjustment date arrives at an exponential rate,*

$$X_t = X_0 e^{-dt} + D \int_0^t e^{-du} C_{t-u} du \quad (13)$$

with suitable positive constants d and D .

Proof. See the Appendix. ■

The first part of this theorem states that the state variable governing the portfolio choice and food consumption decisions of a household is the quantity $w_t - \Lambda(\Delta)x_t$ which, in keeping with the intuition of the previous section, we will call net wealth w_t^{net} . Note that we have re-defined the function $\Lambda(\Delta)$, which was also used in the previous deterministic model, in particular in the portfolio choice rule (4). We abuse notation in this way because the two $\Lambda(\cdot)$ functions in the two models play completely analogous roles, and with this notation the portfolio choice rule in the two models is algebraically identical.¹⁷ The reason is that in the current environment, as in the deterministic environment, $\Lambda(\Delta)$ measures the net present value of outstanding commitments.

To see why $\Lambda(\Delta)x_t$ measures outstanding commitments and w_t^{net} is the key state variable, consider a set of consumers who all happened to adjust on a date t , and all of whom have a level of wealth w_t . As we observed above, market completeness implies that these consumers can efficiently share the risk associated with their next adjustment date among themselves. The reason is that their upcoming adjustment dates are independent, thus for them *as a group*, there

¹⁷We will similarly reuse notation for other variables that play identical roles in the two models.

is no aggregate uncertainty about the total amount of money they have to spend on their current home. To calculate the total cost of the current home for this group, note that for each date s after t , the share of individuals who have not adjusted yet is $1 - G(s)$. Hence, this group must make a total payment of $x(1 - G(s))$ if x is the level of commitment consumption chosen on date t for all $s \geq t$. The total present discounted cost of these payments is the price of outstanding commitments for this group at time t :

$$x \cdot \int_t^\infty e^{-r(s-t)} (1 - G(s - t)) ds.$$

Note that the outstanding payment stream is discounted by r because at the group level, there is no uncertainty about how much needs to be paid.

Following the same logic, for a household who adjusted Δ periods ago, after efficiently sharing risks, the outstanding commitment payment is exactly equal to the expected discounted cost of consuming the current house, i.e., x multiplied by

$$\Lambda(\Delta) = \int_t^\infty e^{-r(s-t)} \frac{1 - G(s - t + \Delta)}{1 - G(\Delta)} ds. \quad (14)$$

This expression is a generalization of the previous one. Given that Δ periods have already passed for this cohort by time t , the share of the remaining households who have not adjusted yet by some date $s \geq t$ is $(1 - G(s - t + \Delta))/(1 - G(\Delta))$. Using these new probability weights yields the more general expression for the present discounted cost of commitment payments for a cohort that adjusted Δ periods ago. Since the outstanding commitment payment is given by $\Lambda(\Delta)$, the household sets aside that amount in safe assets to finance future payments as in the deterministic model. The remaining amount, which is precisely w_t^{net} , is allocated to food and savings, explaining the household portfolio choice equation in the theorem. Note that $\Lambda(\cdot)$ is decreasing over time because households with longer home tenures have a shorter expected duration in their current home.¹⁸ Hence, the stochastic model, like the deterministic one, implies that recent home buyers have higher outstanding commitment payments and should therefore hold safer portfolios.

The second part of this theorem shows the existence of a representative consumer, and proves that commitments provide microfoundations for habit formation by aggregating the individual Euler equations. The intuition for these results is the same as in the previous section. In this case, we obtain an analytic characterization of the dynamic of aggregate habit by establishing that relative food consumption across cohorts has an ergodic distribution. When the population is started from that ergodic distribution, we first prove that habit can be expressed as a time

¹⁸More precisely, a sufficient condition for $\Lambda(\cdot)$ to be decreasing is that the hazard rate $\lambda(\cdot)$ is increasing, as we show in the Appendix.

average of past levels of food consumption, and then make use of a general argument which shows that a time average of food consumption can always be expressed equivalently as a time average of total consumption C_s . Note that the stochastic aggregate model also admits a general equilibrium interpretation as in the previous section.

The final part of the theorem considers the special case when the arrival dates are completely unpredictable. In this case, the habit rule can be written as a simple exponential time average, regardless of the initial distribution of relative food consumption. This result underscores the strong connection between the commitments model and habit formation. The representative consumer's preferences and the habit dynamics of our aggregated 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.

3.4 Extensions: Heterogeneity Across Households

Both models above assumed that households are identical in all respects except for their adjustment dates – i.e., that there is no heterogeneity along any other dimension. This subsection describes why the portfolio choice and aggregation results extend when we allow for heterogeneity in the relative preference for commitments (μ), wealth levels, labor income risk, as well as the frequency of adjustment dates. The heterogeneity across these dimensions can also be correlated across households without affecting the results. The robustness of the theory is especially important because such heterogeneity is prevalent in the data, as we will see below.

To begin, consider the introduction of a risky labor income stream. Maintaining the complete markets assumption, risky labor income that is perfectly correlated with the stock market can easily be introduced.¹⁹ Households who have labor income risk will adjust their stockholdings up to the point where their total risk exposure is given by (11). If we interpret wealth to include human wealth, then the dynamics of wealth and consumption are unaffected. Hence, heterogeneity in labor income risk only affects the portfolio choice rule, and has no effect on wealth or consumption dynamics. Permitting the heterogeneity to be correlated with adjustment dates does not affect this reasoning.

Wealth heterogeneity is easy to handle because we have assumed homogenous utility. If two households are identical in all respects except for their level of wealth, their consumption and investment decisions will be exactly proportional on all dates. Since any household can be broken into other households with smaller wealth levels, the model is consistent with any cross-sectional

¹⁹More generally, any risky labor income stream can be introduced provided that the asset space is sufficiently rich that agents can hedge this risk fully.

wealth distribution.

Finally, consider heterogeneity in μ , and suppose that it is correlated with the adjustment frequency. To see why the results still go through intuitively, consider an economy with two types of agents, those with a low preference for commitments μ_L and those with a high preference μ_H . Assume that households in the low group adjust less frequently than households in the high group (formally, $\lambda_L(\cdot) < \lambda_H(\cdot)$). By complete markets, there exists a representative consumer with habit-type preferences for the whole population. Moreover, for the subpopulations of H and L agents, the equivalent of Theorem 3 holds (with group specific values for the parameters in the aggregation). This immediately implies that the portfolio choice result extends. To show that aggregation extends, we need to verify that total habit for the union of the two groups can be represented using past total consumption C_s .

By the argument of Theorem 3, habit in both groups can be expressed using past values of group-level food consumption. By complete markets, food consumption is perfectly correlated across individuals in different groups, implying that the linear representation of group level habit with group level food consumption can be summed to write aggregate habit as a time average of aggregate food consumption. As noted earlier, such an expression always leads to a representation of total habit using past levels of total consumption.

4 Testing the Model: Estimation Strategy and Data

The theory that commitments are the source of “habit” in the macroeconomy has several predictions that can be tested using microdata. The most substantive predictions relate to the two key features of habit-formation preferences: (1) Commitments should make agents adjust consumption – especially of more committed goods – slowly in response to shocks and (2) Commitments should drive up risk aversion.

There is ample existing evidence for the first prediction. Both the infrequency of moves and vehicle purchases in our data and the formal studies of consumption behavior discussed in section 3 indicate that commitment consumption adjusts very sluggishly. Since commitment consumption is a significant fraction of total consumption, it follows that commitments must make total consumption adjust sluggishly. In addition, the papers discussed in the literature review and other studies of consumption support the prediction that adjustment of more committed goods is more sluggish. For instance, Gruber (1997) and Gruber (1998) find that unemployment spells induce large changes in food consumption but that less than 5% of the unemployed move.

We therefore focus on the prediction about commitments and risk aversion here. One important way in which the amplification of risk aversion by commitments manifests itself is in the

form of a more conservative portfolio allocation. As shown by equations (4) and (11), both the deterministic and stochastic adjustment models developed above imply the following relationship between stock holdings and commitment consumption for a given household i :

$$\alpha_{i,t}w_{i,t} = \xi + \frac{\pi}{\gamma\sigma^2}\Lambda(\Delta_{i,t}) \cdot x_{i,t} + \theta \cdot w_{i,t} + \eta_{i,t}. \quad (15)$$

The disturbance term η introduced in this equation captures background income risk that shifts household portfolios, as well as measurement error and other unobserved heterogeneity in the data that are not explicitly incorporated in our model. The theory predicts that an exogenous increase in commitments should reduce the fraction of wealth held in risky assets (stocks) relative to safe assets (bonds) by the agent. To test whether this prediction holds empirically, we use housing consumption as the commitment good x . We focus on housing because it constitutes a large fraction of the average individual’s budget (see Table 1), and because moves between homes are sufficiently infrequent that housing does indeed appear to be a significant commitment.

In the case of housing, (15) implies that holding wealth fixed, total commitment is 1) rising with the value of the home, x , and 2) falling with home tenure, which is inversely related to Δ . We focus on variation in home tenure rather than home value to obtain variation in a household’s degree of commitment, for two reasons. First, although we have assumed that the commitment good is not an asset in our model, in practice, houses themselves are large assets whose prices can fluctuate. The riskiness of housing can affect portfolio allocations for non-commitment reasons, as suggested by Cocco (2005). Second, variation in home values is tightly linked with households’ idiosyncratic preferences, making it very difficult to find a powerful source of variation in property value that is not plagued by endogeneity or omitted variable biases.

What is the intuition for the connection between home tenure and the degree of commitment? People who just bought a new house see the long stream of mortgage payments they have to make as a fixed obligation, compelling them to “play it safe” in the asset markets in order to be sure they can pay the bills. This is because recent home buyers are more likely to stay for a few years rather than moving immediately after purchase. For example, if a couple buys a first home after marriage, they are likely to anticipate living there until the birth of one or two children generates a need for a larger home. The claim that recent homebuyers have longer expected durations is consistent with the empirical results of Sinai (1997), who estimates the duration dependence of the hazard rate of moving, controlling for unobserved heterogeneity across households in baseline hazard rates.²⁰ Once a family has lived in a house for several years, it is more likely to have

²⁰It is important to distinguish within-household and cross-sectional variation in the probability of moving. Our assumption requires (and empirical evidence confirms) that the within-household probability of moving rises with home tenure; this may not be true in the cross section because of heterogeneity in mobility rates.

outgrown the home (either upward or downward), and its expected duration in the home is shorter. Hence, the house is viewed as less of a commitment, and the family should be more willing to take risks in order to later raise commitment and non-commitment consumption.

Since we will focus on home tenure, the key estimating equation for our empirical analysis is

$$\text{stockholding}_{i,t} = \xi + \beta \times \text{hometenure}_{i,t} + h(\text{wealth}_{i,t}) + g(\text{age}) + \text{controls} + \varepsilon_{i,t}. \quad (16)$$

In the empirical analysis, we consider two measures of “stockholding”: dollars held in stocks and the fraction of wealth held in stocks. There are four differences between (15) and this specification. First, $\beta \times (\text{hometenure})$ is a linear approximation of the function $\Lambda(\Delta)$. Second, we have assumed that home values are fixed; if home value varies across households, the estimate of β will reflect a weighted average of the true coefficient. Third, we have allowed stockholding to depend on household wealth through a non-linear function $h(\text{wealth})$, so that the estimates that are not sensitive to the particular linear relationship implied by our stylized model. Finally, we have introduced age and other controls that may proxy for life cycle effects and other unmodelled heterogeneity in the data. We introduce a new error term ε to reflect these changes.

Testing the commitments theory in this specification is formally equivalent to testing whether $\beta > 0$. Because (16) captures the fundamental intuition that larger commitments should amplify risk aversion, the estimation results in this paper serve not only as a test of the particular models developed above, but also provide evidence regarding the general importance of commitments as foundations for habit formation.

4.1 Data and Sample Selection

We use data from the 1990-1996 panels of the Survey of Income and Program Participation to estimate (16). The SIPP collects income, asset, and demographic information from a sample of approximately 20,000 households. Asset data are generally collected once per panel. The main advantages of the SIPP relative to other commonly used datasets on financial characteristics such as the SCF and PSID are its large sample size and detailed information about covariates such as a complete marital history.

The raw data contains information on 99,136 households, of which 53,680 are married. We restrict attention to this “core sample” of married households in our analysis because we will obtain exogenous variation in home tenure from the timing of marital shocks such as marriage and remarriage after divorce or widowhood.²¹

²¹We also examine divorced individuals who did not re-marry as a control group in section 5 below.

Table 2 gives summary statistics (in real 1990 dollars) for the core sample of married households, as well as four other subgroups, which are discussed in greater detail below. The general demographic and wealth characteristics of the core sample appear to be fairly representative of the married U.S. population. Approximately 40% of total wealth is held in the form of home equity and another 25% is held in illiquid assets such as cars and other real estate, leaving approximately 35% in liquid assets for the average household. Of the \$42,555 of liquid wealth, 37% is held in stocks, 23% in interest-bearing savings accounts, 10% in bonds, 18% in IRA assets, and 12% in “other” liquid assets. The relatively small fraction of wealth held in stocks reflects the fact that only 20% of the married households in the data are stock market participants. “Other assets” can be further broken down into checking accounts (7%), US savings bonds (14%), debt owed to the household (36%), and equity in other financial investments (43%).²² Note that total wealth does not include 401k assets, as the 1990-93 SIPP panels do not collect data for this category. In section 5, we describe tests which indicate that the lack of data on 401k’s is unlikely to bias the portfolio-shifting effects we document in non-retirement accounts.

The data understates the skewness of the distribution of certain variables such as income and property values because these variables are topcoded to protect confidentiality.²³ All results reported below are robust to inclusion or exclusion of the topcoded group.

5 Empirical Evidence

The central difficulty in testing the model is that home tenure is endogenous and correlated with other characteristics that affect portfolio choice. Consequently, the orthogonality condition required to obtain a consistent estimate of β using OLS is unlikely to hold. For instance, in the data, education is inversely related to home tenure, i.e., more educated people are more mobile. If education can only be measured imperfectly and better educated individuals also tend to have lower (unobserved) background income risk, there would be a correlation between ε and home tenure in (16). A similar problem plagues a panel study that tracks the portfolio allocation of a household over time. People tend to buy houses when they obtain secure jobs (e.g., when they get tenure in academia). But this is precisely the time that they may also compensate for their reduced labor income risk by holding riskier assets, again violating orthogonality in a fixed-effects OLS regression.

²²Since 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 category.

²³For 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.

Consistent with this intuition, we find that OLS estimates of β are very sensitive to the inclusion of controls. With a minimal set of controls (age spline, total wealth spline, and year dummies), we obtain an estimate of $\beta = \$20$ (s.e. = \$176) in the group of married stockholding homeowners. However, once we include a rich set of controls (see Table 3, specification (4) for the list), we obtain $\beta = \$220$ (s.e. = \$104). The fact that higher commitment (shorter home tenure) is associated with less stockholding once controls are included suggests that the violations of the OLS orthogonality condition work against finding evidence for the theory.²⁴

In view of these difficulties, we use an instrumental variables approach to estimate β . We begin with a graphical overview that sketches our identification strategy and then discuss a series of regression estimates to fill in the details.

5.1 Graphical Overview

Motivated by the model, where moves are driven by shocks, we use a set of shocks that cause moves as instruments for home tenure after establishing that they are uncorrelated with ε . The shocks we use are all related to the timing of changes in marital status, holding fixed age and wealth: age at first marriage, age at first marriage termination (divorce or spouse’s death), and an indicator variable for remarriage.²⁵

Figure 1a depicts the first-stage relationship between marital shocks and home tenure. The details of how this and other figures are constructed are described later; here we give a short summary of their message. For simplicity, we combine the three instruments into one by defining a new variable, the “age at most recent marital shock,” which equals the age at first marriage among once-marriers and age at first termination among remarriers. The figure plots the residual relationship (in deciles) between the age at most recent marital shock and home tenure, conditioning on a large set of observables such as wealth, age, and demographics. It is clear that there is a strong negative effect of the age at most recent marital shock on home tenure. The intuition is that those whose marital status has changed more recently are more likely to have moved more recently, and therefore have shorter current home tenure.

To use the marital shocks as instruments, we must first ensure that they do not have direct effects on portfolio choice outside the home tenure channel, once we condition on variables such as age and wealth. For instance, we might worry that more educated or informed people tend to marry later and also have different portfolio preferences. Figure 2 illustrates how we test whether

²⁴We confirmed that endogeneity is also a problem in a panel analysis by examining trends in income and wealth around home purchases using the 1996 SIPP panel, where asset data were collected twice.

²⁵More precisely, we use age and remarriage data for the “reference person” in the SIPP, who is an owner or renter of record and is typically the head of the household.

this key orthogonality condition is satisfied. To construct this figure, we focus on a “control group” of households that are either 1) long married homeowners (married more than 40 years), 2) married renters, or 3) divorced/widowed homeowners who did not remarry. As we describe in greater detail below, the theory suggests that there should be no first-stage relationship between marital shocks and home tenure for these three groups. Figure 2a shows that this prediction holds true in the data: the age at most recent marital shock instrument is unrelated to home tenure for the households in the control group. Consequently, in these control groups, the commitments channel is effectively shut down, and any correlation between marital shocks and portfolio allocations would constitute evidence that the key exclusion restriction is violated. Fortunately, as the horizontal fitted line in Figure 2b demonstrates, the marital shock instrument is completely unrelated to the share of wealth held in stocks in the control groups, providing strong evidence that marital shocks have no direct effects on portfolio choice.

Having established that marital shocks can be used as valid instruments, we examine the reduced-form effect of marital shocks on portfolio choice among households that are not in the control group defined above. Figure 1b illustrates that in this “treatment group,” the age at most recent marital shock has a strong negative effect on the share of wealth held in stocks. Figures 1 and 2 together imply that home tenure has a *positive* causal effect on stockholding, confirming that recent home buyers have amplified risk aversion, as the theory predicts.

The rest of this section formalizes the ideas behind these figures using regressions. The next subsection documents the first-stage relationship depicted in Figure 1a for each of marital shock instruments. The second subsection formally tests the exogeneity of these instruments using the three control groups, essentially by comparing reduced form estimates similar to those in Figures 1b and 2b. The final subsection reports two-stage least squares estimates and overidentification tests. We address other potential biases such as the endogeneity of wealth and potential non-commitment reasons for a link between home tenure and portfolio choice in section 6.

5.2 First Stage Effects of Marital Shocks on Home Tenure

We now examine the first stage relationship between marital shocks and home tenure in greater detail. We focus initially on a set of households who are candidates to be affected by marital shock instruments based on our theory. In section 5.2 below, we explain theoretically and confirm empirically that there is no first stage association between the timing of marital shocks and home tenure among renters, long-married homeowners, and divorcees and widows. Here, we restrict attention to the remaining households, who are married homeowners whose current marriage duration is less than forty years.

We also initially restrict attention to the 24.5% of households in this sample who are stockholders. Only stockholders can respond on the intensive margin by changing stockholding in response to changes in commitment levels, and we show later that changes in commitment appear to have little effect on stock market participation rates. We label the set of recently-married homeowners who own stocks the “treatment” group. Table 2 reports summary statistics for the treatment group, which unsurprisingly is considerably wealthier than the average household in the population. To demonstrate that our results are not driven by endogenous sample selection, we later report estimates on the entire sample of all married households in the data.

Consider the age at first marriage instrument. Specification (1) of Table 3 reports estimates of ϑ in the following specification:

$$\text{hometenure} = \delta + \vartheta \times \text{age at 1st marriage} + \text{controls} + \nu. \tag{17}$$

This specification has “few controls”: year dummies, a 10 piece linear spline (partitioned by deciles) for age, and a 10 piece linear spline for total household wealth.²⁶ The sample in this regression consists of once-married households in the treatment group. The reason we exclude twice-married households here is that age at first marriage has little predictive power for their home tenure, as we will see shortly. In the once-married treatment group, holding fixed age and wealth, marrying one year later is estimated to reduce the number of years spent in the current home by 0.2 years (with a t-statistic of 9.86). This is not surprising, given that more than 40% of households in our data bought their current home within 5 years after their first marriage.

Now consider the age at first marriage termination variable, which is defined only for individuals whose first marriage ended. Specification (2) replicates specification (1) using the age at first marriage termination variable for twice-married households in the treatment group. The estimates reveal that age at termination of the first marriage is a strong predictor of home tenure in this group, while age at first marriage is statistically much weaker. This is as we would expect based on our theory: only the most recent move-inducing shock should determine current home tenure, making the timing of first marriage less relevant for those who have had another marital shock since then.

The age at first marriage and age at first termination instruments capture “within-group” variation in two disjoint groups – once and twice-married couples. Specification (3) in table 3 explores the “across-group” variation between these two groups, by investigating the link between

²⁶We always begin by reporting estimates of a specification with few controls to alleviate concerns that the inclusion of endogenous regressors is spuriously generating our results. In this case, only wealth is potentially endogenous, and we address wealth endogeneity in the section 6.

an indicator for remarriage and home tenure.²⁷ Column (3) shows that remarried households have lived in their current home for 3 years less on average than once-married households. The t-statistic for this coefficient exceeds 12.

To summarize the first-stage results for the entire treatment group, we combine the three instruments by defining a new variable, age at most recent marital shock = max(age at first marriage, age at first termination). Specification (4) in Table 3 shows that conditional on total wealth and age, a one year increase in the age at the most recent marital shock reduces current home tenure by 0.25 years on average (t-statistic of 16) in the treatment group. Specification (5) shows that this coefficient is virtually unchanged with “full controls”: 10 piece linear splines for liquid wealth, home equity, property value, and age; linear controls for unsecured debt, business equity, vehicle equity, education, income, and the number of children at home; and industry, occupation, and year fixed effects.

Finally, specification (6) replicates (5) for the “core sample” that consists of all 53,680 married households in the data. The age at most recent shock remains a very powerful predictor of home tenure in this group, although the coefficient is slightly smaller, which is consistent with the existence of control groups within the core sample where the marital shock variable does not predict home tenure.

5.3 Exogeneity Tests for Marital Shock Instruments

The marital shock variables must not have direct effects on portfolio choice to be valid instruments for home tenure. Stated formally, the marital shock variable must be orthogonal to the error term in the estimating equation (16):

$$E[\text{marital shock} \times \varepsilon] = 0. \tag{18}$$

We test this exclusion restriction by conducting placebo tests on a set of “control groups” where the theory predicts that marital shocks should have no effect on current home tenure. Stated formally, we predict (and later confirm in the data) that in each control group

$$\text{hometenure} = \delta_c + 0 \times \text{marital shock} + h_c(\text{wealth}) + \text{controls} + \nu_c$$

²⁷In the remarriage dummy regressions, we exclude 10 observations that contain households with reported wealth about \$5 million. Inclusion of these households in this regression blows up the standard errors in the levels estimates (but does not change the point estimates), because these outliers drive the means. Note that all other results reported in this paper remain statistically significant whether these 10 observations are included or not. Moreover, in a shares specification, these households can be included without affecting the results, because they receive much less weight.

where the subscript c denotes “control.” Under the hypothesis that the true relationship between portfolio risk and home tenure in the control group is given by (16), if (18) holds, the implied reduced form relationship for the control group is

$$\text{stockholding} = \text{const} + 0 \times \text{marital shock} + l_c(\text{wealth}) + \text{controls} + \varepsilon + \beta\nu_c. \quad (19)$$

Assuming that $E[\varepsilon \times \text{marital shock}]$ is equal in the treatment and the control group, it follows that (18) holds if and only if estimation of (19) gives a zero coefficient on the marital shock variable. The assumption underlying our tests of exogeneity is thus a standard “common trends” assumption across the treatment and control groups:

$$E[\varepsilon \times \text{marital shock} \mid \text{treatment}] = E[\varepsilon \times \text{marital shock} \mid \text{control}]. \quad (20)$$

Intuitively, our test requires that the *direct* effect of marital shocks on portfolio choice – if there is one – should be the same in both our treatment and control groups. Conditional on this assumption, if we find no evidence of a link between portfolios and the timing of marital shocks in the controls, we can clearly conclude that there is no such link in the treatment group either.

The three control groups we examine are as follows. 1) *Long-married homeowners*. Our model suggests that a marital shock should affect commitment consumption until some other shock leads to another move. For individuals who have been married for a very long time, the probability that some other moving shock (e.g., job related) has caused them to move is quite high. Hence, we expect no first-stage relationship between current home tenure and age at first marriage for this group. 2) *Married renters*. Renting involves much less of a commitment than home ownership: in the data, the median home tenure for renters is only 3 years. Since consumption of rental units is not sticky over time, there should be no first-stage effect of the age at most recent marital shock on home (apartment) tenure for renters. 3) *Divorced/widowed homeowners who did not remarry*. Remarriers tend to move around the time of their second marriage.²⁸ This suggests that many homeowners who do not remarry do not move when their marriage ends. Consequently, age at first termination should not predict current home tenure in this group.

Tables 4a and 4b report results of the exogeneity tests. All regressions include the full set of controls described in Table 3; results are similar in the few controls specifications. In these and all other tables, standard errors are robust to arbitrary heteroskedasticity of error terms. We begin by testing the exogeneity of the combined instrument, the age at most recent marital shock. Columns (1)-(3) of Table 4a report results of an exogeneity test on the “pooled control

²⁸To confirm this, we regress hometenure on both age at second marriage and age at first termination. In this regression, only age at second marriage predicts current home tenure.

group,” which pools households in the three control groups described above. Column (2) shows that the effect of age at most recent marital shock on home tenure in the pooled control group is insignificant and very small. This first-stage estimate confirms our hypothesis that marital shocks should not predict home tenure in the control groups.

Given this result, we can test the exclusion restriction for this instrument by regressing stockholding on the age at most recent marital shock in the pooled control group. Column (3) shows that the coefficient on age at most recent marital shock in this regression is \$29, with a robust standard error of \$76. Column (1) helps in interpreting the magnitude of this coefficient by reporting an identical regression using the *treatment* group, where we established the existence of a powerful link between marital shocks and home tenure in the previous section. In the treatment group, having a marital shock one year later is predicted to *reduce* stockholding by \$658 with a standard error of \$241. Given that the 95% confidence intervals for the reduced-form coefficients in the treatment and control do not overlap, it is clear that marital shocks appear to have a very different effect in the treatment group, where they affect home tenure. In addition, the precisely estimated zero coefficient in the control group directly indicates that there is no direct link between portfolio allocation and the timing of marital shocks in that group. Hence, provided that the treatment and control groups are sufficiently similar in the sense of (20) – an issue that we will revisit below – we conclude that the age at most recent marital shock instrument is a valid instrument for home tenure.

Columns (1)-(3) measured stockholding by dollars held in stocks (in levels). Specifications (4) and (5) show that similar conclusions are obtained if stockholding is measured by the share of total wealth held in stocks instead. The estimates for the control group are again statistically insignificant and an order of magnitude smaller than the highly significant estimates in the treatment group. This is not surprising: As we discuss in greater detail in Section 5.4 below, the only difference between levels and shares specifications is the way in which the data is weighted, and the results reported in this paper hold across all wealth levels.

The remaining exogeneity tests report separate results for each of three control groups. Columns (6) and (7) focus on married renters, and show that age at most recent marital shock has an insignificant and small effect in the first stage, while the effect in the reduced form is an insignificant *positive* \$690 with a standard error of 401. The positive point estimates, which are consistent with Beckerian theories of marriage (Becker, 1973), imply that if anything, the correlation between age at first marriage age and the error term ε is positive. This would bias the two-stage least squares estimate of β upward, making it more difficult to find evidence supporting the theory, which predicts $\beta < 0$. Comparing the reduced-form estimate in column (6) with

column (1) reveals that the key coefficients are statistically distinguishable in the treatment and renters control group.

In Table 4b we break down the most recent marital shock instrument and examine the effects of each of its components in the remaining two control groups. Specifications (1)-(3) report an exogeneity test for the age at first marriage instrument using the control group of long-married homeowners. Again, the instrument is unrelated to home tenure in this group in the first stage, and the corresponding reduced form effect on stockholding is insignificant and positive. This is in contrast with the results in the treatment group, where we find a negative effect of age at first marriage on stockholding significant at the 10% level. The null hypothesis that the coefficient on age at first marriage in columns (1) and (3) are equal can be rejected with $p < 0.05$.²⁹

Finally, columns (4)-(6) in Table 4b focus on the third control group, non-remarried divorcees, and test the exogeneity of the age at first termination and age at first marriage instruments separately. As expected, for both instruments, the estimates in the first stage as well as in the reduced form are small and insignificant. This contrasts sharply with the reduced form results in the treatment groups for the two instruments – the once-married group for the age at first marriage variable (column 1) and twice-married for the age at first termination (column 4) – where we again find a negative effect significant at the 10% level that is statistically distinguishable from the control group estimates.

Given that there are three control groups and three instruments, there are many more exogeneity tests (not reported) that can be performed using analogous methods. There is no case where the candidate instrument fails the test. In other words, none of the three instruments predicts a strong link between marital shocks and stockholding in any of the control groups.

In summary, the control groups provide strong evidence that the key orthogonality condition (18) holds for all three marital shocks, provided that the “common trends” assumption in (20) holds. To help assess whether this assumption is likely to hold, Table 2 compares summary statistics for the control groups and the treatment group. There are clearly some differences between each of the controls and the treatment – e.g., long-married homeowners are richer and older, renters tend to be poorer and younger, and divorcees tend to be older and have mid-level wealth. However, the treatment group is “sandwiched” in the middle between the different control groups on most dimensions. Our conclusions could only be explained away if the common trends assumption is violated for all three instruments and all three control groups in exactly the same way, which seems implausible.

²⁹Since the coefficients are estimated on disjoint samples, we can assume that the covariance between the two coefficient estimates is zero when conducting this t test.

Before proceeding, we explain the construction of Figures 1 and 2 in greater detail. Figure 1a non-parametrically shows the joint distribution of hometenure and marital shocks in the data, conditioning on the full set of covariates described above. Each point in this figure represents a decile of the treatment group (920 observations). These points are constructed using the following steps. First, raw home tenure residuals are computed for each observation of the treatment group. This is done by regressing home tenure on all the covariates used in specification (4) of Table 4 except the marital shock variable, and computing residuals from this regression. Raw residuals for the age at most recent marital shock variable are computed using an analogous regression. We then break the marital shock residual distribution into deciles, and compute the mean age at marital shock residual and home tenure residual for each decile. The figure plots the mean hometenure residual against the mean marital shock residual, by deciles of the marital shock residual. Figures 1b and 2a-2b were constructed in exactly the same way, except for changes in the sample definition and the dependent variable.

5.4 Two-Stage Least Squares Estimates

Tables 5a and 5b report two-stage least squares (TSLS) estimates of the effect of home tenure on portfolio choice. These estimates can be interpreted as the causal effect of changes in commitment on portfolio choice in view of the preceding analysis. We begin by using the “combined instrument” of the age at the most recent marital shock (marriage, divorce, or spouse death) to create exogenous variation in home tenure. We then perform over-identification tests by breaking up the instrument into its three components, and finally discuss additional robustness checks.

In specification (1) of Table 5a, we estimate the effect of home tenure on stockholding (in levels), which corresponds to β in the key estimating equation (16). This first specification has a minimal set of controls: an age spline, year fixed effects, and a total wealth spline. Wealth is the only potentially endogenous variable in this regression, but its endogeneity is unlikely to bias the estimate of the hometenure coefficient for reasons discussed in section 6 below. The estimates indicate that a one year increase in home tenure causes a \$3000 increase in stockholding on average, consistent with the ratio of the reduced-form and first-stage estimates reported earlier; for comparison, the average wealth in this sample is \$243,621. The 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, property value, and age; controls for unsecured debt, business equity, vehicle equity, education, income, and the number of children at home; and year, occupation, and industry dummies. Under this specification, a one year increase in home

tenure, is still estimated to cause roughly a \$3000 increase in stockholding, and remains significant at the 1% level. The fact that the estimate is virtually unchanged despite the inclusion of a large set of controls indicates that controlling for observed heterogeneity does not affect the results, and suggests that the estimates are not likely to be sensitive to unobservable heterogeneity either.

Where is the \$3000 per year coming from? Specification (3) reports estimates of the regression that corresponds to (2) with safe assets (bonds + money market + CDs + savings accounts) as the dependent variable. It shows that a one year increase in home tenure causes a \$1,558 increase in safe asset holding. This \$1,558 increase in safe assets can be further broken down into a \$1,267 reduction in bondholding (municipal + corporate) and a \$292 reduction in savings accounts, CDs, etc. Hence, households increase their risk exposure by selling bonds and buying stocks as they spend more time in their house, as the model predicts. The discrepancy between the \$3000 shift out of stocks per year and the \$1,558 shift into safe assets is accounted for by an increase in “other assets,” in particular “other financial equity” and “debt owed to the household.” There is no significant association between home tenure and all other forms of wealth. Since the risk properties of “other assets” are unclear, we can be sure that there is a shift of *at least* \$1,558 from safe to risky assets for every additional year that a family spends in a given house.

Columns (4)-(6) repeat specification (2) using the three components of our combined instrument, age at first marriage, age at termination of first marriage, and the remarriage indicator. All three instruments imply that a rise in home tenure causes a rise in stockholding. In addition, the coefficient on home tenure with each of the instruments is of the same order of magnitude and is statistically indistinguishable from the coefficient obtained with the combined instrument. Our instrument set thus easily passes standard overidentification criteria. Any alternative story that undermines the conclusions must be quite complicated, since it would have to explain the findings for all three instruments in both treatment and control groups.

One important concern with the specifications in Table 5a is that they use an endogenously selected sample (recently-married stockholding homeowners), and that selective inclusion into this sample could bias the estimates. To allay this concern, specification (1) in Table 5b replicates (2) from Table 5a for the core sample that includes the universe of *all* 53,535 married households in the dataset. The estimates indicate that a one year increase in home tenure causes a \$568 increase in stockholding for the average household in the core sample (statistically significant at the 5% level). The reason that this coefficient is smaller than for the treatment group is that non-stockholders have no way to respond to increases in commitment, since they cannot cut back further on their exposure to risk. Since stockholders comprise 24.5% of the core sample, the estimate of the hometenure coefficient in column (1) implies that a one year increase in home tenure

causes stockholders to hold an additional $\$568/.245 = \$2,318$ in stocks.³⁰ In the terminology of the instrumental variables literature, the implied effect of the treatment (in this case, a 1 year increase in home tenure) on the treated (in this case, stockholders) is a \$2,318 increase in stockholding. The fact that this estimate has the same magnitude as in all the specifications in Table 5a indicates that our earlier results were not spuriously generated by sample selection bias.

Table 5a uses the level of stocks or safe assets (in dollars) as the dependent variable. Stockholding can also be measured in shares (fraction of wealth held in stocks). When wealth is held fixed, levels ($\$stocks$) and shares ($\$stocks/\$wealth$) specifications are identical, except for the way in which the data are weighted. The levels specifications weight the wealthy more heavily, while the shares specifications weight the wealthy less heavily.³¹ Thus far, we have focused on the levels specifications since these are most relevant for understanding the behavior of portfolios and consumption in the aggregate.

However, to show that our results are robust to the weighting procedure, specification (2) of Table 5b replicates (2) from Table 5a in shares. In this regression, all the monetary variables, including the splines, are included as shares of total wealth.³² The estimates imply that a one year increase in home tenure causes a 0.36% increase in the share of stocks in total wealth for the average household in the treatment group. This estimate is again significant at the 1% level. To assess the magnitude of this coefficient, recall that the mean wealth in the treatment group is \$243,621, implying that the mean increase in stockholding is \$877. The reason that this estimate is smaller than the estimates from the levels regressions (columns 1 and 2) is that the shares regression places greater weight on low wealth households, and low wealth households appear to change their portfolios less than rich households in response to changes in their level of commitment consumption.

To make this point clearer, specification (3) replicates (2) for the subsample of households in the treatment group that have wealth above the mean. In this high-wealth group, a one year increase in home tenure is estimated to cause nearly twice as large an increase in the portfolio share of stocks (0.64%). An analogous regression with the share of safe assets as the dependent variable (not reported) shows that the portfolio share of safe assets falls by 0.54% contemporaneously.

³⁰This calculation assumes, of course, that there is no response on the participation margin. We find no robust association between home tenure and the fraction of households participating in the stock market. This may be because certain households face fixed transaction or informational costs that prevent them from participating, irrespective of the degree of commitment they face.

³¹The levels regressions are likely to suffer from heteroskedasticity in the error terms, but we correct our standard errors for such heteroskedasticity using the standard Huber-White method.

³²Since shares are ill defined for negative-wealth households, we drop the 14 households in the treatment group that report negative total wealth.

Hence, there is a particularly large real shift in the risk composition of the portfolios of high-wealth households over their home tenures.³³

Finally, specification (4) sheds further light on the mechanism through which higher commitments lead to a safer portfolio allocation. It reports estimates of a linear probability model for a bondholder dummy, with the same independent variables as in specification (2). Changes in commitment induce a strong, statistically significant change in the probability of having a bonds account: a one year increase in home tenure causes a 0.7 percentage point reduction in the probability of owning bonds; for comparison, the average probability of owning bonds in the treatment group is 19.7%.

The results above are also robust to many other specification checks that are not reported in the table. First, since we do not have data on the portfolio composition of retirement accounts, 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. To the extent that investment behavior does not vary across different types of retirement accounts, this finding indicates that the lack of data on 401k holdings is not problematic either. 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, 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 again find similar results. Finally, we experimented with clustering of standard errors to allow for various types of serial correlation in the data, and found that clustering does not have a significant effect on the precision of the estimates.

6 Other Potential Biases

6.1 Measurement Error and Endogeneity in Wealth

Testing whether $\beta > 0$ in (16) requires that we hold wealth fixed when we compare households with different housing commitments. This is particularly important because wealth is an extremely powerful predictor of stockholdings in our data. There are two difficulties with “holding wealth fixed” in practice: 1) Wealth measures in the data are likely to suffer from measurement error;

³³One reason that the rich might be more responsive to commitment consumption than the poor is decreasing relative risk aversion (see Carroll, 2000 and Guiso, Japelli and Terlizzese, 1996 for evidence for DRRRA utility). Households with lower risk aversion will invest a larger fraction of their marginal dollar of net wealth in risky assets, as shown by (4).

and 2) The cross-sectional variation in wealth is endogenous. This subsection addresses these concerns.

If the marital shock instruments are correlated with unobserved wealth, measurement error in wealth can induce a nonzero correlation between the instrument and the error term, violating the key orthogonality condition (18). The evidence from the control groups largely addresses this concern by showing that (18) does appear to hold.

Nonetheless, as a robustness check, we now present more direct evidence that marital shock instruments are not picking up unobserved wealth effects. Column (1) of Table 6 estimates an equation analogous to our reduced-form specifications in Table 4 with minimal controls on the treatment group, replacing the dependent variable with the total value of the first two cars currently owned. The coefficient of age at most recent marital shock is -\$12.37 and the standard error is \$13.62, implying that the marital shocks have little to do with current car consumption, once we control for age and total (observed) wealth. The coefficient of the remarriage indicator variable is also small and insignificant. Column (2) shows that similar estimates are obtained when we include the full set of controls. The fact that the instruments do not predict vehicle consumption make it unlikely that they are proxying for unobserved wealth, since wealthier people presumably have more expensive cars.

Columns (3) and (4) repeat this exercise for property value as the dependent variable. The results in column (3) are potentially worrisome, because the effect of age at marital shock on housing consumption is a statistically significant \$504.41. However, once we include controls in column 4, this coefficient becomes small and insignificant. The control that knocks out the marital shock coefficient is years of education, which has a large positive effect on property value. This is intuitive: Education is positively correlated with unobserved wealth (human capital) and age at marriage, so once we hold control for education, marital shocks no longer predict correlates of unobserved wealth.

The second concern is that observed wealth itself may be correlated with the error term because it is endogenous. We now formalize this problem and explain how the estimating equation (16) with the flexible wealth control $h(\text{wealth})$ that we used above was designed to address it. Wealth endogeneity could be a problem because in (15), a violation of the orthogonality condition

$$E[\text{wealth} \times \eta] = 0 \tag{21}$$

makes the estimate of the wealth coefficient θ in (15) inconsistent. However, obtaining a consistent estimate of θ is not central to testing to commitments model. We therefore pursue an estimation strategy that is sure to estimate β consistently at the expense of obtaining an estimate of θ .

Define the (unobserved) function

$$h(w) = E[\eta|w] - \theta \cdot w. \tag{22}$$

The role of $h(w)$ is to isolate the component of the error term that is correlated with any function of observed wealth. Introducing the new error term $\varepsilon = \eta - E[\eta|w]$, equation (15) can be re-written as

$$\text{stockholding}_{i,t} = \xi + \beta \times \text{hometenure}_{i,t} + h(\text{wealth}_{i,t}) + \text{controls} + \varepsilon_{i,t} \tag{23}$$

which is exactly our key estimating equation (16). The advantage of this specification over (15) is that both wealth and $h(\text{wealth})$ are orthogonal to the error term by construction: We defined ε precisely so that $E[\varepsilon|\text{wealth}] = 0$ is satisfied. As a result, specification (23) has no endogeneity problem associated with wealth.

The difficulty in implementing this procedure is that the function $h(\cdot)$ is unknown. We dealt with this problem above by attempting to capture the form of $h(\cdot)$ flexibly using 10-piece linear splines for total wealth or separate splines for its components (e.g., liquid wealth, home equity). The estimates of β are also very similar if a simple linear wealth control is used, suggesting that our main conclusions are not sensitive to the specification of $h(\cdot)$. In addition, the control groups evidence suggests that our wealth controls are sufficiently rich to purge the endogeneity bias. If wealth endogeneity accounts for the results in the treatment group, we should observe the same spurious effect of the instrument in the control groups as well, assuming that the two groups are sufficiently similar. Precisely stated, the control groups evidence rules out correlations between both wealth and ε and the marital shock instruments and ε if the joint distribution of wealth, the instrument and ε are identical in the treatment and controls.

An alternative approach to addressing potential biases from wealth endogeneity is motivated by the observation that β is estimated inconsistently only if there is a correlation between the instrument set for home tenure and observed wealth. The sign of the correlation between wealth and the instrument determines the direction in which the estimate of β will be biased. Column (5) investigates these correlations by replicating the few-controls reduced form specification with total wealth as the dependent variable. It shows that the age at marital shock and remarriage indicator instruments have opposite-signed correlations with wealth. Hence, potential wealth endogeneity biases generated by these instruments must go in opposite directions. But the estimates in Table 5 indicate that both of these instruments give positive, significant estimates of home tenure on stockholding, implying that wealth endogeneity cannot be driving the results. Column (6) in Table 6 shows that the addition of an education control makes the coefficients on both instruments much smaller and insignificant. Hence, once we condition on education, the link between the marital

shock instruments and wealth is broken, making it even less likely that our estimates of β are driven by wealth endogeneity.

6.2 Alternative Channels

The evidence we shown of a causal link between home tenure and portfolio choice supports the commitments theory only if other reasons for such a link are ruled out. In this subsection, we discuss two alternative channels through which home tenure might affect portfolios and argue that neither one explains our findings.

Cocco (2005) shows that in a life cycle model of portfolio choice, background home price risk has a small but non-negligible effect on stockholding. Can home price risk be directly responsible for our findings? To answer this question, first recall that property value does not vary with the instruments in our specifications with controls (specification 6.4). Hence, our results are not a consequence of variation in home price risk associated with variation in home value. Home price risk could also be directly related to home tenure. Households with long home tenure, who are more likely to move in the near future, may also be more concerned about home price risk, because they will be directly exposed to it when they sell their houses. But this effect is likely to work against our results. Home price risk is essentially uncorrelated with the stock market (Cocco, 2005) and the optimal response to an increase in such background risk is to hold *less* stocks (Gollier, 2001, pp 126-129).

Another potential concern is that households accumulate home equity as home tenure rises, directly affecting the portfolio composition of household wealth, or by changing the borrowing constraints that households face.³⁴ We address this issue in two ways. First, our results are essentially unchanged if we control for home equity, implying that contemporaneous movements along this dimension cannot be responsible for our findings. Second, for the wealthy subsample (households above the sample mean), there is no association between home tenure and total home equity (including second houses), and our results continue to hold in this group. Since the levels regressions effectively place more weight on wealthier households, there is no association between the level of home equity and home tenure in the full sample. For example, a regression of total home equity on the age at most recent marital shock with a minimal set of controls in the treatment group yields a coefficient of -\$20 and a standard error of 156. Similarly, mortgage outstanding is also unrelated to the timing of marital shocks. These findings also rule out an explanation based on changes in borrowing constraints associated with home tenure, at least for

³⁴For instance, Constantinides, Donaldson and Mehra (2002) propose a resolution of the equity premium puzzle based on borrowing constraints that vary over the life cycle.

wealthy households.

7 Conclusion

This paper has shown that consumption commitments – goods such as houses, cars, furniture, and service contracts whose consumption cannot be freely adjusted in the short run – amplify households’ risk aversion and make total consumption respond slowly to shocks. In an economy populated with agents with heterogeneous commitments, the aggregate dynamics of portfolios and total consumption coincide precisely with those that arise from a representative consumer model with habit-formation preferences. Hence, consumption commitments provide neoclassical foundations for “habit” in the aggregate.

We tested the theory’s prediction that an exogenous increase in commitment should induce a household to hold a less risky portfolio. We showed that home tenure acts as a proxy for a household’s commitment level, and that marital shocks produce variation in home tenure that is exogenous to portfolio decisions. Consistent with the model’s prediction, IV estimates show that households hold much safer portfolios when they have moved into their current home recently. In other words, families take less risk when they are unlikely to move in the near future and are more constrained by the stream of financial obligations (e.g., mortgage payments) associated with their house.

The model of habit proposed here also generates many other predictions that could be tested and used to further refine the commitments theory. For instance, it implies that consumption goods that are easier to adjust – such as food – should exhibit weaker evidence of habit formation in the aggregate than more broadly defined categories of consumption. These and other time- and cross-sectional variation in asset price and 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, that is, Δ , as opposed to calendar time. The Bellman equation for the maximization problem between two reset dates is

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

which yields, using Ito's lemma

$$\rho V_t = \max_{f, \alpha} \left\{ \left(\frac{f_t^{1-\gamma}}{1-\gamma} + \mu \frac{x_t^{1-\gamma}}{1-\gamma} \right) + \frac{dV}{dt} + \frac{dV}{dw} [(r + \alpha_t \pi) w_t - c_t] + \frac{1}{2} \frac{dV^2}{d^2 w_t} (\alpha_t \sigma w_t)^2 \right\}. \quad (24)$$

We guess that the value function is of the form

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

where φ_t and $\Lambda(t)$ are deterministic functions to be determined. Here 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 (24) yields the consumption rule

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

or equivalently

$$f_t = \varphi_t^{-\frac{1}{\gamma}} (w_t - \Lambda(t)x_t) \quad (25)$$

and the investment rule

$$\alpha_t = \frac{\pi}{\gamma \sigma^2} \left(1 - \Lambda(t) \frac{x_t}{w_t} \right). \quad (26)$$

Plugging these back into the Bellman equation and simplifying yields

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

In order for this equation to hold, we need that

$$w_t - \Lambda(t)x_t = w_t - \frac{x_t}{r} - \frac{d\Lambda(t)}{dt} \frac{x_t}{r_t}$$

is satisfied. Equivalently, $d\Lambda(t)/dt = r\Lambda(t) - 1$, which, coupled with the terminal condition $\Lambda(T) = 0$ has the solution

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

Note that $\Lambda(t)$ as defined here is the present discounted value of a cash-flow of 1 every period up to the next reset date. In words, the consumer invests $\Lambda \cdot x$ dollars in a separate, safe account to be able to finance commitment consumption until the next reset date comes.

Substituting the optimal consumption and investment rules as well as the formula for Λ into the Bellman equation (24) implies after some calculations

$$\frac{\dot{\varphi}_t}{\varphi_t} = -\gamma\varphi_t^{\frac{-1}{\gamma}} + \rho - (1 - \gamma) \left\{ \frac{\pi^2}{2\gamma\sigma^2} + r \right\}. \quad (27)$$

This is a differential equation. Denoting $L = \rho - (1 - \gamma) \left\{ \frac{\pi^2}{2\gamma\sigma^2} + r \right\}$, the solution can be written as

$$\varphi_t = \left(\frac{\gamma}{L} + e^{\frac{L}{\gamma}t} K \right)^\gamma \quad (28)$$

where K is a constant that we determine below.

On a reset date, the household chooses the new level of commitment consumption by maximizing the value function

$$V_0(w_0) = \max_x \varphi_0 \frac{(w_0 - \Lambda(0)x)^{1-\gamma}}{1-\gamma} + \mu \frac{1 - e^{-\rho T}}{\rho} \cdot \frac{x^{1-\gamma}}{1-\gamma}.$$

Denoting $E = (1 - e^{-\rho T})/\rho$, the solution of this program is

$$x = w_0 \cdot \frac{(\varphi_0 \Lambda(0))^{-1/\gamma}}{\varphi_0^{-1/\gamma} \Lambda(0)^{(\gamma-1)/\gamma} + (\mu E)^{-1/\gamma}} = w_0 \cdot \kappa \quad (29)$$

where we denoted the constant multiplying w_0 by κ .

Because adjustment dates are fully forecastable, food consumption is continuous on an adjustment date (a value matching condition). Formally,

$$\varphi_T^{-1/\gamma} w_T = \varphi_0^{-1/\gamma} (w_T (1 - \Lambda(0)\kappa))$$

where the left hand side is food consumption right before adjustment, and the right hand side is food consumption immediately after adjustment, using the new, endogenously chosen level of net wealth. After some manipulations and substituting for κ this equation leads to

$$\varphi_T^{1/\gamma} = \varphi_0^{1/\gamma} + \frac{\Lambda(0)^{(\gamma-1)/\gamma}}{(\mu E)^{-1/\gamma}}$$

and substituting in the formula for φ_t given in (28) yields

$$K = \frac{\lambda(0)^{(\gamma-1)/\gamma}}{(\mu E)^{-1/\gamma} (e^{\frac{L}{\gamma}T} - 1)}.$$

This expresses K with exogenous parameters. Using the formula for K we can express the function φ from (28) with exogenous parameters only.

We have solved for all endogenous variables in the model using exogenous parameters. Our proposed value function with the implied optimal policies satisfies the Bellman equation, verifying that our guess was correct. Now revisit our assumption that t measures the time elapsed since the last reset date. Noting that the functions φ and Λ only depend on the time elapsed since the last reset date Δ , we can define

$\psi_{\Delta(t)} = \varphi_{\Delta(t)}^{-\frac{1}{\gamma}}$, and then formulas (25), (26) and (29) show that the optimal policy rules are as claimed in the theorem.

To solve for the dynamics of food consumption, first observe that the implied dynamics for the net wealth of a household between adjustments is

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

Substituting in the ODE for Λ yields

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

Turning to food consumption, we can use the above expression for the dynamics of wealth to write

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

and substituting in the differential equation (27) yields

$$\frac{df_t}{dt} = \left\{ \frac{\pi^2}{2\gamma^2\sigma^2} + \frac{1}{\gamma} \left(\frac{\pi^2}{2\sigma^2} + r - \rho \right) \right\} dt + \frac{\pi}{\gamma\sigma} dz. \quad (30)$$

Proof of Proposition 1

Let $H_t(s)$ be the share of food consumption that adjusted less than s periods ago on date t . Then $H_t(\cdot)$ is a distribution function; let $h_t(s)$ denote the associated density. If $y(q)$ is the relative food consumption of cohort q , then $h_t(s) = y(q_0(t-s))$ where $q_0(u)$ is the cohort of people who adjusted on date u . This implies that $H_{t+T}(s) = H_t(s)$, so we will use the notation $H_\tau(s)$. Note that $h_{t+u}(s) = h_t(s-u)$. The assumption that $H_t(\cdot)$ has a density is justified because $Y(\cdot)$ has a density.

By (25), the net wealth of a consumer who consumes f today and last adjusted Δ periods ago is

$$f \cdot \varphi_{\Delta}^{1/\gamma} = w^{\text{net}}.$$

Therefore, aggregate food consumption as a share of aggregate net wealth is

$$b(\tau) = \frac{1}{\int_0^T \varphi_s^{1/\gamma} dH_\tau(s)}$$

which depends on calendar time only through τ . From equation (29), a household that adjusts on date t sets its new level of commitment consumption to be

$$f_t \varphi_0^{1/\gamma} \frac{\kappa}{1 - \Lambda(0)\kappa}.$$

It follows that during a short dt time period, new commitments amount to

$$dt \cdot h_\tau(0) F_t \varphi_0^{1/\gamma} \frac{\kappa}{1 - \Lambda(0)\kappa}.$$

Current commitments can be written as a sum of new commitments undertaken during the last T periods, and commitments that were in place at date zero and have not been changed since. Using the above expression for new commitments then leads to

$$X_t = \varphi_0^{1/\gamma} \frac{\kappa}{1 - \Lambda(0)\kappa} \int_{t-T}^T h_s(0) F_s ds + k(t) X_0.$$

Here, the second term represents commitments that have not been adjusted since date zero. As a result, $k(t)$ is a deterministic function with $k(t) = 0$ if $t > T$. Importantly, the weights in the first term depend only on τ . Using the expression for $b(\tau)$, we can also rewrite aggregate commitments as a weighted average of past levels of net wealth:

$$X_t = \varphi_0^{1/\gamma} \frac{\kappa}{1 - \Lambda(0)\kappa} \int_{t-T}^T \frac{h_s(0)}{\int_0^T \varphi_v^{1/\gamma} dH_s(v)} W_s^{\text{net}} ds + k(t) X_0.$$

Again, the weights, denoted by $a(\cdot)$ in the proposition, only depend on τ .

When the distribution of $H_t(s)$ is uniform, all cohorts have equal levels of food consumption. As a result, $H_t(s)$ no longer depends on t , and both $a(\cdot)$ and $b(\cdot)$ are constants.

Proof of Proposition 2

We begin with a lemma that shows how to express a time average of food consumption using a time average of total consumption. We will use this lemma again when we discuss the model with stochastic adjustment dates.

Lemma 1 *If commitment consumption is of the form*

$$X_t = \int_0^t j(t-s) F_s ds + k(t) X_0$$

for some functions $j(\cdot)$ and $k(\cdot)$ then we can write

$$X_t = o(t) X_0 + \int_0^t \zeta(t-s) C_s ds$$

where the functions $\zeta(\cdot)$ and $o(\cdot)$ uniquely solve the integral equations

$$\zeta(u) = j(u) - \int_0^u \zeta(v) j(u-v) dv \tag{31}$$

$$o(t) = k(t) - \int_0^t \zeta(t-s) k(s) ds \tag{32}$$

with initial conditions $\zeta(0) = j(0)$, $o(0) = k(0)$.

Proof. Consider the process

$$\tilde{X}_t = o(t) X_0 + \int_0^t \zeta(t-s) C_s ds.$$

We will show that $\tilde{X}_t = X_t$ for all $t \geq 0$. First note that

$$\begin{aligned}\tilde{X}_t &= o(t)X_0 + \int_0^t \zeta(t-s) [F_s + X_s] ds \\ &= o(t)X_0 + \int_0^t \zeta(t-s)F_s + \zeta(t-s) \left[\int_0^s j(s-u)F_u du + k(s)X_0 \right] ds \\ &= o(t)X_0 + \int_0^t F_s \left[\zeta(t-s) + \int_0^{t-s} j(u)\zeta(t-s-u) du \right] ds + X_0 \int_0^t \zeta(t-s)k(s)ds.\end{aligned}$$

$X_t = \tilde{X}_t$ will hold if

$$j(t-s) = \zeta(t-s) + \int_0^{t-s} j(u)\zeta(t-s-u) du$$

or with $t-s = u$

$$\zeta(u) = j(u) - \int_0^u \zeta(v)j(u-v)dv$$

as well as

$$o(u) = k(u) - \int_0^u \zeta(u-v)k(v)dv.$$

Substituting in $u = 0$ gives $\zeta(0) = j(0)$ and $o(0) = k(0)$. The integral equation for $\zeta(u)$ then yields a unique solution, which can be used to determine $o(\cdot)$. By the above argument, a pair of functions that solve these equations also give $X_t = \tilde{X}_t$, which is the desired representation. ■

Now turn to the proof of Proposition 2. By Proposition 1, we have $j(u) = a/b$ for $0 \leq u \leq T$, and $j(u) = 0$ for $u > T$. From the lemma, we must have

$$\frac{a}{b} = \zeta(u) + \frac{a}{b} \int_0^u \zeta(s)ds \quad (33)$$

for $0 \leq u < T$, and

$$0 = \zeta(u) + \frac{a}{b} \int_{u-T}^u \zeta(s)ds \quad (34)$$

for $u \geq T$, as well as $\zeta(0) = a/b$. Differentiating the first equation for $\zeta(\cdot)$ with respect to u leads to an ODE

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

The solution for $0 \leq u < T$ can be written as

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

(34) also implies that for $u > T$

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

As a result, $\zeta(u)$ cannot have the same sign on any interval of length larger than T . For example, if it were positive on $[u-T, u]$ then $\zeta(u)$ would have to be negative. The equation also implies that

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

which shows that as long as $aT < b$, $\zeta(u)$ goes to zero geometrically, because

$$\frac{aT}{b} \cdot \max_{[(k-1)T, kT]} |\zeta(v)| > \max_{[kT, (k+1)T]} |\zeta(v)|.$$

Because $k(t) = 0$ if $t > T$, formula (32) implies that as long as $\zeta(\cdot)$ goes to zero geometrically, so will $o(\cdot)$.

To establish that the weights are bounded, first differentiate (34) in u :

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

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

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

One solution of this ODE is $\beta(u) = \exp\left(\frac{a}{b}u\right)$. By Gronwall's lemma, we can then bound the absolute value of the particular solution we are interested in by $K_2 \cdot \exp\left(\frac{a}{b}u\right)$ with some positive constant K_2 . It follows that $|\zeta(u)| < K_2 \cdot \exp\left(\frac{a}{b}u\right) \cdot \exp\left(-\frac{a}{b}u\right) = K_2$ which shows that $\zeta(u)$ is indeed bounded.

Finally, note that the parameters a and b are derived from the underlying parameters of the model. To demonstrate that the last part of the proposition has content, we now show that there are underlying parameters for which $aT < b$ 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 $aT < b$ continues to hold.

Proof of Theorem 3

Risksharing arrangements. By the argument in the main text, the household seeks to insure against two types of risk: 1) Payment risk, which is the randomness of total payments associated with the uncertain length of home tenure, and 2) food consumption risk, associated with the possibility that moving induces a jump in food consumption. The first of these can be insured by pooling payment risk across consumers. By doing so, each household swaps its uncertain payment stream for a deterministic one that equals the expected present value of outstanding housing obligations. But this implies that a move induced by an exogenous shock brings about a loss in net wealth (as well as the opportunity to adjust), because the consumer needs to finance the outstanding payment obligations of a new house. This is the second risk source mentioned above: without further insurance, the loss in net wealth would lead to a cut in food consumption.

By pooling payment risk, the state variables for the household problem become net wealth w_t^{net} , commitments x , and the time elapsed since last adjustment Δ . To deal with food consumption risk, the consumer will buy Arrow-Debreu securities that pay off on adjustment dates. Importantly, perfect insurance against such risk will be achieved by purchasing, on each date, a security that pays off if the next adjustment date happens immediately. Besides holding a time-varying amount of this security, no additional insurance is needed, because such "immediate insurance" securities together with the safe asset dynamically complete the market for adjustment date risk, and all such insurance is priced at expected value.

Formally, assume that on each date t , the consumer buys insurance that pays ωw_t^{net} in the event when the consumer has to move on date t , and zero otherwise. Note that ω may depend on Δ and other state variables. The cost of that insurance on date t is $\lambda(\Delta)dt \cdot \omega w_t^{\text{net}}$ because it is priced at expected value.

Bellman equation and optimal policy. With the above notation, we now turn to the Bellman equation of the household problem. The value function can be written as a function of the state variables w_t^{net} , x and Δ . Then dynamic optimization implies

$$\begin{aligned} \rho V_{\Delta}(w^{\text{net}}, x) &= \max_{f, \alpha, \omega} \frac{f^{1-\gamma}}{1-\gamma} + \mu \frac{x^{1-\gamma}}{1-\gamma} + \frac{dV}{d\Delta} + \frac{dV}{dw^{\text{net}}} [(r + \tilde{\alpha}\pi)w^{\text{net}} - f - \lambda(\Delta)\omega w^{\text{net}}] + \\ &\quad \frac{1}{2} \frac{d^2V}{d^2w^{\text{net}}} \alpha^2 \sigma^2 (w^{\text{net}})^2 + \lambda(\Delta) [V_0(w^{\text{new}} - \Lambda(0)x^{\text{new}}, x^{\text{new}}) - V_{\Delta}(w^{\text{net}}, x)]. \end{aligned} \quad (35)$$

Here $\tilde{\alpha}$ is the share of stocks in net wealth, so that the share of stocks in total wealth is given by $\alpha = \tilde{\alpha}w^{\text{net}}/w$. The first five terms in this Bellman equation correspond to the terms in the deterministic counterpart (24), with the only difference being that food consumption insurance leads to an outflow of $\lambda(\Delta)\omega w^{\text{net}}$ each period. The last term represents the continuation value associated with an adjustment date. The variable w^{new} stands for total wealth on adjustment, after insurance payments are settled.

Throughout the rest of the argument, we can think about the wealth used for commitment payments as wealth held in a separate account. Housing payments flow from that account, but between adjustments it does not require any cash inflows. We therefore abstract from these payments below.

We solve the Bellman equation following similar steps as in the deterministic adjustment case. We guess that the value function has the form

$$V_{\Delta}(w^{\text{net}}, x) = \varphi_{\Delta} \frac{(w^{\text{net}})^{1-\gamma}}{1-\gamma} + \mu \frac{x^{1-\gamma}}{1-\gamma} \cdot \int_0^{\infty} e^{-\rho s} \frac{1 - G(s + \Delta)}{1 - G(\Delta)} ds \quad (36)$$

where the first term is the value from food consumption before the next adjustment date plus total value after the next adjustment date; and the second term is the value of housing consumption before the next adjustment date. Our approach will be to first solve for all endogenous quantities as a function of the endogenous φ_{Δ} . This will transform the Bellman equation into an ODE for φ_{Δ} . We then solve that ODE.

Substituting our guess into (7) and taking the first order condition in f and $\tilde{\alpha}$ yields the optimal policy rules

$$f = \varphi_{\Delta}^{-1/\gamma} w^{\text{net}} \quad (37)$$

and

$$\tilde{\alpha} = \frac{\pi}{\gamma \sigma^2} w^{\text{net}}. \quad (38)$$

On an adjustment date, the optimal new level of commitment consumption x^{new} , given the new level of total wealth w^{new} , is the solution to maximizing (36):

$$\max_{x^{\text{new}}} \varphi_0 \frac{(w^{\text{new}} - \Lambda_0 x^{\text{new}})^{1-\gamma}}{1-\gamma} + \mu \frac{(x^{\text{new}})^{1-\gamma}}{1-\gamma} \cdot \int_0^{\infty} e^{-\rho s} [1 - G(s)] ds.$$

Denoting $\int_0^{\infty} e^{-\rho s} [1 - G(s)] ds = E$, we find

$$x^{\text{new}} = w^{\text{new}} \cdot \kappa$$

where κ is

$$\kappa = \frac{(\varphi_0 \Lambda_0)^{-1/\gamma}}{\varphi_0^{-1/\gamma} \Lambda_0^{(\gamma-1)/\gamma} + (\mu E)^{-1/\gamma}}.$$

Next we determine the extent of food consumption insurance by solving for the optimal choice of ω . Since the role of this insurance is to avoid surprise jumps in food consumption, we can pin down ω by equating food consumption just before and just after an adjustment date.³⁵

$$\varphi_\Delta^{-1/\gamma} w^{\text{net}} = \varphi_0^{-1/\gamma} (w^{\text{new}} - \Lambda(0)x^{\text{new}}) = \varphi_0^{-1/\gamma} (w^{\text{net}}(1 + \omega)(1 - \Lambda_0 \kappa))$$

Here the left hand side is food consumption just before adjustment. To understand the right hand side, note that upon adjustment, the household earns ωw^{net} from the insurance policy, but it has to put down a share $\Lambda(0)\kappa$ of its new wealth to the separate account associated with housing payments. Substituting the formula for κ into this equation yields

$$\omega = \frac{\varphi_\Delta^{-1/\gamma}}{\varphi_0^{-1/\gamma}} \frac{1}{1 - \Lambda_0 \kappa} - 1 = \frac{\varphi_\Delta^{-1/\gamma}}{\varphi_0^{-1/\gamma}} \frac{\varphi_0^{-1/\gamma} \Lambda_0^{(\gamma-1)/\gamma} + (\mu E)^{-1/\gamma}}{(\mu E)^{-1/\gamma}} - 1$$

which expresses ω with exogenous variables and the function φ .

We now turn to substitute the optimal policy rules into the Bellman equation. This requires evaluating the continuation value of the household on an adjustment fate (i.e., calculating the last term in (7)):

$$V_0(w^{\text{new}}, x^{\text{new}}) = \varphi_0 \frac{(w^{\text{net}}(1 + \omega)(1 - \Lambda_0 \kappa))^{1-\gamma}}{1 - \gamma} + \mu E \cdot \frac{(w^{\text{net}}(1 + \omega)\kappa)^{1-\gamma}}{1 - \gamma}$$

and using the formulas for ω and κ we obtain

$$V_0(w^{\text{new}}, x^{\text{new}}) = \frac{(w^{\text{net}})^{1-\gamma}}{1 - \gamma} \varphi_\Delta^{(\gamma-1)/\gamma} \left[\varphi_0^{1/\gamma} + (\mu E)^{1/\gamma} (\Lambda_0)^{(\gamma-1)/\gamma} \right].$$

Substitution of these formulas into the Bellman equation (7) leads, after algebraic manipulations, to

$$\rho - (1 - \gamma)(r + \frac{1}{2} \frac{\pi^2}{\gamma \sigma^2}) + \gamma \lambda(\Delta) = \frac{\varphi'_\Delta}{\varphi_\Delta} + \gamma \varphi_\Delta^{-1/\gamma} \left[1 + \lambda(\Delta) \left(\varphi_0^{1/\gamma} + (\mu E)^{1/\gamma} (\Lambda_0)^{(\gamma-1)/\gamma} \right) \right]. \quad (39)$$

Solving the ODE. The last formula is an ordinary differential equation for the function φ_Δ . Our final task is to prove that this ODE has a solution. To gain some intuition, first assume that $\lambda(\Delta)$ is the constant function $\lambda(\Delta) = \bar{\lambda}$. In this case, φ also has to be a constant $\bar{\varphi}$, because Δ contains no information about future adjustment dates. As a result, the ODE collapses into an expression for $\bar{\varphi}$:

$$\bar{\varphi}^{-1/\gamma} = \frac{\rho - (1 - \gamma)(r + \frac{1}{2} \frac{\pi^2}{\gamma \sigma^2})}{1 + \bar{\lambda} \Lambda_0^{(\gamma-1)/\gamma} (\mu E)^{1/\gamma}} \quad (40)$$

a well defined, positive number, given that we assume $\gamma > 1$. Importantly, this formula implies that $\bar{\varphi}$ is increasing in $\bar{\lambda}$.

Now go back to the general ODE. Recall the notation $\lambda(0) = \lambda_L$ and $\lambda(T) = \lambda_H$, and that for $s > T$, $\lambda(s) = \lambda_H$ is a constant. Define the values φ_L and φ_H that would prevail if λ was constant at levels λ_L respectively λ_H ; in other words, φ_L and φ_H are defined by (40) with $\lambda = \lambda_L$ and $\lambda = \lambda_H$.

³⁵Importantly, maximizing the objective function directly in ω gives the same result.

We need to find an initial condition for the general ODE (39). We do so by making use of a terminal condition. As noted above, for $s > T$, we have $\lambda(s) = \lambda_H$ a constant. As a result, for $s > T$ we must have $\varphi_s = \varphi_H$ a constant. This terminal condition would pin down the path of the solution to the ODE, except for the complication that the value of φ_0 appears in equation (39). This implies that we cannot solve the ODE backwards without knowing the value of φ_0 .

We deal with this problem by showing that there exists an initial condition such that the dynamic system achieves the desired terminal condition. First assume that we start the ODE from the initial value φ_L , using the actual $\lambda(\Delta)$ function as the control. Since φ'_Δ is decreasing in λ in formula (39), as λ rises in the beginning, φ_Δ will have to fall. And then φ_Δ will never reach the level of φ_L again, because at that point, the derivative would be negative. In particular, at time T , $\varphi_T \leq \varphi_L$.

By a very similar argument, one can show that when started from the initial value φ_H , the path of φ_Δ will always be weakly above φ_H , and so $\varphi_T \geq \varphi_H$. Finally, note that $\varphi_L \leq \varphi_H$ because by assumption $\lambda_L \leq \lambda_H$. The continuity of the solution path as a function of the initial condition implies that for some intermediate initial value φ_0 , the path of φ_Δ hits φ_H exactly at time T . We have found a solution that satisfies the terminal condition, and this gives the solution to the household problem.

We can derive the time path of food consumption for the household the same way as in the deterministic adjustment case, to find that

$$f_t = f_s \exp \left\{ \frac{\pi}{\gamma\sigma} (z_t - z_s) + \frac{1}{\gamma} \left(\frac{\pi^2}{2\sigma^2} + r - \rho \right) (t - s) \right\}. \quad (41)$$

Since this does not depend on adjustment dates, we conclude that the food consumption of all households is perfectly correlated. This verifies that the optimal policy achieves perfect risksharing across consumers in the population.

Finally, we show, as claimed by the theorem, that $\Lambda(\cdot)$ is a decreasing function. Note, the fact that $\lambda(\cdot)$ is increasing means that $\log(1 - G(s))$ is concave and decreasing. As a result, $(1 - G(s + \Delta))/(1 - G(\Delta))$ is a decreasing function of Δ . But then the weights in (14) are all decreasing, which shows that $\Lambda(\cdot)$ is decreasing.

Aggregation. Now consider the population of consumers. Let $H_t(s)$ be the share of aggregate food consumption consumed by households who adjusted less than s periods ago. A household who consumes f_t today and last adjusted Δ periods ago has net wealth given by

$$f \cdot \varphi_\Delta^{1/\gamma} = w^{\text{net}}.$$

Therefore aggregate food consumption as a share of aggregate net wealth can be written as

$$\frac{1}{\int_s \varphi_s^{1/\gamma} dH_t(s)}. \quad (42)$$

Thus the ratio of aggregate food consumption to aggregate net wealth will be a constant (that is, independent of t) if the distribution $H_t(s)$ does not depend on t .

Ergodic distribution. We show that such a steady state cross-sectional distribution $H(s)$ exists. The evolution of $H_t(s)$ over time is given by

$$\frac{d(1 - H_t(s))}{dt} = h_t(s) - \int_s^\infty \lambda(u) h_t(u) du \quad (43)$$

where we assumed that $H_t(\cdot)$ is differentiable with density $h_t(\cdot)$. This assumption will hold if the initial distribution $H_0(\cdot)$ is smooth. The intuition behind the equation is as follows. During a short dt time period, $1 - H_t(s)$, the share of food consumption that adjusted more than s periods ago, changes for two reasons. First, households who adjusted $s - dt$ periods ago are added to this group. Second, a consumer who adjusted u periods ago adjusts with probability $\lambda(u)$, leaving this group. An ergodic distribution would correspond to $dH_t(s)/dt = 0$. Substituting in, we find the condition

$$h(s) = \int_s^\infty \lambda(u) h(u) du$$

for an ergodic distribution $H(s)$. Differentiating in s yields the differential equation

$$h'(s) = -\lambda(s) f(s)$$

which can be solved to give

$$h(s) = K_3 \cdot \exp \left\{ - \int_0^s \lambda(u) du \right\}.$$

Because $\lambda(\cdot)$ is nondecreasing, the integral of this density is finite, and K_3 will be determined by setting the integral to one. Using (43), it is easy to verify that the resulting $H(s)$ distribution is truly ergodic. If the population is started from that initial distribution, it will always have the same cross-sectional distribution over time.

Habit representation. A household that adjusts on date t will set its new level of commitment consumption to equal

$$f_t \cdot \varphi_0^{1/\gamma} \frac{\kappa}{1 - \Lambda_0 \kappa}.$$

Hence during a short dt time period, new commitments as a share of aggregate food consumption are

$$dt \cdot \varphi_0^{1/\gamma} \frac{\kappa}{1 - \Lambda_0 \kappa} \cdot \int_s^\infty \lambda(s) dH(s), \quad (44)$$

which is a constant. Because the aggregate food consumption to net wealth ratio is also a constant b , it follows that every day, new commitments equal a fixed proportion of aggregate net wealth. If that proportion is a then we have

$$X_t = \int_0^t \frac{a}{b} F_s \cdot (1 - G(t - s)) ds + X_0 k(t) \quad (45)$$

where $k(t)$ is the share of aggregate commitment consumption that has not adjusted since date zero. Formally, we have

$$k(t) = \int_0^\infty x(\Delta) \frac{1 - G(t + \Delta)}{1 - G(\Delta)} d\Delta$$

where $x(\Delta)$ is the share of commitment consumption by people who at date zero have not adjusted for Δ periods. Applying Lemma 1 immediately gives a representation of commitments as a time average of past total consumption.

In the special case when $\lambda(\Delta)$ is a constant, we have $G(s) = 1 - e^{-\lambda s}$, and from (45) we need to use Lemma 1 with $j(u) = (a/b)e^{-\lambda u}$. The integral equations (31) and (32) are satisfied with $\zeta(u) = a \cdot e^{-(\lambda + a/b)u}$ and $o(u) = e^{-(\lambda + a/b)u}$. Thus with a choice of $d = \lambda + a/b$ and $D = a/b$ we get equation (13). In this special case, the adjustment process has no memory, and the representation holds for any initial distribution.

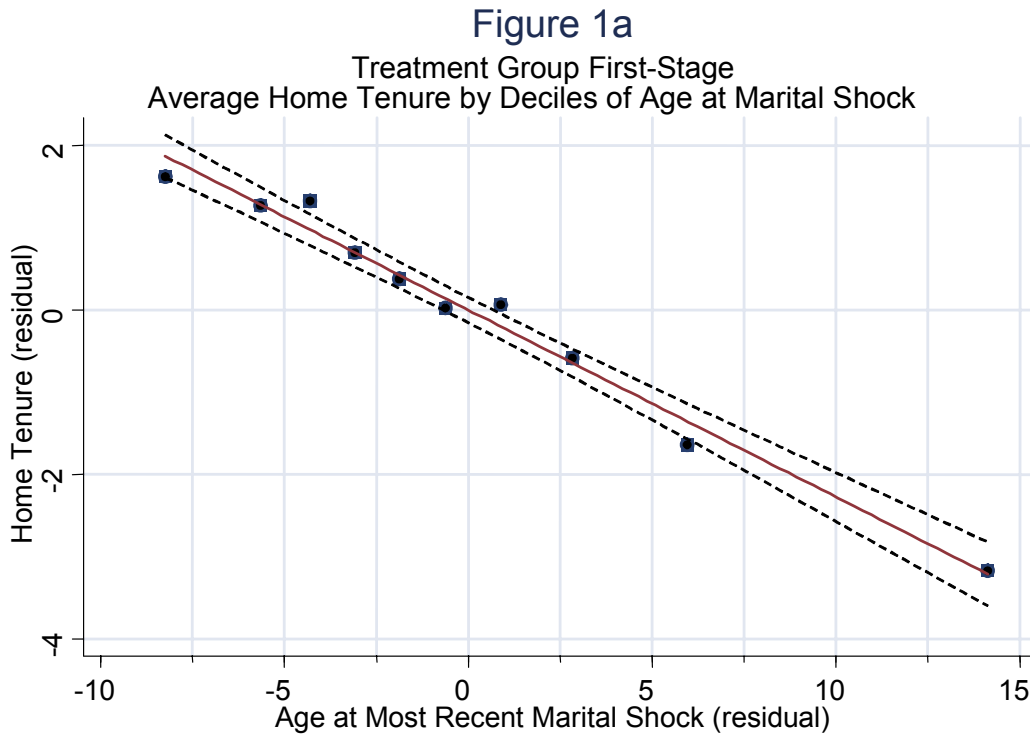
References

- Attanasio, O. (2000). Consumer Durables and Inertial Behavior: Estimation and Aggregation of (S,s) Rules for Automobile Purchases, *Review of Economic Studies* **67**: 667–696.
- 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.
- Boldrin, M., Christiano, L. J. and Fisher, J. D. M. (2001). Habit Persistence, Asset Returns and the Business Cycle, *American Economic Review* **91**: 149–166.
- Calvo, G. (1983). Staggered Prices in a Utility Maximizing Framework, *Journal of Monetary Economics* **12**: 383–398.
- 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. Working paper, John Hopkins University.
- Carroll, C. D., Overland, J. and Weil, D. N. (2000). Saving and Growth with Habit Formation, *American Economic Review* **90**: 341–355.
- Chen, X. and Ludvigson, S. (2004). Land of addicts? An Empirical Investigation of Habit-Based Asset-Pricing Models. Working paper, New York University.
- Chetty, R. (2003). Consumption Commitments, Unemployment Durations, and Local Risk Aversion, *NBER Working Paper* **10211**.
- Christiano, L. J., Eichenbaum, M. and Evans, C. L. (2004). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy*, forthcoming.
- Cocco, J. (2005). Portfolio Choice in the Presence of Housing. *Review of Financial Studies*, forthcoming.
- 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.
- Cox, J. C., Ingersoll, J. E. and Ross, S. A. (1985). An Intertemporal General Equilibrium Model of Asset Prices, *Econometrica* **53**: 363–384.
- Dybvig, P. H. (1995). Dusenberry's Ratcheting of Consumption: Optimal Dynamic Consumption and Investment Given Intolerance for any Decline in Standard of Living, *Review of Economic Studies* **62**: 287–313.
- Dynan, K. E. (2000). Habit Formation in Consumer Preferences: Evidence from Panel Data, *American Economic Review* **90**: 391–406.
- Eberly, J. (1994). Adjustment of Consumers' Durables Stocks: Evidence from Automobile Purchases, *Journal of Political Economy* **102**: 403–436.
- Eichenbaum, M., Hansen, L. P. and Singleton, K. J. (1988). A Time Series Analysis of Representative Agent Models of Consumption and Leisure under Uncertainty, *Quarterly Journal of Economics* **103**: 51–78.
- Ferson, W. E. and Constantinides, G. (1991). Habit Persistence and Durability in Aggregate Consumption: Empirical Tests, *Journal of Financial Economics* **29**: 199–240.
- Flavin, M. (2001). Owner-Occupied Housing in the Presence of Adjustment Costs: Implications for Asset Pricing and Nondurable Consumption. Working Paper, UCSD.
- Flavin, M. and Nakagawa, S. (2003). A Model of Housing in the Presence of Adjustment Costs: A Structural Model of Habit Persistence. Working Paper, UCSD.
- Fratantoni, M. C. (2001). Homeownership, Committed Expenditure Risk, and the Stockholding Puzzle, *Oxford Economic Papers* **53**: 241–259.
- Fuhrer, J. C. (2000). Habit Formation in Consumption and Its Implications for Monetary-Policy Models, *American Economic Review* **90**: 367–390.
- 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.
- Gruber, J. (1997). The Consumption Smoothing Benefits of Unemployment Insurance, *American Economic Review* **87**: 192–205.
- Gruber, J. (1998). Unemployment Insurance, Consumption Smoothing, and Private Insurance: Evidence from the PSID and CEX, *Research in Employment Policy* **1**: 3–32.
- Guiso, L., Japelli, T. and Terlizzese, D. (1996). Income Risk, Borrowing Constraints and Portfolio Choice, *American Economic Review* **86**: 158–172.
- Heaton, J. (1993). The Interaction between Time-Nonseparable Preferences and Time-Aggregation, *Econometrica* **61**: 353–385.
- Jagannathan, R. and Wang, Y. (2004). Consumption Risk and the Cost of Equity Capital. Working paper, Kellogg School of Management.
- Koren, M. and Szeidl, A. (2003). Portfolio Choice with Illiquid Assets. CEPR Working Paper 3795.
- Li, R. (2003). Household Portfolio Choice with Precommitment. Working paper, University of Toronto.
- Ljungqvist, L. and Uhlig, H. (2000). Tax Policy and Aggregate Demand Management under Catching Up with the Joneses, *American Economic Review* **90**: 356–366.
- Lupton, J. P. (2003). Household Portfolio Choice and the Habit Liability: Evidence from Panel Data. Working paper, Federal Reserve Board.
- Lustig, H. and Nieuwerburgh, S. V. (2003). Housing Collateral, Consumption Insurance and Risk Premia. Working paper, University of Chicago and NYU Stern.
- Martin, R. F. (2003). Consumption, Durable Goods and Transaction Costs. International Finance Discussion Papers 756, Board of Governors of the Federal Reserve System.
- McCallum, B. T. and Nelson, E. (1999). Nominal Income Targetting in an Open-Economy Optimizing Model, *Journal of Monetary Economics* pp. 553–578.


- Merrill, S. R. (1984). Home Equity and the Elderly. In H. Aaron and G. Burtless (editors), Retirement and Economic Behavior. Brookings Institution, Washington DC.
- Olney, M. (1999). Avoiding Default: The Role of Credit in the Consumption Collapse of 1930, *Quarterly Journal of Economics* pp. 319–335.
- Parker, J. A. and Julliard, C. (2004). Consumption Risk and the Cross-Section of Expected Returns. *Journal of Political Economy*, forthcoming.
- Piazzesi, M., Schneider, M. and Tuzel, S. (2003). Housing, Consumption and Asset Pricing. Working paper, University of Chicago.
- Postlewaite, A., Samuelson, L. and Silverman, D. (2004). Consumption Commitments and Preferences for Risk. Working Paper, Penn Institute for Economic Research.
- Ravina, E. (2004). Keeping Up with the Joneses: Evidence from Micro Data. Working Paper, Northwestern University.
- Reis, R. (2003). Inattentive Consumers. Working paper, Harvard University.
- Shea, J. (1994). Should We Test the Life Cycle-Permanent Income Hypothesis with Food Consumption Data?, *Economics Letters* **45**: 63–68.
- Shore, S. H. and Sinai, T. (2004). Household Risks and the Demand for Housing Commitments. Working Paper, The Wharton School, University of Pennsylvania.
- Sinai, T. (1997). Taxation, User Cost, and Household Mobility Decisions. Working Paper, The Wharton School, University of Pennsylvania.
- Sundaresan, S. M. (1989). Intertemporally Dependent Preferences and the Volatility of Consumption and Wealth, *Review of Financial Studies* **3**: 73–89.
- Taylor, J. (1979). Staggered Wage Setting in a Macro Model, *American Economic Review* **69**: 108–113.
- Venti, S. F. and Wise, D. A. (1989). Aging, Moving, and Housing Wealth. in D. Wise (editor) The Economics of Aging, University of Chicago Press.
- Venti, S. F. and Wise, D. A. (2000). Aging and Housing Equity, *NBER Working Paper* **7882**.
- Yogo, M. (2003). A Consumption-based Explanation of Expected Stock Returns. Working paper, Harvard University.



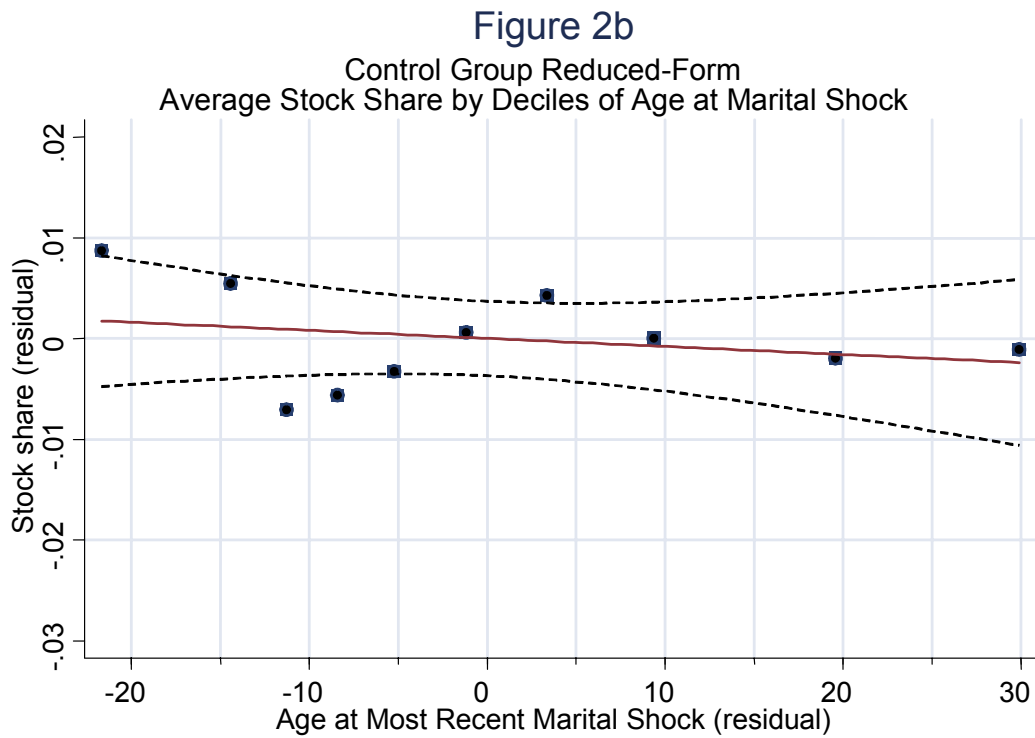
NOTE – This figure is a residual plot showing the data underlying the first-stage relationship between the timing of marital shocks and home tenure for the treatment group (see Table 2 for definition). Each point in this figure represents 10% of the treatment group (920 observations). These points are constructed using the following steps: First, raw home tenure residuals are computed for each observation of the treatment group. This is done by regressing home tenure on all covariates listed in Table 4A, column 4 (shares specification, full controls) except the marital shock variable, and computing residuals from this regression. Second, raw residuals for the age at most recent marital shock variable are computed using an analogous regression specification. Finally, we break the marital shock residual distribution into deciles, and compute the mean age at marital shock residual and home tenure residual for each decile. The figure plots the mean home tenure residual against the marital shock residual by deciles of the marital shock residual. The best-fit (OLS) line for these ten points and the 95% confidence interval for this line are also shown.



NOTE – This figure is a residual plot showing the data underlying the reduced-form relationship between the timing of marital shocks and the share of stocks in a household’s total wealth for the treatment group. As in Figure 1a, each point represents 10% of the treatment group (920 observations). The x-coordinates are constructed exactly as in Figure 1a. The y coordinates are computed by taking means of stock share residuals by decile of the marital shock residual from a regression equivalent to that in Table 4A, column 4, except that the marital shock variable is omitted. The figure plots the mean stock share residual against the marital shock residual by deciles of the marital shock residual. The best-fit (OLS) line for these ten points and the 95% confidence interval for this line are also shown.



NOTE – This figure replicates Figure 1a for the “pooled control” group (see Table 2 for definition), showing that there is no first-stage relationship between home tenure and the timing of the most recent marital shock for this group. Each point represents 10% of the pooled control group (560 observations). See notes to Figure 1a for details on how these points are constructed. The relative scale of the x and y axes is the same as in Figure 1a, making all slopes directly comparable across the two figures.



NOTE – This figure replicates Figure 1b for the “pooled control” group (see Table 2 for definition), showing that there is no reduced-form relationship between stockholding and the timing of the most recent marital shock for this group. Each point represents 10% of the pooled control group (560 observations). See notes to Figure 1b for details on how these points are constructed. The relative scale of the x and y axes is the same as in Figure 1b, making all slopes directly comparable across the two figures.

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

NOTE – Source: BLS tabulations from Consumer Expenditure Survey, 2000. Take-home pay defined as gross income net of taxes and mandatory insurance/pension contributions. National Income and Product Accounts (NIPA) definition of “nondurables and services” includes all categories in this table except vehicles and furniture.

TABLE 2
DESCRIPTIVE STATISTICS

Variable	Core Sample	Treatment	Controls		
	Married Households	Recent-Mar Homeowners	Long-Married Homeowners	Married Renters	Divorced Homeowners
Total household wealth	119,179 (231,567)	243,621 (396,717)	348,465 (466,744)	93,675 (183,199)	210,415 (330,228)
Home equity	48,699 (60,029)	79,748 (67,946)	104,400 (71,933)	0	82,816 (67,495)
Property Value	95,047 (74,849)	135,114 (80,623)	115,101 (76,177)	0	107,656 (72,177)
Liquid wealth	42,555 (184,140)	113,866 (348,152)	197,061 (429,298)	62,573 (149,175)	99,926 (298,108)
Stocks	15,888 (154,028)	59,978 (304,602)	104,884 (382,517)	30,653 (93,287)	49,329 (276,169)
Savings accts, CDs, money mkt	9,847 (21,736)	16,177 (27,103)	31,758 (38,113)	11,168 (22,180)	19,138 (26,698)
Municipal and corporate bonds	3,915 (28,982)	9,129 (44,927)	22,323 (74,871)	5,976 (34,016)	13,691 (55,678)
IRA wealth	7,734 (23,040)	17,971 (35,562)	23,438 (40,483)	8,110 (20,077)	10,523 (25,823)
Other assets	5,172 (53,566)	10,610 (104,125)	14,658 (58,262)	6,666 (62,145)	7,245 (37,662)
Other real estate	12,214 (48,264)	25,693 (72,822)	29,711 (72,211)	14,572 (49,913)	16,089 (54,591)
Business equity	8,027 (47,081)	13,306 (62,030)	6,476 (46,877)	9,372 (51,108)	5,011 (38,839)
Vehicle equity	7,685 (07,499)	11,009 (08,744)	10,816 (08,653)	7,158 (06,976)	6,572 (06,730)
Unsecured debt	4,428 (18,908)	4,856 (18,388)	1,380 (10,450)	6,136 (16,622)	2,985 (47,769)
Annual income	43,464 (34,578)	66,349 (46,571)	41,600 (28,515)	48,637 (32,149)	33,495 (30,070)
Years of education	12.86 (3.18)	14.60 (2.59)	13.29 (2.92)	14.43 (2.7)	13.46 (2.8)
Number of children at home	0.59 (1.00)	0.61 (.97)	0.01 (.15)	0.59 (1.01)	0.15 (.5)
Age	47.62 (15.23)	46.53 (11.25)	70.57 (6.44)	42.03 (14.94)	61.87 (15.07)
Age at first marriage	23.31 (5.19)	24.21 (5.36)	22.98 (3.8)	24.37 (5.13)	22.96 (5.32)
Age at termination of 1st marriage	31.24 (10.13)	34.80 (11.18)			
Home tenure (years)	13.63 (11.99)	11.35 (9.49)	25.01 13.79	5.92 (7.71)	20.49 (16.02)
Sample size	55,288	9,310	2,244	1,010	2,716

NOTE – This table reports means and standard deviations of variables for treatment and control groups. The data source is the Survey of Income and Program Participation 1990-96 Asset Topical Modules. All monetary variables are in real 1990 dollars. Liquid wealth is defined as total wealth minus all home equity, business equity, and vehicle equity. Other assets comprise money owed to respondent, savings bonds, checking accounts, and equity in other investments. The core sample consists of all married households in the dataset. The treatment sample consists of once or twice-married stockholding homeowners whose current marriage duration is less than 40 years. The long-married homeowners sample consists of once or twice-married stockholding homeowners who have been married for more than 40 years. The married renters sample includes once or twice-married stockholding renters. The divorced homeowners sample includes once-married, currently single stockholding homeowners. For age at termination of first marriage, sample size is 11,942 in column 1 and 1,967 in column 2.

TABLE 3
FIRST-STAGE: EFFECT OF MARITAL SHOCKS ON HOME TENURE

	(1)	(2)	(3)	(4)	(5)	(6)
	Age at 1st mar Once-married	Age at Term Twice-Married	Remarriage Pooled Treat	Combined Instrument	Combined Inst. Full controls	Combined Inst. Entire sample
Dependent Variable: hometenure hometenure hometenure hometenure hometenure hometenure						
Age at first marriage	-0.206 (0.021)	-0.148 (0.072)				
Age at termination of first marriage		-0.178 (0.036)				
Remarriage indicator			-2.981 (0.231)			
Age at most recent marital shock				-0.253 (0.016)	-0.218 (0.016)	-0.169 (0.00707)
Years of education					-0.120 (0.041)	-0.118 (0.0159)
no. of children in home					0.021 (0.079)	-0.0865 (0.0316)
Business equity					-1.04e-6 (1.29e-6)	3.73e-7 (7.70e-7)
Vehicle equity					-8.35e-6 (9.87e-6)	-1.14e-5 (5.74e-6)
Unsecured debt					-1.04e-7 (3.22e-6)	-2.88e-6 (1.36e-6)
Annual income					4.90e-7 (1.83e-6)	1.00e-6 (1.24e-6)
Age spline	x	x	x	x	x	x
Total wealth spline	x	x	x	x		
Liquid wealth spline					x	x
Home equity spline					x	x
Property value spline					x	x
Year fixed effects	x	x	x	x	x	x
Occup. fixed effects					x	x
Industry fixed effects					x	x
Observations	7319	1856	9165	9175	9175	55535

NOTE – Heteroskedasticity robust standard errors are reported in parentheses. Columns 1-5 use the treatment group (once or twice married stockholding homeowners who current marriage duration is less than 40 years). Column 1 includes only once-married individuals; column 2 only twice-married individuals. Column 3 includes the entire treatment group except 10 households with wealth above \$5 million; see text for further details. Columns 4 and 5 include the entire treatment group. Column 6 includes all currently married homeowners. Age at most recent marital shock is defined as max(age at first marriage, age at termination of first marriage). Remarriage indicator is 0 for once-married reference person and 1 for twice-married. All splines are 10-piece linear splines partitioned by deciles of the relevant variable.

TABLE 4A
EXOGENEITY TESTS FOR AGE AT MOST RECENT MARITAL SHOCK INSTRUMENT

Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	<u>Treatment</u>	<u>Pooled control groups</u>		<u>Treatment</u>	<u>Pooled cntrl</u>	<u>Renters control</u>	
	Red-form	1st stage	Red-form	Red-form	Red-form	1st stage	Red-form
	stocks	htenure	stocks	stock share	stock share	htenure	stocks
Age at most recent marital shock	-658.313 (241.788)	0.0094 (0.0130)	29.711 (76.827)	-0.00079 (0.00024)	-0.000081 (.000138)	-0.062 (0.059)	690.953 (401.695)
Years of education	-1,005.297 (369.164)	-0.299 (-0.075)	-2,188.84 (561.30)	-0.00182 (0.00069)	-0.00122 (0.00083)	-0.147 (0.111)	582.856 (872.52)
Children at home	945.478 (1,427.904)	0.64 (-0.428)	2,919.61 (1,182.79)	0.00291 (0.00177)	0.0120 (0.0052)	0.220 (0.183)	4,097.37 (1,754.38)
Business equity	-0.012 (0.031)	2.51e-6 (3.96e-6)	-0.022 (0.041)	0.18647 (0.15471)	-0.470 (0.770)	8.70e-6 (3.55e-6)	-0.096 (0.048)
Vehicle equity	-0.275 (0.220)	-2.34e-5 (2.49e-5)	-0.345 (0.196)	0.02142 (0.15796)	-0.474 (0.770)	-3.14e-5 (3.32e-5)	0.050 (0.282)
Unsecured debt	-0.105 (0.144)	3.74e-6 (1.87e-6)	0.007 (0.022)	0.00302 (0.00437)	0.00237 (0.00412)	-6.60e-6 (7.15e-6)	0.026 (0.071)
Annual income	-0.186 (0.031)	-7.54e-6 (6.23e-6)	-0.446 (0.100)	-0.01166 (0.00476)	-0.00096 (0.00143)	-6.61e-6 (7.12e-6)	-0.229 (0.114)
Age spline	x	x	x	x	x	x	x
Liquid wealth spline	x	x	x	x	x	x	x
Home equity spline	x	x	x	x	x		
Property value spline	x	x	x	x	x		
Year fixed effects	x	x	x	x	x	x	x
Occup. fixed effects	x	x	x	x	x	x	x
Industry fixed effects	x	x	x	x	x	x	x
Observations	9222	4815	5970	9208	5959	1002	1010

NOTE – Heteroskedasticity robust standard errors are reported in parentheses. Columns 1 and 4 use the treatment group (once or twice married stockholding homeowners whose current marriage duration is less than 40 years). Columns 2, 3 and 5 use pooled control group of long-married homeowners + married renters + divorced homeowners (see Table 2 for the exact definition). Columns 6 and 7 include currently married stockholding renters who married once or twice. In columns 4 and 5, all monetary variables including the dependent variable and the splines are normalized by total wealth; observations with negative total wealth are omitted. Age at most recent marital shock is defined as max(age at first marriage, age at termination of first marriage). All splines are 10-piece linear splines partitioned by deciles of the relevant variable.

TABLE 4B
ADDITIONAL EXOGENEITY TESTS

	(1)	(2)	(3)	(4)	(5)	(6)
	Once married	Long-married homeowners		Twice married	Divorced homeowners	
	<u>Treatment</u>	<u>Control</u>		<u>Treatment</u>	<u>Control</u>	
	Red-form	1st stage	Red-form	Red-form	1st stage	Red-form
Dependent variable:	stocks	hometenure	stocks	stocks	hometenure	stocks
Age at first marriage	-805.127 (452.026)	0.0897 (0.0967)	701.122 (651.596)	-266.135 (610.08)	-0.0466 (0.0497)	-24.846 (230.233)
Age at termination of first marriage				-569.127 (304.93)	-0.033 (0.033)	-18.324 (137.485)
Years of education	-1142.926 (364.697)	-0.293 (0.115)	-483.202 (647.720)	-345.447 (1046.12)	-0.266 (0.115)	-1644.023 (591.835)
Children at home	1539.194 (1814.173)	-2.906 (1.375)	22483.9 (8157.66)	-1098.103 (1606.53)	1.171 (0.502)	2313.10 (1882.158)
Business equity	-0.0331 (0 .0361)	1.24e-6 (5.40e-6)	0.0524 (0 .0580)	0.0411 (0.0322)	8.69e-6 (6.29e-6)	-0.0173 (0.0400)
Vehicle equity	-0.1392 (0.2624)	3.23e-6 (3.57e-5)	0.267 (0.301)	-0.656 (0.344)	-2.06e-5 (3.98e-5)	-0.294 (0.240)
Unsecured debt	-0.164 (0.1951)	-1.50e-5 (1.54e-5)	0.0345 (0.1088)	0.0144 (0.0369)	3.20e-6 (1.44e-6)	0.0051 (0.0131)
Annual income	-0.1412 (0.0425)	-1.13e-5 (1.02e-5)	-0.395 (0.143)	-0.3075 (0.0699)	4.62e-7 (7.99e-6)	-0.264 (0.102)
Age spline	x	x	x	x	x	x
Liquid wealth spline	x	x	x	x	x	x
Home equity spline	x	x	x	x	x	x
Property value spline	x	x	x	x	x	x
Year fixed effects	x	x	x	x	x	x
Occup. fixed effects	x	x	x	x	x	x
Industry fixed effects	x	x	x	x	x	x
Observations	7352	2096	2113	1870	2401	2421

NOTE – Heteroskedasticity robust standard errors are reported in parentheses. Column 1 includes once-married and column 4 includes twice-married households in the treatment group. Sample in columns 2 and 3 consists of currently married stockholding homeowners, married once or twice, whose current marriage has lasted for at least 40 years. Columns 5 and 6 include once-married, currently single stockholding homeowners. All splines are 10-piece linear splines partitioned by deciles of the relevant variable.

TABLE 5A
EFFECT OF HOME TENURE ON PORTFOLIOS: TWO-STAGE LEAST SQUARES ESTIMATES

Instrument:	(1)	(2)	(3)	(4)	(5)	(6)
	Stocks	Stocks	Safe assets	Stocks	Stocks	Stocks
	Age at most recent marital shock			Age at 1st mar	Age at 1st term	Remarriage
Dependent variable:	stocks	stocks	safe assets	stocks	stocks	stocks
Home tenure (years)	3,190.471 (1,162.787)**	3,081.985 (1,132.402)**	-1,558.397 (523.774)**	5375.072 (3,127.94) ⁺	4055.708 (2,149.67) ⁺	1480.733 (621.94)*
Age at first marriage					304.67 (731.083)	
Business equity		-0.008 (0.032)	0.001 (0.014)	-0.035 (-0.038)	0.071 (0.037) ⁺	-0.012 (0.026)
Vehicle equity		-0.251 (0.226)	0.312 (0.099)**	-0.167 (-0.26)	-0.43 (-0.357)	-0.286 (0.132)*
Unsecured debt		-0.104 (0.145)	-0.048 (0.034)	-0.169 (-0.198)	0.029 (0.04)	0.009 (0.036)
Annual income		-0.187 (0.031)**	0.119 (0.022)**	-0.142 (0.044)**	-0.313 (0.070)**	-0.12 (0.026)**
Years of education		-629.015 (393.731)	-121.083 (292.721)	-457.472 (-526.026)	487.005 (1289.25)	-544.10 (328.46) ⁺
Children in home		897.070 (1,436.548)	723.017 (569.442)	1522.979 (-1860.642)	-1376.64 (1846.70)	-243.95 (801.73)
Age spline	x	x	x	x	x	x
Total wealth spline	x					
Liquid wealth spline		x	x	x	x	x
Home equity spline		x	x	x	x	x
Property value spline		x	x	x	x	x
Year fixed effects	x	x	x	x	x	x
Occup. fixed effects		x	x	x	x	x
Industry fixed effects		x	x	x	x	x
Observations	9175	9175	9175	7319	1856	9165

NOTE – Heteroskedasticity robust standard errors are reported in parentheses: + denotes significance at 10%; * denotes significance at 5%; ** denotes significance at 1%. In columns 1-3 home tenure is instrumented by age at most recent marital shock = max(age at first marriage, age at termination of first marriage) in the treatment group (see Table 2 for definition). In column 3, the dependent variable (safe assets) is defined as the sum of total wealth held in savings accounts, CDs, money market, and municipal and corporate bonds. Columns 4-5 split the treatment group into two subsamples. In column 4 home tenure is instrumented by age at first marriage in the once-married subsample. In column 5 home tenure is instrumented by age at termination of first marriage in the twice-married subsample. Column 6 instruments for home tenure using the remarriage indicator (1 for twice-married and 0 for once-married individuals) in the entire treatment group, excluding the 10 households with wealth above \$5 million for reasons discussed in the text. All splines are 10-piece linear splines partitioned by deciles of the relevant variable.

TABLE 5B
TWO-STAGE LEAST SQUARES ESTIMATES: ROBUSTNESS CHECKS

	(1)	(2)	(3)	(4)
	Entire sample	Shares	Shares	Bond market
	Full controls	Full controls	High wealth	participation
Dependent variable:	stocks	stk share	stk share	bondholder
Home tenure (years)	567.599 (289.938)*	0.0036 (0.0011)**	0.0064 (0.0019)**	-0.0077 (0.0030)**
Business equity	-0.011 (0.011)	0.2320 (0.1591)	-0.1119 (0.2238)	-5.29e-8 (7.11e-8)
Vehicle equity	-0.182 (0.082)*	0.0648 (0.1615)	-0.4085 (0.2611)	0.0015 (0.0005)**
Unsecured debt	-0.027 (0.031)	0.0026 (0.0049)	0.0501 (0.0449)	-3.26e-7 (2.49e-7)
Annual income	-0.127 (0.025)**	-0.0129 (0.0050)*	-0.1344 (0.0234)**	2.73e-7 (1.04e-7)**
Years of education	-300.853 (136.202)*	-0.0009 (0.0008)	0.0012 (0.0015)	0.0152 (0.0019)**
Children in home	155.408 (261.691)	0.0029 (0.0018)	-0.0033 (0.0042)	0.0031 (0.0044)
Age spline	x	x	x	x
Total wealth spline				
Liquid wealth spline	x	x	x	x
Home equity spline	x	x	x	x
Property value spline	x	x	x	x
Year fixed effects	x	x	x	x
Occup. fixed effects	x	x	x	x
Industry fixed effects	x	x	x	x
Observations	53535	9161	2876	9175

NOTE – Heteroskedasticity robust standard errors are reported in parentheses: + denotes significance at 10%; * denotes significance at 5%; ** denotes significance at 1%. Home tenure is instrumented by age at most recent marital shock = max(age at first marriage, age at termination of first marriage) in all specifications. Column 1 includes the core sample of all currently married homeowners. All other columns use the treatment group (once or twice married stockholding homeowners whose current marriage duration is less than 40 years). Column 2 includes individuals with positive wealth only, while column 3 in addition restricts wealth to be above the treatment sample mean of \$250,000. In both 2 and 3, all monetary variables including the dependent variable and the splines are normalized by total wealth. In column 4, the dependent variable is a dummy variable for bond market participation. All splines are 10-piece linear splines partitioned by deciles of the relevant variable.

TABLE 6
ADDITIONAL SPECIFICATION TESTS: MISMEASUREMENT AND ENDOGENEITY OF WEALTH

	(1)	(2)	(3)	(4)	(5)	(6)
	Cars	Cars	Property value	Property value	Wealth effects of	Wealth effects with
	Few controls	with controls	Few controls	with controls	instruments	education control
Dependent variable:	car value	car value	property val.	property val.	total wealth	total wealth
Age at most recent marital shock	-12.37 (13.62)	-21.73 (13.68)	504.41 (126.628)	156.72 (125.957)	2656.81 (950.592)	1095.51 (951.980)
Remarriage indicator	307.48 (229.90)	358.22 (230.23)	-3420.35 (1,985.36)	-639.18 (1,940.24)	-30156.72 (13,332.16)	-10481.02 (13,516.28)
Years of education		63.39 (38.45)		3,395.195 (317.748)		22,655.914 (1,499.139)
Children in home		-5.97 (97.89)		165.221 (842.862)		
Income		0.020 (0.0023)		0.296 (0.025)		
Age spline	x	x	x	x	x	x
Total wealth spline	x	x	x	x		
Year fixed effects	x	x	x	x	x	x
Occup. fixed effects		x		x		
Industry fixed effects		x		x		
Observations	9310	9310	9222	9222	9310	9310

NOTE – Heteroskedasticity robust standard errors are reported in parentheses. Sample in all columns is the treatment group (once or twice married stockholding homeowners whose current marriage duration is less than 40 years). Dependent variable in columns 1 and 2 is the current total value of up to two cars owned by the household. Age at most recent marital shock is defined as max(age at first marriage, age at termination of first marriage). Remarriage indicator is 1 for twice-married and 0 for once-married individuals in this sample. All splines are 10-piece linear splines partitioned by deciles of the relevant variable.