Moving Costs, Nondurable Consumption and Portfolio Choice

Nancy L. Stokey*

University of Chicago

September 27, 2007

Abstract

The substantial adjustment cost for housing transactions affects a homeowner's decision about when to move. It also affects her consumption of nondurables and her portfolio, at dates when a new house is purchased and over time intervals between moves. These decisions are studied using a calibrated dynamic model. For reasonable transaction costs, portfolios display substantial swings during intervals between moves and make large jumps when a new house is purchased. This suggests that a measure of housing wealth will be useful in empirical studies of portfolios at the micro level. The transaction cost has only a modest effect on nondurable consumption, however. Consequently, it can explain very little of the equity premium puzzle. These conclusions are robust to the value assumed for the elasticity of substitution between housing and nondurables, a parameter that over a broad range has a surprisingly weak impact on the consumer's behavior.

*I am grateful to Monika Piazzesi, Robert Lucas, Narayana Kocherlakota, and Robert Shimer for helpful comments.

For most individuals housing accounts for large fractions of both consumption and wealth. But housing is important for another reason as well. Moving typically entails substantial adjustment costs, so individuals adjust their consumption of housing services infrequently. As age, wealth, family size, and other household characteristics change, the consumer must decide whether and when to sell her current house and buy a new one, incurring the adjustment cost. She must also make decisions about nondurable consumption and her portfolio of financial assets. Between moves the size and direction of the latter adjustments are influenced by the fact that housing is fixed, and when she sells one house and buys another her nondurable consumption and portfolio take discrete jumps.

This paper studies the behavior of an infinitely lived consumer making these decisions. The model focuses on changes in the consumer's wealth as the driving variable, ignoring life cycle effects. This approach, which allows the use of a time-invariant Bellman-type equation, highlights some of the main forces at work and makes the problem tractable.

A calibrated version of the model is simulated, and the quantitative results are reasonable. With an adjustment cost of 8%, the consumer allows her wealth (permanent income) to rise or fall by about 50% between moves.

The consumer's portfolio of financial assets displays broad swings between moves and large jumps at the time of a move. Portfolio choice in the model depends on the local risk aversion of the consumer's value function for wealth, and the shape of this function is distorted by the presence of the transaction cost. Thus, with a transaction cost, risk aversion depends on the consumer's ratio of total wealth to housing wealth, although absent the transaction cost it is constant. As in Grossman and Laroque (1990), risk aversion in the value function is lower when the ratio of housing wealth to total wealth is near a threshold that triggers a move and higher when that ratio is at the level it assumes just after a move. Thus, the share of wealth held in the risky asset changes as the consumer's wealth rises or falls, increasing by eleven percentage points between its post-transaction level and its level just before a move to a bigger house. It then jumps down when the new house is purchased.

These wide swings suggest that including a measure of housing will be useful in empirical studies of cross-section or panel data on portfolios. The way that housing affects these decisions is subtle—it is non-monotone—but the calibrated model suggests that the effects are substantial.

Between moves nondurable consumption rises and falls with the consumer's wealth, with the size of the change depending on the elasticity of substitution between housing and nondurables and the elasticity of intertemporal substitution. But nondurable and total consumption are remarkably similar to what they would be in the absence of a transaction cost. They are also remarkably insensitive to the value assumed for the elasticity between nondurables and housing. This insensitivity may explain why empirical estimates of that parameter vary over such a wide range.

The decades-old hypothesis that adjustment costs for housing explain the equity premium puzzle is also examined. Given the insensitivity of total consumption to the transaction cost, the conclusion here is not surprising: the adjustment cost works in the right direction, but for reasonable parameter values the effect is small. Even with a very low elasticity of substitution, the most favorable case, the transaction cost can explain only a modest fraction of the puzzle.

Finally, it is interesting to note that the theoretical model here produces a value function that is strictly concave. Thus, the non-concavities found in other models arise from additional features, not from adjustment costs alone.

The rest of the paper is organized as follows. Section 1 contains a brief review of the related literature. Preferences are described in section 2 and a model without transaction costs is studied briefly in section 3. The model with transaction costs is set out in section 4, the calibration is described in section 5, and the simulations are presented in section 6 Section 7 discusses the model's predictions about the equity premium puzzle, and section 8 concludes.

1. Related literature

There is a sizable literature, going back two decades, asking whether including durable goods can improve the fit of asset pricing models. Early attempts assumed that consumption of durables is flexible in the sense that there are no adjustment costs. In this group are the papers by Dunn and Singleton (1986) and Eichenbaum and Hansen (1990). They found that including durables does little to improve the fit of the model.

The theoretical paper by Grossman and Laroque (1990) provided a framework for studying the behavior of an individual who consumes only one good, housing services, and faces adjustment costs for changing her level of consumption. They showed that the adjustment cost affects the consumer's portfolio choice in a systematic way. Specifically, consumers who have recently adjusted their housing stock, and hence anticipate a long interval of time before another adjustment, are more risk averse than those who anticipate making an adjustment in the near future. Their model does not include nondurable consumption, however, so it is difficult to calibrate and provides no predictions about the behavior of standard Euler equations.

Several subsequent papers have further explored the implications of adjustment costs. Marshall and Parekh (1999) study a model in which the adjustment cost applies to total consumption, not just housing. They find that even small values for this adjustment cost induce much smoother consumption behavior, and hence are quite successful in explaining the equity premium puzzle. However, it is not clear what those adjustment costs represent, or what data could be used to estimate or calibrate them. Fukushima (2005) looks at a model with discretionary and precommitted consumption goods, where in any period only a randomly chosen fraction of households are allowed to adjust the precommitted component of their consumption. The precommiteed component makes consumption smoother, but the ad hoc nature of the adjustment mechanism makes the model difficult to match to panel data.

Other work has explored the potential for adjustment costs to produce value functions with regions that are not concave. Chetty and Szeidl (2007) examine the effect of an adjustment cost in a setting with one large informational shock and one decision. They find that the ex post value of the consumer, as a function of realized wealth, has alternating convex and concave portions. Similarly, Vereshchagina (2007) studies a model with a one-time housing adjustment decision and obtains a similar result. But transaction costs do not necessarily produce non-concavities. As will be seen below, the model here has a value function that is strictly concave.

Flavin and Nakagawa (2004) study a model similar to the one here that also includes life cycle effects and house price risk, and nests a habit persistence model as well. Using data from the PSID, they estimate a very low elasticity of substitution between housing and nondurables. They also find that while the habit persistence model can be rejected, the adjustment cost model cannot be. Using aggregate (NIPA) data, Siegel (2004) studies a similar model, and also finds evidence that adjustment costs are important. However, using a different measure of housing consumption in the PSID, he estimates a much higher elasticity of substitution between housing and nondurables. The present paper is closely related to these two, simplifying the model to make it analytically tractable and amenable to calibration and simulation.

Martin (2003) looks at nondurable consumption around the time of a housing adjustment. Using PSID data, he distinguishes households that are likely to make upward and downward adjustments in their housing from those that are unlikely to adjust. He finds evidence those likely to move to a larger house reduce their consumption of nondurables, and those likely to move to a smaller house raise their consumption of nondurables. The model here predicts such behavior if the elasticity of substitution between housing and nondurables is smaller than the elasticity of intertemporal substitution.

Finally, two recent papers find evidence of state-dependent risk aversion. Using data from the Survey of Income and Program Participation, Chetty and Szeidl (2004) find that households who have moved recently, and hence are unlikely to move again in the near future, choose less risky portfolios than those with longer tenures at their current residence. Using a sample of homeowners from the PSID, Kullmann and Siegel (2005) find that lower ratios of net worth to housing wealth are correlated with lower stock market participation and reduced holdings of stocks and other risky assets. The model here predicts that risk aversion varies with the ratio of total wealth to housing wealth, but the relationship is not monotone. Risk aversion is high just after a housing transaction, when the wealth ratio has been adjusted to a target value. It then falls as the consumer's wealth increases or falls, with housing constant. But risk aversion is not monotone in housing tenure either, since total wealth does not change monotonically with tenure. [1]

2. Preliminaries

There are two consumption goods, housing services H and a single composite nondurable C. The flow of housing services H reflects both size and quality, including features like location, lot size, and other attributes. The consumer has CES preferences over the two goods,

$$U(C,H) = \begin{cases} \left[\omega C^{1-\zeta} + (1-\omega)\, H^{1-\zeta}\right]^{1/(1-\zeta)}, & \zeta \neq 1, \\ C^{\omega} H^{1-\omega}, & \zeta = 1, \end{cases}$$

[1] In addition, many papers have studied other channels—like house price risk—through which housing affects portfolio choice and nondurable consumption. For example, see Flavin and Yamashita (2002), Campbell and Cocco (2005), Cocco (2005), Fukushima (2005), Lustig and Nieuwerburgh (2006), Piazzesi, Schneider, and Tuzel (2007), and Fillat (2007).

where $\omega \in [0, 1)$ is the relative weight on nondurables, and $1/\zeta$ is the elasticity of substitution, and her intertemporal utility function is

$$\mathrm{E}_0 \left[\int_0^\infty e^{-\rho t} \frac{\{U[C(t), H(t)]\}^{1-\theta}}{1-\theta} dt \right], \tag{1}$$

where θ is the coefficient of relative risk aversion and ρ is the rate of time preference.

The consumer's only income is the return on her portfolio. She holds two assets, a safe one with a constant rate of return r, and a risky one with mean return $\mu > r$ and variance $\sigma^2 > 0$. Define the ratio

$$\gamma \equiv (\mu - r)/\sigma^2,$$

so $\gamma > 0$ is the inverse 'price' of risk. We will assume that the return r on the safe asset is also the interest rate on mortgages.

The price of the nondurable is normalized to one. The purchase price of housing is constant, and housing units can be chosen so that this price is also one. The direct cost of housing then has three components: interest at the rate r, depreciation at the rate δ, and maintenance at the rate m. Hence the (flow) cost of housing services is $p_h = r + \delta + m$. Housing may also have an indirect cost because it enters the portfolio constraint. That constraint will be discussed below.

Let Q denote the consumer's total wealth, H the value of her house, and A her holdings of the risky asset. Then $Q - A$ denotes wealth in the safe asset, including housing.

Under the parameter restrictions that will be used here the consumer always chooses $A > 0$, so we do not need to impose a lower bound on A. But we will impose an upper bound. The constraint has two parts.

First, there is an exogenously given minimal equity $\epsilon \in (0, 1]$ that an owner must hold in her house. Since the mortgage interest rate is the same as the return on the safe asset, this is equivalent to requiring that owner hold safe assets equal to ϵH. For $\epsilon = 0$ the consumer can be interpreted as a renter.

In addition, if the consumer is sufficiently risk tolerant she may want to short the safe asset, i.e., to buy the risky asset on margin. We will allow her to do so but will limit the size of such holdings—which is like imposing a margin requirement—and assume that the minimal equity in her house cannot be used as collateral. Specifically, we will require

$$A \in [0, a_{ss}(Q - \epsilon H)], \tag{2}$$

where $a_{ss} \geq 1$ reflects the size of the margin requirement. If $a_{ss} = 1$, the consumer cannot buy the risky asset on margin.

Given C, H, A, Q, the change in the consumer's total wealth over a short interval of time dt is

$$\begin{aligned} dQ &= [r(Q - A - H) + \mu A - (\delta + m) H - C] dt + \sigma A dz \qquad (3) \\ &= [rQ + (\mu - r) A - p_h H - C] dt + \sigma A dz, \end{aligned}$$

where z is a Wiener process. If $Q - A \geq H$ the consumer owns her house outright, and if the inequality is strict she has additional wealth invested at the risk-free rate.

The following parameter restrictions will be used throughout.

ASSUMPTION 1:

$$\begin{array}{llll} 0 < r < \mu, & \sigma^2 > 0, & \zeta > 0, & 0 \leq \omega \leq 1, \\ 0 \leq \epsilon \leq 1, & a_{ss} \geq 1, & \rho > 0, & \theta > 0, \quad \theta \neq 1, \\ \delta, m \geq 0, & \rho + (1 - \theta)\delta > 0. & & \end{array}$$

The case $\theta = 1$, which represents logarithmic utility, can be treated along similar lines. The last restriction will be used in section 4.

3. The frictionless model

A useful benchmark for comparisons is the model with no transaction cost. In this case the consumer's problem is to choose (C, H, A) to maximize (1) subject to the budget constraint (3) and the portfolio constraint (2), given initial wealth $Q_0 > 0$.

Since the objective function is homogeneous of degree $(1-\theta)$ in (C, H, A, Q) and the constraints are homogeneous of degree one, the optimal ratios H/Q, A/Q, etc. are constant over time. Hence the consumer's problem can be written as

$$W(Q_0) = \max_{\substack{c \geq 0, h \in [0, 1/\epsilon] \\ a \in [0, a_{ss}(1-\epsilon h)]}} \mathrm{E}_0 \left[\int_0^\infty e^{-\rho t} \frac{[u(c) h Q(t)]^{1-\theta}}{1-\theta} dt \right] \tag{4}$$

$$\text{s.t.} \quad \frac{dQ}{Q} = \left[r + a(\mu - r) - (p_h + c) h \right] dt + a\sigma dz,$$

where $c \equiv C/H$ is the ratio of nondurable consumption to housing services, $h \equiv H/Q$ is the ratio of housing to wealth, $a \equiv A/Q$ is the portfolio share in the risky asset, and $u(c) \equiv U(c, 1)$ is the intensive form of the CES aggregator.

For any fixed (c, h, a), total wealth Q is a geometric Brownian motion with constant drift and variance. Hence $\mathrm{E}_0 \left[Q(t)^{1-\theta} \right] = Q_0^{1-\theta} e^{\Gamma(c,h,a;\theta)t}$, where

$$\Gamma(c, h, a; \theta) \equiv (1-\theta) \left[r + (\mu - r) a - (p_h + c) h - \theta \frac{1}{2} (\sigma a)^2 \right]. \tag{5}$$

Consequently, if $\rho > \Gamma$ the value function in (4) has the form $W(Q_0) = Q_0^{1-\theta} w^*$, where

$$w^* \equiv \max_{\substack{c \geq 0, h \in [0, 1/\epsilon] \\ a \in [0, a_{ss}(1-\epsilon h)]}} \frac{[u(c) h]^{1-\theta}}{1-\theta} \frac{1}{\rho - \Gamma(c, h, a; \theta)}. \tag{6}$$

The next assumption insures that Γ satisfies the required condition.

Assumption 2: If $0 < \theta < 1$,

$$\rho > (1-\theta) \times \begin{cases} \left[r + (\mu - r) a_{ss} - \theta a_{ss}^2 \sigma^2 / 2 \right], & \text{if } \theta < \gamma / a_{ss}, \\ \left[r + (\gamma/\theta)(\mu - r)/2 \right], & \text{if } \theta \geq \gamma / a_{ss}. \end{cases}$$

The following proposition characterizes the solution. The proposition first describes the solution for renters, consumers with $\epsilon = 0$. Although a renter faces a short sale constraint, that constraint does not involve her housing choice. Thus, there is a certain type of separation between her consumption and portfolio decisions, even if her portfolio constraint binds. For a buyer, a consumer with $\epsilon > 0$, the solution is the same as the renter's if her (tighter) portfolio constraint is satisfied for the renter's choices. If that constraint is not satisfied, the buyer's consumption and portfolio choices are intertwined in a more complicated way.

PROPOSITION 1: Let Assumptions 1 and 2 hold.

(a) For $\epsilon = 0$ the unique solution to the problem in (6) is

$$\begin{aligned} a_R(\theta) &= \min\{\gamma/\theta,\ a_{ss}\}, \\ c_R &= \left(\frac{\omega p_h}{1-\omega}\right)^{1/\zeta}, \\ h_R(\theta) &= \frac{1}{c_R + p_h}\frac{1}{\theta}\left\{\rho - (1-\theta)\left[r + \sigma^2\gamma a_R(\theta) - \theta\frac{1}{2}\sigma^2 a_R^2(\theta)\right]\right\}. \end{aligned} \tag{7}$$

Moreover, h_R is strictly concave in θ, reaching a maximum where

$$\frac{r-\rho}{\sigma^2} = \left(\frac{\theta}{2}a_R - \frac{\gamma}{\theta}\right)\theta a_R.$$

(b) For $\epsilon > 0$, the unique solution to the problem in (6) is as in (7) if $a_{ss}\left[1 - \epsilon h_R(\theta)\right] \geq \gamma/\theta$. Otherwise

$$\begin{aligned} [c_B(\theta,\epsilon), h_B(\theta,\epsilon)] &= \arg\max_{c,h} \frac{[u(c)h]^{1-\theta}}{1-\theta}\frac{1}{\rho - \Gamma[c, h, a_{ss}(1-\epsilon h)]}, \\ a_B(\theta,\epsilon) &= a_{ss}\left[1 - \epsilon h_B(\theta,\epsilon)\right]. \end{aligned} \tag{8}$$

In this case

$$h_B(\theta,\epsilon) < h_R(\theta) \qquad \text{and} \qquad c_B(\theta,\epsilon) > c_R.$$

PROOF: See the Appendix.

For renters the share of wealth in the risky asset $a_R(\theta)$ is strictly positive and depends only on γ/θ and a_{ss}. For those who are sufficiently risk averse the solution is interior, at $a_R = \gamma/\theta$, while for those who are sufficiently risk tolerant the constraint binds and the solution is $a_R = a_{ss}$.

For renters the ratio c_R of nondurable consumption, and hence also the expenditure share of housing $p_h/(p_h + c_R)$, depend only on the rental price p_h and the parameters ω and ζ, and not on $\theta, \rho, \gamma, \sigma^2$.

For renters the ratio of total expenditure to wealth, $(c_R + p_h) h_R(\theta)$, depends on $\theta, \rho, \gamma, \sigma^2$ but not on p_h, ω or ζ. The function $h_R(\theta)$ has an inverted U-shape. If $r = \rho$ and a_{ss} is large, then $h_R(\theta)$ peaks at $\hat{\theta} = 2$. The peak occurs at a lower value if $\rho > r$.

The hump shape is the result of two opposing forces. Consumers with lower θ are more risk tolerant, so they hold portfolios with higher expected rates of return. Hence they also choose lower ratios of expenditure to wealth. But consumers with higher θ have have a stronger incentive to smooth consumption over time. Hence they also prefer lower ratios of expenditure to wealth. The first force predominates for values of θ below a certain threshold, and the second for θ above that threshold, leading to the hump shape for $h_R(\theta)$.

For a portfolio-constrained owner, the house/wealth ratio is lower than for a renter. For this consumer housing services have an extra cost at the margin, the incremental portfolio distortion. Hence she chooses a lower ratio of housing to nondurables.

4. The model with transaction costs

Suppose the consumer must pay a transaction cost of λH when she adjusts her housing, where $\lambda > 0$. Then she will adjust her housing only occasionally, by discrete amounts, and her budget constraint has two parts. At dates when she adjusts her housing her wealth falls by the amount of the transaction cost. At all other times the durable depreciates deterministically and wealth grows stochastically.

In addition to voluntary housing adjustments, suppose that moves may be required for exogenous reasons. Job changes that involve relocating to a new city and changes in family size are two possible interpretations of these moves. Assume that this shock is Poisson, with a constant arrival rate κ.

Define the stopping time T_X as the arrival of the next exogenous relocation shock, and define the stopping time T_A as the time the consumer chooses for the next adjustment in case an exogenous has not occurred. The time of the consumer's next housing adjustment is the minimum of these two, $T' \equiv T_A \wedge T_X$.

With a transaction cost for housing two state variables are needed, Q and H. But the consumer's value function $V(Q, H)$ is, as before, homogeneous of degree $(1-\theta)$ in the state variables, and the policy functions for C, A, and H' are homogeneous of degree one. Hence a normalized form of the problem can be written in terms of a single state variable, a ratio. It is convenient to use $q = Q/H = 1/h$. The Bellman equation is then

$$v(q_0) = \sup_{\{c(t),a(t)\},T_A,q'} \mathrm{E}_0 \left\{ \int_0^{T'} e^{-\eta t} \frac{u[c(t)]^{1-\theta}}{1-\theta} dt + e^{-\eta T'} \left(\frac{q(T') - \lambda}{q'} \right)^{1-\theta} v(q') \right\} \quad (9)$$

$$\begin{aligned} \text{s.t.} \quad dq &= \{[r + \delta + (\mu - r)\, a]\, q - (p_h + c)\}\, dt + \sigma a q dz, \\ a &\in \left[0, a_{ss}\left(1 - \frac{\epsilon}{q}\right)\right], \qquad\qquad t \in [0, T'), \\ T' &= T_A \wedge T_X, \\ q' &\geq \epsilon, \end{aligned}$$

where $v(q) \equiv V(q, 1)$,

$$\eta \equiv \rho + (1-\theta)\,\delta,$$

and as before $c = C/H$ and $a = A/Q$. Assumption 1 insures $\eta > 0$. A solution consists of a value function $v(q)$ defined on $\mathbf{R}_+$ satisfying (9), and policy functions $\{c(t), a(t)\}, T_A, q'$ that attain the maximum. As shown in the Appendix, under Assumptions 1 and 2 v is well defined.

Two properties of the solution are immediate from (9). First, the maximizing value for q' does not depend on the state $q(T')$ when the adjustment is made. Define

$$M \equiv \max_{q'} \frac{v(q')}{q'^{1-\theta}}, \quad (10)$$

and let S denote the return point, the optimal value for q'. Thus, M is the optimized value for an individual with net wealth $Q = 1$ when she buys a house, and S is the wealth/house ratio she chooses.

In addition, the stopping time chosen by the consumer has the form $T_A = T(b) \wedge T(B)$, where $T(\beta)$ denotes the first time the stochastic process q reaches β, and $0 \leq b < B < +\infty$ are optimally chosen thresholds. Thus, the state has an *inaction region*, the open interval (b, B). While the state remains inside this interval the consumer does not sell her house voluntarily, although the exogenous moving shock may force her to do so. The consumer immediately adjusts her housing if q is outside the interval (b, B). Hence the value function outside the inaction region has the form

$$v(q) = (q - \lambda)^{1-\theta} M, \qquad q \notin (b, B). \tag{11}$$

After an initial transaction, if required, the state remains inside the interval (b, B).

To characterize the value function v, the critical points b, S, B, and the policy functions c and a, we can use the fact that inside the inaction region the value function satisfies the Bellman-type equation

$$\begin{aligned}(\eta + \kappa)\, v(q) \quad = \quad & \max_{c, a \in [0, a_{ss}(1-\epsilon/q)], q'} \left\{ \frac{u(c)^{1-\theta}}{1-\theta} + m(q) v'(q) \right. \\ & \left. + \frac{1}{2} s^2(q) v''(q) + \kappa (q - \lambda)^{1-\theta} \frac{v(q')}{(q')^{1-\theta}} \right\},\end{aligned} \tag{12}$$

where

$$\begin{aligned} m(q) \quad &\equiv \quad [r + \delta + (\mu - r)\, a]\, q - (p_h + c)\,, \\ s^2(q) \quad &\equiv \quad (\sigma a q)^2\,, \end{aligned}$$

are the instantaneous drift and variance for q under the optimal policies $a(q)$ and $c(q)$. (See Stokey 2007, Ch. 9 for a more detailed discussion.)

The interpretation of (12) is fairly standard. The first term on the right is the current utility flow from consumption. The second and third, which come from an application of Ito's lemma, are the expected 'capital gain' from changes in the state variable. To interpret the final term, subtract $\kappa v(q)$ from both sides. Then the final term on the right, which is negative, is the expected net loss from the exogenous moving shock. The remaining term on the left side $\eta v(q)$ is the 'current return' on the value v.

The optimal policies for the portfolio and nondurable consumption are found by maximizing the term in braces in (12). Hence the optimal portfolio is

$$a(q) = \min\left\{\frac{\gamma}{-qv''/v'}, a_{ss}\left(1 - \frac{\epsilon}{q}\right)\right\}. \tag{13}$$

The expression on the right is exactly analogous to the one for the problem with no transaction costs in (7). The only difference is that here the relative risk aversion of the value function, $-qv''/v'$, varies with q. With no transaction cost the transaction cost the value function has the form $W(Q) = Q^{1-\theta}w^*$, so $-QW''/W' = \theta$.

The condition for nondurable consumption is

$$\frac{d}{dc}\left[\frac{u(c(q))^{1-\theta}}{1-\theta}\right] = v'(q). \tag{14}$$

The term in square brackets is instantaneous utility, as a function of nondurable consumption only, when housing services are fixed at unity. Nondurable consumption $c(q)$ increases with wealth, but the extent to which it varies depends on the substitution elasticity $1/\zeta$ and the intertemporal elasticity $1/\theta$. Lower elasticities imply a weaker response for nondurable consumption.

Optimal choice of the boundaries b and B requires that value matching and smooth pasting conditions hold. That is, both v and v' must be continuous. From (11) we see that this requires

$$\begin{aligned} \lim_{q\downarrow b} v(q) &= (b-\lambda)^{1-\theta} M, \\ \lim_{q\uparrow B} v(q) &= (B-\lambda)^{1-\theta} M, \end{aligned} \tag{15}$$

$$
\begin{aligned}
\lim_{q\downarrow b} v'(q) &= (1-\theta)\,(b-\lambda)^{-\theta}\,M, \\
\lim_{q\uparrow B} v'(q) &= (1-\theta)\,(B-\lambda)^{-\theta}\,M,
\end{aligned}
$$

where M, defined in (10), is the optimized value for a consumer with unit wealth (after the transaction cost is paid) who is buying a new house. In addition, the return point S satisfies

$$
\begin{aligned}
v(S) &= S^{1-\theta}M, \qquad (16) \\
v'(S) &= (1-\theta)\,S^{-\theta}M.
\end{aligned}
$$

Although an analytic solution is not available, it is not difficult to compute solutions computed numerically, and we turn next to the simulations.

5. Calibration

The model has thirteen parameters: $(\mu, \sigma, r, a_{ss}, \epsilon)$ describing asset markets, $(\delta, \lambda, m, \kappa)$ for housing, and $(\rho, \theta, \zeta, \omega)$ describing preferences. Parameters about which there is better information will be fixed throughout the analysis. For the others I will choose benchmark values and conduct sensitivity experiments.

The asset returns will be fixed throughout at $\mu = 0.077$, $\sigma = 0.1655$, and $r = 0.015$. These values are fairly standard. (See Kocherlakota, 1996, and Mehra and Prescott, 2006.)

The short sale parameter will be fixed at $a_{ss} = 1.20$ throughout, which allows some scope for the consumer to buy the risky asset on margin. Minimum down payments for homeowners are typically 10-15%. The upper end of this range will be used here, $\epsilon = 0.15$. In the simulations below the portfolio constraint involving these two parameters almost never binds.

Since most people maintain their houses rather than allowing them to depreciate, I will set $\delta = 0$ throughout.

Smith, Rosen, and Fallis (1988) estimate the monetary cost of selling a house to be 8% - 10% of the value of the unit. This figure includes agents' commissions, legal fees, taxes and other transaction costs, and moving costs. In addition there are costs that are harder to measure, such as the time cost of search, the psychic cost of disruption, and so on. I will use a conservative figure, $\lambda = 0.08$, for the benchmark and experiment with other values.

A key element in the model is the ratio of total wealth to housing wealth. The model includes only tangible wealth, while in fact the bulk of 'total wealth'—in the sense of what generates income—is intangible wealth, human capital. Capital's share in national income is about 1/3, so total wealth is about 3 times the stock of physical capital. Residential structures are about 40 - 50% of total private fixed capital.[2], so total physical capital is about 2.0 - 2.5 times the housing stock. Multiplying these two ratios suggests a figure of about 6.0 - 7.5 for the ratio of total wealth to housing wealth.

The total wealth/housing ratio in the model is sensitive to the maintenance cost m. That parameter is set at $m = 0.04$, a value that produces average wealth/housing ratios in the appropriate range. This figure for maintenance does not seem unreasonable, since it should be interpreted broadly. Thus, it includes property taxes, heating, and other costs that are difficult to adjust and proportional to the value of the house.

Little direct evidence is available on the hazard rate for exogenous moves. I will use $\kappa = 0$ for the benchmark and experiment with a positive value.

There are four preference parameters, ρ, θ, ζ, and ω. For the rate of time preference, $\rho = 0.025$ will be used throughout. This figure is fairly standard.

There is less agreement about the elasticity of intertemporal substitution, with values for θ in the range of $[1, 10]$ all having their advocates. With asset returns fixed at their market values, this parameter is important in determining the average

[2] See Davis and Heathcote (2005, Table 7), who use NIPA data for 1948-2001.

growth rate of consumption and wealth. The value $\theta = 4$ produces a growth rate close to the historical average of 2%, so it will be used for the benchmark and sensitivity experiments will be conducted with other values.

There is even less consensus about the elasticity of substitution between housing and nondurables. Using data from a policy experiment that involved low-income renters in two cities, Hanushek and Quigley (1980) estimate price elasticities of $\varepsilon = 0.45$ and 0.64. Siegel (2004) obtains two estimates based on homeowners in the PSID over the period 1978-1997, using the self-reported value of the owner occupied house. Aggregating across households and using only the time series information, the estimated elasticity is 0.53. Using the household level information and limiting the sample to households that own stocks, the estimated elasticity is in the range $[1.23, 1.54]$.

Flavin and Nakagawa (2004) also use data from the PSID, for 1975-1985, but they employ a different measure of housing to sidestep the problem of price variation across cities. They obtain an elasticity of substitution of $\varepsilon = 0.13$. Using NIPA data on real rents and the aggregate expenditure share of housing over the period 1936-2001, Piazzesi, et. al. (2007) estimate the elasticity to be in the range $[1.05, 1.25]$. Using CEX data for 27 cities in 2003, a simple regression of the expenditure share of housing on the relative price of housing leads to an estimated elasticity of $\varepsilon = 0.45$.[3]

The value $\varepsilon = 0.5$ will be used as the benchmark, and sensitivity experiments conducted with values of $\varepsilon = 0.15$, 1.0, and 1.25.

The weight parameter ω will be calibrated using the expenditure share of housing. Aggregate data from NIPA suggest an expenditure share of about 20% over the period 1960-2005, with relatively little variation. Data from the CEX suggests a somewhat

[3] This estimate excludes Anchorage, which is an extreme outlier. The price data, from Aten (2005, Tables 3 and 4), are for 2003. The expenditure shares, from Tables 21 - 24 of the Consumer Expenditure Survey, are for 2003-2004.

higher figure, around 33%. I will use an intermediate value, calibrating ω in each simulation so that the expenditure share of housing is 30%.

Table 1 displays the benchmark parameters.

Table 1

$\mu = 0.077$	$r = 0.015$	$\rho = 0.025$	$\zeta = 2$
$\sigma = 0.1655$	$a_{ss} = 1.20$	$\theta = 4$	$\lambda = 0.08$
$\delta = \kappa = 0$	$\epsilon = 0.15$	$m = 0.04$	$\omega = 0.184$

Another figure that can be used to check the predictions of the model (or to calibrate κ) is the average length of residence. For persons 15 years and older who live in owner occupied housing, this figure is 11.3 years.[4]

6. Quantitative results

a. Frictionless model.—

Figure 1 displays results for the model with no transaction costs ($\lambda = 0$) and risk aversion $\theta \in [0.75, 5.0]$. The preference parameter ω is calibrated to give housing an expenditure share of 30%, and the benchmark values are used for the other parameters.

Figure 1a shows the portfolio share in the risky asset. The short sale constraint binds for consumers who are sufficiently risk tolerant, those with $\theta < \theta^c \approx 1.9$.

Figure 1b shows the ratio of expenditures to wealth, $(p_h + c)\, h$. The curve is single peaked, as Proposition 1 predicts for $\epsilon = 0$.

Figure 1c shows the average (long run) growth rate for income, consumption and wealth. It declines with θ over most of the range, with a kink at θ^c. For $\theta < \theta^c$ the growth rate declines with θ because the expenditure flow increases, while the portfolio

[4]Calculated from Figure 4 in Schachter and Kuenzi (2002), which is based on Census (SIPP) data for 1996.

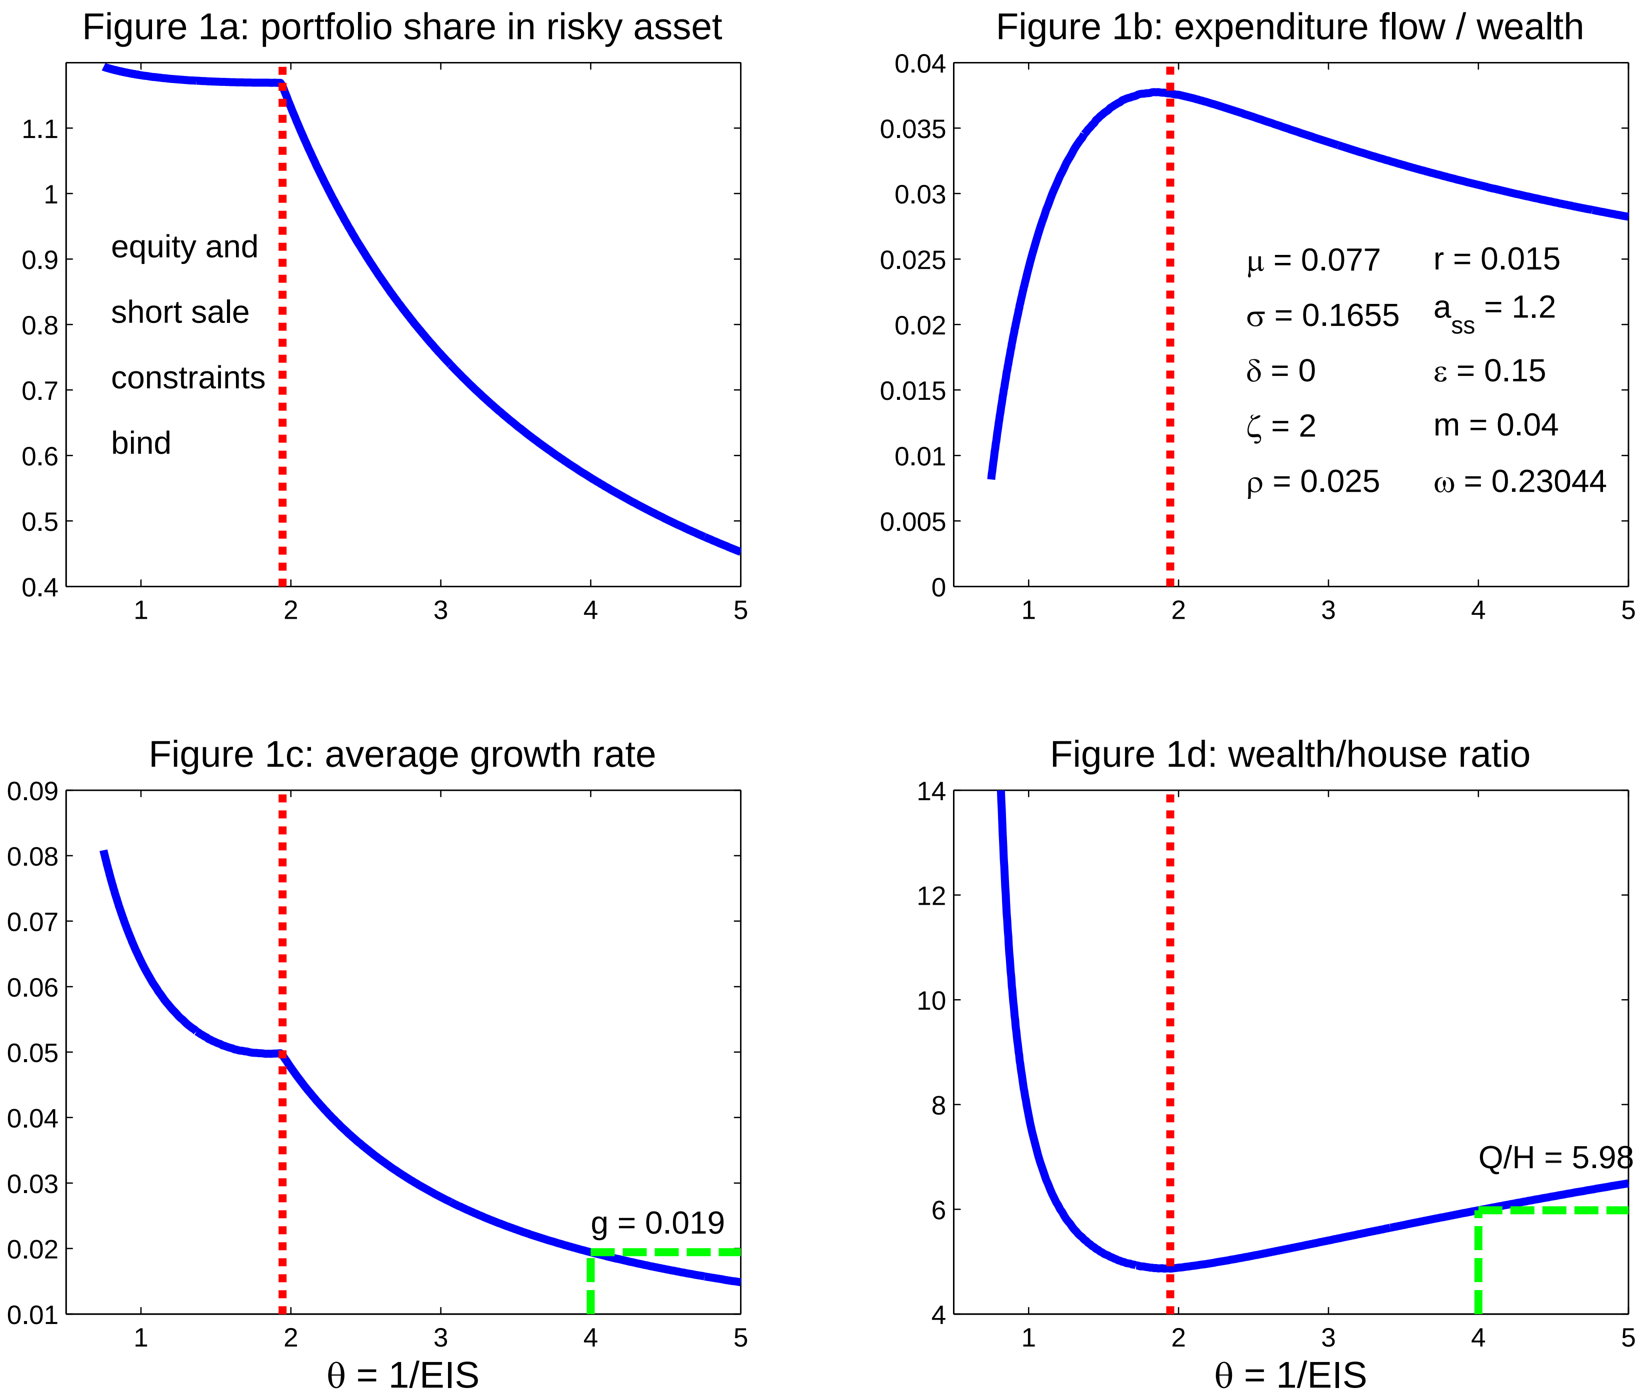
Figure 1a: portfolio share in risky asset
equity and
short sale
constraints
bind
Figure 1b: expenditure flow / wealth
μ = 0.077
r = 0.015
σ = 0.1655
a_ss = 1.2
δ = 0
ε = 0.15
ζ = 2
m = 0.04
ρ = 0.025
ω = 0.23044
Figure 1c: average growth rate
g = 0.019
θ = 1/EIS
Figure 1d: wealth/house ratio
Q/H = 5.98
θ = 1/EIS

allocation is almost constant, constrained at $(1 - \epsilon h)\, a_{ss}$. For $\theta > \theta^c$ expenditures decrease with θ, but there is also a portfolio reallocation toward the safe asset, which reduces the return on the portfolio. The latter effect swamps the former, so the growth rate continues to decline. For $\theta = 4$, the growth rate is $g = 1.9\%$.

Figure 1d shows the ratio of total wealth to housing wealth, which is roughly a mirror image of the expenditure flow/wealth ratio in Figure 1b. For $\theta = 4$ the ratio is $Q/H = 6.0$.

b. Benchmark model.—

Figures 2 - 4 display results for the benchmark calibration. The portfolio constraint involving housing equity and short sales does not bind.

Figure 2 shows the value function and its first two derivatives, as well as the value function for a consumer who faces no transaction costs. (The same value for ω is used here for the consumer who faces no transaction cost.) Both value functions are smooth and concave, their first derivatives are smooth and convex, and their second derivatives are smooth and concave. The transaction cost does not create kinks or nonconvexities.

The adjustment thresholds, indicated with dotted lines, are wealth/house ratios of $b = 3.4$ and $B = 9.5$, and the ratio chosen when a new house is purchased, indicated with a small open circle, is $S = 6.2$. Thus, an upward adjustment is made when wealth has increased by about 52% and a downward adjustment when it has fallen by 45%. The long run average, which is 6.4, is higher than the (constant) ratio of 5.4 chosen by a consumer who faces no transaction cost. The transaction cost makes housing more expensive so less is consumed, producing a higher ratio of total wealth to housing wealth.

Figures 3 and 4 describe the consumer's behavior between housing transactions. Since there is no depreciation, the consumer's flow of housing services is constant

Figure 2: Value function and its derivatives

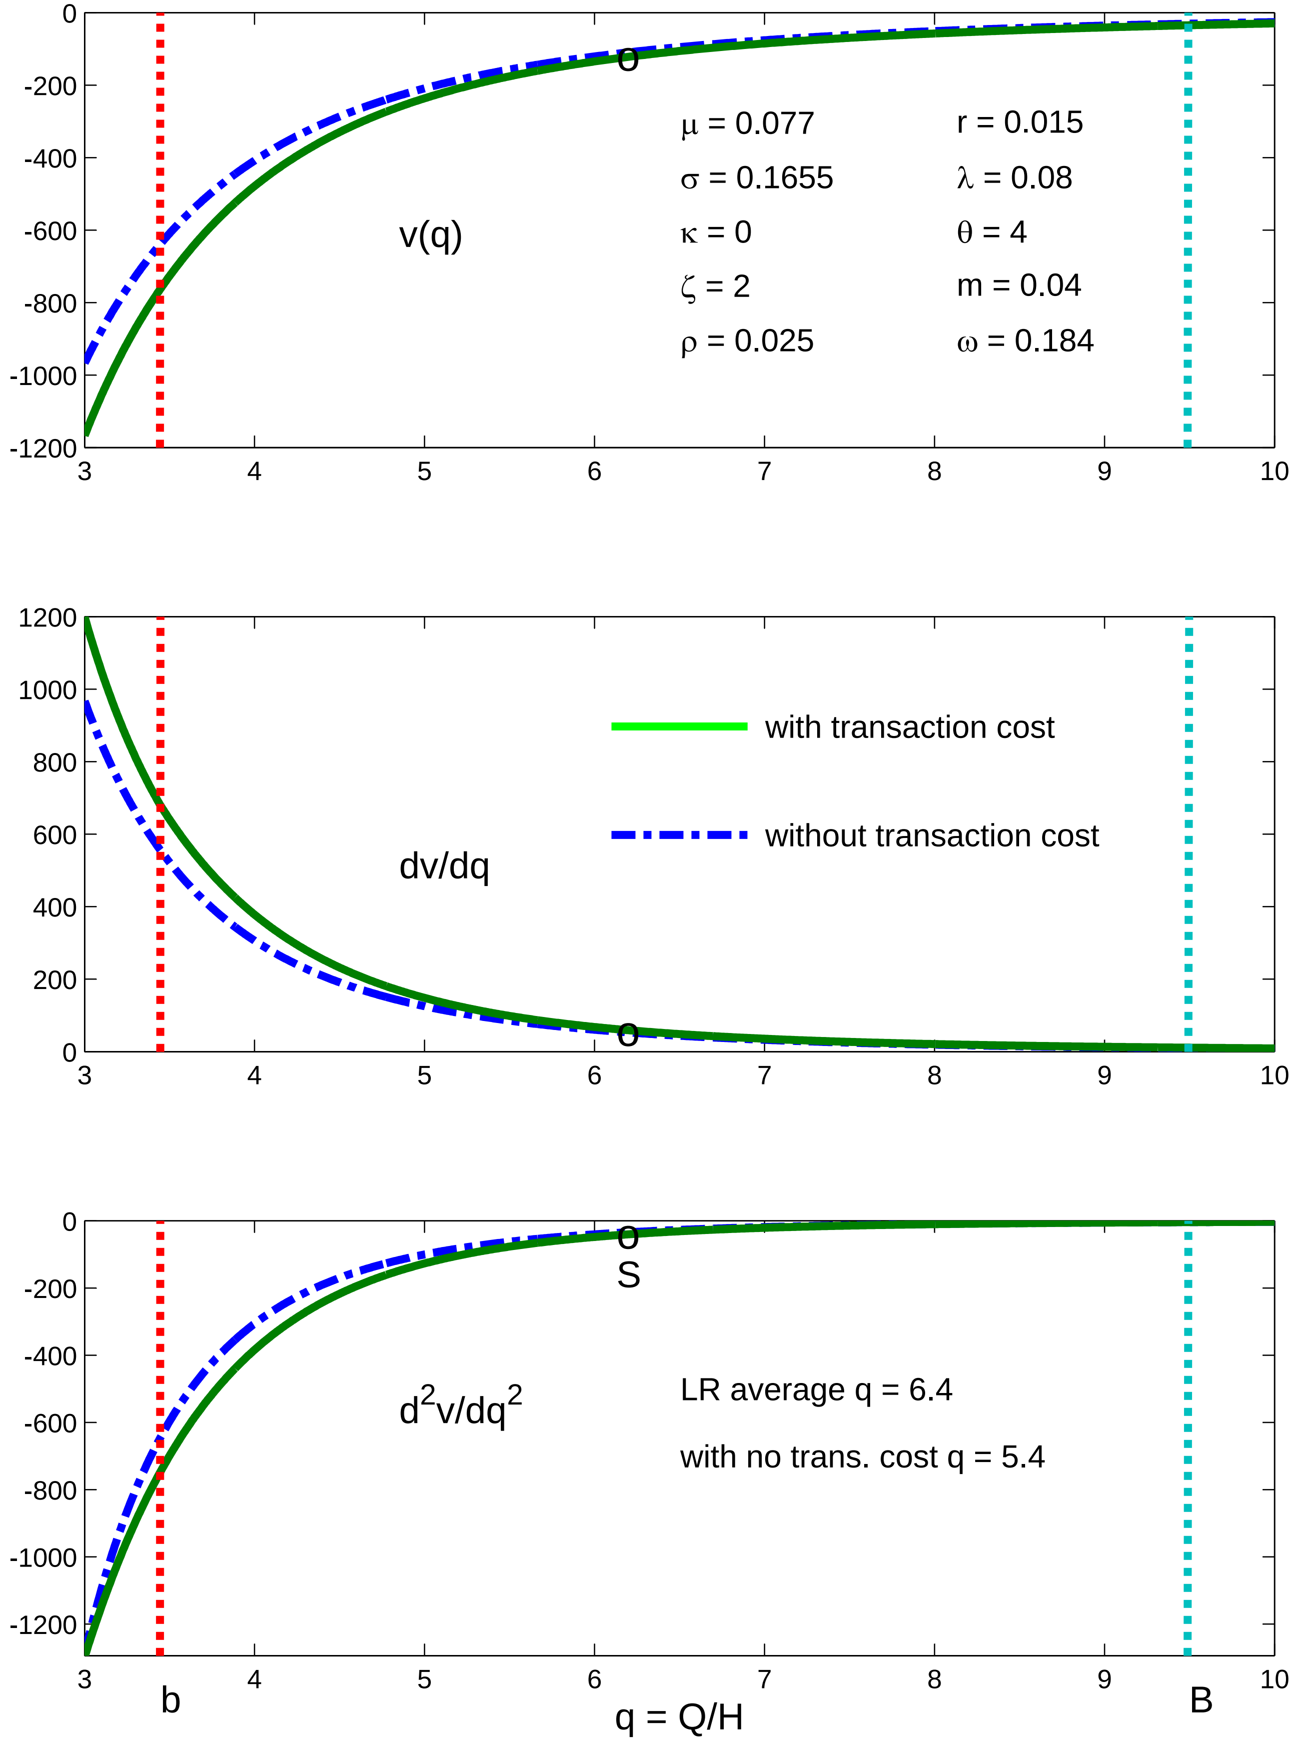

over such intervals. Long run averages are calculated using the density function for q following a start at $q = S$. (The density, not shown, has a fairly symmetric tent shape, with a peak at S.)

Figure 3a shows the share of her wealth that the consumer holds in risky assets. This function is U-shaped, reflecting the fact—first noted by Grossman and Laroque—that the consumer is more risk tolerant when she is close to the adjustment thresholds, and more risk averse in the middle of the inaction region. The fairly high risk aversion coefficient used here, $\theta = 4$, means that the consumer puts only 53% - 64% of her wealth in the risky asset. The long run average is 55%, a little lower than the 57% chosen by a consumer facing no transaction cost.

Figure 3b shows nondurable consumption relative to housing wealth. It moves linearly with total wealth, rising 69% or falling 51% relative to its level just after the most recent housing adjustment. Its long run average is 13%, which is a little higher than the 11% for a consumer who faces no transaction costs. The transaction cost induces the consumer to shift her consumption mix toward nondurables.

The average return on the portfolio (not shown) is 4.9%, and the average growth rate of consumption, income and wealth (they are all the same) is 2.0%.

Figure 4a shows how total expenditure and its two components change with wealth, for the benchmark consumer and for one who faces no transaction cost. Since the consumer's preferences are homothetic, in the frictionless world expenditures on housing and nondurables—and hence their sum—increase in proportion to wealth. Thus, the three dashed lines, for the consumer who faces no transaction cost, are rays from the origin. With a transaction cost, the consumer's expenditures on nondurables increase as her wealth increases, but her housing expenditure is constant. In this normalized model it is simply $p_h = r + m = 0.055$. With housing fixed, the consumer who faces a transaction cost substitutes into nondurables. Thus, her nondurable expenditure increases more strongly with wealth than for the consumer who faces

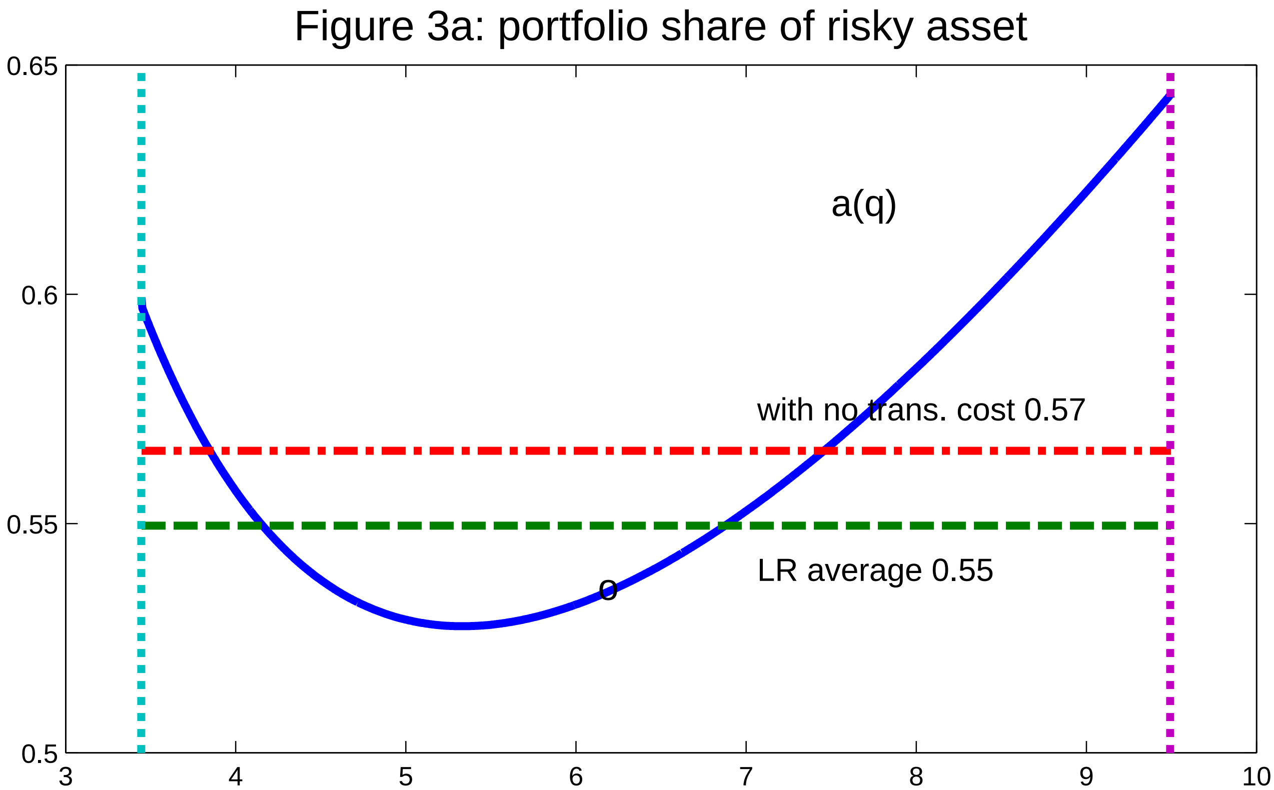
Figure 3a: portfolio share of risky asset
0.65
0.6
0.55
0.5
3
4
5
6
7
8
9
10
a(q)
with no trans. cost 0.57
LR average 0.55
o

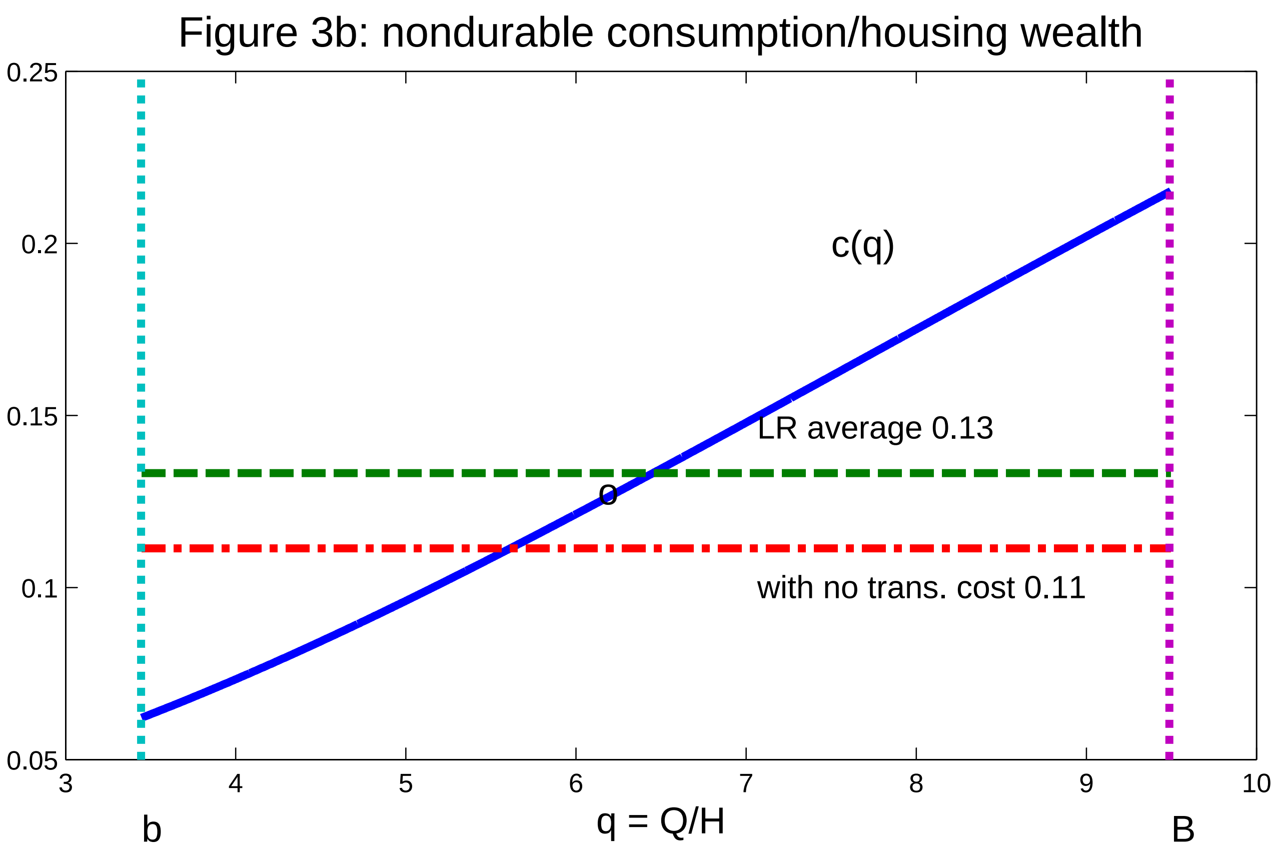
Figure 3b: nondurable consumption/housing wealth
0.25
0.2
0.15
0.1
0.05
3
4
5
6
7
8
9
10
c(q)
LR average 0.13
o
with no trans. cost 0.11
b
q = Q/H
B

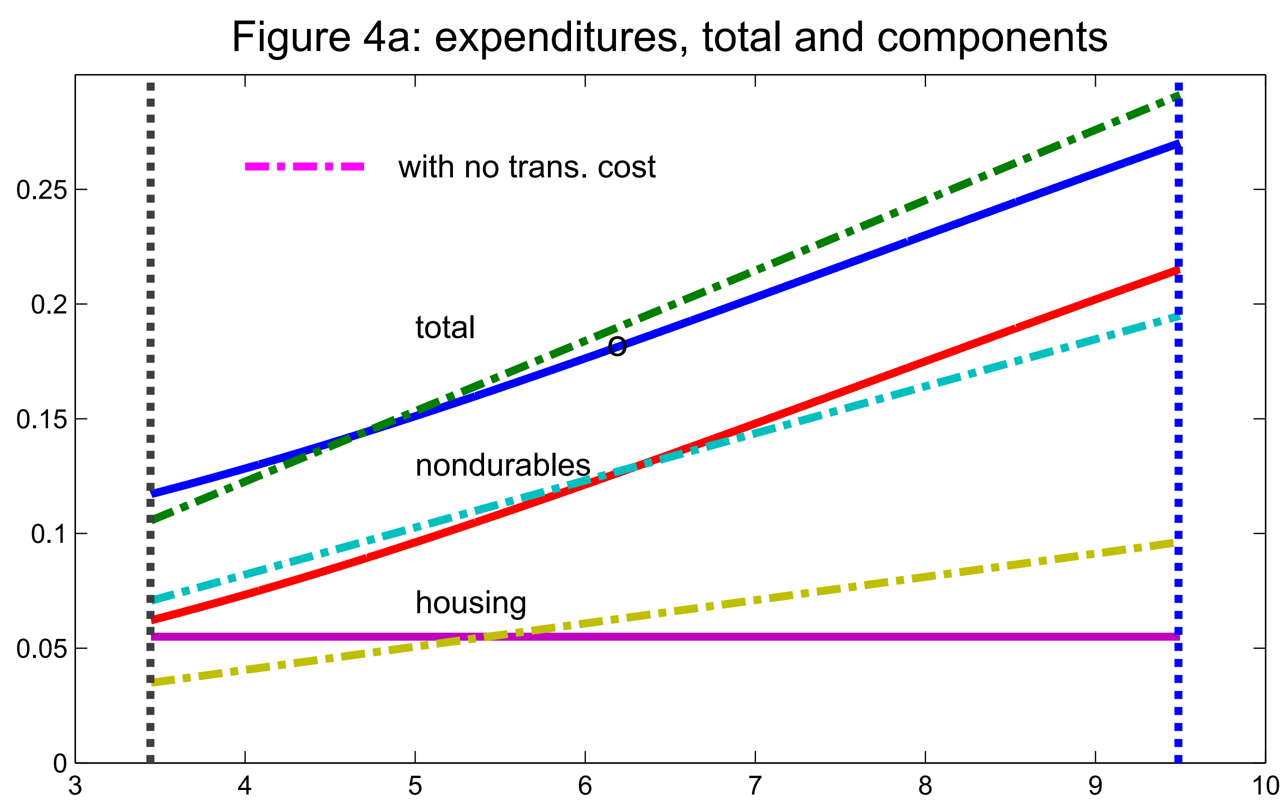
Figure 4a: expenditures, total and components
with no trans. cost
total
nondurables
housing
0
0.05
0.1
0.15
0.2
0.25
3
4
5
6
7
8
9
10

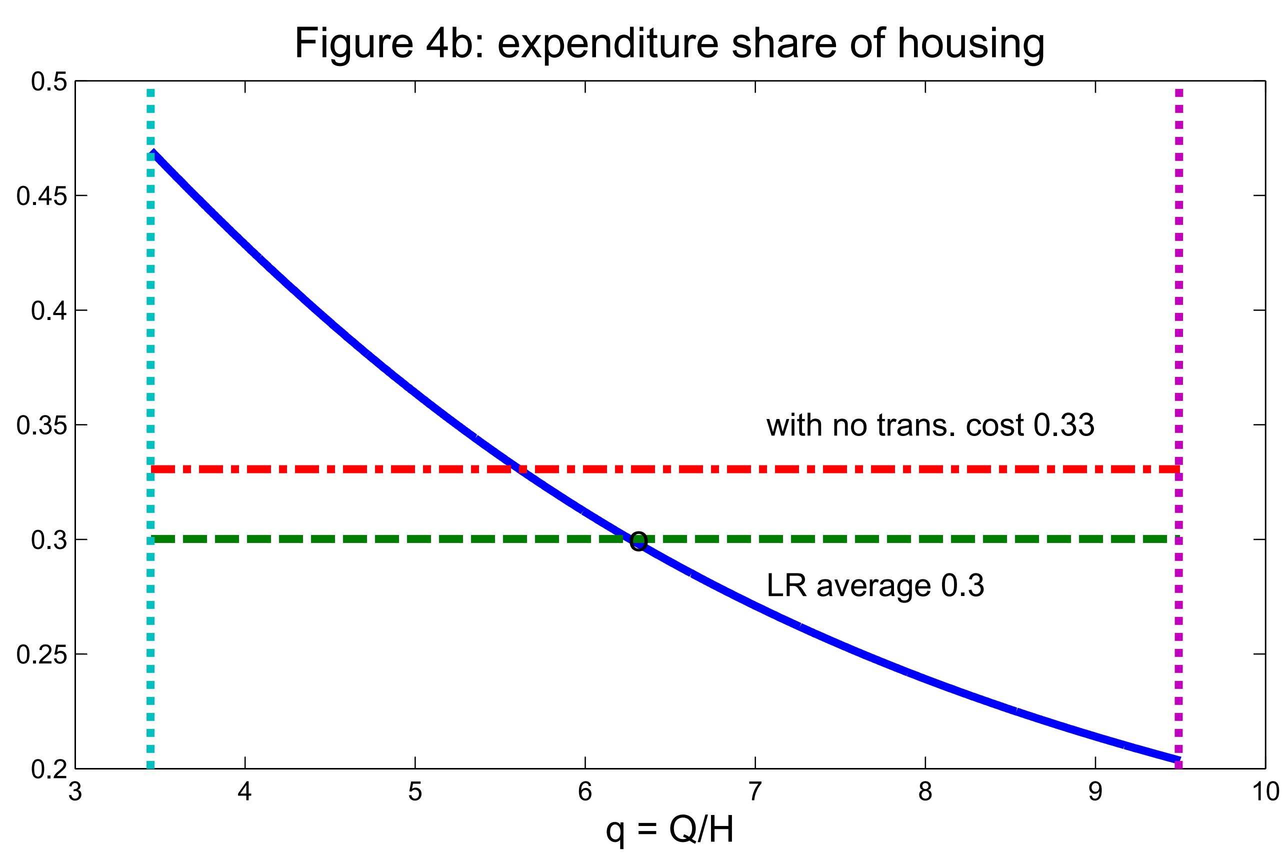
Figure 4b: expenditure share of housing
with no trans. cost 0.33
LR average 0.3
0.2
0.25
0.3
0.35
0.4
0.45
0.5
3
4
5
6
7
8
9
10
q = Q/H

no transaction cost. Nevertheless, her total expenditure increases less strongly with wealth than it would if the transaction cost were absent.

Figure 4b shows the housing share of total expenditure, which falls from about 47% at the lower threshold to about 20% at the upper threshold. The long run average value, which by construction is 30%, is slightly lower than the 33% for a consumer who faces no transaction costs (and has the same preference parameters).

On average the consumer spends about 60% of her income, slightly less than a consumer who faces no transaction costs.

Table 2 describes housing transactions. The fraction of downward adjustments is small, only 9%. Since the consumer's wealth grows, on average, only a (relatively rare) sequence of bad portfolio returns induces her to downsize her house.

Table 2

	at b	at B
probability of adjustment ($q = S$)	0.09	0.91
new/old house	0.54	1.51
new/old nondurable consumption	1.11	0.89
change in portfolio share	−0.06	−0.11

For transactions at the lower threshold the value of the new house is about 54% of the value of the one being sold, nondurable consumption rises by about 11%, and the portfolio share in the risky asset falls by 6 percentage points. For transactions at the upper threshold, the value of the new house is about 51% higher than the value of the one being sold, nondurable consumption falls by about 11%, and the portfolio share in the risky asset falls by 11 percentage points.

For a consumer who has just transacted, with $q = S$, the expected time to the next adjustment is 20.8 years. This figure is the expected value for completed durations, so it is not comparable to the cross section average of 11.3 years in the data.

c. Sensitivity analysis: ε.—

As noted above, there is conflicting evidence about the elasticity of substitution between housing and nondurables. To explore the effect of this parameter, we will next compare results for elasticities of $\varepsilon = 1/\zeta = 0.15$, 0.50, 1.0, and 1.25. In each case ω is adjusted to keep the average expenditure share for housing at 30%.

A higher elasticity allows the consumer to substitute more easily into nondurables as her wealth increases, reducing the incentive to pay the transaction cost associated with a housing adjustment. Thus, as shown in Table 3, the inaction region gets wider as the elasticity of substitution increases—b falls and B rises.

Table 3

ε	b	B	$\mathrm{E}[T]$	g	$\Pr(B)$
0.15	3.9	9.0	14.9	1.98	0.87
0.50	3.4	9.5	20.8	1.97	0.91
1.00	3.1	10.0	26.6	1.96	0.93
1.25	2.9	10.2	28.7	1.96	0.94

The wider inaction region leads to longer expected times between adjustments, with the expected duration rising from 14.9 years to 28.7 years. The long run growth rate g does not change much as the substitution elasticity varies, remaining at about 2.0%. The probability that the next adjustment is at the upper threshold increases slightly with the elasticity, rising from 0.87 to 0.94. As ε increases, the widening of the inaction region reduces the probability of a sequence of low returns sufficiently long and severe to induce a downward housing adjustment.

Figure 5a shows the portfolio policies. All four have a U shape. The U is flatter for higher elasticities, but the functions are quite similar except for the lowest elasticity, $\varepsilon = 0.15$, where it displays much sharper fluctuations.

Figure 5b shows nondurable consumption. In all cases it increases with wealth, but higher elasticities lead to larger adjustments—steeper slopes. The behavior of

Figure 5a: Portfolio shares, various subst. elasticities

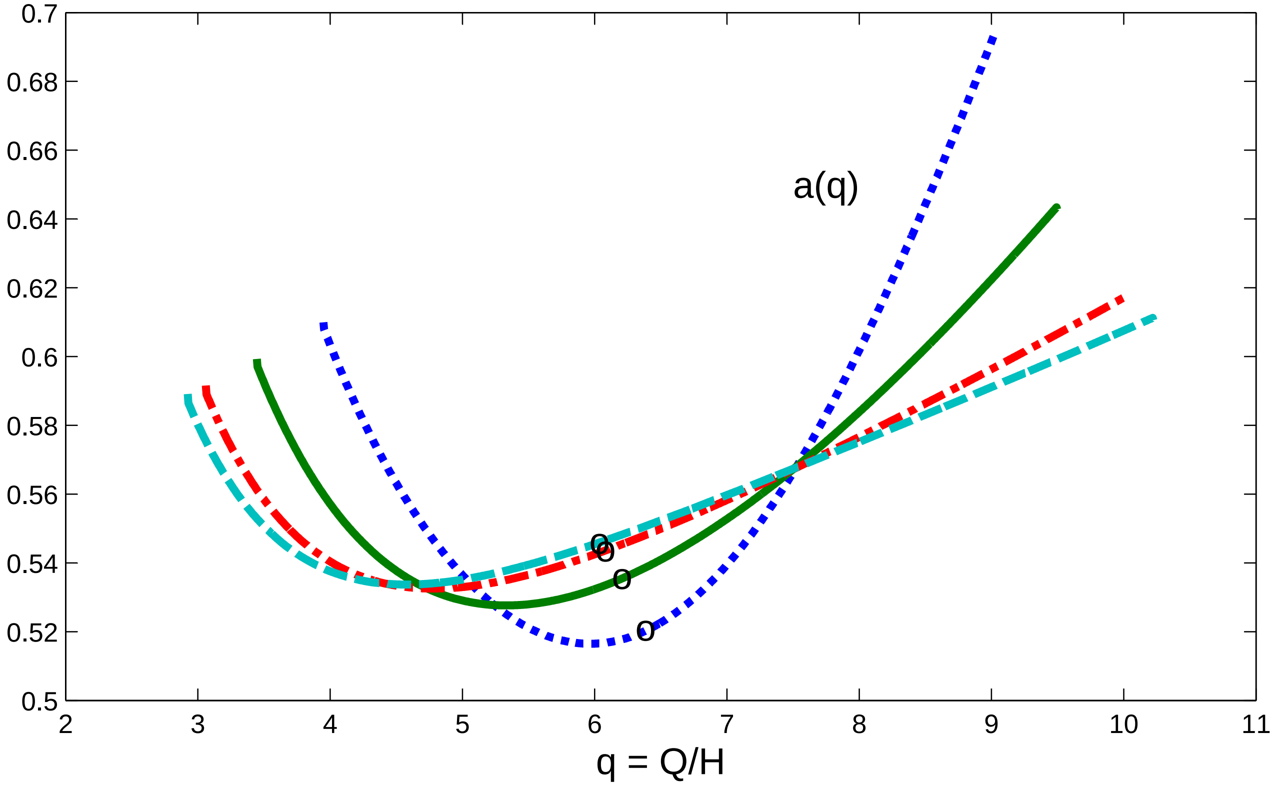

Figure 5b: Nondurable cons., various subst. elasticities

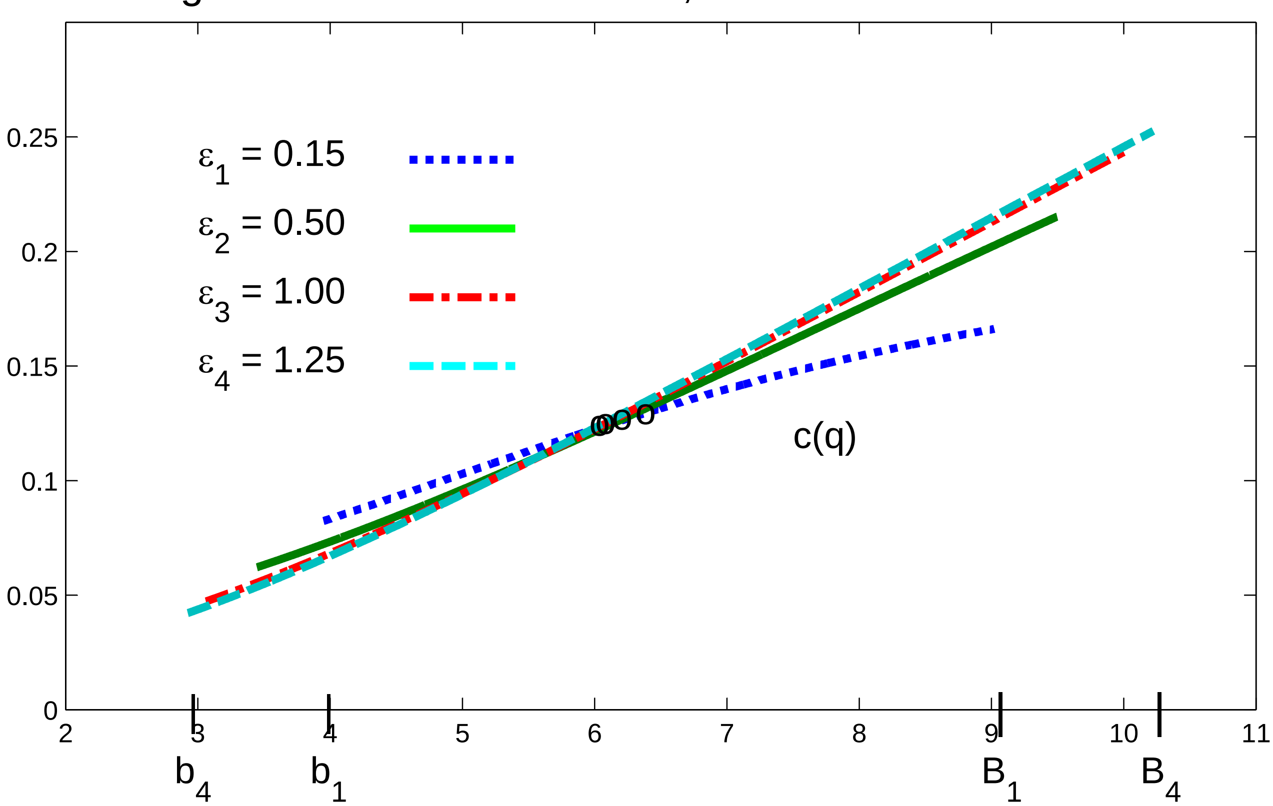

nondurable consumption is remarkably similar for the three higher elasticities. Only for $\varepsilon = 0.15$ is the curve significantly flatter.

These figures suggest why the elasticity of substitution is so difficult to estimate: within a broad range it has remarkably little effect on behavior.

Next consider the changes when a transaction is made. We will focus on the upper threshold, since most transactions occur there. These changes are displayed in Table 4.

Table 4, transactions at B

ε	$\hat{H}/H$	$\hat{C}_{ND}/C_{ND}$	$a(S) - a(B)$
0.15	1.4	1.09	−0.17
0.50	1.5	0.90	−0.11
1.00	1.6	0.84	−0.07
1.25	1.7	0.83	−0.06

The ratio of new to old house values after an adjustment, $\hat{H}/H$, increases slightly with the elasticity of substitution, from 1.4 to 1.7. This pattern is a straightforward result of the widening of the inaction region.

The ratio of new/old nondurable consumption, $\hat{C}_{ND}/C_{ND}$, falls with the elasticity. Recall that since $1/\theta = 0.25$, for $\varepsilon = 0.25$ (not shown), preferences are additively separable between housing and nondurables. In this case nondurable consumption is unchanged after a housing transaction, $\hat{C}_{ND}/C_{ND} = 1$. For $\varepsilon = 0.15$ the ratio exceeds one, and the consumer increases her nondurable consumption when she purchases a larger house. For the higher elasticities the ratio is less than one, and the consumer reduces her nondurable consumption after purchasing a larger house, with the size of the reduction increasing with the elasticity. Thus, for elasticities exceeding $1/\theta$ the consumer behaves like someone who is 'house poor,' although she is not liquidity constrained, as that term suggests.

When a transaction is made at the upper threshold, the consumer reduces the share of her portfolio in the risky asset. This shift simply reverses, in a single jump, the increases that occurred gradually as the consumer's wealth accumulated since the last housing transaction. The jump is larger for lower elasticities. For $\varepsilon = 0.15$, the consumer reduces her risky asset holdings by 17 percentage points. For $\varepsilon = 1.25$, the reduction is only 6 percentage points.

d. Sensitivity analysis: $\lambda, m, \kappa, \theta$.—

A higher transaction cost λ has two effects: it increases the incentive to avoid a housing adjustment and it increases the cost of housing. The first effect tends to widen the inaction region and increase the expected duration. The second effect shifts the consumption mix toward nondurables. Thus, it increases the ratio of total wealth to housing, both at the transaction point and on average, and it tends to shift both thresholds upward. Thus, the two effects work in opposite directions at the lower threshold and in the same direction at the upper threshold.

Figure 6 displays these effects. An increase in λ widens the inaction region, and the effect is greater at the upper end. A higher transaction cost also makes the consumer more risk averse, but has virtually no effect on nondurable consumption for a given ratio of total wealth to housing. Instead it simply shifts the likely range for that ratio. The expected duration between adjustments rises from 16.0 years for $\lambda = 0.04$ to 20.8 and 23.8 at the higher values.

A higher maintenance cost m also makes housing more expensive, shifting consumption toward nondurables and increasing the ratio of total wealth to housing wealth. Thus, an increase in m shifts the curves in Figures 3a and 3b to the right, raising the critical values b, S, B.

A positive hazard rate for exogenous moves, $\kappa > 0$, means that the consumer is sometimes forced to sell her house and pay the transaction cost. This makes housing less attractive and moves more frequent. As with an increase in the maintenance

Figure 6a: Portfolio shares, various transaction costs

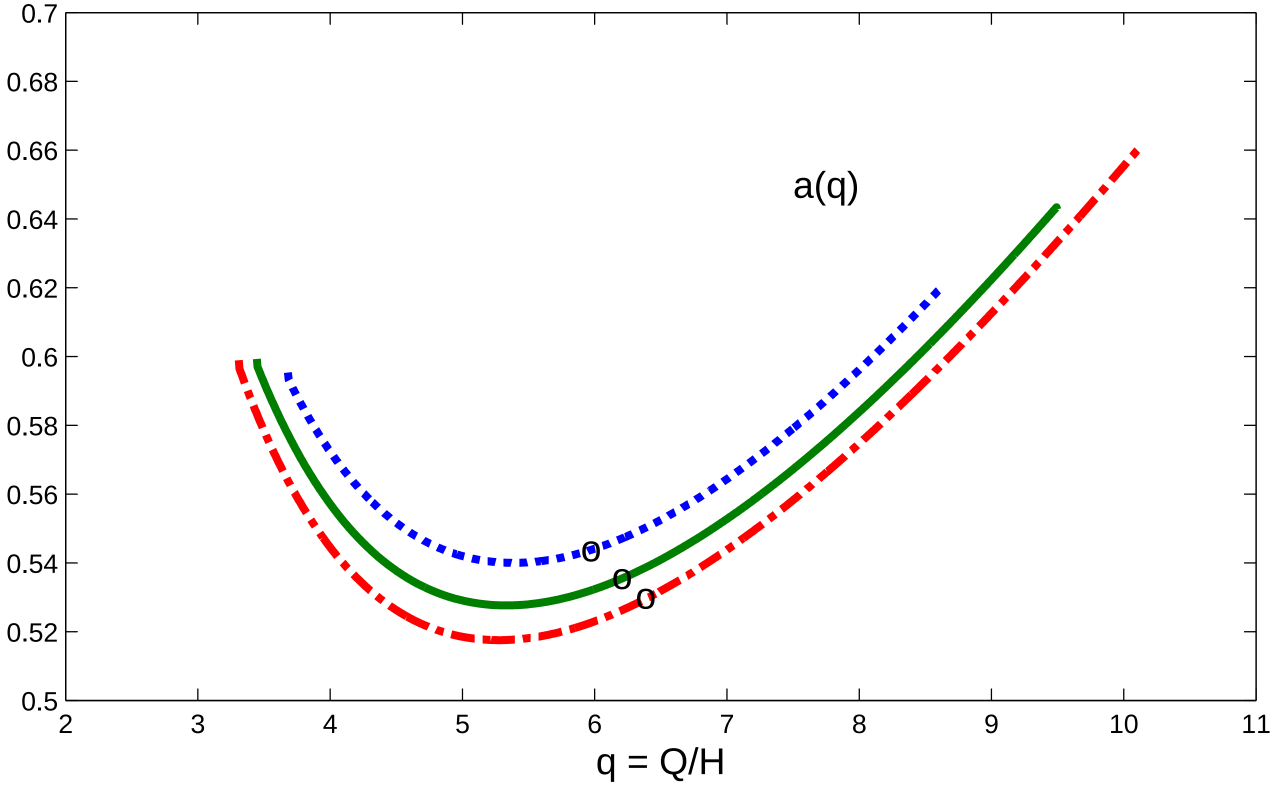

Figure 6b: Nondurable cons., various transaction costs

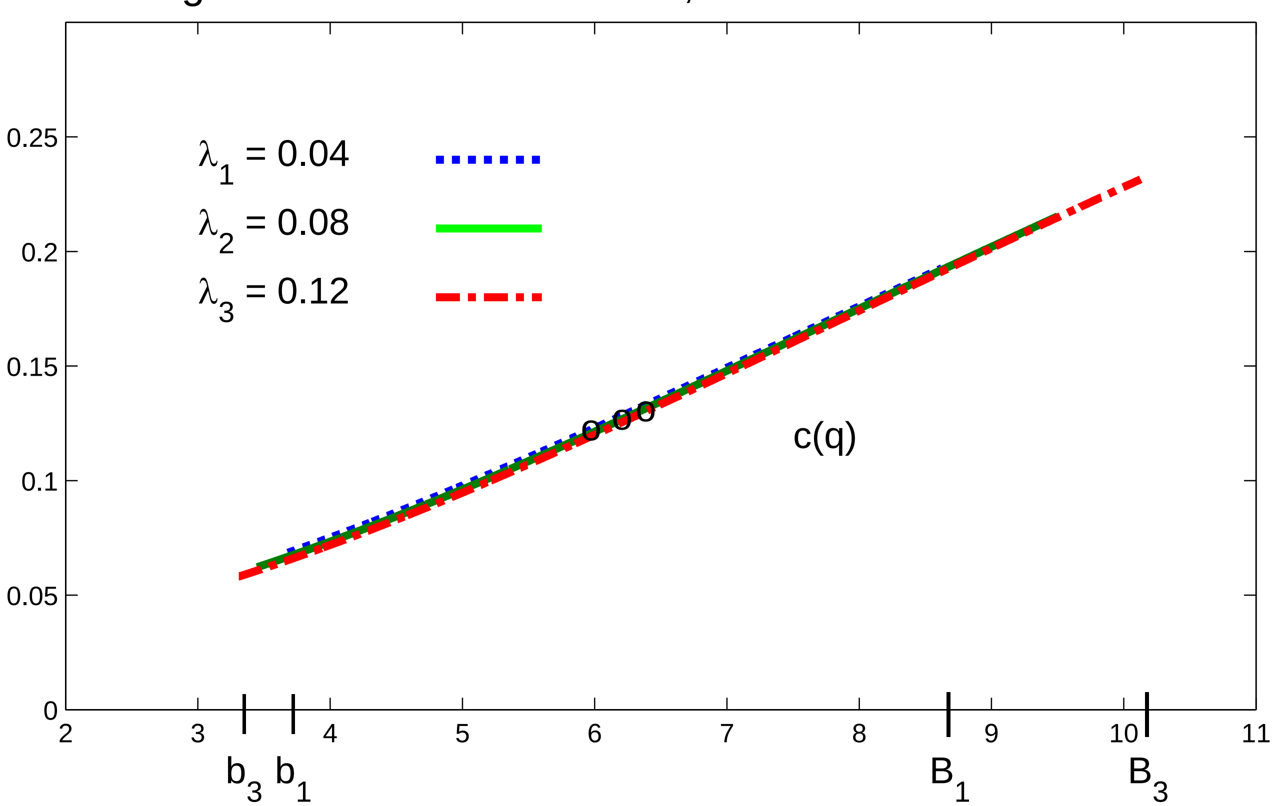

cost, the critical values b, S and B increase, and the curves describing the portfolio and nondurable consumption shift to the right. For a hazard rate of $\kappa = 0.01.$, the average ratio of total wealth to housing rises from 6.4 to 7.0, the expected duration falls from 20.8 to 19.6 years.

Figure 7 shows the effects of changing the risk aversion parameter θ. The most dramatic effect is, as expected, on the portfolio, with more risk tolerant consumers holding more of the risky asset. For $\theta = 2$ the short sale constraint comes into play, constraining the consumer when she is near either transaction threshold. A reduction in θ increases the consumer's willingness to substitute intertemporally, shifting the lower threshold b downward. Changes in θ also produce large changes in the average growth rate, as suggested by Figure 1c.

7. The equity premium puzzle

A standard exercise in finance[5] uses the Euler equation

$$\mathrm{E}_t \left[\frac{U'(X_{t+s})}{U'(X_t)} e^{-\rho s} \left(1 + r^j_{t+s}\right) \right] = 1, \tag{17}$$

which hold for any asset j, to conclude that

$$\mu^j - r = \theta \mathrm{Cov}\left(\frac{dX_t}{X_t}, r^j_t \right), \tag{18}$$

where X_t is total consumption, r^j_{t+s} is the instantaneous return on asset j at $t+s$, μ^j is the expected return on asset j, r is the risk-free rate, and θ is the coefficient of relative risk aversion. This is the equation commonly used to back out an estimate of the risk aversion parameter θ, using data on consumption growth and asset returns.

The equity premium puzzle noted by Mehra and Prescott (1985) is a puzzle because the covariance of consumption growth with asset returns is low, while the excess return on risky assets is high. Thus, a large value of θ is needed to justify the

[5]See Mankiw and Zeldes (1991) for more detail.

Figure 7a: Portfolio shares, various risk aversion

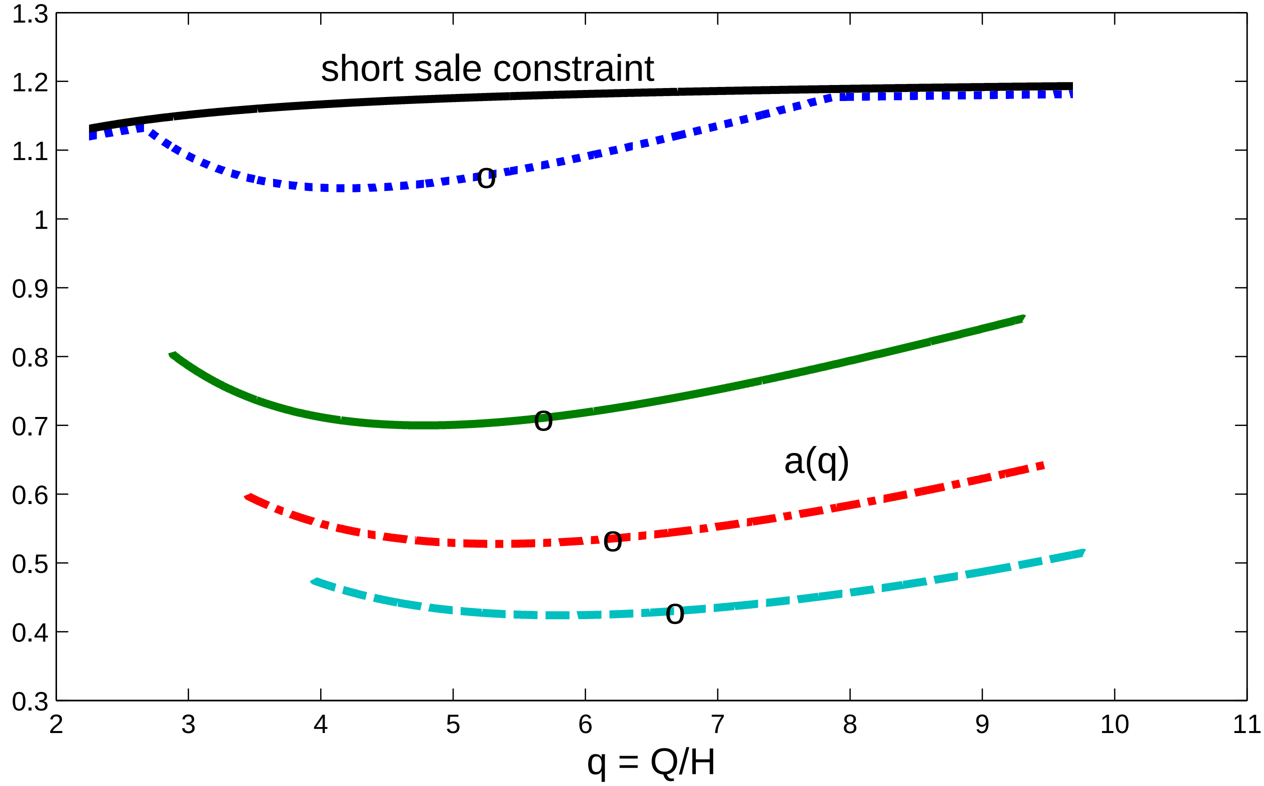

Figure 7b: Nondurable cons., various risk aversion

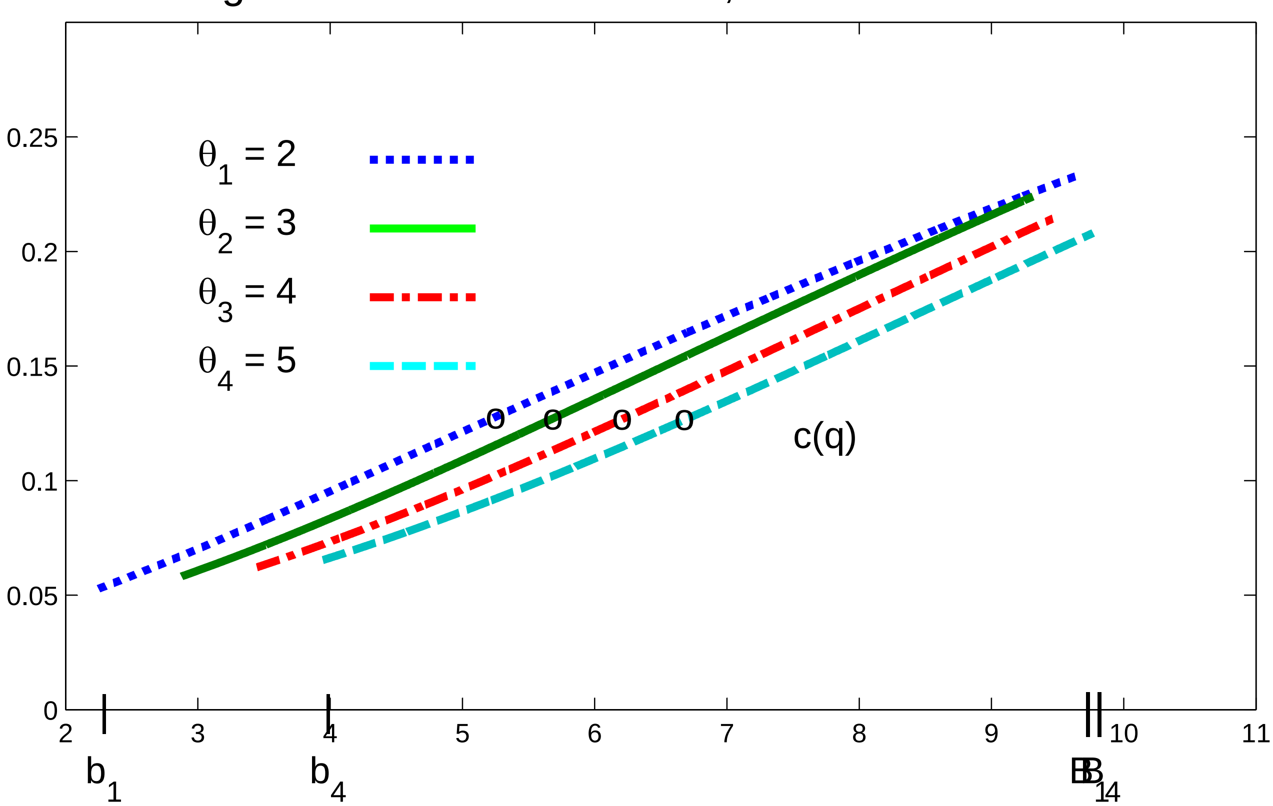

excess return on the left side of (18). While it is usually labeled as a puzzle about the excessively high return on equity, it can as well be viewed as a puzzle about the excessive smoothness of consumption.

But the relationship in (17) is derived in a frictionless model. If some components of consumption are costly to adjust, then they will vary less than predicted by the frictionless model, and hence the covariance of consumption growth and asset returns will be lower. Thus, a transaction cost for housing offers a potential explanation for smooth consumption. With an adjustment cost, housing consumption is constant over long intervals, and for elasticities $\varepsilon < 1/\theta$ nondurable consumption is also smoother. The question then is quantitative: are these effects large enough to explain the puzzle? Using the model here we can calculate the magnitude of the error an econometrician would make if he estimated θ using the (misspecified) frictionless model.

Let $r_t^a = \mu dt + \sigma dz_t$ denote the instantaneous return on the model's risky asset. First note that with no transaction cost, as in section 3, consumption growth tracks growth in wealth, so

$$\frac{dX_t}{X_t} = \frac{dQ_t}{Q_t} = \left[(1-a^*)\, r + \mu a^* - x^*\right] dt + \sigma a^* dz_t,$$

where x^* is the (constant) ratio of consumption expenditures to wealth and a^* is the (constant) portfolio share in the risky asset. The term in square brackets is not stochastic, so

$$\begin{aligned} \operatorname{Cov}\left(\frac{dX_t}{X_t}, r_t^a\right) &= \frac{1}{dt}\mathrm{E}\left[(\sigma a^* dz_t)(\sigma dz_t)\right] \\ &= a^*\sigma^2. \end{aligned}$$

Since $a^* = (\mu - r)/\sigma^2\theta$, the econometrician using (18) would obtain an estimate (neglecting sampling error) of θ. In the absence of a transaction cost the model is correctly specified, and the estimate of θ is correct.

Now suppose there is a transaction cost, and consider a consumer who is using the thresholds b, S, B and the policy functions $c(q)$ and $a(q)$. To compute the covariance that would be obtained using long time series, we must average over q values inside the interval (b, B) using the stationary density and also take into account the discrete jumps that occur at the boundaries.

First consider expenditure growth inside the inaction region. If the consumer's wealth is q_t, then her consumption expenditure is

$$X_t = p_h + c(q_t),$$

and the increment to her wealth is

$$dq_t = h(q_t)dt + a(q_t)q_t\sigma dz_t,$$

where

$$h(q_t) \equiv r\left[1 - a(q_t)\right]q_t + \mu a(q_t)q_t - \left[p_h + c(q_t)\right]$$

is the expected return on her portfolio less consumption expenditures. Thus, inside the inaction region expenditure growth is

$$\frac{dX_t}{X_t} = \frac{c'(q_t)}{p_h + c(q_t)}\left[h(q_t)dt + a(q_t)q_t\sigma dz_t\right].$$

As before the first term is not stochastic. Thus, averaging across wealth levels with the stationary distribution $\psi(q)$, we obtain

$$\mathrm{Cov}\left(\frac{dX_t}{X_t}, r_t^a\right) = \sigma^2\int_b^B \frac{c'(q)}{p_h + c(q)}a(q)q\psi(q)dq + M, \tag{19}$$

where M is the contribution of the jump terms.

Next consider jumps. An adjustment at B means the consumer is purchasing a larger house. Although her nondurable consumption may fall, depending on the calibration (cf. Table 2), her total expenditure always increases. The reverse occurs after an adjustment at b. Let $J_B > 0$ and $J_b < 0$ denote the expenditure jumps at the two thresholds. Next, note that the jump at B occurs only if $dz > 0$, and the jump at b only if $dz < 0$. Hence $M \geq 0$, and the jump terms can only add to the (positive) first term in (19). Thus, setting $M = 0$ gives a lower bound on the covariance, and using (19) we can place an upper bound on the value $\hat{\theta}$ that an econometrician using (18) would obtain (neglecting sampling error).

For the benchmark calibration the estimated value would be $\hat{\theta} \leq 4.6$, while the true value is $\theta = 4.0$. A lower elasticity of substitution between housing and nondurables increases the size of the error, and for the very low elasticity of $\varepsilon = 0.15$ the estimated value is $\hat{\theta} \leq 7.1$ Larger transaction costs also increase the error, but only slightly. With an elasticity of $\varepsilon = 0.15$ and a transaction cost of $\lambda = 0.15$, the estimated value is still $\hat{\theta} \leq 7.1$. Thus, while the effect is in the right direction it is too small to explain much of the equity premium puzzle.

Figure 8 displays total expenditures as a function of total wealth, for a consumer with a house normalized to size one, for various scenarios. The dotted line is for a consumer who faces no transaction cost, with ω calibrated to give housing a share of 30% in total consumption. The substitution elasticity does not matter for this consumer. The other three curves describe consumers who face a transaction cost of 8% and have elasticities as indicated. Even for the very low elasticity, $\varepsilon = 0.15$, the transaction cost has has a modest effect on total expenditure. A higher transaction cost, 15% instead of 8%, narrows the inaction region for this consumer, but otherwise has almost no impact.

8. Conclusions

Adjustment costs for housing are substantial, so it is not surprising that they have significant effects on portfolios as well. Empirical work thus far confirms this conclusion, with evidence that anticipation of an impending housing adjustment, length of tenure in the current house, or the ratio of housing wealth to total wealth affect the more flexible aspects of behavior like nondurable consumption and portfolios. The model here provides a basis for refining empirical frameworks of this type. It also

Figure 8: Total expenditures, various elasticities

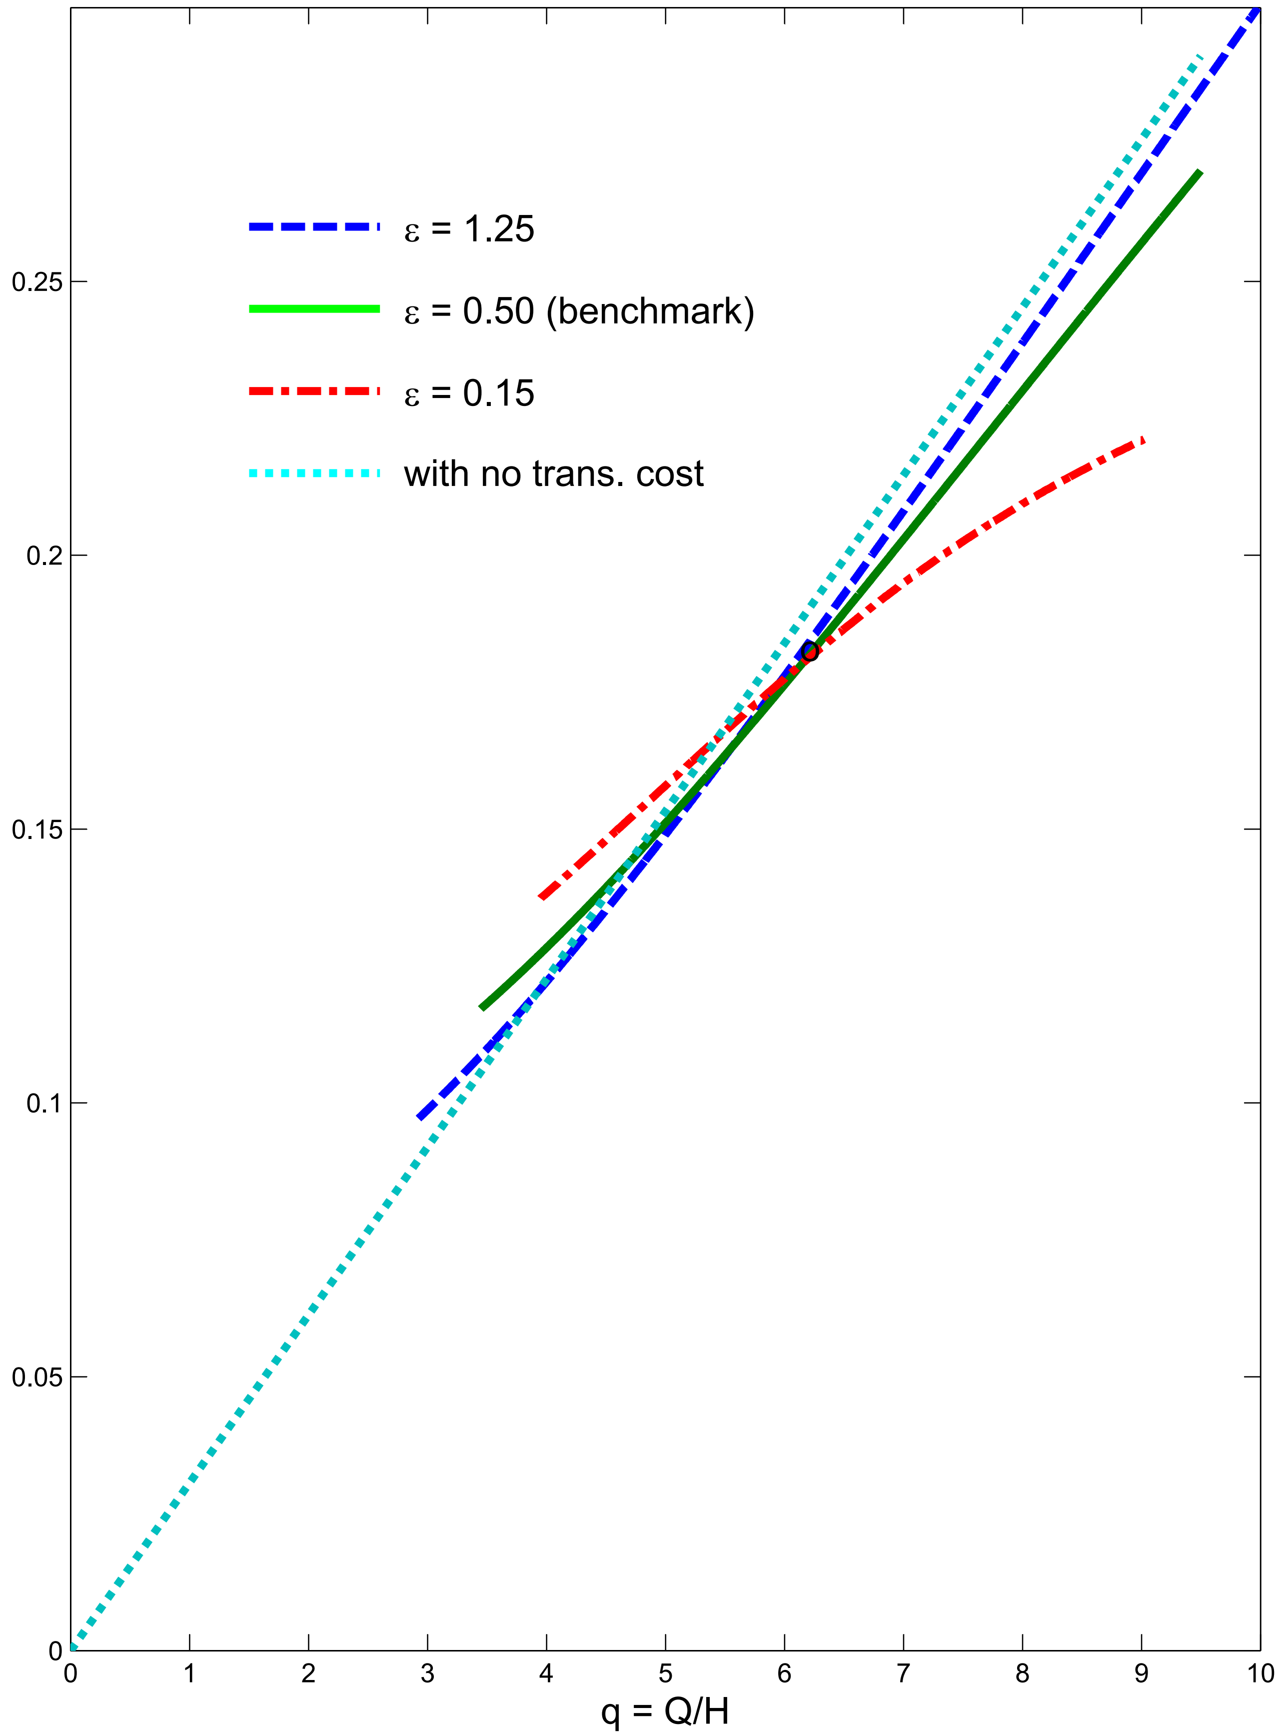

offers a rich set of predictions for other aspects of consumer behavior, such moving probabilities and the ratio of new to old house values when a family moves.

Adjustment costs for housing have surprisingly little impact on consumption of nondurables, however, and hence they provide little help in resolving the equity premium puzzle. Even with a very low substitution elasticity, there is simply too much response in nondurable consumption. Increasing the expenditure share for 'housing,' by including transportation, furniture, and other consumption components that are closely linked to housing, might help. But it would then be difficult to justify a low substitution elasticity between the broad 'housing' good and the remaining set of nondurables.

It is interesting that the elasticity of substitution between housing and nondurables plays a relatively minor role. The fact that changing that parameter over a wide range has so little effect on the simulation results may explain why empirical studies have produced such a wide range of estimates.

It is also interesting that the model here produces a value function that is strictly concave over its entire domain. Thus, it is not adjustment costs *per se* that produce the non-concavities in Chetty and Szeidl (2007) and Vereshchagina (2007).

The model here excludes several important features: labor income, life cycle considerations, and housing price risk. Extending the model to include them is an interesting avenue for further research.

APPENDIX

PROOF OF PROPOSITION 1: First we must show that (6) has a finite maximum. If $0 < \theta < 1$, we must show that $\rho > \Gamma$ for all feasible (c, h, a), so that utility does not diverge to $+\infty$. If $\theta > 1$, we must show that $\rho > \Gamma$ for at least one feasible (c, h, a), so that utility does not diverge to $-\infty$.

To this end we will first show that if

$$\rho > \Gamma(c,h,a;\theta), \begin{cases} \text{all feasible } (c,h,a), & \text{if } 0<\theta<1, \\ \text{some feasible } (c,h,a), & \text{if } \theta>1, \end{cases} \tag{20}$$

then the optimal portfolio $\alpha(h;\theta)$ is

$$\alpha(h;\theta) = \min\left\{\frac{\gamma}{\theta}, a_{ss}(1-\epsilon h)\right\}. \tag{21}$$

Then we will show that under Assumptions 1 and 2, (20) holds.

(i) Suppose (20) holds. Since a appears in (6) only as an argument of Γ, the optimal portfolio solves

$$\max_{a\in[0,a_{ss}]} \frac{1}{1-\theta}\frac{1}{\rho-\Gamma(c,h,a;\theta)}.$$

Hence the objective is to maximize Γ if $0<\theta<1$ and to minimize Γ if $\theta>1$. Note that

$$\begin{aligned} \Gamma_a(c,h,a) &= (1-\theta)\left[(\mu-r)-\theta a\sigma^2\right] \\ &= (1-\theta)\,\sigma^2(\gamma-\theta a). \end{aligned}$$

If $0<\theta<1$, then $\Gamma_{aa}<0$, so Γ is concave. Since $\Gamma_a(c,h,0)>0$, there cannot be an optimum at $a=0$. If $\gamma/\theta > a_{ss}(1-\epsilon h)$, then

$$\Gamma_a(c,h,a_{ss}(1-\epsilon h);\theta) = (1-\theta)\,\sigma^2\left[\gamma-\theta a_{ss}(1-\epsilon h)\right] > 0,$$

so the solution is at a corner, $\alpha(h;\theta) = a_{ss}(1-\epsilon h)$. Otherwise the solution is interior and satisfies $\Gamma_a = 0$. Hence the optimal portfolio is as in (21).

If $\theta>1$, the objective is to minimize $\Gamma(c,h,a)$. In this case Γ is convex, and the preceding argument holds with a sign change. Hence (21) also holds for $\theta>1$.

(ii) Next we will show that Assumptions 1 and 2 insure (20) holds. Suppose $0<\theta<1$. Then

$$\frac{d}{dh}\Gamma(c,h,\alpha(h;\theta);\theta) \quad = \quad \Gamma_h + \Gamma_a\alpha'(h;\theta)$$

$$
\begin{aligned}
&= -(1-\theta)(p_h+c) + (1-\theta)\sigma^2[\gamma - \theta\alpha(h;\theta)]\,\alpha'(h;\theta) \\
&\leq -(1-\theta)(p_h+c) \\
&< 0, \qquad (22)
\end{aligned}
$$

where the second line uses the fact that $\gamma/\theta \geq \alpha(h;\theta)$ and $\alpha'(h;\theta) \leq 0$. Hence for any $\epsilon \geq 0$,

$$
\rho > \Gamma(0,0,\alpha(0;\theta);\theta) \geq \Gamma(c,h,\alpha(h;\theta);\theta) \geq \Gamma(c,h,a;\theta), \quad \text{all feasible } (c,h,a),
$$

where the first inequality uses Assumption 2, the second uses (22) and the fact that $\partial\Gamma/\partial c < 0$, and the third uses the fact that $\alpha(h;\theta,\epsilon)$ maximizes $\Gamma(c,h,a;\theta)$.

If $\theta > 1$, then since $r > 0$, for $a = 0$ and all c, h sufficiently small,

$$
\rho > 0 > (1-\theta)[r - (p_h+c)h] = \Gamma(c,h,0).
$$

Proof of part (a): Suppose $\epsilon = 0$. The optimal portfolio a_R is as in (21). And for any expenditure flow $E > 0$, the optimal consumption mix solves

$$
\max_{c,h} \; u(c)h \qquad \text{s.t. } (p_h+c)h = E,
$$

so

$$
c_R = \arg\max_c \frac{u(c)}{p_h+c}.
$$

For the CES preferences here, the solution is (7).

Since a_R does not involve h, maximizing (6) with respect to h implies

$$
\begin{aligned}
0 &= \frac{1-\theta}{h_R} + \frac{\Gamma_h}{\rho-\Gamma} \\
&= (1-\theta)\left(\frac{1}{h_R} - \frac{p_h+c_R}{\rho-\Gamma}\right),
\end{aligned}
$$

or

$$
(p_h+c_R)h_R = \rho - \Gamma[0,0,\alpha_R(\theta);\theta] + (1-\theta)(p_h+c_R)h_R,
$$

or

$$\theta\left(p_h+c_R\right)h_R=\rho-(1-\theta)\left[r+(\mu-r)\,a_R(\theta)-\theta\frac{1}{2}\sigma^2a_R^2(\theta)\right],$$

as claimed. Then

$$\left(p_h+c_R\right)h_R=\frac{\rho-r}{\theta}+r-(1-\theta)\left(\frac{1}{\theta}\gamma\sigma^2a_R-\frac{1}{2}\sigma^2a_R^2\right)$$

so

$$\begin{aligned}\left(p_h+c_R\right)h_R'(\theta) &= -\frac{1}{\theta^2}\left(\rho-r-\gamma\sigma^2a_R\right)-\frac{1}{2}\sigma^2\alpha_R^2-(1-\theta)\left(\frac{1}{\theta}\gamma-\alpha_R\right)\sigma^2a_R' \\ &= \frac{1}{\theta^2}\left[r-\rho+\left(\gamma-\frac{1}{2}\theta^2\alpha_R\right)\sigma^2a_R\right],\end{aligned}$$

establishing the last claim.

Proof of part (b): The solution in (7) solves the problem with $\epsilon>0$ if and only if a_R,h_R satisfies the tighter portfolio constraint. Otherwise the tighter portfolio constraint binds, so $h_B(\theta,\epsilon)<h_R(\theta)$. Using (21), the consumer's problem is as in (8). The conditions for a maximum are

$$\begin{aligned}0 &= \frac{(1-\theta)\,u'(c)}{u(c)}+\frac{\Gamma_c}{\rho-\Gamma}, \\ 0 &= \frac{1-\theta}{h}+\frac{\Gamma_h-\Gamma_a\epsilon a_{ss}}{\rho-\Gamma},\end{aligned}$$

or

$$\begin{aligned}\frac{u'(c)}{u(c)} &= \frac{h}{\rho-\Gamma}, \\ \frac{1}{h} &= \frac{c+p_h+\left[\gamma-\theta a_{ss}\left(1-\epsilon h\right)\right]\sigma^2\epsilon a_{ss}}{\rho-\Gamma}.\end{aligned}$$

Combining these two gives gives

$$\begin{aligned}\frac{1}{c+p_h+\left[\gamma/\theta-a_{ss}\left(1-\epsilon h\right)\right]\theta\sigma^2\epsilon a_{ss}} &= \frac{u'(c)}{u(c)} \\ &= \frac{\omega}{\omega c+(1-\omega)\,c^{\zeta}},\end{aligned}$$

so

$$c_B(\theta,\epsilon) = \left[\frac{\omega}{1-\omega}\left\{p_h + \left[\gamma/\theta - a_{ss}\left(1-\epsilon h\right)\right]\theta\sigma^2\epsilon a_{ss}\right\}\right]^{1/\zeta}.$$

For a renter the term on the right in square brackets is zero, while for a constrained buyer it is positive. Hence $c_B(\theta,\epsilon) > c_R$. ■

PROOF THAT v is well defined: It suffice so show that for any fixed initial condition (Q_0, H_0) with $Q_0 > \lambda H_0$, the value $V(Q_0, H_0)$ is finite.

The transaction cost cannot raise utility, so clearly $V(Q_0, H_0)$ is bounded above by $Q_0^{1-\theta}w^*$. In addition, under Assumption 1 it is possible to choose a feasible strategy for which expected utility is bounded below. For example, set $A \equiv 0$, so the risky asset is not held; let $C = cH$, where $c > 0$ is small; set $H = hQ$ when a transaction is made, where $h > 0$ is small; choose a long period T between transactions; and make the first transaction at date 0. During intervals when no transaction is made wealth grows at a constant rate

$$\frac{dQ}{Q} = \left[r - \left(p_h + c\right)h\right]dt.$$

For c and h sufficiently small, this growth rate is positive. Hence $W_{n+1} > W_n$, all n. where W_n is wealth after the n^{th} housing transaction. Over each interval with no transaction, utility is $vW_n^{1-\theta}$, where v is a constant. Hence lifetime utility under this strategy is

$$\sum_{n=0}^{\infty} e^{-\rho nT}vW_n^{1-\theta} \geq v\left(Q_0 - \lambda H_0\right)^{1-\theta}\frac{1}{1-e^{\rho T}},$$

where the right side is finite.

REFERENCES

[1] Aten, Bettina H. 2005. Report on interarea price levels, working paper 2005-11, Bureau of Economic Analysis, November.

[2] Campbell, John Y. and Joao R. Cocco. 2005. How do house prices affect consumption? Evidence from micro data, NBER working paper 11534.

[3] Chetty, Raj and Adam Szeidl. 2004. Consumption commitments: neoclassical foundations for habit formation, NBER working paper 10970.

[4] Chetty, Raj and Adam Szeidl. 2007. Consumption commitments and risk preferences, *Quarterly Journal of Economics*, 122: 831-877

[5] Cocco, J., 2005. Portfolio choice in the presence of housing, *Review of Financial Studies* 18, 535-567.

[6] Dunn, K., Singleton, K., 1986. Modeling the term structure of interest rates under non-separable utility and durability of goods, *Journal of Financial Economics* 17, 27-55.

[7] Eichenbaum, Martin S. and Lars P. Hansen. 1990. Estimating models with intertemporal substitution using aggregate time series data. *Journal of Business and Economic Statistics* 8, 53-69.

[8] Fillat, Jose L. Housing as a measure for long-run risk in asset pricing, working paper, University of Chicago.

[9] Flavin, Marjorie and Shinobu Nakagawa. 2004. A model of housing in the presence of adjustment costs: a structural interpretation of habit persistence, NBER working paper 10458.

[10] Flavin, Marjorie and Takashi Yamashita, 2002. Owner-occupied housing and the composition of the household portfolio, *American Economic Review* 92, 345-362.

[11] Fukushima, Kenichi. 2005. Asset pricing implications of precommitted consumption, working paper, U. of Minnesota.

[12] Grossman, Sanford and Guy Laroque. 1990. Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption good, *Econometrica* 58: 25-51.

[13] Hanushek, Eric A. and John M. Quigley. 1980. What is the price elasticity of housing demand? *Review of Economics and Statistics* 62: 449-454.

[14] Kocherlakota, Narayana R. 1996. The equity premium: it's still a puzzle, *Journal of Economic Literature*, 34: 42-71.

[15] Kullman, Cornelia, and Stephan Siegel. 2004. Real estate and its role in household portfolio choice, Working paper, UBC.

[16] Lustig, H., Van Nieuwerburgh, S., 2006. Housing collateral, consumption insurance and risk premia, working paper, UCLA..

[17] Mankiw, N. Gregory and Stephen P. Zeldes. 1991. The consumption of stockholders and nonstockholders, *Journal of Financial Economics* 29: 97-112.

[18] Marshall, David A. and Nayan G. Parekh. 1999. Can costs of consumption adjustment explain asset pricing puzzles? *Journal of Finance* 54: 623-654.

[19] Martin, Robert F. 2003 Consumption, durable goods, and transaction costs, International Finance Discussion Paper No. 756, Federal Reserve Board, Washington, DC.

[20] Mehra, Rajnish and Edward C. Prescott. 1985. The equity premium: a puzzle, *Journal of Monetary Economics* 15: 145-61.

[21] Mehra, Rajnish and Edward C. Prescott. 2007. The equity premium: ABCs. In *Handbook of the Equity Risk Premium*, Rajnish Mehra, ed., Elsevier, forthcoming.

[22] Piazzesi, Monika, Martin Schneider, and Selale Tuzel. 2007. Housing, consumption and asset pricing, *Journal of Financial Economics*, 83: 531-569.

[23] Schachter, Jason P. and Jeffrey J. Kuenzi. 2002. Seasonality of moves and the duration and tenure of residence: 1996, Population Division working paper 69, U.S. Census Bureau, Washington, DC, December.

[24] Siegel, Stephan. 2004. Consumption-based asset pricing: durable goods, adjustment costs, and aggregation, working paper, Graduate School of Business, Columbia University.

[25] Smith, L.B., K.T. Rosen, and G. Fallis. 1988. Recent developments in economic models of housing markets, *Journal of Economic Literature* XXVI, 29-64.

[26] Stokey, Nancy L. 2007. *Brownian Models in Economics*, Princeton University Press, forthcoming.

[27] Vereshchagina, Galina. 2007. Preferences for risk in a dynamic model with consumption commitments. University of Iowa working paper.