PERUVIAN ECONOMIC ASSOCIATION

On the cyclicality of durable consumption responses

Ken Miyahara

Working Paper No. 210, September 2025

ON THE CYCLICALITY OF DURABLE CONSUMPTION RESPONSES*

Ken Miyahara

September 2, 2025

Abstract

Building on canonical asset pricing and portfolio choice frameworks with illiquid durable goods (Grossman and Laroque (1990), Flavin and Nakagawa (2008), Stokey (2009)), I propose a heterogeneous agent portfolio choice model to assess the cyclicality of aggregate durable consumption responses. The model features idiosyncratic utility switching costs that allow it to match the distribution of durable adjustments *sizes* in PSID data. By leveraging a structural mapping between adjustment hazards and the cross-sectional distribution of wealth, the framework provides a robust method for estimating fundamental adjustment frictions directly from observed behavior. We find that the model predicts procyclical and non-linear durable demand responses and a disproportionate decline in upward adjustments during recessions. The main result demonstrates that policy, such as fiscal stimulus payments, is significantly more potent during economic booms than in recessions. This asymmetry highlights the state-dependent nature of stabilization policies and their varying effectiveness across the business cycle.

JEL Classification Numbers: E2, G11, G51

Key Words: Lumpy adjustment, portfolio choice, generalized hazard function.

*Miyahara: University of Chicago (email: kmiyahara@uchicago.edu). I thank George Constantinides and Lars Peter Hansen for their guidance and support throughout this project. The Matlab code used draws heavily from Greg Kaplan's teaching codes.

1 Introduction

Micro-level lumpy behavior has significant aggregate implications. Prior research on capital investment dynamics argues that a key consequence of micro-level lumpiness is the history dependence of aggregate impulse responses (Bachmann, Caballero, and Engel, 2006). Specifically, investment responses to further shocks are larger during economic booms than during recessions. Similar state-dependent behavior have been documented for durable goods demand by Berger and Vavra (2015). Building on these research questions, this paper investigates the cyclicality of durable consumption responses using a novel heterogeneous-agent portfolio choice model.

Figure 1, which uses data from the US Panel Study of Income Dynamics (PSID), illustrates empirical patterns of aggregate durable demand. The first panel shows a significant decrease in the frequency of housing adjustments during the Great Financial Crisis, while the second panel reveals that this decline coincided with a shift in the composition of households, with downsizing becoming relatively more frequent. This observation that the frequency and composition of durable adjustments vary with the business cycle suggests that aggregate demand responses are state-dependent, with weaker responses likely during recessions.

These observations raise several important questions for macroeconomic analysis. How can we accurately predict these swings in the frequency and composition of durable adjustments? How responsive is the economy to further shocks in these different states? And how can we rigorously discipline the answers to these questions with observed durable demand behavior?

This paper addresses these questions by proposing a tractable heterogeneous-agent portfolio choice model. The model is built on the empirical observation that the size of durable adjustments—the change in value between a newly acquired durable and the old one—is highly heterogeneous across agents. The estimation strategy introduces three innovations relative to the existing literature. First, it targets the observed distribution of durable adjustment sizes from PSID data. Second, it exploits the structural mapping between adjustment

(a) Frequency of Housing Adjustment (PSID) (b) Fraction of House Upsizing (PSID)

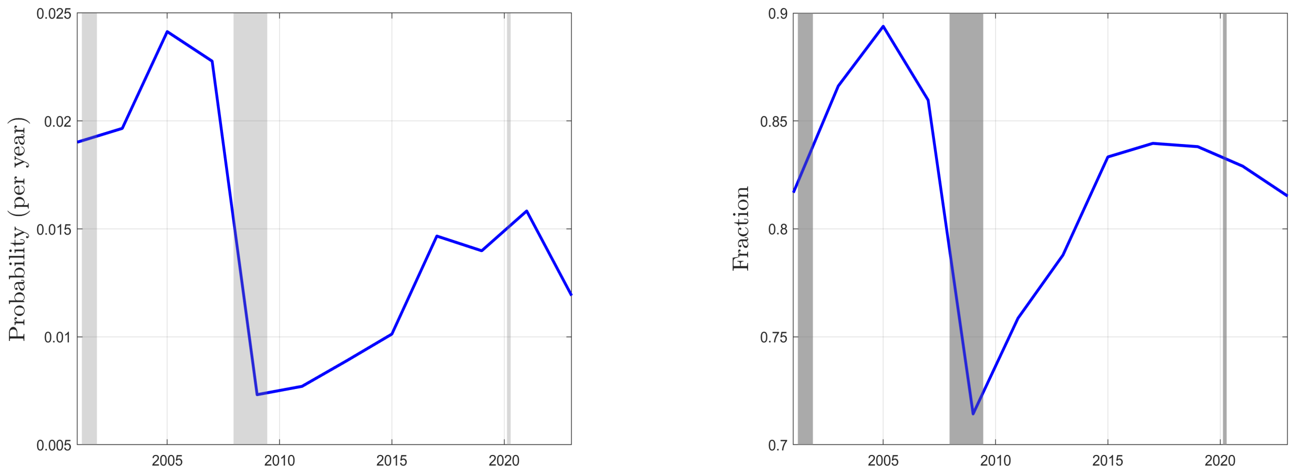

Note: Own calculations using PSID data. Probabilities are empirical frequencies of occurrence in the sample.

Figure 1: Housing Adjustment in the US (1999-2023)

hazards (the probabilities of adjustment per unit of time) and the cross-sectional distribution of wealth. Third, it infers microeconomic frictions purely from observed *adjustment* behavior, avoiding the need to impute desired sizes of adjustment using potentially noisy wealth and income data as a proxy for the present value of total wealth.

Model summary. The model develops a continuous-time portfolio choice framework where households derive utility from both non-durable consumption and durable goods. Agents face transaction costs proportional to the durable value and idiosyncratic utility switching costs, which create a disincentive for frequent durable adjustments. While adjustment opportunities arrive randomly, agents optimally choose when to exercise them to maximize utility. These switching costs represent a variety of non-monetary frictions, such as search costs for finding a new home, amenity-related attachments or the psychological disutility of relocation.

The model successfully reproduces key features of the empirical distribution of durable adjustment sizes observed in the PSID data. It accurately captures the bimodal shape of the distribution and substantial heterogeneity, an empirical feature that is difficult to reconcile with models featuring only proportional transaction costs. In our framework, this heterogeneity arises endogenously from variation in idiosyncratic switching costs. Households facing higher switching costs wait longer and execute larger adjustments when they do move, whereas those with lower costs adjust more frequently and by smaller amounts. This mechanism allows the model to match the observed dispersion in housing mobility without invoking heterogeneity in preference parameters or credit conditions.

Estimation Strategy. The estimation of the switching cost distribution is performed using a novel approach that maximizes the fit between the model's predicted distribution of durable adjustment sizes and its empirical counterpart from the PSID data.

We estimate the adjustment hazard function because it can be mapped to a distribution of adjustment sizes. After an estimated hazard function is obtained, Proposition 2 shows how to recover the fundamental distribution of switching costs. In brief, Proposition 1 presents a novel method by which hazard functions can be mapped into optimal non-durable consumption, portfolio and desired durable value policies. With the full set of policy functions, one obtains the invariant cross-sectional distribution of *desired* adjustment sizes. Finally, combining the adjustment hazard with the distribution of desired adjustments yields the model's prediction for the distribution of actual adjustment sizes.

It is important to note that the only data used in this estimation is the adjustment behavior encoded in the distribution of the size of changes. Our method avoids the strong assumption that unobserved adjustment behavior can be inferred from wealth and income data. Specifically, Berger and Vavra (2015) impute unobservable desired adjustments (called "gaps") using wealth and income data and then match their model's distribution of desired adjustments and hazard function to these imputed moments. Because the desired level of durable consumption is driven by the unobservable present value of financial and human wealth, our method is more robust as it infers frictions from actual adjustment behavior alone.

Results. The estimation results reveal that the model successfully matches the bimodal distribution of durable adjustments observed in the PSID data. By targeting this empirical distribution, the model recovers a distribution of idiosyncratic switching costs that is asymmetric. Relative to a uniform switching cost distribution benchmark, the estimated distribution feature lower costs being more likely for downward adjustments and higher costs being more likely for upward adjustments. When exposed to an aggregate wealth shock, the model's dynamics successfully reproduce the qualitative empirical patterns observed during the Great Recession, showing a significant drop in the frequency of durable adjustments and a disproportionate decrease in the fraction of upward changes. This behavior highlights the endogenous state-dependence of household decisions. The main result of the paper demonstrates that durable demand responses are procyclical and non-linear. Specifically, a further positive wealth shock in a boom state yields a durable demand response that is 60% to 80% larger than a comparable shock in a recession state, as measured by the discounted cumulative impulse response of the first two years after the shock. This cyclicality arises from the interaction between the cross-sectional distribution of households and the estimated, state-dependent adjustment hazards, underscoring the different potency of stabilization policies across the business cycle.

Related Literature. This paper relates to several strands of the literature. It most closely connects to the work on state-dependent effects of economic shocks by Bachmann et al. (2006) and Berger and Vavra (2015), which face a similar challenge: estimating adjustment frictions from data. Pioneering work by Bachmann et al. (2006) calibrate investment frictions to match sectoral level investment responses to identified productivity shocks within a partial equilibrium model of investment dynamics.

On the other hand, Berger and Vavra (2015) use the policy function of a Bewley model to impute unobservable desired adjustment sizes, which they refer to as ajustment "gaps". They then construct hazard functions by measuring the relative frequencies of adjustments within these imputed gap groups. Their estimation strategy involves matching these imputed gap and hazard distributions from the data to the predictions of their model, thereby estimating two parameters of adjustment frictions.

Our paper improves upon their analysis on three fronts. First, we use data on relative adjustment frequencies and sizes to discipline our model, whereas the authors use an adjustment indicator variable to construct imputed hazards.[1] Because our paper focuses on understanding demand responses, we argue that matching the intensive margin of durable demand (sizes of adjustment) is as important as matching the extensive margin (frequency). Second, our estimation approach exploits the fact that adjustment frictions are fully encoded on the adjustment hazard function. This means that the cross-sectional distribution of desired adjustments (”gaps”) depends on adjustment frictions only through the hazard function, which renders the distribution of gaps a redundant moment condition. Third, our estimation is immune to criticisms regarding the soundness of measured wealth data as an approximation of the present value of financial and human wealth.

Our model with adjustment frictions builds on the foundational work on asset pricing and portfolio choice by Grossman and Laroque (1990) and Stokey (2009). These theoretical models have found wide empirical support. Flavin and Nakagawa (2008) study the effects of life-cycle and house price risk on portfolio choice. They find that while a habit persistence model for durable consumption is rejected by PSID data, a model with transaction costs is not. Martin (2003) adds non-durable consumption to a Grossman and Laroque (1990) framework and finds that the model’s prediction for spending behavior on non-durable goods aligns with that in PSID data. The pioneering models of these authors, however, cannot reproduce the variation in the intensive margin of durable demand seen in the data because they feature only fixed transaction costs. This leads to (S,s) adjustment rules with only two possible adjustment sizes: one negative and one positive. To address this limitation, we directly extend Stokey (2009) framework by adding an idiosyncratic utility switching cost to match the rich heterogeneity of adjustment sizes observed in the data. As mentioned before,

[1] In principle, they could contrast imputed desired adjustments with actual adjustments for the sub-sample of agents that adjust. But we do not find such analysis in their paper.

we think that matching this intensive margin is potentially as important as matching the extensive margin for understanding demand responses.

Finally, our paper shares a similar objective with Lippi and Oskolkov (2023) in estimating a distribution of adjustment frictions in a generalized hazard function environment. We extends their result of mapping hazard functions to distributions of desired adjustments to a richer setting where drift and volatility of Brownian shocks is not exogenous but are derived from an optimal control problem. This technical contribution is portable to other environments featuring lumpy behavior and optimal drift and volatility control, and it is clearly indicated by Proposition 1 and Corollary 1.

The remainder of this paper is organized as follows. Section Section 2 lays out the heterogeneous-agent portfolio choice model and derives the key policy functions. Section 3 presents mappings between the model and steady-state data on durable demand. Section 4 describes the data and details our novel estimation strategy. Section 5 presents the model's predictions for procyclical demand responses. Finally, Section 6 offers concluding remarks.

2 Model

This section lays out the individual agent's economic problem under a sequential formulation exploiting the homogeneity of preferences and budget constraints with respect to durable values and arrives to a useful recursive formulation. The model builds on Stokey (2009).

2.1 Preliminaries

Time is continuous and runs from date 0 to infinity. The agent derives utility from non-durable consumption C and durable D, with Cobb-Douglas preferences over the two goods. The agent also enjoys utility from a stochastic process h_t which she considers exogenous to her decision problem. Let me define a flow utility function given by

$$U(C_t, D_t, h_t) \equiv \rho\left[\alpha\log(C_t) + (1-\alpha)\log(D_t)\right] + \frac{1}{2}\frac{1}{\gamma-1}h_t^2, \tag{1}$$

where $\rho > 0$ denotes the discount rate, $\alpha \in (0,1)$ the (desired) expenditure share on non-durables, and $\gamma > 1$ a penalization parameter for the agent's robust control problem.

The agent's only source of income is portfolio returns. She invests in a risk-free asset yielding rate $r > 0$ and allocates a proportion $\theta > 0$ in a risky asset with excess return $r_e > 0$ and volatility $\sigma > 0$. Durable purchases require a down payment share $\epsilon \in (0,1)$, with the remaining $1-\epsilon$ financed at a borrowing rate with credit spread $s > 0$. For simplicity, depreciation and maintenance are assumed to be zero.[2]

Upon adjustment, the agent sells the existing durable of value D and acquires a new durable of value $\hat{D}$, incurring dealer fees fD and receives net sales proceeds $\epsilon\big(D - \hat{D}\big)$. As a result, the jump in financial wealth is $-fD + \epsilon(D - \hat{D})$

Financial wealth evolution. Let W denote financial wealth, the agent's wealth evolves as

$$dW = \big[rW_t + r_e\theta_t W_t - C_t - (1-\epsilon)(r+s)D_t\big]dt + \theta_t W_t \sigma dZ_t + \big[-fD_t + \epsilon(D_t - \hat{D}_t)\big]dN_t, \tag{2}$$

where Z_t is a Wiener process representing idiosyncratic shocks and N_t is an optimally chosen Poisson counter representing adjustment events.

Discussion of assumptions. The innovation component of the risky return (dZ_t) is assumed to be uninsurable idiosyncratic investment risk. Therefore the interpretation of these investment opportunities are those of uncorrelated private businesses. This assumption is in place to ensure meaningful heterogeneity in durable demand (a feature in the data) and

[2] Dropping a maintenance cost is without loss of generality. In this model, the credit spread payment cannot be told apart from a maintenance cost proportional to the value of the durable.

tractability in aggregating individual behavior to macroeconomic variables.

The goal of introducing idiosyncratic risk in the economic environment is to account for persistent swings in the present value of total wealth (permanent income) which triggers a desire to re-balance the value of durable consumption. Therefore an alternative interpretation is that shocks dZ_t constitute idiosyncratic investment risk and proportional shocks to the present value of human wealth.

2.2 Sequential Problem

Adjustment opportunities arrive at random times τ, governed by a Poisson process with exogenous intensity $\kappa > 0$. However, because adjustment incurs costs, the agent may optimally choose not to adjust when given the opportunity. Specifically, at each adjustment date, the agent faces two types of costs: proportional transaction costs fD and a proportional utility switching cost $\exp(-\psi)$ to the continuation value, where ψ is drawn from a distribution F with non-negative support and assumed i.i.d. across agents and over time.

The agent selects a cutoff $\bar{\psi}$ –the maximum switching cost she is willing to "pay"– thereby determining the endogenous jump intensity or adjustment hazard $\lambda_w = \kappa F(\bar{\psi})$. Allowing κ to be arbitrarily large nests a pure fixed cost model as a limiting case.

At each point in time, the agent chooses the non-durable consumption rate C_t, the risky portfolio share θ_t, the desired new durable stock $\hat{D}$ if adjusting, and a adjustment hazard $\lambda_{w,t} = \kappa F(\bar{\psi}_t) \in [0, \kappa]$.

Robustness concerns. Agent's expectation in equation (3) entertains drift distortions to the baseline Wiener process Z_t subject to an entropy penalization included in the flow utility in equation (1). The change of measure amounts to considering that

$$Z_t = h_t\, dt + Z_t^h,$$

where Z_t^h is perceived to be a Wiener process. As a result, she considers an additive distortion of $\sigma\theta W h_t$ to her financial wealth drift and takes this into account when making her portfolio choice. For simplicity, and to focus on the effects of portfolio risk adjustment, the agent considers that jump intensities λ_w are well specified.[3]

In the absence of jumps, this preference specification is equivalent to Kreps and Porteus (1978) and Epstein and Zin (1989) recursive utility with an elasticity of intertemporal substitution (EIS) parameter equal to unity and risk aversion parameter of γ. A sharp characterization of the h_t process is achieved in equation (11) with the problem's recursive formulation below.

Recursive utility equivalence. This paragraph is precise about the connection between recursive utility and a model with robustness concerns in absence of jump processes. It closely follows Hansen and Sargent (2025)

A homogeneous of degree one representation of recursive utility without a durable good is

$$\exp(V_t) = \left[(1 - \exp(-\Delta\rho))C_t^{1-\nu} + \exp(-\Delta\rho)R_t^{1-\nu}\right]^{\frac{1}{1-\nu}}$$

where

$$R_t = \left\{\mathbb{E}_t\left[\exp\left([1-\gamma]\, V_{t+\Delta}\right)\right]\right\}^{\frac{1}{1-\gamma}}.$$

The logarithm of the continuation value process will evolve as

$$dV_t = \mu_t^V\, dt + \sigma_t^V\, dZ_t.$$

Hansen and Sargent (2025) show that in the special case of unit elasticity of intertemporal substitution (as $\nu \to 1$) we obtain

$$0 = -\rho\left(V_t - \log(C_t)\right) + \mu_t^V + \frac{1-\gamma}{2}\left(\sigma_t^V\right)^2.$$

[3]This feature of the model breaks the equivalence with recursive utility explained below but retains its convenient features.

which is equivalent to

$$\rho V_t = \min_{h_t} \left\{ \rho \log(C_t) + \frac{1}{2(1-\gamma)} h_t^2 + \mu_t^V + \sigma_t^V \cdot h_t \right\}.$$

which represents a problem with robustness concerns with drift distortion h_t subject to a quadratic penalization with parameter γ. Therefore recursive utility is equivalent to the standard HJB restriction with a distorted local mean $\mu_t^V + \sigma_t^V \cdot h_t$ and a penalization contribution to flow utility which motivates our definition of flow utility in equation (1). The difference with the specification in equation (10) is that we apply it to the normalized log continuation value function $v \equiv V - \log(D)$ where D is the durable value, that log consumption is replaced by $\alpha \log(C_t)$ where α is the (desired) non-durable expenditure share, and that we do not include a risk adjustment for jump risk. This last feature of our model breaks the equivalence with recursive utility but retains its convenient features of specifying restrictions using the log value function and allowing to disentangle intertemporal substitution to risk aversion, which is key to porftolio choice.

The log value function. As discussed above, working with robustness concerns or recursive utility allows to characterize optimal behavior using the local mean of the *logarithm* of the continuation value V_t.[4] We will henceforth call the log continuation value function simply the log value function. Here we write it in its sequential form for expositional clarity. The agent's log value function satisfies

$$\begin{aligned} V(W, D) = \max_{\{C_t, \theta_t, \bar{\psi}_t\}} E_0 \Big[\int_0^\tau e^{-\rho t} U(C_t, D_t, h_t) dt + e^{-\rho \tau}(1 - F(\bar{\psi}_t)) V(W_\tau, D) \\ + e^{-\rho \tau} F(\bar{\psi}_t) \left\{ \max_{\hat{D}} V[W_\tau - fD + \epsilon\big(D - \hat{D}\big), \hat{D}] - E[\psi | \psi < \bar{\psi}_t] \right\} | W_0 = W \Big], \end{aligned} \tag{3}$$

The first term accumulates flow utility until an adjustment opprtunity arrives. The h_t term entering the flow utility and the underlying perceived wealth evolution is taken as

[4] In this sentence the continuation value is viewed a stochastic process dependent on the state w_t.

exogenous by the agent and solves a utility minimization problem subject to a quadratic entropy penalization detailed in equation (11). The second term is the expected present value of *status quo* where $F(\bar{\psi}_t)$ is the conditional probability of acquiring a new durable. The third term is the expected present value of re-balancing the durable good. The first term in the continuation value includes the choice of durable value $\hat{D}$, the payment of dealer's fees fD and the net proceeds of the transaction $\epsilon(D - \hat{D})$ which are proportional to the agent's equity in the durable ϵ. The second term in the continuation value accounts for the expected switching costs paid given a cutoff policy of $\bar{\psi}_t$. As described above, the switching costs is proportional to the continuation value, $\exp(-\psi)$, and it becomes additive because we characterize the dynamics of the log value function.

Optimal adjustment hazard. The optimal cutoff $\bar{\psi}^*$ equates the switching cost to the benefit of adjustment

$$\bar{\psi}^* = y(W, D) \equiv \max_{\hat{D}} V(\hat{W}(\hat{D}), \hat{D}) - V(W, D), \tag{4}$$

where $\hat{W}(\hat{D}) = W - fD + \epsilon\big(D - \hat{D}\big)$ is the post-adjustment financial wealth and y is the benefit of adjustment. This condition highlights the state-dependence, that agents with greater benefits from adjustment will select higher hazard rates.

Discussion of assumptions. What economic forces do the idiosyncratic switching costs ψ represent? These costs capture non-monetary frictions—such as the cognitive costs of moving, the search cost of finding a new home, or amenity-related attachments—that lead households to prefer their current durable good beyond its market valuation. This assumption is directly motivated by empirical patterns in the PSID data. There, households display substantial heterogeneity in the size of durable adjustments: some transition to slightly larger homes, while others leap to significantly larger ones. This pattern is difficult to reconcile with standard models featuring only fixed transaction costs. In contrast, the present framework accommodates such heterogeneity in a natural and disciplined way. Variation in observed adjustment sizes arises from underlying differences in switching costs across households. Agents facing higher switching cost ψ wait longer to adjust, and as a result, make larger, more discrete changes when the gains outweigh the friction. As a result, this mechanism allows for richer, empirically grounded modeling of durable consumption dynamics.

Why is it important to include robustness concerns? As mentioned above, including a robust control problem allows to work with the logarithm of the value function which transforms multiplicative adjustments, due to homogeneity, to additive adjustments. This feature allows great simplification in the estimation procedure. A constant relative risk aversion (CRRA) utility specification also allows for this transformation when working with a risk aversion parameter of unity, log utility, but it yields extremely high risky asset portfolio shares.

Sequential Problem: Homothetic Formulation

Thanks to the assumption of no idiosyncratic labor income risk, the agent's state space can be reduced to a single dimension: the ratio of financial to durable wealth $w \equiv W/D$. This feature of the model allows for enormous numerical simplification which I will harvest in the estimation procedure.

The value function is homogeneous of degree 1 in the value of the durable D which motivates the normalization $v(w) \equiv V(w,1) = V(W,D) - \log(D)$ where V is the log value function. Let $c \equiv C/D$ represent the ratio of non-durable to durable consumption, the wealth-to-durable ratio w evolves as

$$dw_t = \mu_w(w_t, c_t, \theta_t)dt + \sigma_w(w_t, \theta_t)dZ_t + \left[\hat{w}_t - w_t\right]dN_t, \tag{5}$$

$$\mu_w(w, c, \theta) = rw + r_e\theta w - c - (1-\epsilon)(r+s), \tag{6}$$

$$\sigma_w(w, \theta) = \theta w \sigma, \tag{7}$$

where $\hat{w} \equiv \hat{W}/\hat{D} = (w - f + \epsilon)D/\hat{D} - \epsilon$ by simple manipulation of definitions. Since there exists a one-to-one mapping between $\hat{w}$ and $\hat{D}$, we formulate the adjustment choice in terms of $\hat{w}$. The log value function $V(W, D)$ normalized by $\log(D)$ satisfies

$$v(w) = \max_{\{c_t,\theta_t,\bar{\psi}_t\}} E_0 \Big[\int_0^\tau e^{-\rho t} U(c_t, 1, h_t) dt + e^{-\rho\tau}(1 - F(\bar{\psi}_t)) v(w_\tau) \\ + e^{-\rho\tau} F(\bar{\psi}_t) \left(\max_{\hat{w}} \{ v(\hat{w}) + \log(w - f + \epsilon) - \log(\hat{w} + \epsilon) \} - E[\psi | \psi < \bar{\psi}_t] \right) | \, w_0 = w \Big], \tag{8}$$

where the noticeable difference is the addition of the amount $\log(\hat{D}/D) = \log(w - f + \epsilon) - \log(\hat{w} + \epsilon)$ to the continuation value. This is because $V(\hat{W}, \hat{D})$ is normalized by $\log(D)$ instead of $\log(\hat{D})$. [5]

2.3 Recursive Formulation

The recursive formulation offers a more compact and analytically tractable characterization of the agent's decision problem. Let

$$y(w) \equiv \max_{\hat{w}} \{ v(\hat{w}) + \log(w - f + \epsilon) - \log(\hat{w} + \epsilon) - v(w) \} \tag{9}$$

denote the benefit of adjustment (also defined in equation (4)), then the normalized log value function $v(w)$ satisfies the following Hamilton-Jacobi-Bellman (HJB) equation

$$\rho v(w) = \max_{c,\theta} \min_{h^*} \left\{ U(c, 1, h^*) + [\mu_w(w, c, \theta) + \sigma_w(w, \theta) \, h^*] \, v'(w) + \frac{[\sigma_w(w, \theta)]^2}{2} v''(w) \right\} + H[y(w)], \tag{10}$$

where μ_w, σ_w are defined in equation (6) and equation (7) and U is defined in equation (1). The term h^* is the drift distortion of the underlying Wiener process Z_t which creates a wealth drift distortion of $\sigma_w \, h^*$. The distortion faced by the agent is state dependent and is given

[5] To see why $\log(\hat{D}/D)$ constitutes the correct addition to the continuation value notice that by definition of v, $V(\hat{W}, \hat{D}) = v(\hat{w}) + \log(\hat{D}) = v(\hat{w}) + \log(\hat{D}/D) + \log(D)$.

by

$$h^* \equiv \text{argmin}_h \left\{ U(c(w), 1, h) + \sigma_w \left[w, \theta(w) \right] \, h \, v'(w) \right\}, \tag{11}$$

where $c(w), \theta(w)$ are the (optimal) policy functions. The impulse control term H encodes the value of optimally exercising the option to adjust and is defined as

$$H(y) \equiv \max_{\bar{\psi}} \kappa F(\bar{\psi}) \left[y - E\big(\psi \,|\, \psi < \bar{\psi}\big) \right], \qquad \text{with} \qquad \lambda_w(w) = H'(y(w)) = \kappa F(y(w)).$$

It is useful to note that the impulse control term nests a pure proportional cost model as in Grossman and Laroque (1990) for the case of switching costs ψ equal to zero arriving with probability one, i.e. $F(0) = 1$, and $\kappa \uparrow \infty$.

Looking at the structure of the continuation value upon adjustment in $y(w)$, it is evident that the target ratio $\hat{w}$ is the same across agents. Finally, non-durable consumption and portfolio rules $c(w), \theta(w)$, together with $\lambda_w(w)$, fully characterize the agent's optimal behavior.

2.4 Policy Functions

This subsection highlights the main properties of policy functions which will then be used to map the model to aggregate variables.

Optimality conditions. The recursive formulation allows to characterize the agent's behavior using first order conditions derived from the HJB equation (10). At every wealth ratio

Figure 2: Policy functions

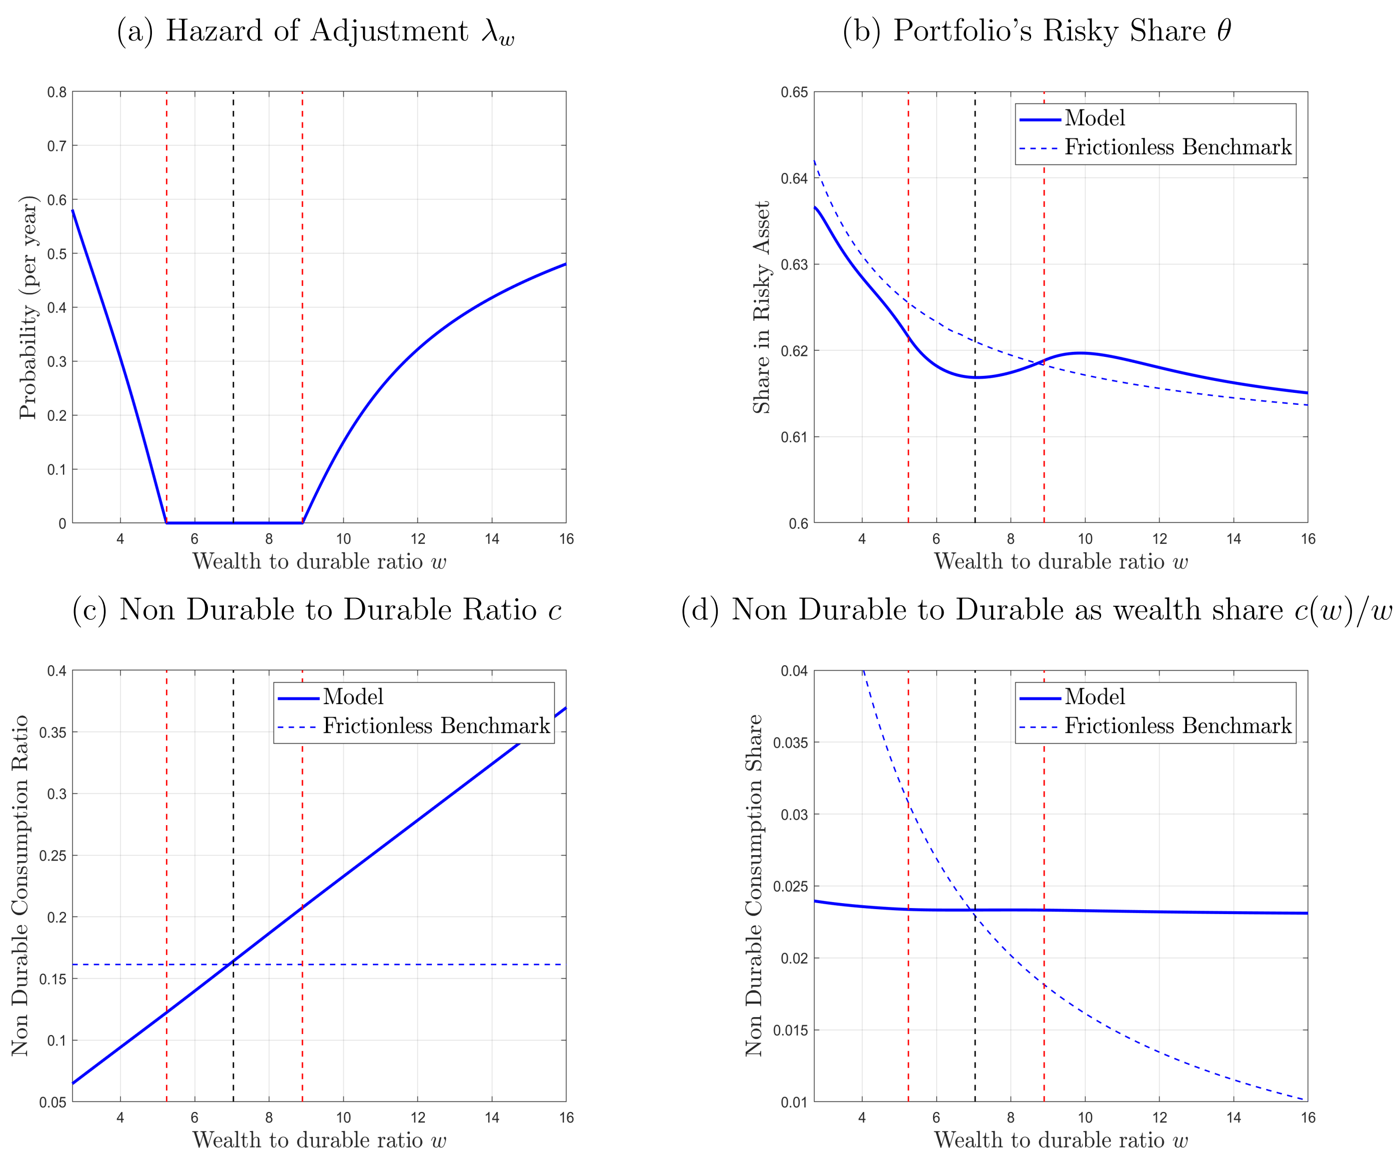

Note: The model corresponds to the benchmark calibration in Table 1

w policies c, θ, λ_w and drift distortion h^* solve

$$c(w) = \rho\, \alpha\, \frac{1}{v'(w)}, \tag{12}$$

$$\theta(w) = \frac{r_e\, v'(w)}{\sigma^2\, w\, \{(\gamma - 1)\, [v'(w)]^2 - v''(w)\}}, \tag{13}$$

$$h^*(w) = (1 - \gamma)\, \sigma\, \theta(w)\, w\, v'(w), \tag{14}$$

$$\lambda_w(w) = \kappa F\, [y(w)] \tag{15}$$

In the estimation procedure described below, I will exploit the fact that policies c, θ and drift distortion h^* only depend on the value function and adjustment friction parameters $\{\kappa, F\}$ through the marginal value function v' (and its first derivative) whereas the adjustment hazard depends directly on these primitive objects. It might appear that $y(w)$ depends on the value function because v directly appears and $\hat{w}$ does appear too. However, these objects depend only on the marginal value function as well. Notice that

$$v(\hat{w}) - v(w) = \int_w^{\hat{w}} v'(w)\, dw, \quad \text{and} \quad \hat{w} = \left\{ w : v'(w) = \frac{1}{w + \epsilon} \right\}. \tag{16}$$

The first expression is straightforward and the second simply states that $\hat{w}$ solves the maximization problem embedded in the value of optimally adjusting, $y(w)$, in equation (9).

Lemma 1. Policies c, θ and drift distortion h^* only depend on the value function and parameters $\{\kappa, F\}$ through the marginal value function v'.

In what follows I use a simple calibration to illustrate the features of household's behavior.

Adjustment hazard. The first panel of Figure 2 displays the endogenous hazard function λ_w, interpreted as the yearly probability of durable adjustment. The agent's policy features a region of pure inaction –bounded by the two dashed red lines– where adjustment is strictly suboptimal. In this region, the utility gains from changing the durable stock are insufficient to outweigh the transaction cost fD, leading the agent to remain inactive. Mathematically, this region corresponds to wealth ratios such that the benefit of adjustment is non positive, $y(w) \leq 0$.

Outside the inaction region, the probability of adjustment becomes positive and increases with the benefits $y(w)$. Notably, even in these regions, adjustment remains probabilistic due to the presence of idiosyncratic switching costs. Upon adjustment, the wealth ratio jumps to a target level $\hat{w}$, shown by the dashed black line, which lies in the interior of the inaction region.

A distinctive feature of the hazard function is its asymmetry: the incentive to downsize after negative financial shocks increases more rapidly than the incentive to upscale following positive shocks. This asymmetry arises from the concavity of log flow utility. Marginal utility declines steeply with increases in non-durable consumption so the relative gain from up-sizing is muted compared to the benefit from downsizing.

Portfolio choice. Turning to the portfolio policy θ on the second panel of Figure 2, the agent selects risky asset shares that are strictly below the frictionless benchmark, depicted by the dashed blue line, when it is located within the inaction region and above when it holds high wealth ratios outside the inaction region. This arises because, although the agent faces frictions in adjusting durables, she retains full flexibility over financial portfolio choices. In the frictionless model, relative risk aversion is constant and the risky share decreases because the investor maintains a constant risky share over *total* wealth including the equity portion of illiquid durable wealth D. In our model, relative risk aversion becomes endogenous and state-dependent: it is lower near the boundaries of the inaction region, where the agent is close to adjusting, and highest at the target wealth $\hat{w}$.[6]

A crucial question in the household finance literature is how adjustment frictions alter portfolio choice. Flavin and Nakagawa (2008) analyze a model similar to ours, but without switching costs, to study the effects of life-cycle and house price risk on portfolio choice. They

[6] Note that what matters is the expected time until an adjustment and not the adjustment hazard *per se*. The adjustment hazard is constant and equal to zero within the pure inaction region.

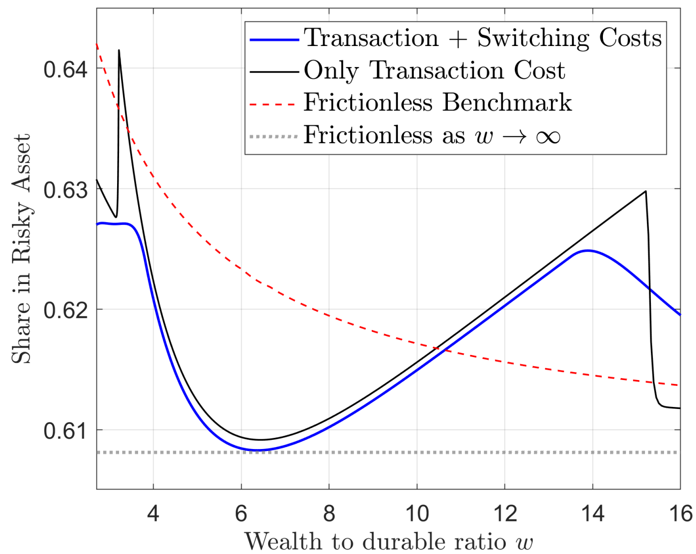

Figure 3: Portfolio choice for a proportional transaction cost of 5%

find that while a habit persistence model for durable consumption is rejected by PSID data, a model with transaction costs is not. Given this, it is reassuring that the portfolio choice predictions of our model do not significantly depart from those of the standard transaction cost model. Figure 3 provides such a robustness check, using a typical calibration from the literature with a transaction cost of 5% of the house value.

Non-durable consumption. Figure 2's third and fourth panel present the non-durable consumption ratio and as a share of the wealth ratio. In the absence of adjustment frictions, the agent would continuously rebalance the durable stock to maintain a constant expenditure ratio $c^* = \alpha(r+s)/(1-\alpha)$, depicted by the dashed blue line.[7] Note that such value coincides with the policy at the target wealth ratio $\hat{w}$.

With frictions, however, the agent tolerates deviations from this benchmark. In the third panel we can observe that for low values of w, non-durable consumption falls below the frictionless level due to lower permanent income, the constant mortgage expenditures and the inability to adjust the durable stock. Interestingly, in the fourth panel we can see that the optimal non-durable consumption policy is close to a fixed share out of wealth ration w.

[7] For the result on the frictionless problem see equation 6 in Stokey (2009)

This is driven by the choice of EIS parameter equal to 1.

3 Mapping the Model to the Steady-State Data

The model generates sharp predictions about household durable adjustment behavior, which can be contrasted empirically. This section derives the model's implications for the frequency and size distribution of durable adjustments in a stationary environment i.e. absent aggregate shocks.

Aggregation and Stationary Distribution. A key object for linking the model to the data is the stationary distribution of wealth-to-durable ratios $m_w(w)$ which satisfies a stationary Kolmogorov Forward (KF) equation

$$\lambda_w(w)m_w(w) + \frac{d}{dw}\big(\mu_w[w, c(w), \theta(w)]\, m(w)\big) = \frac{1}{2}\frac{d^2}{dw^2}\big(\sigma_w[w, \theta(w)]m_w(w)\big), \qquad w \neq \hat{w}, \tag{17}$$

with boundary conditions $\lim_{w\downarrow 0} m_w(w) = 0$, $\lim_{w\uparrow\infty} m_w(w) = 0$, $\lim_{w\uparrow\hat{w}} m_w(w) = \lim_{w\downarrow\hat{w}} m_w(w)$ and $\int_0^\infty m_w dw = 1$. The equation balances inflows and outflows of probability density at each point in the wealth ratios space. The left-hand side represents exits from state w due to endogenous adjustment (first term) and drift (second term), while the right-hand side captures arrivals due to stochastic portfolio shocks.

A critical modeling choice is the interpretation of the stochastic process dZ_t in equation (5). To ensure meaningful heterogeneity across agents, I treat the risky asset as a purely idiosyncratic investment—such as ownership in a private firm or entrepreneurial venture. See discussion in Section 2.1.

From States to Observables: Frequency and Size Distribution. To connect the state variable w with observed durable adjustment behavior, we define the log size of the desired durable change as

$$\Delta d = \log(\hat{D}/D) = \log(w - f + \epsilon) - \log(\hat{w} + \epsilon) \equiv g(w),$$

where the map g is monotonic and invertible. In other words, because the target wealth ratio $\hat{w}$ is constant across agents, each w maps uniquely to a desired adjustment size Δd.

Using this bijection, we transform the adjustment hazard and stationary distribution from $w-$space to $\Delta d-$ space

$$\lambda = \lambda_w \circ g^{-1}, \qquad m = m_w \circ g^{-1}. \tag{18}$$

Figure 4

(a) Hazard of adjustment λ (b) Distribution of desired adjustments m

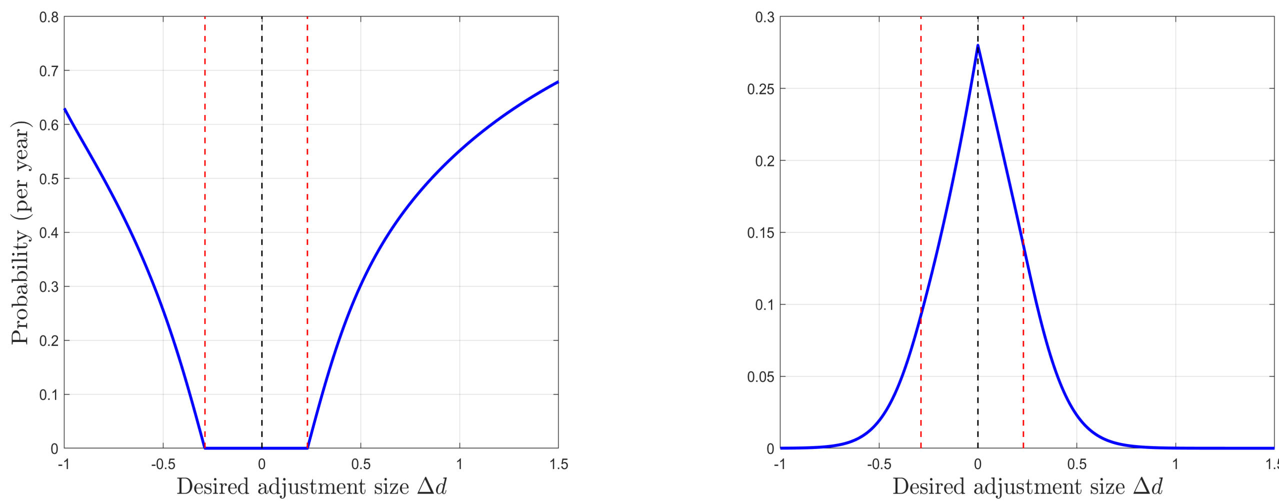

Note: The model corresponds to the benchmark calibration in Table 1

Figure 4 shows the hazard and cross-sectional distributions of desired adjustment sizes. In the first panel, the hazard features a pure inaction region showing that small adjustments are sub-optimal because of the proportional transaction cost f. Importantly, the inaction barriers are asymmetric. Before taking action, the agent waits to have a larger negative desired adjustment size than it does for a positive adjustment. This is because of the concavity of utility and the relative magnitude of the transaction cost f for a low wealth and a high wealth agent. This feature is also depicted in the second panel where it is clear that the right boundary crosses the distribution at a higher density than the left boundary which highlights that it is further from the target $\Delta d = 0$.

The (total) frequency of durable adjustment is then given by

$$N = \int_{-\infty}^{\infty} \lambda(\Delta d)\, m(\Delta d)\, \mathrm{d}\Delta d, \tag{19}$$

which aggregates across the population by weighting each agent's hazard rate by the density of agents with that desired adjustment. We similarly define upward and downward adjustment frequencies as

$$N_{+} = \int_{0}^{\infty} \lambda(\Delta d)\, m(\Delta d)\, \mathrm{d}\Delta d, \quad N_{-} = \int_{-\infty}^{0} \lambda(\Delta d)\, m(\Delta d)\, \mathrm{d}\Delta d. \tag{20}$$

Finally, the implied distribution of adjustment sizes among those who actually adjust is given by

$$q(\Delta d) = \frac{\lambda(\Delta d) m(\Delta d)}{N} \tag{21}$$

which we compare to empirical data in Figure 5.

4 Data and Estimation

This section describes the data, a benchmark calibration and the estimation algorithm employed to uncover primitive adjustment frictions.

4.1 Data Description

The empirical analysis relies on data from the Panel Study of Income Dynamics **?**, with particular focus on housing and mobility behavior reported biennially. The primary variable of interest is the self-reported value of owner-occupied housing.

Following Berger and Vavra (2015), I restrict the sample to homeowners whose household

Figure 5: Distribution of adjustment sizes in PSID Data and Model

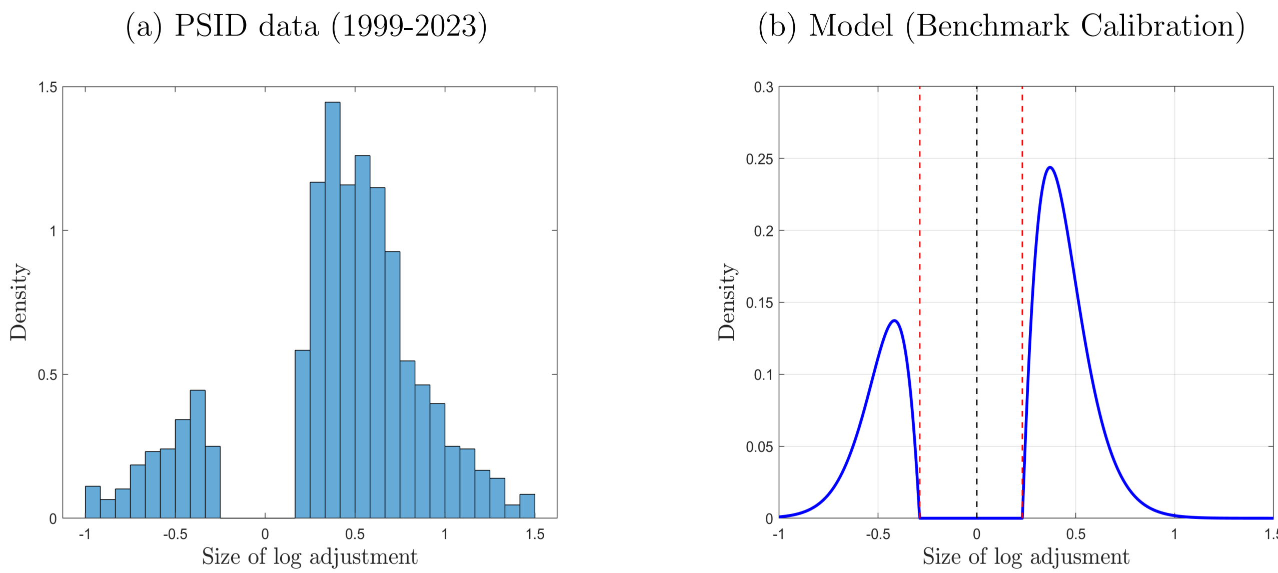

Note: The PSID cleaning procedure is explained in the main text. See model calibration in Table 1

heads are aged 65 or younger. Durable adjustment events are identified using a stringent definition to ensure that observed changes reflect true behavioral responses rather than passive asset revaluation or measurement error. Specifically, a household is classified as adjusting its durable stock if:

1. It reports a residential move relative to the previous PSID wave,
2. It reports having sold its prior home, and
3. The reported change in home value is between 0.20 and 1.5 log points for increases and 0.3 and 1 log points for decreases.

To isolate active up-sizing or down-sizing from mechanical price movements, all housing values are normalized using the National Income and Product Accounts (NIPA) housing services price index. I follow this procedure exactly as described in Berger and Vavra (2015) with the exception of using an asymmetric threshold to classify valid adjustment. I employ a larger threshold for negative adjustments because such behavior is inherent of the economics of inaction due to transaction costs and it is a featured shared with benchmark models as Stokey (2009) and Grossman and Laroque (1990).[8]

The first panel of Figure 5 displays the empirical distribution of housing adjustment sizes (log changes in normalized value) observed between 1999 and 2023. The distribution is markedly heterogeneous and exhibits a bimodal structure. The heterogeneity in adjustment sizes is at odds with a fixed transaction cost model. In particular, the presence of both small and large adjustments suggests a role for unobserved heterogeneity in adjustment costs.[9] The distribution is bimodal which is a hallmark of the fixed transaction costs in the housing market, small house adjustments are not worth the monetary moving costs.

The second panel of Figure 5 displays the distribution of adjustment sizes generated by the structural model. The model captures the broad heterogeneity in observed adjustment behavior. Notably, the model reproduces the empirical bimodality, which emerges naturally from the idiosyncratic switching cost mechanism embedded in the hazard function.

4.2 Benchmark Calibration

To illustrate the model's mechanisms, and to give the estimation algorithm an initial guess for the distribution of switching costs F and the arrival of adjustment opportunities κ, I adopt a benchmark calibration based on Stokey (2009), with a modification to the transaction cost f which I take to be 0.5% instead of Stokey's 8%. The key reason for this large difference in transaction costs is that we match different targets. I match positive adjustments starting at 20% (to make my study comparable to Berger and Vavra (2015)) whereas the Stokey (2009) matches the average housing tenure.

At this point, I allow for the possibility that arrival rates and switching costs are inherently different for home value increases and decreases. Denote κ_u, F_u the arrival rate and

[8] In Stokey (2009), the author estimates boundaries $\underline{\Delta d} = -0.71$ and $\overline{\Delta d} = 0.46$, see the second column of their Table 2.

[9] In general, some unobserved heterogeneity (UH) is necessary to match this empirical fact. In this paper I model this UH as idiosyncratic switching costs. Further research should explore whether heterogeneity in risk aversion, access to different financial instruments or other sources of differences across agents has meaningful potential to confront this pattern.

distribution of switching cost to up-size and κ_d, F_d the analogous to down-size.

Table 1 summarizes the full set of baseline parameters. Of note, the arrival rates of adjustment opportunities is set to once per day. This high arrival frequency ensures that any inaction observed in the model is not driven by infrequent opportunities, but rather by endogenous optimal inaction due to switching costs.

Welfare. Regarding the level of switching costs, they are more easily understood using a money metric. Here I use a compensating variation metric which asks *what lump-sum transfer would an agent require to be as well off as in a world without the switching costs?*. An agent in the 25 percentile of the wealth-to-durable w distribution would require a transfer of 2% of her home value whereas one at the 75 percentile would require a 3% transfer. As a tail measure, the agent in the 95 percentile of the wealth-to-durable distribution would require a 4% transfer. This amounts to 8,000 USD, 12,000 USD and 16,000 USD respectively assuming a median value of a home of 400,000 in 2025 USD. Are these welfare costs big or small? A point of reference is the welfare cost stemming from proportional transaction costs f. In my model a transaction cost of $f = 0.50\%$ commands a mean compensating variation of 16,000 USD (4% of home value). Therefore the magnitude of switching costs is smaller than the one transaction costs and nonetheless allows to match the rich heterogeneity in adjustment sizes in the data.

For comparison, Stokey (2009) obtains a 3.75% of home value welfare cost for a $f = 0.15\%$ transaction cost.[10] It is important to note that our models are not nested, Stokey (2009) uses a CRRA utility framework whereas I use a case of recursive utility with EIS equal to unity.

4.3 Estimation

In this subsection I will lay out a framework to estimate the switching cost primitives $\{\kappa_i, F_i\}_{i=\{u,d\}}$ by minimizing the distance between the model's distribution of adjustments to

[10] In particular, she obtains a 0.5% cost out of total wealth and total wealth is 7.5 times the value of the home.

Table 1: Benchmark Calibration Parameters

Parameter	Value	Description
ρ	2.6%	Discount rate
γ	3.6	Risk aversion or entropy penalization
α	0.88	Expenditure share in non-durables (desired)
r	1%	Risk-free interest rate
r^e	6%	Risk premium
σ	16.55%	Idiosyncratic risk
s	1.2%	Mortgage rate spread
ϵ	15%	Home equity requirement
f	0.5%	Transaction cost
$\kappa_u = \kappa_d$	1/90	Arrival of adjustment opportunities (1 per day)
Ψ_u	0.01	Upper bound of random utility cost: $F_u \sim U([0, \Psi_u])$
Ψ_d	0.03	Upper bound of random utility cost: $F_d \sim U([0, \Psi_d])$

the one measured in PSID data. The method is inspired by Lippi and Oskolkov (2023) but contains novel insights that are portable to other environments.

4.3.1 Marginal Valuation and an Inverse Mapping

One alternative for estimation is to parameterize the distribution of switching costs F_i by choosing a finite support and probabilities, then computing policy functions and obtaining the distribution of desired adjustments $m(\Delta d)$ and the hazard function $\lambda(\Delta d)$. The product of these yields the model prediction for the distribution of adjustments $q(\Delta d)$ to be contrasted with data. This method can prove complex as the level of the value function changes with each iteration and the choice of the support must be updated as well as probabilities.

A second alternative is to directly choose a hazard function and derive the distribution of desired adjustments, obtaining adjustments distribution q. This is the approach taken in Lippi and Oskolkov (2023). However, their method is not directly applicable. In their case, drift and exposure (to the Wiener process) are constant parameters which allows for a straightforward mapping between the hazard function and the cross-sectional distribution of desired adjustments. In this paper, drift and risk exposure are endogenous policies. Can we find a similar map between the adjustment hazard function λ and the distribution of desired adjustments m?

The answer is positive and it is possible because of a sharp link between the hazard function and the marginal (log) value function. Hansen and Souganidis (2025) show that marginal valuation can be decomposed into contributions due to (1) flow utility, (2) jump intensity and (3) the marginal value upon a jump. In this paper, only the third channel affects marginal valuation because non-durable consumption and hazards (jump intensities) are chosen optimally and thus vanish due to the Envelop Theorem. This is summarized in the next proposition.

PROPOSITION 1. **Marginal Valuation.** The marginal value function $\varphi \equiv v'$ satisfies

$$[\rho - r + \lambda_w(w)]\,\varphi(w) = \frac{\lambda_w(w)}{w - f + \epsilon} + [\mu_w(w,c,\theta) + \sigma_w(w,\theta)\,h^*]\,\varphi'(w) + \frac{[\sigma_w(w,\theta)]^2}{2}\varphi''(w), \tag{22}$$

where $c, \theta, \lambda_w, h^*$ satisfy equations (12) to (15).

Proof. Differentiating equation (10) with respect to w, I use the Envelope Theorem and only differentiate w where it appears directly. The impulse control term yields

$$H'(y(w))\,y'(w) = \lambda_w(w)\left[\frac{1}{w - f + \epsilon} - \varphi(w)\right]$$

where I have used the optimality of the hazard and the fact that terms involving the optimal target $\hat{w}$ are additive and vanish through differentiation. The derivatives of the drift and exposure term yield

$$(r + r_e\theta + \sigma\theta h^*)\,\varphi(w) + (\theta\sigma)^2 w\varphi(w) = r\varphi(w)$$

where I have used the first order conditions of portfolio choice θ. The remaining terms are immediate. ■

Proposition 1 tells us that the marginal value function shares the same perceived drift and exposure to small shocks than the value function v but it does not feature jump dynamics. Instead it includes a flow term $\lambda_w(w)/(w - f + \epsilon)$ contributed by the derivative of the continuation value and a discount rate adjustment $\lambda_w(w) - r$ contributed by the derivative of the continuation value and the derivative of the drift and exposure terms. From a Feynman-Kac or asset pricing point of view, the marginal value function is the present value of flows $\lambda_w(w)/(w - f + \epsilon)$ with discount rate $\rho - r$ and (state dependent) death rate $\lambda_w(w)$. Formally,

$$\varphi(w) = \mathbb{E}\left\{\exp\left[-(\rho - r)t - \int_0^t \lambda_w(w_t)\right] \frac{\lambda_w(w_t)}{w_t - f + \epsilon} \mid w(0) = w\right\}$$

where $\mathbb{E}$ integrates over a perceived wealth evolution that does not feature jumps

$$dw_t = \left[\mu_w(w_t, c_t, \theta_t) + \sigma_w(w_t, \theta_t) h_t^*\right] dt + \sigma_w(w_t, \theta_t) dZ_t.$$

COROLLARY 1. Given λ_w, the marginal value function φ is the solution to a non-linear ordinary differential equation (ODE).

Proof. Using Lemma 1 we can write c, θ, h^* and therefore μ_w and σ_w as functions of φ. As a result equation (22) is a non-linear ODE with inhomogeneous term and adjusted discount rate governed by λ_w. ■

Corollary 1 provides a powerful method to obtain consumption and portfolio policies consistent with the hazard function and optimality. Additionally, because these policies are all that is required to obtain the distribution m_w of wealth ratios, we obtain a mapping between the hazard function and the cross-sectional distribution of wealth m_w. The final step to map hazards and distribution λ_w, m_w to the distribution of adjustments is the target wealth level $\hat{w}$. This variable is also obtained from the marginal value function and satisfies equation (16) reproduced here for convenience

$$\varphi(\hat{w}) = \frac{1}{\hat{w} + \epsilon}.$$

This step completes the mapping from the hazard function λ_w to the distribution of adjustment sizes q. Estimation of λ_w is then straightforward. The process is summarized here for convenience:

1. Guess λ_w
2. Solve for φ using equation (22)
3. Obtain policies c, θ using marginal value φ
4. Compute m_w using equation (17) and policies λ_w, c, θ
5. Using equation (16), obtain $\hat{w}$ and transformed hazard function and distribution λ, m
6. Compute distribution of adjustment sizes q using λ, m in equation (21)
7. Compare q with q_{data}. If they are close enough, stop, else continue with 1.

The fact that the marginal value function only depends on the optimization problem through the hazard function λ_w stems from the fact that the continuation value $v(\hat{w})$ vanishes from marginal valuation because the multiplicative scaling due to homogeneity $\hat{D}/D$ in equation (8) becomes additive from working with the log value function in a framework with robustness concerns. Such a simplification is impossible in a CRRA utility framework with risk aversion parameter $\gamma > 1$.

Recover primitives κ, F. After having obtained an estimate for the hazard function how do we recover the primitive adjustment frictions $\{\kappa_i, F_i\}_{i=\{u,d\}}$ where we allow for different arrival rates and switching cost distributions for upward and downward adjustments. The method is summarized in the following proposition

Proposition 2. **Recovering adjustment frictions.** Given marginal value function $\{\varphi_j\}$ and hazard function $\{\lambda_{w,j}\}$ represented as vectors of length N, the primitives $\{\kappa_i, \{F_{i,j}, \psi_{i,j}\}\}_{i=\{u,d\}}$ satisfy

$$\kappa_u = \lambda_N, \qquad \kappa_d = \lambda_1, \tag{23}$$

$$F_{i,j} = \frac{\lambda_{w,j}}{\kappa_i}, \qquad \psi_{i,j} = y_j, \quad i \in \{u, d\}, \tag{24}$$

$$y_j = \int_{w_j}^{w_{\hat{j}}} \varphi_j dw_j + \log(w_j - f + \epsilon) - \log(w_{\hat{j}} + \epsilon) \tag{25}$$

where $w_{\hat{j}}$ is the last wealth ratio such that $\varphi_j > (w_j + \epsilon)^{-1}$.

The proof is analogous to the one in Lippi and Oskolkov (2023). Proposition 2 shows how to recover parameters for the arrival rate of adjustment opportunities $\{\kappa_i\}$, the support of the switching cost distribution $\{\psi_{i,j}\}$ and the cumulative density function $\{F_{i,j}\}$ for upward and downward adjustments.

Results. The hazard function is estimated minimizing the relative entropy of the data and model distributions of adjustment sizes. Formally,

$$\lambda_w \equiv \text{argmin}\ \mathbb{E}\left[\log\left(\frac{q(\Delta d; \lambda_w)}{q_{\text{data}}(\Delta d)}\right)\frac{q(\Delta d; \lambda_w)}{q_{\text{data}}(\Delta d)}\right] \tag{26}$$

where the expectation is with respect to the data distribution and some smoothing is applied to eliminate the case of dividing by zero.

Figure 6 shows the estimation results. The left panel depicts the distribution of adjustment sizes in PSID data and the predicted distribution from two models: the benchmark calibration of Table 1 and the estimated model. Notably, the estimation matches well the right fat tail that the benchmark calibration misses. This is due to the shape of the estimated hazard in the right panel. The estimated hazard is flatter and lower for large wealth ratios which explains the thickness of the right tail.

Figure 6: Estimation Results: Distribution of Adjustment Sizes and Hazard

(a) Distribution of Adjustment Sizes

(b) Hazard Function λ

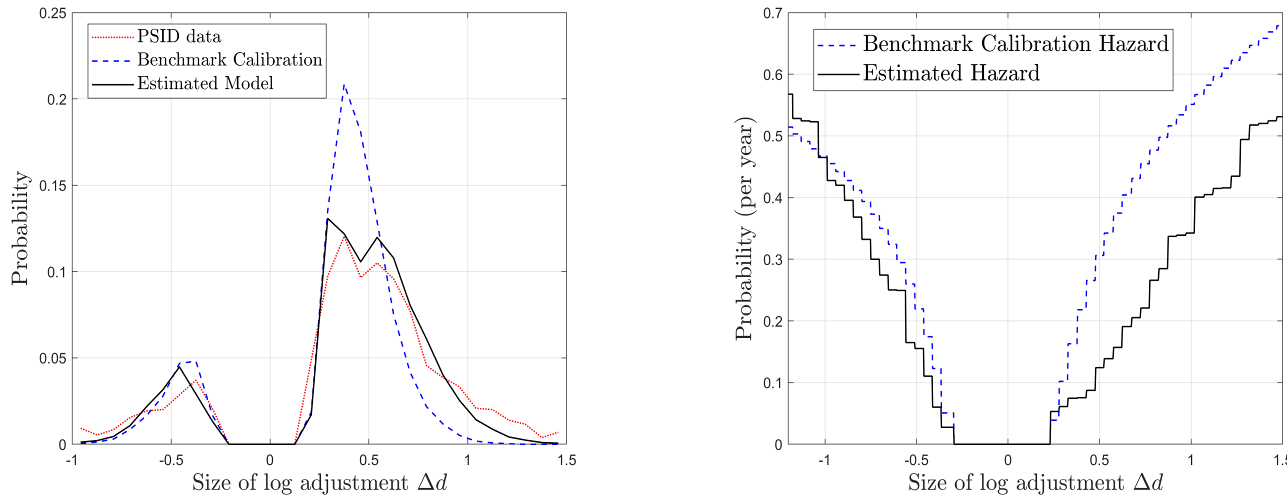

Note: See benchmark model calibration in Table 1. The estimated hazard is obtained following the algorithm detailed in the main text.

Recovering primitives. Using Proposition 2 we can recover the distribution of switching costs and the arrival rate of adjustment opportunities. The left panel of Figure 7 shows the estimated distribution of switching costs for upward adjustments and the uniform distribution of the benchmark calibration. Our estimation result shows that low switching costs have a positive but low probability. Consequently, agents endogenously wait for large upward desired adjustments, which are more valuable, in the face of large switching costs. This feature is borne out in the data since the right tail of upward adjustment is thick, see the first panel of Figure 5. The right panel of Figure 7 depicts the distribution of switching costs for downward adjustments. In this case, the estimates indicate that small switching costs are more likely which shows up in the data as a spike at adjustments of size -0.5 and smaller and a thinner left tail.

The estimated opportunity to adjust up arrives in a year with a probability of 0.98 and the opportunity to adjust down with a 0.58 probability. These figures stem from $\kappa_u = 3.68$

and $\kappa_d = 0.88$.[11] It is important to be careful with results regarding κ_i. The aim in this paper is to estimate the relative magnitude κ_u/κ_d which informs the relative mass of up and downward adjustment. To have a meaningful discussion about their absolute magnitude it is important to match the total frequency of adjustments in the data since that observable statistic directly connects to the arrival rate of opportunities to adjust through equation (19). Any attempt to match the low frequency in the data with a size of return shocks as large as $\sigma = 0.16$ requires a very wide inaction region (Stokey, 2009) and therefore needs to ignore small adjustments sizes as the one used in this study and the literature (Berger and Vavra, 2015). Because we want to relate as close as possible to such empirical literature, I leave the attempt to reconcile the size of return shocks σ with the frequency of adjustments to future research.

Figure 7: Estimation Results: Switching Cost Distributions

(a) CDF F_u of upward adjustment (b) CDF F_d of downward adjustment

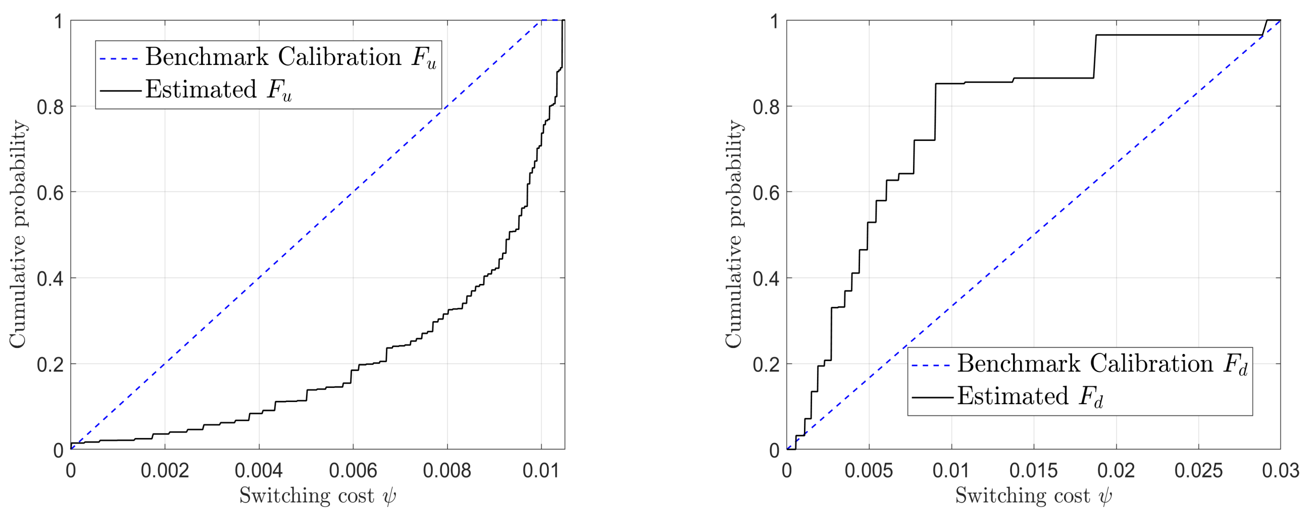

Note: See benchmark model calibration in Table 1. The estimated hazard is obtained following the algorithm detailed in the main text.

Untargeted moments. The model has several predictions for aggregate behavior which were not targeted. These are the mean wealth to durable ratio, the mean wealth growth, the

[11]To convert Poisson arrival rates to probabilities we use the Poisson CDF $1 - \exp(-\kappa)$.

proportion of upward adjustments and the frequency of adjustments. Additionally, we show the welfare costs entailed by the switching costs.

Table 2 presents the aforementioned moment predictions and shows the good fit of the model with the exception of the frequency of adjustments. As mentioned in the previous paragraph there is an economic tension between the size of returns risk σ and the width of the inaction region. A wider inaction region helps in matching the low frequency of adjustments but leads to ignoring smaller adjustment sizes of the type studied in this paper and the related literature. Further research to reconcile these facts is necessary.

Turning to the welfare consequences of switching costs, we find that the compensating variation is smaller for low wealth agents, relative to the benchmark calibration, and larger for high wealth agents. This is consistent with the estimated distribution of costs where low costs are more likely for downward adjustments and less likely for upward adjustments.

Table 2: Estimated Model Predictions for Untargeted Moments and Welfare Costs

Variable	**Value**	**Reference and Source**
Mean wealth	8.2	7.5 Stokey (2009)
Mean wealth growth	2.1%	2% Stokey (2009)
Proportion of upward adjustments	0.84	0.84 PSID data
Frequency of adjustments	0.040	0.015 PSID data
Welfare cost of switching costs p25	7,700 USD	–
p50	9,700 USD	–
p75	12,500 USD	–
p95	37,000 USD	–

5 Responses to Aggregate Shocks

In Figure 1, we observed that the frequency of housing adjustments declined and maintained a low level after the great recession. We also note, by inspecting the right panel, that this decline saw a disproportionate decrease in the frequency of upward adjustments. In this section, we use the estimated model to rationalize this observation and, at the same time, obtain predictions for the relative magnitude of demand responses in booms and recessions.

5.1 Frequency and Fraction of Upward Adjustments.

I consider an economy at steady-state and study the dynamics that follow after an aggregate once-and-for-all proportional shock to agents' wealth. In particular, I study the response of the frequency of housing adjustment and the fraction of upward adjustments. For illustration, and to relate to the literature on state-dependent economic responses, I will call a *boom* state one where every agent receives a proportional 15% increase in their wealth-to-durable ratio and a *recession* state one where such a shock is a negative 15%. These magnitudes correspond to a 25% fluctuation in the risky asset return together with a 60% share of wealth on the risky asset.

Figure 8: Distribution of Wealth After an Aggregate Shock

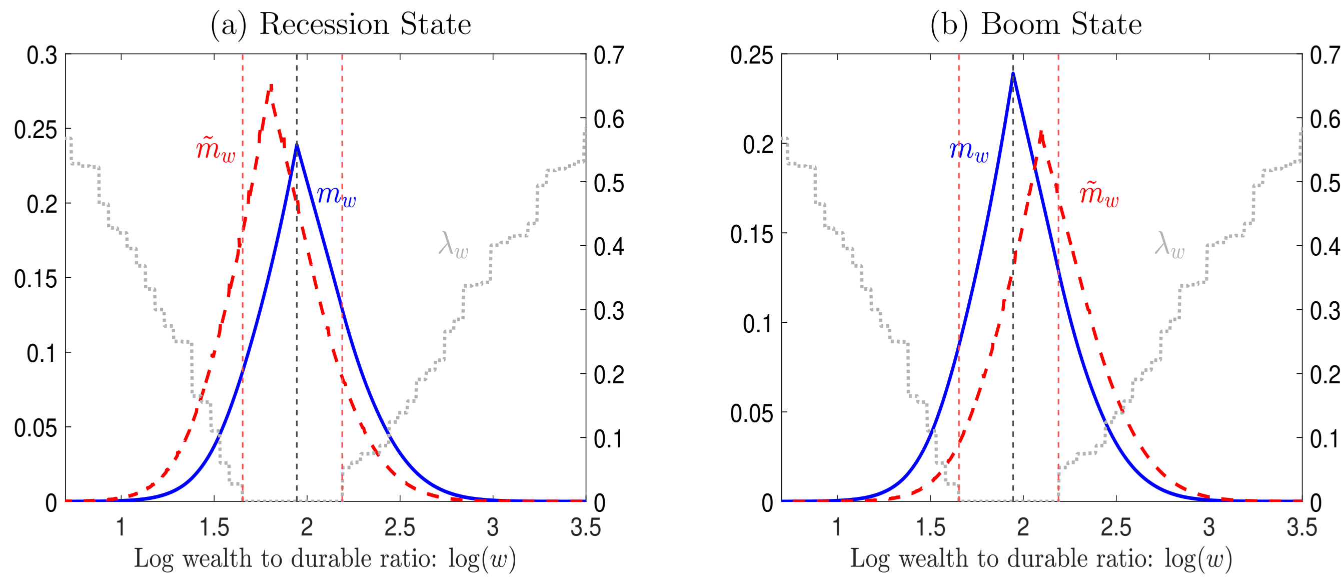

Note: Density is in the left scale and yearly probability of adjustment is in the right scale. This graph uses the estimated model. For visual convenience I depict the distribution as a function of log wealth so that the aggregate shock amounts to a horizontal displacement. The apparent change in the level of the density is not a mistake but a feature of using uneven grids for plotting and numerical methods.

Figure 8 depicts these two states. In the left panel, the red dashed line indicates the distribution of wealth ratios in a recession state whereas the thick blue line shows the steady state distribution of wealth. We can observe that the recession state features a larger mass of agents within the pure inaction region and an increased share of agents with negative desired adjustments outside the inaction region. Therefore, it is not obvious that the model obtains an absolute decrease in the frequency. Conversely, in the boom state, a group of agents that were inside the inaction region have been moved toward positive desired adjustments.

For simplicity, we develop the analysis of the frequency and proportion of upward adjustments in a partial equilibrium setup assuming that firms use the steady-state policy λ_w.[12] As a consequence, the model dynamics lies in the distributional dynamics $\tilde{m}(t, w)$, prompted by starting from a distribution different from the steady state, and its interaction with steady state policy λ_w. Formally, the distribution of wealth ratios $\tilde{m}(t, w)$ solves a partial differential equation which is a time varying version of the Kolmogorov Forward equation (17) with initial condition $\tilde{m}(0, w) = m_w [w \cdot \exp(\pm 0.15)]$ where m_w is the invariant distribution (see Cavallo et al. (2024)).

Figure 9: Frequency and Fraction of Upward Adjustments

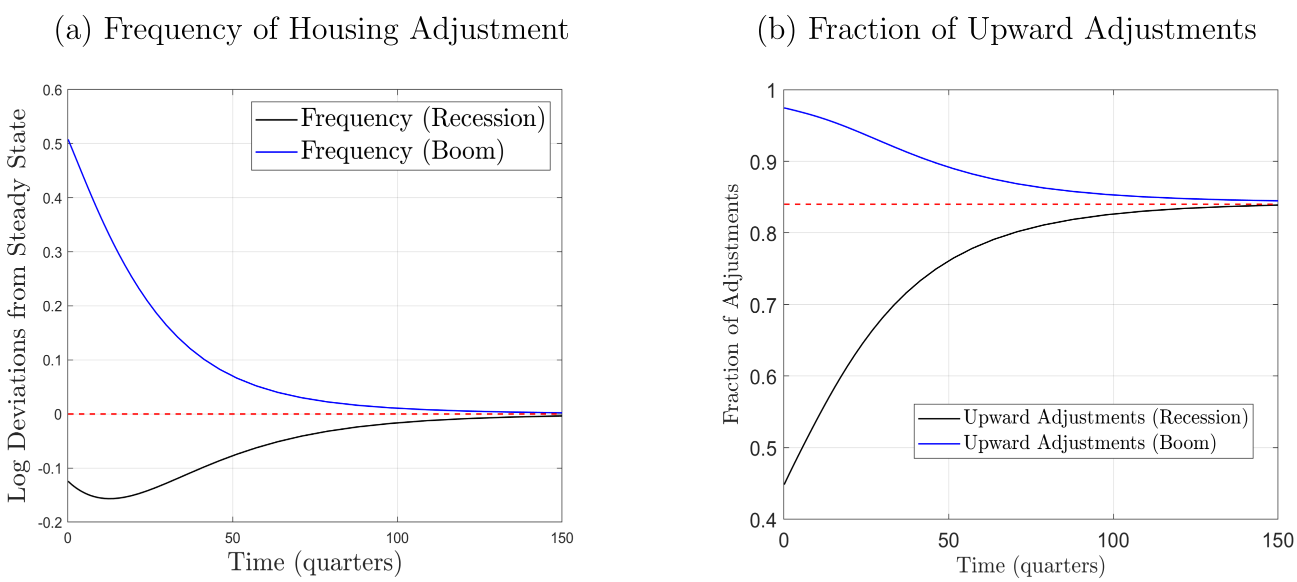

Note: The model uses the estimated hazard function.

[12]This analysis is close in spirit to Cavallo, Lippi, and Miyahara (2024) which applies it to a price setting setup.

Using the steady state policy λ_w and the distribution $\tilde{m}_w(t, w)$ we are able to compute the response of the frequency and the proportion of upward adjustments following equation (19) and equation (20), Figure 9 plots these responses. Importantly, both figures replicate the qualitative patterns of Figure 1 in PSID data generating a decrease in the frequency with a disproportionate decrease in the frequency of upward adjustments after the great recession. Quantitatively, the drop in frequency is milder in the model and the drop in the fraction of upward adjustments is stronger.

As mentioned, the frequency response in the first panel of Figure 9 features two counteracting forces. First the aggregate shock locates many agents within the pure inaction region, decreasing the frequency of adjustments. Second, the aggregate shock prompts some agents that were inactive to execute a downward adjustment, increasing the frequency of adjustments. The fact that the aggregate response is negative indicates that the first force dominates. As a robustness check, Figure 11 of Appendix A shows that an array of recession shock sizes generate a decrease in the frequency of adjustments. Interestingly, the effect is non linear and the largest negative impact response is achieved at a recessionary shock of size 10% where the inaction channel is strongest.

The second panel, instead indicates that most action in steady state comes from positive adjustments. Specifically, 0.84 of adjustments correspond to upward modifications in house values. However, the aggregate shock markedly tilts the composition of adjustments to an almost equal fraction at the onset of the shock. Again, this is because agents that were wishing to adjust upward now find themselves in the inaction region and others that were inactive now have large incentives to downsize. This observation hints that housing demand responses to fiscal stimulus during recessions would require large stimulus packages since they must affect the incentives of agents that currently want to downsize to up-sizing and the relatively wide inaction region dampens much of the policy impulse.

5.2 Cyclical Demand Responses

We now turn to the central question of the paper. Does durable demand responds more or less strongly to further economic stimulus during recessions and booms? This question is of substantial economic interest since it sheds light into the trade offs faced by policy makers when considering the benefits and costs of stabilization policies such as fiscal stimulus checks.

We consider an economy in the recession state and ask what is the additional durable demand response prompted by a further positive wealth shock of size $\delta > 0$. That is we compare two dynamical systems, the recession dynamics characterized by state distribution $\tilde{m}_r(t, w)$ with initial condition $\tilde{m}_r(0, w) = m_w(w \cdot \exp(0.15))$ (where m_w is the invariant distribution) and the shocked recession system $\tilde{m}_{r,\delta}(t, w)$ with initial condition $\tilde{m}_{r,\delta}(0, w) = \tilde{m}_r(0, w \cdot \exp(-\delta))$.

Definition 1. **Impulse Response Function.** We define the durable demand impulse response function (IRF) at horizon t of a wealth shock of size $\delta > 0$ to a state $s \in \{\text{recession}, \text{boom}\}$ as

$$IRF_s(t; \delta) \equiv \int_{-\infty}^{\infty} \Delta d \, \lambda(\Delta d) \left[\tilde{m}_{s,\delta}(t, g^{-1}(\Delta d)) - \tilde{m}_s(t, g^{-1}(\Delta d))\right] \mathrm{d}\Delta d \tag{27}$$

where $g(w) = \log(w - f + \epsilon) - \log(\hat{w} + \epsilon)$ is the invertible map from wealth ratios w to desired adjustment sizes Δd used in equation (18). The IRF is expressed in units of durable demand per unit of time.

The first term of the integrand represents the size of desired durable demand, the second term is the hazard function that represents a probability (per unit of time) that such demand will be realized, and the third term computes the effect of the further wealth shock δ and not the total effect from the current state plus the additional shock.

We also define a discounted cumulative impulse response to summarize the potency to stimulate demand of a given shock δ up to time T.

Definition 2. **Discounted Cumulative Impulse Response.** We define the discounted cumulative impulse response (DCIR) of a wealth shock of size $\delta > 0$ to a state $s \in \{\text{recession}, \text{boom}\}$ as

$$DCIR_s(T, \delta) \equiv \int_0^T \exp(-\rho t)\, IRF_s(t; \delta)\, dt, \tag{28}$$

where ρ is the agent's discount rate. The cumulative response is expressed in units of durable demand per unit of time.

We choose to discount the cumulative impulse responses given the high persistency of the demand responses and the idea that the CIR will be evaluated by an impatient social planner. Below, we choose to study the DCIR up to a 2 year horizon in an attempt to capture responses relevant for a policy maker with a economic stabilization objective. The literature sometimes studies the instantaneous impact response (Bachmann et al. (2006) and Berger and Vavra (2015)) but that measure is misleading in the case of short-lived shocks.

Panel (a) of Figure 10 depicts the impulse response function expressed as percentage point of annual durable demand. The demand response in a boom is higher at impact and remains larger for about 6 years. Inspecting panel (b) we can see that the demand response during a boom is predicted to be 55% larger during a recession. Regarding levels, panel (c) and (d) show the discounted cumulative responses after the first two years in percentage points (pp) of annual steady state demand, the $DCIR(2, \delta = 1\%)$. During a recession, a 1% wealth shock brings a 2 year cumulative increase in durable demand of about 65 pp of annual demand whereas during a boom, the response is of 42 pp.

Panels (c) and (d) are designed to reveal non-linearities in demand responses. They depict the discounted 2 year cumulative response normalize by the wealth shock size. Therefore units are percentage points of annual durable demand per a 1% shock to wealth, akin to a benefit-to-cost ratio. Panel (c) shows that responses during recessions are almost linear but slightly concave providing a slightly worse benefit-to-cost ratio for demand stabilization at

Figure 10: Durable Demand Responses

(a) $IRF_s(t, \delta = 1\%)$

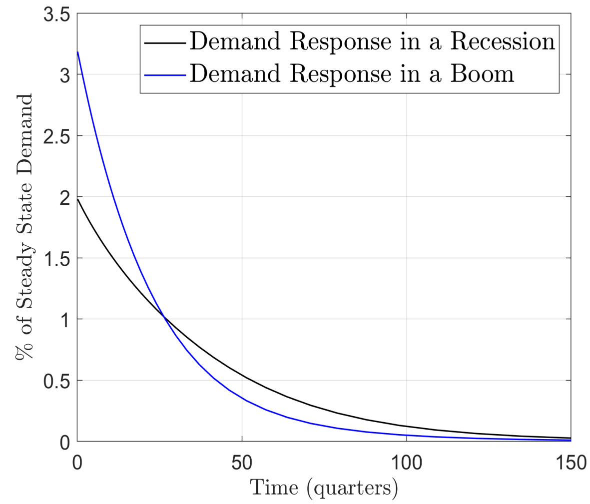

(b) Relative Demand Response $\frac{DCIR_b(2,\delta)}{DCIR_r(2,\delta)}$

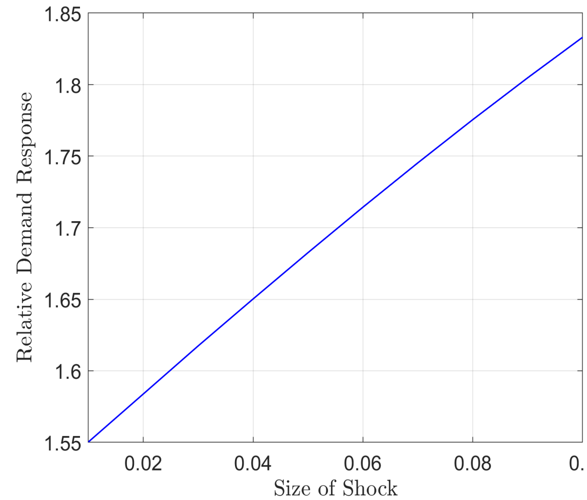

(c) Normalized CIR in Recessions $\frac{DCIR_r(2,\delta)}{\delta}$

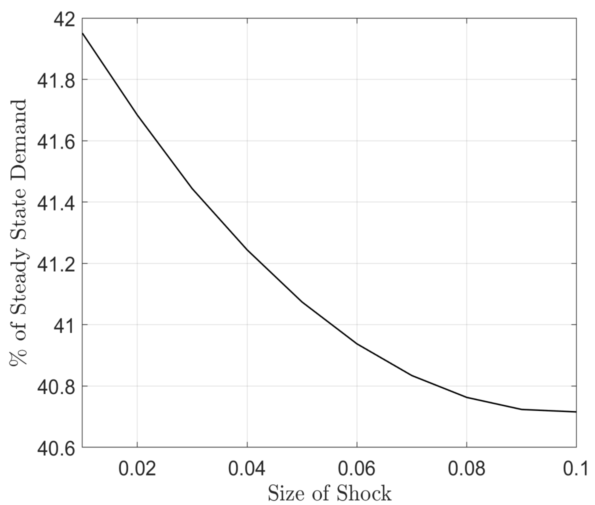

(d) Normalized CIR in Booms $\frac{DCIR_b(2,\delta)}{\delta}$

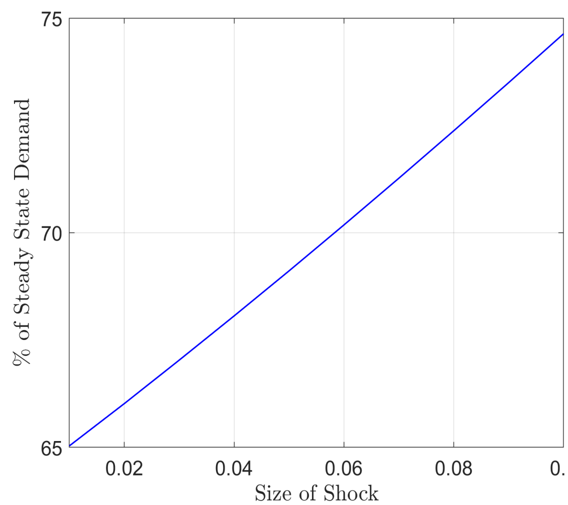

Note: The models use the estimated hazard function.

large shock sizes. Panel (d), on the contrary, shows that demand responses during a boom are convex in the size of the shock and a policy expanding wealth by 10% yields an additional 10 pp of stimulus per each percentage point of the wealth shock. Namely, a full additional unit of annual demand. As way of summary, panel (b) shows the relative demand response and indicates that the relative expansionary potency of wealth shocks is increasing in the size of the shock from 55% more expansionary to 85%. [13]

Why are responses so different? The reason lies in the state of the economy and its ability to spur durable demand. Looking at the left panel of Figure 8, the red dashed line has a large share of agents located well within the inaction region where transaction costs incentivize households to wait for larger swings in wealth before making a move. Furthermore, in the recession state, the group of agents with positive incentives to adjust (to the right of the inaction region) is small and located at low adjustment hazard rates (around a 5% annual probability). These two features make it so that a further positive shock (i) keeps most agents still within the inaction region and (ii) shifts a small group of households to slightly higher hazard rates. On the contrary, inspecting the right panel of Figure 8 the boom state has its mode very close to the edge of the inaction region and features a large group of households with positive incentives to adjust (to the right of the inaction region) and located at higher adjustment hazard rates (around a 15% annual probability). A further positive shock to the boom state then (i) pushes a substantial group of agents to the positive adjustment region and (ii) displaces the agents that already had positive incentives to adjust to even higher desired adjustments sizes with higher adjustment hazards. In sum, it is the interaction of the agents' distribution of wealth with endogenous hazard functions that yield these asymmetric and state-dependent effects.

[13] Interestingly, Bachmann et al. (2006) find a similar figure for the relative magnitude of investment responses.

6 Concluding Remarks

This paper investigates the cyclicality of durable consumption responses through the lens of a novel heterogeneous-agent portfolio choice model. By focusing on the observed micro-level heterogeneity in durable adjustment sizes, we are able to provide a more disciplined and empirically grounded approach to estimating adjustment frictions. The proposed methodology leverages a structural mapping between adjustment hazards and the stationary distribution of desired durable changes, a technical contribution that is portable to other environments featuring lumpy behavior.

Our findings yield several economically relevant insights. First, the model successfully rationalizes the procyclical behavior of the total frequency and the fraction of upward adjustments observed in the data. It then strongly predicts procyclical aggregate durable consumption responses. We demonstrate that this cyclicality is a direct consequence of the interaction between the cross-sectional distribution of households and their endogenous, state-dependent adjustment probabilities. Second, we show that stabilization policies aimed at stimulating durable demand are significantly more potent during economic booms than during recessions. Responses are also non-linear, with a positive wealth shock in a boom state yielding a durable demand response that is 60% to 80% larger than a comparable shock in a recession state. This highlights a crucial consideration for policymakers, as it suggests that a naive approach to stimulus may have wildly different benefit-to-cost ratios depending on the state of the economy. Finally, our work contributes to the literature by providing a robust framework for estimating fundamental adjustment frictions directly from observed adjustment behavior, circumventing the need for imputation and potentially inaccurate measures of wealth. This framework, which includes switching costs, does not alter the portfolio choice and asset pricing predictions relative to the standard transaction cost model, which aligns with empirical support from several prior studies.

Future research should explore the implications of different sources of household heterogeneity, such as inattention and information frictions (Alvarez, Guiso, and Lippi (2012), Gabaix (2017)), and their impact on aggregate dynamics. Another potentially important direction is analyzing the effect of mortgage rates on durable demand responses. A complication in such analysis is that the current mortgage rate becomes an additional state variable. Finally, a more formal method for estimation can be devised using maximum likelihood estimation because parameters (hazard functions λ) can be mapped into distributions of actual adjustments q which assign a likelihood to observations on adjustments.

References

Alvarez, F., L. Guiso, and F. Lippi (2012, May). Durable consumption and asset management with transaction and observation costs. *American Economic Review 102*(5), 2272–2300.

Bachmann, R., R. J. Caballero, and E. M. Engel (2006, June). Aggregate implications of lumpy investment: New evidence and a dsge model. Working Paper 12336, National Bureau of Economic Research.

Berger, D. and J. Vavra (2015, January). Consumption dynamics during recessions. *Econometrica 83*, 101–154.

Cavallo, A., F. Lippi, and K. Miyahara (2024, December). Large shocks travel fast. *American Economic Review: Insights 6*(4), 558–74.

Epstein, L. G. and S. E. Zin (1989). Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica 57*(:), 937–969.

Flavin, M. and S. Nakagawa (2008, March). A model of housing in the presence of adjustment costs: A structural interpretation of habit persistence. *American Economic Review 98*(1), 474–495.

Gabaix, X. (2017, Dec). Behavioral inattention. NBER Working Papers 24096, National Bureau of Economic Research, Inc.

Grossman, S. J. and G. Laroque (1990). Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods. *Econometrica 58*(1), 25–51.

Hansen, L. P. and T. J. Sargent (2025, September). *Risk, Uncertainty, and Value.* QuantMFR.

Hansen, L. P. and P. Souganidis (2025). Stochastic responses and marginal valuation. *Manuscript, University of Chicago*.

Kreps, D. M. and E. L. Porteus (1978). Temporal Resolution of Uncertainty and Dynamic Choice. *Econometrica 46*(1), 185–200.

Lippi, F. and A. Oskolkov (2023, November). A structural model of asymmetric lumpy investment. *Manuscript, Einaudi Institute of Economics and Finance*.

Martin, R. F. (2003). Consumption, durable goods, and transaction costs. International Finance Discussion Papers 756, Board of Governors of the Federal Reserve System (U.S.).

Stokey, N. L. (2009). Moving costs, nondurable consumption and portfolio choice. *Journal of Economic Theory 144*(6), 2419–2439.

A Appendix

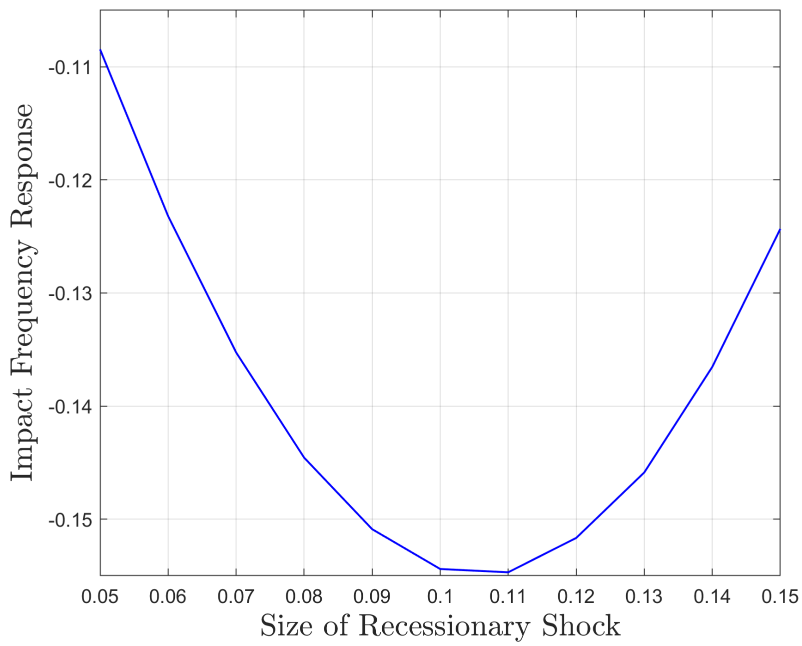

Figure 11: Frequency Response at Impact

Note: Levels are log deviations from steady state frequency of adjustments.