

Consumption Dynamics During Recessions

David Berger
Northwestern University

Joseph Vavra*
University of Chicago and NBER

9/16/2014

Abstract

Are there times when durable spending is less responsive to economic stimulus? We argue that aggregate durable expenditures respond more sluggishly to economic shocks during recessions because microeconomic frictions lead to declines in the frequency of households' durable adjustment. We show this by first using indirect inference to estimate a heterogeneous agent incomplete markets model with fixed costs of durable adjustment to match consumption dynamics in PSID microdata. We then show that aggregating this model delivers an extremely procyclical Impulse Response Function (IRF) of durable spending to aggregate shocks. For example, the response of durable spending to an income shock in 1999 is estimated to be almost twice as large as if it occurred in 2009. This procyclical IRF holds in response to standard business cycle shocks as well as in response to various policy shocks, and it is robust to general equilibrium. After estimating this robust theoretical implication of micro frictions, we provide additional direct empirical evidence for its importance using both cross-sectional and time-series data.

JEL Classification: E21, E32, D91

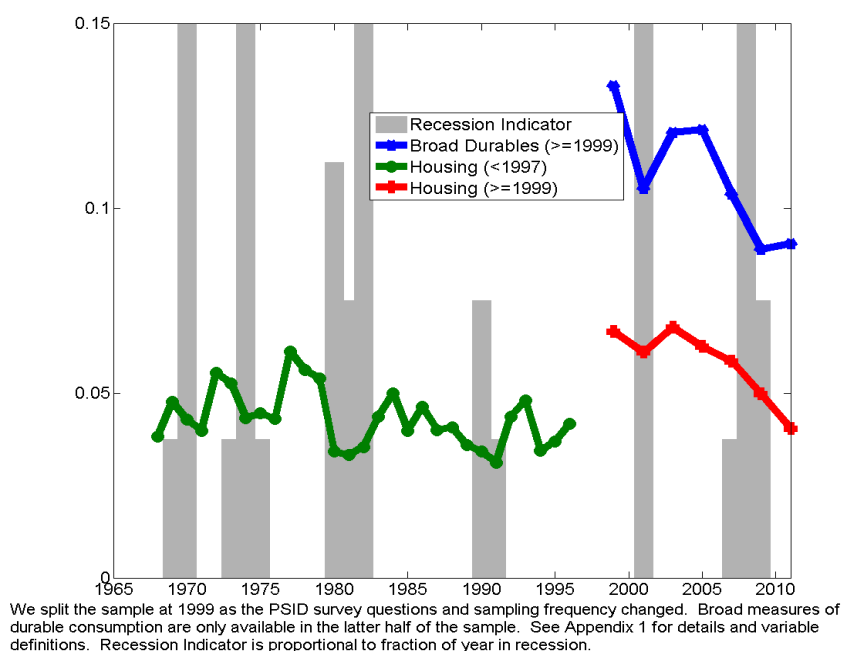
Keywords: Durables, Fixed Costs, Consumption, Non-linear Impulse Response, Indirect Inference

*We would like to thank the Editor and anonymous referees as well as Ian Dew-Becker, Steve Davis, Eduardo Engel, Jonathan Heathcote, Erik Hurst, Loukas Karabarbounis, Giuseppe Moscarini, Aysegul Sahin, Tony Smith and numerous seminar and conference participants. We would like to thank David Argente for excellent research assistance. Correspondence: joseph.vavra@chicagobooth.edu.

1 Introduction

Does the response of aggregate durable spending to a given change in policy depend on the state of the business cycle? In this paper, we argue that microeconomic durable frictions lead to sluggish macro responses during recessions.

Figure 1: Frequency of Durable Adjustment



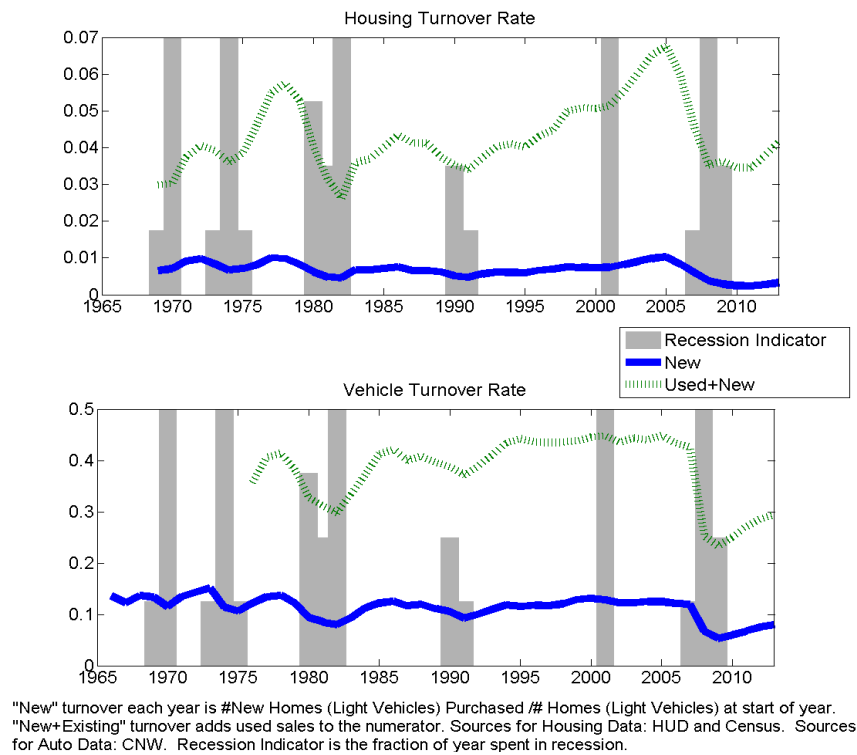
We begin by using various data to show that while durable adjustment is always infrequent, households are particularly unlikely to adjust their durable holdings during recessions. Figure 1 shows the frequency of durable adjustment in PSID data across time.¹ We show this both for a broad measure of durables, which is only available beginning with the PSID redesign in 1999, as well as for housing, which is available for a longer time-series. Panel logit regressions imply that recessions lead to a significant decline in the probability of broad durable adjustment of approximately 20% and a decline in the probability of buying/selling

¹See Appendix 1 for a detailed description of the data construction for these and subsequent figures. Broad durables include both housing and vehicles while housing includes only housing adjustment. Frequencies are annual.

a house of approximately 15%. In addition to this time-series result, we find a strong relationship between local business cycles and durable adjustment: a two-standard deviation increase in state unemployment lowers the probability of broad durable adjustment in the PSID by 30%-40% after controlling for various combinations of state, year and household fixed effects. See Appendix 1 for formal regression results.

Aggregate durable turnover shows a similar pattern: Figure 2 shows various measures of durable sales in a year divided by initial stocks. The first panel shows the behavior of new and used vehicle sales (as measured by CNW market research) and the second panel shows the behavior of new and existing housing sales (as measured by Census and HUD). While it is well-known that new durable purchases are highly cyclical, it is less widely documented that used durables exhibit similar patterns. These facts reinforce each other so that the probability that a randomly chosen house or car changes hands falls dramatically in recessions.²

Figure 2: How Frequently Does the Durable Stock Change Hands?



These microeconomic adjustment patterns have important implications for business cycle dynamics. In particular, infrequent and lumpy durable adjustment at the household level

²New (New+Existing) house turnover is 19% (22%) lower in recession years than non-recession years. Similarly, New (New+Existing) vehicle turnover is 11% (14%) lower. See Appendix 1 for description of our data.

leads aggregate durable expenditures to become much less responsive to shocks or unanticipated policy changes during recessions. Why is there a cyclical link between micro lumpiness and aggregate responsiveness? Declines in wealth and income during recessions lead fewer households to adjust their durable holdings upwards and more households to adjust them downwards. However, the presence of depreciation means that the number of increases declines more quickly than the number of decreases grows. Thus, during recessions, fewer households adjust their durable holdings, which sharply reduces the elasticity of aggregate durable expenditures to aggregate shocks.

Understanding the behavior of broad durable expenditures is crucial for understanding recessions. Consumer durables and residential investment respectively accounted for 24% and 33% of the total decline in real GDP between 2007-2009 so that declines in broad durable spending account for more than half of the recession.³ From 1960-2013 both components of GDP were highly cyclical and volatile, with reductions in consumer durable spending (residential investment) accounting for 26.6% (58.3%) of real GDP changes during recessions.⁴ Leamer [2007] shows that residential investment and durable spending are the two most importance components in explaining "Weakness in GDP" going into recessions prior to 2007-2009. Thus, in a pure accounting sense, stabilizing broad durable expenditures would substantially moderate the business cycle, and indeed, a number of policy interventions during the Great Recession were specifically designed to stimulate durable demand.⁵

We argue for the quantitative and empirical relevance of procyclical durable responsiveness in five steps:

- 1) We use indirect inference to estimate a heterogeneous agent incomplete markets model with fixed costs of durable adjustment to match household behavior in PSID. In particular, we use a novel "gaps" based approach that maximizes the fit between model and data along the dimensions which are most important for explaining durable adjustment. This procedure is extremely successful as our model is able to explain 72-86% of observed variation in household adjustment probabilities. In addition, our estimated model matches a variety of facts that are not directly targeted.

- 2) After arguing that our estimated model matches micro consumption dynamics, we explore its implications for aggregate dynamics. We begin the macro analysis with a series of aggregate shocks in partial equilibrium. We start with a partial equilibrium analysis because it allows us to explore a more empirically realistic baseline model and provide more sensitivity

³This is the change in components of real chained GDP divided by the change in total real chained GDP from 2007q4 to 2009q2.

⁴This is the average contribution to percent change in real gross domestic product from BEA Table 1.1.2 calculated over NBER recession quarters.

⁵For example, the Cash for Clunkers and First Time Home Buyers credit.

analysis relative to what is feasible in general equilibrium. In addition, it allows us to explore the implications of business cycles for household dynamics in a model that perfectly replicates the aggregate behavior of income and wealth. We show that the response of aggregate durable expenditures to a variety of shocks is highly procyclical. In particular, we allow for shocks to income, wealth, taxes, interest rates and subsidies to durable adjustment. In all cases there is substantial state-dependence so that the same shock has much smaller effects if it occurs in a recession than if it occurs during an expansion.⁶

3) As discussed above, the procyclical impulse response in our model is driven by variation across time in the distribution of households' durable holdings together with the probability they adjust. We next show that we can directly test for this reduced form implication in PSID data, and we show that the data strongly supports this theoretical implication.

4) While steps 1-3 provide evidence for procyclical responsiveness in partial equilibrium, a large literature argues that general equilibrium can undo these effects. To assess this, we next add general equilibrium to our model and show that our conclusions are robust.⁷ The key reason that GE is not particularly important in our framework is that households can save in both illiquid wealth and liquid wealth. If households can only save in only one asset so that $Y = C + I$ as in Khan and Thomas [2008], then lumpy investment behavior necessarily induces violations of consumption smoothing. With two sources of savings so that $Y = C + I_k + I_d$ this is not the case.

5) Finally, we provide additional reduced form evidence that the response of durable spending to economic shocks is indeed procyclical. In particular, we exploit geographical variation to show that the response of MSA-level automobile spending to identified wealth shocks strongly interacts with local economic conditions: the response of auto spending to wealth shocks is much higher in MSAs experiencing local booms than in MSAs experiencing local recessions.⁸

Thus, a variety of structural, reduced-form and time-series evidence supports the conclusion that durable expenditures respond less strongly to shocks during recessions. However, it is important to note that our results are about the relative effectiveness of durable stimulus over the business cycle, so they do not on their own imply that durable stimulus is ineffective during recessions. What they do imply is that policy makers will get less bang-for-the-buck from policies designed to stabilize durable expenditures during recessions than suggested by

⁶Importantly, our model implies a state-dependent IRF but not an asymmetric IRF.

⁷One disadvantage of GE is that the set of permissible exogenous shocks is more limited. For example, we can no longer introduce exogenous shocks to interest rates since these are determined endogenously. For this reason, we focus on TFP shocks in general equilibrium.

⁸In the appendix, we also show that our model with fixed costs of durable adjustment is a substantially better fit to aggregate time-series evidence than are existing models with durable consumption.

linear VAR evidence. Indeed, in Berger and Vavra [2014] we provide evidence using non-linear VARs that durable spending multipliers are substantially lower during recessions than those implied by linear VARs. In addition, Kaplan and Violante [2014] argue that policies designed to stimulate non-durable spending are likely to become more effective during recessions, so such policies may be relatively more attractive for stabilization.

For most of the paper, we focus on analyzing a broad measure of durable spending that encompasses both consumer durables and housing. We focus on this broad notion of durables since procyclical responsiveness should apply to any purchase which is long-lived and illiquid. These are important characteristics at the broadest level of durable aggregation. In addition, both consumer durable spending and residential investment are very large and have similar cyclical patterns.⁹ Nevertheless, focusing on this broad notion of durables forces us to abstract from some institutional features that may be important for housing but not for autos (or vice versa). For this reason, we consider several robustness checks that focus separately on different durable components and show that our conclusions remain.

There is a long line of literature studying models with durable consumption.¹⁰ In pioneering work, Eberly [1994] estimates (S,s) triggers for household auto purchases based on the stylized model of Grossman and Laroque [1990]. She then interacts these estimated triggers with estimates of the household wealth distribution to explain the aggregate time-series for U.S. auto purchases. We expand on this approach in several important ways. Since her work is based on Grossman and Laroque [1990] she imposes a single (S,s) trigger. In addition, she must exclude liquidity constrained households from her analysis. We show that in our model, which allows for binding borrowing constraints as well as (S,s) triggers that vary with income and wealth, these assumptions matter. In contrast to our model, the Grossman and Laroque [1990] model has very little predictive power for most households' durable adjustment. In a similar stylized model, Bar-Ilan and Blinder [1992] argue that (S,s) models should lead durable spending to depend on the past history of durable purchases and thus the distribution of households' current gaps.

Attanasio [2000] and Bajari, Chan, Krueger, and Miller [2013] use alternative estima-

⁹It is important to note that the stock of housing is somewhat larger than the stock of consumer durables, but that this is due to slightly lower depreciation rates. The average level of real consumer durable expenditure is slightly larger than the average level of real residential investment.

¹⁰See Mankiw [1982], Bernanke [1985], and Caballero [1990] for studies of durables and the PIH hypothesis. Bertola and Caballero [1990], Grossman and Laroque [1990] and Caballero [1993] provide analytical models of durable consumption with fixed costs. Leahy and Zeira [2005], Luengo-Prado [2006], and Browning and Crossley [2009] study the role of durable wealth for explaining non-durable consumption. There is also a large body of work studying various aspects of durable consumption over the life-cycle including Dunn [1998], Krueger and Fernandez-Villaverde [2010], and Diaz and Luengo-Prado [2010].

tion procedures to try to understand automobile and housing demand. In contrast to our approach, they estimate reduced form policy functions rather than solving the households' dynamic programming problem. This approach makes the results less suitable for analyzing policy changes that might alter the estimated reduced form relationships.

Iacoviello and Pavan [2013] build an incomplete markets model with fixed costs of housing adjustment and aggregate shocks. In contrast to our paper, they perform a simple calibration exercise for the parameters of the model and do not explore its ability to explain micro dynamics. In addition, they focus on entirely different aggregate questions. While our model is infinite horizon, they instead build a life-cycle model, and computational considerations then require an annual rather than quarterly frequency. As such, their model is less suited for examining business cycle dynamics and they instead focus on explaining secular changes in aggregate volatility.

Kaplan and Violante [2014] study the implications of illiquid wealth holdings such as durables for the behavior of non-durable consumption and show that they are able to explain the response of non-durable consumption to one-time fiscal rebate payments. In addition, they briefly show that illiquidity can potentially lead to state-dependent consumption dynamics. We view our work as highly complementary to their own but it is distinct in several ways. For the most part, we focus on the implications of illiquidity for durable expenditures rather than for non-durable spending because durable spending is substantially more important for understanding business cycle behavior. Since our motivation is understanding how micro consumption dynamics influence aggregate business cycles, our model also features a variety of aggregate shocks and we explore the implications of general equilibrium.

Finally, our paper is closely related to Caballero, Engel, and Haltiwanger [1995], Caballero, Engel, and Haltiwanger [1997] and Bachmann, Caballero, and Engel [2013], which argue for time-varying responsiveness arising from lumpy firm behavior. Besides the obvious difference that we study households rather than firms, there are several distinctions between our analyses. In Caballero, Engel, and Haltiwanger [1995] and Caballero, Engel, and Haltiwanger [1997] they impute capital and employment gaps and explore their aggregate implications. Their gap imputation relies on an assumption that firms' reset targets follow a random walk, while our procedure requires no such assumption. Bachmann, Caballero, and Engel [2013] build a quantitative GE model of firm investment and targets various aggregate time-series facts to address concerns that these early papers were not robust to general equilibrium and lacked quantitative realism. However, they do not test their model implications in micro data.¹¹

¹¹The literature on firm lumpiness must also contend with issues that are not present in our household environment. In particular, it can make a large quantitative difference whether these models are calibrated to match firm vs. establishment moments, and it is not clear what level of aggregation corresponds to an

To summarize, our analysis overlaps in part with many papers, but we believe we are the first paper to jointly explore the micro and macro implications of household durable adjustment in an estimated, structural GE model. We believe the synthesis of microdata, structural modeling and general equilibrium is important for providing an accurate assessment of the impact of policy changes.

2 Model and Estimation

2.1 Model Description

Our baseline model for estimation is a standard incomplete markets model with the addition of household durable consumption subject to fixed costs of adjustment. Households maximize expected discounted utility of a consumption aggregate, and they are subject to idiosyncratic earnings shocks as well as borrowing constraints. In this section, we describe the partial equilibrium version of the model with no aggregate shocks, and in the following sections we discuss the addition of aggregate shocks: first in partial and then in general equilibrium.

Households solve:

$$\max_{c_t^i, d_t^i, a_t^i} E \sum \beta^t \left(\frac{[(c_t^i)^v (d_t^i)^{1-v}]^{1-\gamma} - 1}{1-\gamma} \right)$$

s.t.

$$c_t^i = wh\eta_t^i(1-\tau) + (1+r)a_{t-1}^i + d_{t-1}^i(1-\delta_d) - d_t^i - a_t^i - A(d_t^i, d_{t-1}^i)$$

$$a_t^i \geq -(1-\theta)d_t^i; d_t^i \geq 0$$

$$\log \eta_t^i = \rho_\eta \log \eta_{t-1}^i + \varepsilon_t^i \text{ with } \varepsilon_t^i \sim N(0, \sigma_\eta),$$

where c_t^i , d_t^i and a_t^i are household i 's non-durable consumption, durable stock, and liquid assets, respectively. The parameter β is the quarterly discount factor, v is the relative weight on non-durable consumption in period utility, and $1/\gamma$ is the intertemporal elasticity of substitution.¹² η_t^i represents shocks to idiosyncratic labor earnings, h is a household's

economic decision maker. In contrast, for household level durable adjustment, the correct level of aggregation does not have any such ambiguity.

¹²Piazzesi and Schneider [2007] provides some evidence in favor of the Cobb-Douglas period utility function. Note the Cobb-Douglas utility function also means we can normalize the service flows from durables to be equal to the stock without loss of generality.

fixed¹³ hours of work while w and r are the aggregate wage and interest rate, δ_d is the depreciation rate of durables and τ is a proportional payroll tax. Finally, $A(d_t^i, d_{t-1}^i)$ is the fixed adjustment cost that households face when adjusting their durable stock. We assume that A takes the form

$$A(d, d_{-1}) = \begin{cases} 0 & \text{if } d = [1 - \delta_d (1 - \chi)] d_{-1} \\ F^d (1 - \delta_d) d_{-1} + F^t wh\eta_t^i & \text{else.} \end{cases}$$

Following Bachmann, Caballero, and Engel [2013] $0 \leq \chi \leq 1$ is a "required maintenance" parameter. Positive values of χ represent the fact that some maintenance is required to continue enjoying the flows from durable consumption, e.g., fixing a flat tire on a car or fixing a broken furnace in a house.¹⁴ When a household adjusts its durable stock, it must pay fixed adjustment costs that take two forms. First, they lose a fixed fraction of the value of their durable stock. These costs correspond to brokers fees, titling costs, etc. Second, households face some time cost of adjusting their durable holdings. These costs correspond to, e.g., the time involved in searching for a new house or in researching which car to purchase. We allow for this general specification because these two adjustment costs may interact differently with the business cycle. The opportunity cost of time is procyclical so that time costs will tend to generate countercyclical durable adjustment. Conversely, fixed costs that are proportional to the stock of durables have the most bite when income is low and tend to generate procyclical durable adjustment. Estimating a specification with both costs allows the data to inform their relative importance.

Given these assumptions, the infinite horizon problem can be recast recursively as

¹³Endogenizing hours complicates the model and does not affect our main conclusions.

¹⁴In previous versions of this paper, we considered an adjustment cost function that allowed households to endogenously choose the amount of maintenance between 0 and 1 without paying the fixed adjustment cost. This led to similar results but substantially increases the computational burden of the model, which makes estimation infeasible.

$$\begin{aligned}
V(a_{-1}, d_{-1}, \eta) &= \max [V^{adjust}(a_{-1}, d_{-1}, \eta), V^{noadjust}(a_{-1}, d_{-1}, \eta)] \\
&\quad \text{with} \\
V^{adjust}(a_{-1}, d_{-1}, \eta) &= \max_{c,d,a} \frac{[c^v d^{1-v}]^{1-\gamma}}{1-\gamma} + \beta E_\varepsilon V(a, d, \eta') \\
&\quad \text{s.t.} \\
c &= wh\eta(1-\tau) + (1+r)a_{-1} + d_{-1}(1-\delta_d) - d - a - F^d(1-\delta_d)d_{-1} - F^t wh\eta \\
a &> -(1-\theta)d \\
\log \eta' &= \rho_\eta \log \eta + \varepsilon \quad \text{with } \varepsilon \sim N(0, \sigma_\eta) \\
V^{noadjust}(a_{-1}, d_{-1}, \eta) &= \max_{c,a} \frac{[c^v d^{1-v}]^{1-\gamma}}{1-\gamma} + \beta E_\varepsilon V(a, d_{-1}(1-\delta_d(1-\chi)), \eta') \\
&\quad \text{s.t.} \\
c &= wh\eta(1-\tau) + (1+r)a_{-1} - \delta_d \chi d_{-1} - a \\
a &> -(1-\theta)d \\
\log \eta' &= \rho_\eta \log \eta + \varepsilon \quad \text{with } \varepsilon \sim N(0, \sigma_\eta)
\end{aligned}$$

We now turn to a discussion of how we estimate the parameters of the model. The computational solution of the model is discussed in Appendix 4.

2.2 Estimation

To decrease computational burden, our estimation procedure proceeds in two steps: we first calibrate some subset of parameters for which we have reliable external evidence. We then estimate the remaining parameters using an indirect inference procedure, which we describe shortly.

2.2.1 Calibration and Model Restrictions

We calibrate several parameters of our model in standard ways but have explored the robustness of our conclusions to changes in these parameters. We set $r = 0.0125$, which delivers an annual interest rate of approximately 5%, and we set the discount factor $\beta = 0.98$. In our benchmark model we set $\gamma = 2$. We normalize $w = 1$ and set $h = 1/3$. We calibrate the idiosyncratic productivity process to match the persistence and variance of annual labor earnings in PSID data which yields a persistence of idiosyncratic earnings of 0.975 and a standard deviation of 0.1, and we set the payroll tax equal to 5% to reflect a combination

of the statutory rate with phaseouts for high income.¹⁵ We calibrate the depreciation rate of durables to match data from the BEA, weighted by the relative size of the housing and consumer durable stocks. That is, we set $\delta_d = \delta_H^{BEA} \frac{H^{BEA}}{H^{BEA} + CD^{BEA}} + \delta_{CD}^{BEA} \frac{CD^{BEA}}{H^{BEA} + CD^{BEA}}$, which delivers a quarterly value of 0.018. For simplicity, we abstract from growth, but this does not affect our conclusions.¹⁶

In the general formulation of our model, durables serve a dual role: they provide direct utility to households, but they also serve as collateral against which households can borrow. For most of the analysis that follows, we will shut down this second channel by setting $\theta = 1$. However, in Appendix 3 we estimate a version of the model with $\theta = 0.20$ so that households need only pay a 20% down payment to purchase new durable holdings. We show that this version of the model delivers similar results both for micro and macro durable dynamics.

There are two main reasons that we choose to make the model with $\theta = 1$ our benchmark: First, when $\theta < 1$ and there is no adjustment costs on a , the model implies that households can costlessly adjust their durable equity. In other words, such a parameterization implies that households can costlessly refinance, which is clearly counterfactual. Since it is infeasible to solve a more realistic model with liquid assets, semi-liquid durable equity, and illiquid durables we concentrate on the case with no refinancing rather than the case with costless refinancing as our benchmark. Second, if collateral constraints become looser during expansions this will tend to amplify all of our results since when down-payment requirements are low households can rapidly adjust their durable holdings in response to shocks. In contrast, when down-payment requirements are large, households must save a larger amount of liquid assets before increasing their durable holdings. By shutting this channel down, our quantitative conclusions are thus relatively conservative. Setting $\theta = 1$ in our benchmark model also makes our results more comparable to the model in Kaplan and Violante [2014], which rules out collateralized borrowing against illiquid assets.

In addition to exploring the role of collateral constraints, Appendix 3 also explores a second important empirical extension of our model. In particular, we consider the role of rental markets for our analysis and provide evidence that introducing rental markets has little quantitative effect on our results. While rental markets are not particularly

¹⁵Since the tax is fixed across time, and hours are exogenous this plays essentially no role in our analysis. Using a higher value to match overall income taxes (or excluding taxes from the model entirely) yields nearly identical results along all dimensions. We only include the tax so that we can perform simple policy experiments with temporary and permanent tax changes in the following section.

¹⁶Varying the depreciation rate within reasonable ranges or including trend income growth consistent with that over our sample period had negligible effects on our results. For simplicity and comparison with most business cycle models, we abstract from growth and focus on deviations around a steady-state.

important for consumer durables, they play a large role in housing markets. It would be desirable to build a model with separate consumer durables and housing, but this is technically infeasible. We disallow rental markets in our benchmark model for three reasons: 1) Consumer durable spending represents more than half of total broad durable spending from 1960-2013, and rental markets are not important for consumer durables. 2) Introducing rental markets increases the computational burden of the problem substantially by adding an additional choice.¹⁷ 3) The indirect inference procedure we describe next is based on the "gap" between a households' current durable holdings and those it would hold if it temporarily faced no adjustment costs. Defining gaps in a world with rental markets is not straightforward.¹⁸ Focusing on a benchmark model without rentals and restricting the empirical analysis to homeowners obviates all of these issues. Nevertheless, Appendix 3 shows that the introduction of rental markets does not alter our conclusions.¹⁹

2.2.2 Estimation Procedure

The remaining parameters of our model are the proportional fixed cost of durable adjustment, F^d , the time cost of durable adjustment F^t , the non-durable weight in utility v , and the level of required maintenance χ . In addition, we also estimate a measurement error parameter, σ_ϵ , which allows for all variables in the data and model to be reported with some error. We assume that the reported value of a variable \widehat{Z} is the true value Z plus some percentage measurement error: $\widehat{Z} = Z(1 + \widehat{\epsilon})$ with $\widehat{\epsilon} \sim iid N(0, \sigma_\epsilon)$.

We estimate these parameters using a "gap" based indirect inference procedure. First note that in models with fixed adjustment costs, we can always define a gap $x = \log d^* - \log(d_{-1})$, where d^* is the choice of d that solves the maximization problem in V^{adjust} . Intuitively, x measures the difference between the stock of durables that a household inherits at the start of a period and the stock of durables that a household would choose if it adjusted today. However, since the household does face adjustment costs, its actual choice of durables today may or may not be equal to d^* . If $V^{adjust} > V^{noadjust}$ then the household will choose to adjust and set $d = d^*$ and otherwise, the household will choose to not adjust and will set $d = d_{-1}(1 - \delta_d(1 - \chi))$. The larger the (absolute) value of x the more likely it is that the gains from adjusting exceed the fixed adjustment cost. Thus, the adjustment hazard $h(x)$ will be increasing in the (absolute) size of the gap. This implies that measuring household

¹⁷It also requires estimating the relative value of renting versus owning.

¹⁸In the data it is also not obvious how to define durable stocks for households that simultaneously rent apartments while owning vehicles.

¹⁹In addition, homeownership rates are fairly stable across time with only mild procyclicality. Furthermore, for the small changes that are observed, homeownership rates rise somewhat more quickly in expansions than they fall in booms, so that accounting for rental markets in the data would amplify our conclusions.

gaps is essential for understanding households' durable adjustment decisions.

In addition, the distribution of gaps $f(x)$ plus the adjustment hazard $h(x)$ also determines aggregate durable expenditures at a particular point in time. Aggregate durable expenditures will be given by the amount that a given household purchases when adjusting times its probability of adjusting. This implies that aggregate durable expenditures are given by $I_D = \int x h_t(x) f_t(x) dx$, where $h_t(x)$ is the probability of adjusting at time t as a function of the gap and $f_t(x)$ is the distribution of gaps at time t .²⁰ Given that the distribution of gaps and hazards is critical for understanding both micro and macro adjustment, the goal of our indirect inference procedure is to pick parameters so that distributions of gaps and hazards in the model match those in the data.

The parameters we are estimating crucially affect the demand for durables, and hence the distribution of gaps and probability of adjustment. In particular, F^d affects the width and steepness of the adjustment hazard and F^t affects the symmetry of the hazard since households that are decreasing durables tend to be poorer and have lower opportunity costs of time. The level of durable vs. non-durables and thus the mean gap in the data is affected by v , χ affects the skewness of the gap distribution, and the degree of measurement error affects the level of the hazard (the probability that a household with no observed gap adjusts anyway).²¹

Using superscript m to represent model objects and superscript d to represent data objects, let $f_p^m(x)$ and $h_p^m(x)$ be the distribution of gaps and hazard implied by the model with vector of parameters p . If we knew the distribution of gaps and hazards in the data, $f^d(x)$ and $h^d(x)$, we would then pick p to solve $\min_p \left\{ \int \left[(f_p^m(x) - f^d(x))^2 + (h_p^m(x) - h^d(x))^2 \right] dx \right\}$. That is, we would pick our parameters so that the simulated distribution of gaps and hazards in the model match those in the data. If we observed x in the data, this procedure would be straightforward. The obvious complication with implementing this procedure is that we do not observe x in the data, so we cannot compute $f^d(x)$ and $h^d(x)$.

While we do not directly observe x in the data, this procedure is not hopeless because we can impute x using restrictions implied by our structural model.²² We know that in our model, there is a mapping from observables to d^* and thus x . That is, for a particular set of parameters, we can construct a model-generated function G^m that maps variables z which are observable in both the data and the model to x , which is only observable in the model:

²⁰This intuitive expression ignores maintenance expenditures, but quantitatively in the simulated model these are close to constant across time so that this intuitive expression is highly accurate for capturing changes in durable expenditures numerically.

²¹We have a more formal discussion of identification in Appendix 3.

²²This is analogous to the procedure in Caballero, Engel, and Haltiwanger [1995] and Caballero, Engel, and Haltiwanger [1997], but in those papers they impute gaps using some simple rules of thumb that approximate the true model but are not actually consistent with optimal behavior.

$x^m = G^m(z^m)$. By applying this same function to actual data, we can then impute a gap measure: $x^d = G^m(z^d)$. Thus, imposing structural restrictions from our model allows us to overcome a methodological challenge by imputing unobservable empirical objects from observable empirical objects.

The data requirement for estimation is then that we observe the variables in z , and that we observe households' adjustment decisions so that we can construct $h^d(G^m(z))$. We leave a more complete discussion of the functional form of G^m , as well as a discussion of z for Appendix 2. There we argue that data on a, d, c is required to accurately predict model gaps.²³ In addition, we require panel data on these objects so that we can construct adjustment hazards and control for unmodeled household fixed effects. To our knowledge, the only data sets satisfying this restriction are the PSID²⁴ (from 1999-2011), and the Italian Survey of Household Income and Wealth (SHIW). We concentrate mainly on PSID data but discuss some results for SHIW in Appendix 2. We mention the data for our estimation before formally stating our estimation procedure because it introduces two additional complications: 1) PSID data is self-reported and subject to substantial measurement error. 2) Beginning with the sample redesign, PSID only collects data every other year while our model period is quarterly. We address both of these complications directly by aggregating our model data to the same frequency as the PSID and introducing measurement error when comparing our model objects to their empirical counterparts. With this in mind, we now formally state our estimation procedure:

1) For a given set of parameters p , solve the model and compute $x^m = G^m(z^m)$. 2) Introduce measurement error and aggregate the model to the same frequency as PSID to compute model gaps with sampling error: $\widehat{x}^m = G^m(\widehat{z}^m)$. 3) Compute imputed gaps in the PSID: $\widehat{x}^d = G^m(\widehat{z}^d)$.²⁵ 4) Compute the difference between model simulated hazards and densities and those in the data: $L_p = \int \left[\left(f_p^m(\widehat{x}^m) - f^d(\widehat{x}^d) \right)^2 + \left(h_p^m(\widehat{x}^m) - h^d(\widehat{x}^d) \right)^2 \right] dx$. We then repeat 1-4 with a different set of parameters and minimize L . Finally, we bootstrap standard errors for all model parameters as well as distributions and hazards, but for brevity we leave the discussion to Appendix 2.

In the standard language of indirect inference, our reduced form auxiliary model is given by $f(G^m(\widehat{z}^m))$ and associated hazard $h(G^m(\widehat{z}^m))$. Let $\Gamma(\widehat{z}^m)$ be the joint density of

²³We have tried also using income as an additional element of z and it yielded similar results. See Appendix 3.

²⁴Prior to the PSID sample redesign in 1999, only food consumption was recorded and there was no consistent data on vehicle holdings

²⁵Note that in the data, we only observe empirical objects with measurement error so for notational symmetry we replace z^d with \widehat{z}^d from this point forward, since we only compare model objects with measurement error to the data.

model variables. This joint density together with its evolution encompasses the full structure of the model, but the pdf $f(G^m(\hat{z}))$ summarizes the complicated joint-density of model variables with measurement error $\Gamma(\hat{z}^m)$ in a one-dimensional distribution of gaps. The hazard $h(G^m(\hat{z}^m))$ collapses the time-series evolution of the joint-density of \hat{z}^m into a one-dimensional probability of adjustment as a function of gaps. Thus, our indirect inference estimator in essence collapses high-dimensional joint-densities $\Gamma(\hat{z})$ to more practical one-dimensional functions. Since our reduced form auxiliary model is collapsing some information from the full structural model and is also introducing measurement error and time-aggregation bias, it will in general be a misspecified description of the dynamics of the true model. However, it is important to note that as usual in indirect inference, consistent estimation does not require the auxiliary model to be correctly specified. As long as the reduced form model is computed identically on actual and simulated data we will achieve consistent estimation. We further discuss this point in addition to a discussion of identification of our structural parameters in Appendix 2.

Now that we have formally stated our estimation procedure, we provide some additional discussion in intuitive terms before turning to results. It is important to note that since the distribution of gaps in the model as well as in the data are purely functions of the joint distribution of observables, \hat{z} , our estimation strategy is in some sense trying to make these joint distributions line up with each other. If the joint distribution of observables in the model $\Gamma(\hat{z}^m)$ was able to perfectly match the joint distribution of observables in the data $\Gamma(\hat{z})$, then by construction the distribution of gaps in the model and data would be identical.

However, given that we have few parameters and that $\Gamma(\hat{z}^d)$ is an extremely high dimensional object, a perfect fit is clearly unobtainable.²⁶ Since it is infeasible to perfectly match the joint-distribution of wealth, durable holdings and non-durable consumption, which moments of this distribution are most important to match? Our gap-based indirect inference procedure provides the answer to this question. We should weight moments of $\Gamma(\hat{z}^d)$ by their importance for determining gaps. For example, if our model told us that the ratio of non-durable to durable consumption was extremely important for determining household gaps, while liquid wealth was unimportant, then our estimation strategy would place more weight on matching the former distribution and less weight on matching the latter.

In the following section, we will show that our best fit parameters yield a distribution of gaps in the model that is an extremely good fit to the distribution in the data, which means

²⁶A large literature exists just trying to match the wealth-distribution. Matching $\Gamma(z^d)$ is a vastly more difficult goal since wealth, durable and non-durable expenditures are not independent. For example, existing studies that target the wealth distribution attempt to match $\int f(a)da$ while $\Gamma(z^d) = \int_a \int_d \int_c f_{a,d,c}(a, d, c)dadddc$ is clearly a much more complicated object.

we match the moments of $\Gamma(\hat{z})$ that are important for determining gaps.²⁷ More importantly, we show that our model is very accurate at predicting actual durable adjustment in the data. Since $h^d(G^m(\hat{z}^d))$ is the actual adjustment probability in the data for a household with imputed gap $\hat{x}^d = G^m(\hat{z}^d)$ there is no guarantee that this empirical adjustment probability will correspond to that in the model. This implies that calculating the empirical hazard as a function of imputed durable gaps provides a test for model misspecification: if our structural model is misspecified then our imputed gaps \hat{x}^d will not be particularly useful for explaining observed adjustment probabilities. For example, if G was a random uniform function, $h^d(\hat{x}^d)$ would be completely flat. If imputed gaps are completely random then households with large imputed gaps will be just as likely to adjust as households with gaps of zero. Finding an upward sloping empirical hazard as a function of (absolute) imputed gaps is evidence that our model provides useful predictive power for households' actual durable adjustment decisions in the data. In essence, our estimation procedure is trying to maximize the ability of our model to explain actual durable adjustment patterns but there is no guarantee that we would succeed at this goal.

We now turn to a brief description of our data and then present results showing that our model is a very good fit for both the density of gaps and the empirical adjustment hazard while simpler existing models are unsuccessful at explaining actual durable adjustment.

2.2.3 Data Description

Here we briefly describe the data and sample restrictions for our benchmark estimation. We leave a more detailed description and various robustness descriptions for Appendix 1. Our estimation uses data from the PSID from 1999-2011. Households are interviewed every other year, and are asked a variety of questions about non-durable consumption, wealth, housing, vehicles and income. While it would be desirable to extend the analysis to data before 1999, the previous PSID samples only collected food consumption rather than broad non-durable consumption. In addition, vehicle data was not constructed consistently across time.

The value of c is the sum of all components of food consumption, utilities, transportation expenses, schooling expenses and health services. Our measure of d is the sum of housing and vehicle values and a is the sum of business value, stocks, IRAs, cash and bonds minus the value of outstanding debt. Since our benchmark model does not include rental markets, we restrict our estimation to continuous home-owners in our benchmark results. In Appendix 3

²⁷In the following section, we show that our model is a good fit for various moments of $\Gamma(z)$, which shows that our best fit parameters do not produce unrealistic distributions of observables. This suggests that an alternative estimation procedure directly targeting $\Gamma(z)$ would likely yield similar results. However, by construction, such a procedure would be less accurate at predicting actual household durable adjustment patterns.

we discuss an extension of our model to include rental markets and adjust our PSID analysis accordingly. After constructing measures of c, d , and a per household member we deflate nominal values using NIPA price indices, adjust for household age and remove a household fixed effect.²⁸

We restrict our analysis to households that are in the nationally representative core sample, whose head is less than 65 years of age, and which have non-missing data on c , d and a . See Appendix 1 for additional discussion of our data, cleaning procedures and alternative robustness checks.

2.3 Estimation Results

Table 1 displays our parameter point estimates together with bootstrapped 95% confidence intervals. Our point estimate for the fraction of the value of durables lost when adjusting is 0.0525. This is line with estimates of the size of realtors fees, and it is also similar to values typically used in the literature.²⁹ In the following sections, we show that this fixed cost has important implications for aggregate dynamics. In contrast, we estimate a negligible (and statistically insignificant) time cost of durable adjustment. While not directly targeted, we show that the non-durable share in utility of 0.88 delivers ratios of durable to non-durable expenditures that are consistent with the data. The point estimate for our measurement error parameter implies that measurement error is distributed mean zero with standard deviation of 8%. This implies that a reported value will be within 5% (10%, 15%) of the true value approximately 47% (80%, 94%) of the time. Finally, our estimated maintenance parameter implies that households offset 80% of depreciation each quarter.³⁰

Table I

Parameter	Point Estimate	95% Confidence Interval
F^d (Fixed Cost Stock)	0.0525	(0.043,0.068)
F^t (Fixed Cost Time)	0.001	(0.000,0.004)
ν (Utility Flow Non-Dur)	0.88	(0.875,0.885)
m (Measurement Error)	0.08	(0.06,0.10)
χ (Maintenance)	0.80	(0.75,0.95)

²⁸The age fixed effects removes pure demographic effects, which we do not model. Household fixed effects remove any unmodeled permanent differences across households (which are ex-ante identical in the model).

²⁹Diaz and Luengo-Prado [2010] calibrate a value of 0.05 and Bajari, Chan, Krueger, and Miller [2013] estimates a value of 0.06 in models of housing adjustment. Eberly [1994] uses a transaction cost of 0.05 in her analysis of automobiles.

³⁰This relatively large maintenance value is required to explain the fact that both housing and vehicle adjustment are less frequent than would be expected by the "raw" depreciation numbers.

Given these estimated parameters, how well does our model fit the distribution of gaps and hazard in the data? Figure 3 shows the distribution of gaps in the model, \widehat{x}^m , and imputed gaps in the data, \widehat{x}^d . The shaded areas are bootstrapped 95 percent confidence intervals. Overall, the fit is extremely close with overlapping confidence intervals at all points. Again, this close fit between model and data means that for our best fit parameters the model is able to match $\Gamma(\widehat{z}^d)$ along the dimensions important for explaining gaps. In addition, the estimated densities are moderately negatively skewed due to the presence of depreciation, which we will show has important implications for aggregate dynamics.³¹

Figure 3: Distribution of Gaps in Model and PSID

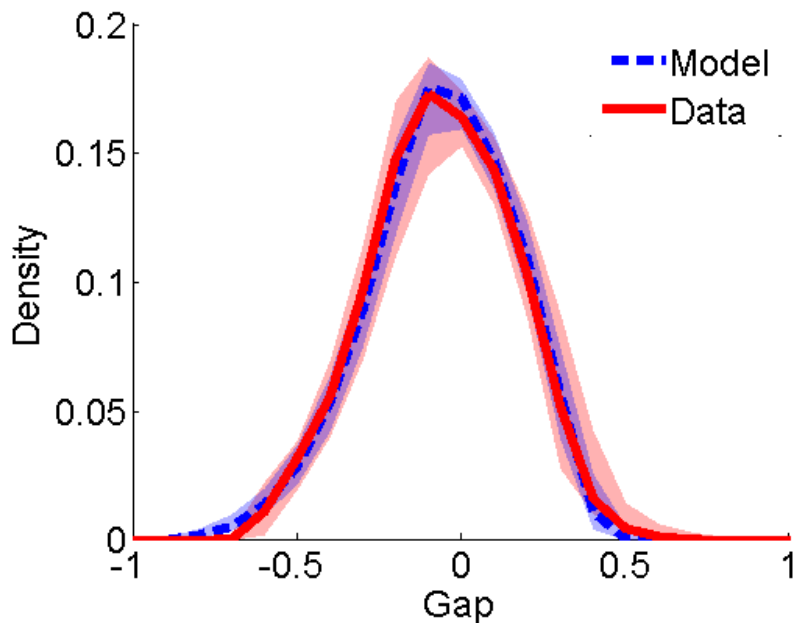



Figure 4 shows the adjustment hazard in the model and in the data. In the model, this is equal to the probability that a household adjusts for a given gap \widehat{x}^m . Note that the hazard in the model does not follow a strict (S,s) rule, jumping from 0 to 1 at some threshold. This is because different state variables can map to the same gap so that sometimes a household with a given gap will choose to adjust and other times it will not.³² In the data, the hazard is equal to the actual empirical probability that a household with an imputed gap \widehat{x}^d chooses to adjust. As stressed in the previous section, this is a very strong test of whether the structural model is well-specified. If the model is very misspecified then the imputed gaps \widehat{x}^d will have little predictive power for empirical durable adjustment.

³¹The skewness of the gap distribution is approximately -.35.

³²Note that if we conditioned on all state variables rather than just the gap, households would follow a strict (S,s) rule.

Figure 4: Adjustment Probabilities in Model and PSID



Overall, our model is extremely successful at predicting actual durable adjustment in the data.³³ We can assess this more formally using several quantitative measures of the fit between the model and hazard. First, we can compare the additional explanatory power of our model versus a Calvo model that just implies all households adjust with the same probability. That is, we compute $R^2 = 1 - \frac{\sum_{\hat{x}} (f^m(\hat{x})[h^m(\hat{x}) - h^d(\hat{x})])^2}{\sum_{\hat{x}} (f^m(\hat{x})[h^m(\hat{x}) - freq])^2}$. This tells us how much of the total variation in the hazard predicted by the model is observed in the empirical hazard.³⁴ The $R^2 = 0.91$. Thus, 91% of the total variation in hazards predicted by our model is observed in actual data. This statistic tells us something about the global fit of the model over the whole distribution of gaps, but we might also be interested in a local measure of fit: given a gap, how well does the model predicted hazard match the empirical hazard? To assess this, we compute $\int \left(\frac{|h^m(\hat{x}) - h^d(\hat{x})|}{h^d(\hat{x})} f^m(\hat{x}) d\hat{x} \right)$. This tells us the average percentage deviation between the model and empirical hazards. We find that the average deviation is .276 so that given a gap, we can on average explain 72% of the observed hazard

³³Clearly the standard errors for the empirical hazard are wider than those for the model but the hazard is strongly upward sloping. Wider standard errors in the data occur because there is some idiosyncratic adjustment in the data not explained by our model and this "noise" interacts with sampling error in regions of the gap distribution with little mass.

³⁴We weight the deviations between $h^d(\hat{x})$ and $h^m(\hat{x})$ by $f^m(\hat{x})$ to account for the fact that more gap mass is close to zero than far out in the tails. That means that we should care more about getting the hazard correct in the middle of the distribution. If we weight all points in the hazard equally rather than weighting by the gap density then we get an $R^2 = 0.98$.

in the data.³⁵ Thus, while our model is not a perfect fit to the empirical hazard, we can explain a very large fraction of observed adjustment probabilities.³⁶

It is important to note that since we only have five parameters, there was no guarantee that any configuration of parameters would be successful at matching the reduced form hazard and density of gaps. In this sense, the predictive power of our model is not driven mechanically by imputing empirical gaps from our model structure. In Appendix 2, we show this more explicitly by performing an identical estimation procedure using the model of Grossman and Laroque [1990], which has served as the basis for many existing empirical studies. We show there that the empirical hazard computed from imputed gaps is nearly flat, and is, if anything downward sloping. This implies that the model actually has modestly negative predictive power: when the model predicts that households in the data should be more likely to adjust than average, they are empirically less likely to adjust than average.

Thus, we view the strong ability of our benchmark model to predict empirical adjustment patterns as its main strength: while the structure used to impute gaps is complicated, given our imputed gap we are highly accurate at predicting when households will adjust.

In addition to the hazard and density, which are explicitly targeted by our indirect inference estimation procedure, we can also assess the model fit along various dimensions which are not directly targeted. Kaplan and Violante [2014] emphasize the importance of "wealthy-hand-to-mouth" consumers for explaining household consumption dynamics. They argue that many households have a large fraction of their wealth in illiquid assets such as durables and that these households may behave quite differently from those with access to liquid wealth. Thus, if we want to take seriously the implications of our model for consumption dynamics, it is important that it implies reasonable numbers of both hand-to-mouth and wealthy-hand-to-mouth households. We define a hand-to-mouth household as one who has liquid assets less than one-half of their monthly labor earnings, and we then define a hand-to-mouth household as wealthy-hand-to-mouth if its durable holdings are greater than the 25th percentile of durable holdings in the population.³⁷ Using this definition in PSID data, 28.7% of households are hand-to-mouth and 17.8% of households are wealthy-hand-to-mouth. While the estimation does not directly target these numbers it

³⁵The average absolute difference (rather than percentage difference) between the model and empirical hazard is .038. In addition, as noted in the previous footnote, we weight deviations by the density of gaps. If we instead weight all points on the hazard equally then we explain 86% of the observed hazard in the data and find an absolute deviation of .032.

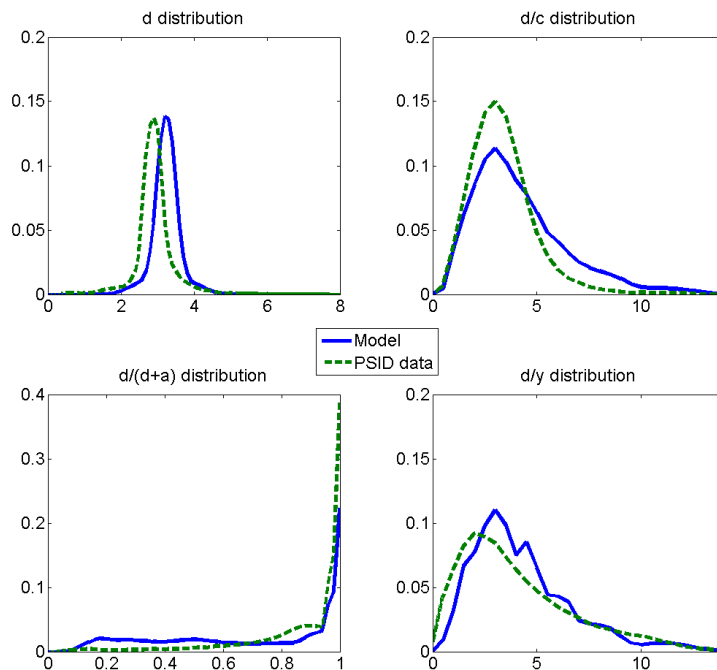
³⁶We do not model life-cycle interactions or shocks to locational preferences that might interact with housing decisions. We suspect that the unexplained portion of durable adjustment is largely driven by these factors, which should be largely independent of the business cycle.

³⁷Using different definitions such as 1/4 of monthly earnings for hand-to-mouth or the 50th percentile for durable holdings produces similar fits between model and data.

produces extremely similar results, with 26.4% of households hand-to-mouth and 18.0% of households wealthy-hand-to-mouth.

In addition, our model produces an average frequency of adjustment which is close to that in PSID data. Using our broad definition of durables that encompasses housing+vehicles, durable stocks in the PSID data have an annual frequency of adjustment of 10.8%. The model implies a frequency of adjustment of only slightly above this at 12.9%.

Figure 5: Durable Holdings in Model and Data



Finally, we can assess the model’s ability to match the overall patterns of durable holdings in the data. In Figure 5 we show four different slices of the durable distribution in $\Gamma(\hat{z})$. We show the unconditional distribution of durable holdings as well as the relationship between durable holdings and non-durable consumption, the relationship between durable holdings and total wealth, and the relationship between durable holdings and income. Overall, the model is a good fit to the data. The only place where the model misses somewhat more substantially is on the mean level of durable holdings. While the distribution of durable holdings around the mean in the model and data are quite similar, the model overstates mean durable holdings relative to the data by roughly 10%. However, this does not have important consequences for any of our results. We can refit a version of the model that explicitly targets mean durable holdings to be equal in the model and data. By construction, this model is a

slightly worse fit for the distribution of gaps and hazards but gives almost identical aggregate results. Again, the fact that our benchmark estimation does not exactly match mean durable holdings in the data implies that mean durable holdings are not particularly important for determining the distribution of gaps and hazards.³⁸

Overall our estimated model is a good fit to household level consumption dynamics both along targeted as well as untargeted dimensions. Bolstered by this good microeconomic fit, we now explore the aggregate implications of our model.

3 Aggregate Implications of Lumpy Durable Purchases

We explore the macroeconomic implications of our model by first exploring the response to a number of shocks in partial equilibrium. The use of partial equilibrium analysis has several advantages: 1) Partial equilibrium is substantially faster to compute than general equilibrium, which allows us to explore the robustness of our results to various extensions such as collateralized borrowing and rental markets and to perform additional sensitivity analysis. 2) In partial equilibrium we can explore a number of aggregate shocks (such as exogenous changes in interest rates) that are more challenging to model in general equilibrium. We will argue that lumpy micro adjustment has important implications for how actual durable spending responds to any shock that changes desired durable holdings, so it is important to explore robustness to a variety of shocks.³⁹ 3) In partial equilibrium, we can pick a sequence of aggregate income and wealth shocks that exactly reproduces the behavior of U.S. GDP and capital across time so that our simulated economy well-approximates the actual U.S. economy.

3.1 Aggregate Income Shocks

We begin by exploring the implications of aggregate income shocks. For brevity, we leave the full model description to Appendix 4 and just discuss the changes in the model relative to the previous section. In the previous section, we assumed that $\log \eta' = \rho_\eta \log \eta + \varepsilon$ with $\varepsilon \sim N(0, \sigma_\eta)$ where ε is an idiosyncratic income shock. We introduce aggregate income shocks by assuming that total household wages $\log y_{tot}$ are the sum of an idiosyncratic component plus an additional aggregate shock:

³⁸This makes sense since as average durable holdings rise both d_{-1} and d^* will tend to rise and gaps are not that affected. Furthermore, time-series movements in aggregate durable expenditures also don't depend much on mean durable holdings since they are determined by changes in d rather than levels of d .

³⁹Since our mechanism applies in a wide-variety of environments and in response to a variety of shocks, we prefer to focus on documenting the generality of our mechanism rather than taking a stand on a particular source of business cycle shocks or focusing on the institutional details of one particular policy change.

$$\log y_{tot} = \log \eta + \log y.$$

As before, we assume that the idiosyncratic component of income, $\log \eta$, follows an AR process with persistence 0.975 and standard deviation of 0.10. We assume that the aggregate component of income, $\log y$, follows an AR process with persistence 0.87 and standard deviation 0.008 to match the behavior of hpfilted GDP from 1960-2013.⁴⁰ This adds one additional aggregate state-variable to the household's problem but solution methods are otherwise unchanged. We solve this model using the parameters previously estimated and then compute impulse response functions to $\log y$ shocks.

Motivated by the evidence in Figure 1, we are particularly interested in whether micro non-linearities in durable adjustment lead durable spending to respond differently to income shocks which occur at different points in the business cycle. To do this, we must first define booms and recessions in our model. We match our model to U.S. data by picking a particular sequence of aggregate income shocks in the model $\log y_{1960q1}, \dots, \log y_{2013q4}$ to exactly reproduce hpfilted US GDP from 1960-2013. Given this sequence of shocks, we can then compute the impulse response of durable expenditures to an *additional* impulse to aggregate income at each date. That is, we feed a sequence of aggregate shocks exogenously into our model for 212 quarters. Then given the history of aggregate shocks up to each date, we compute the full impulse response function to an additional shock at that date. See Appendix 4 for additional discussion of the computation of impulse responses.

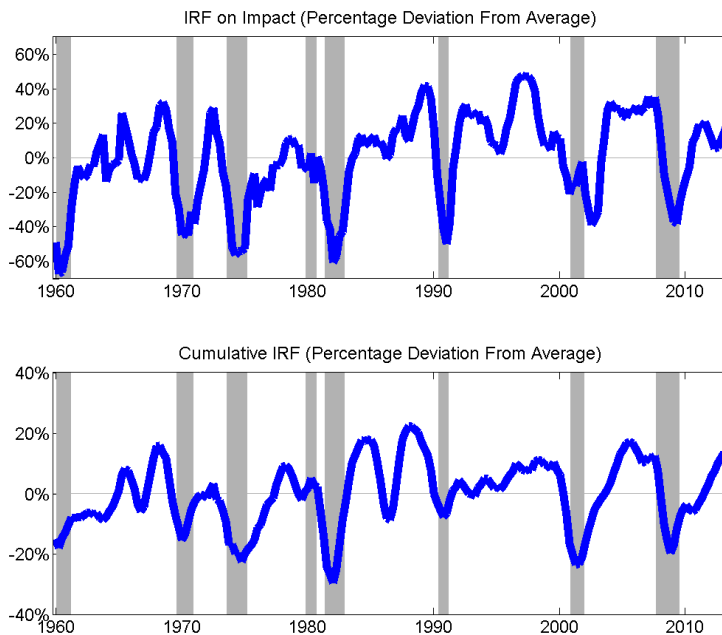
We summarize our state-dependent impulse response in three ways. First, following Bachmann, Caballero, and Engel [2013], we compute the first element of the impulse response function (IRF) for each quarter between 1960q1 and 2013q4. This "Responsiveness Index" provides an estimate of how much durable expenditures will respond to an aggregate shock to income in the quarter in which it occurs. We are particularly interested in the IRF on impact since this has direct relevance for how durable expenditures are likely to respond in the short-run to shocks or stimulus policies during recessions. Second, we report the cumulative response⁴¹ of durable expenditures to the same impulse to income. Figure 6 shows that measured either using either method, durable expenditures are substantially less responsive to income shocks during recessions.⁴²

⁴⁰Calibrating the income shocks to labor compensation yields nearly identical results.

⁴¹The cumulative response is the total area under the impulse response function from 1-8 quarters (after which the IRFs are indistinguishable from zero).

⁴²As we will show more formally when discussing what drives this result, this is evidence of an impulse response that depends on the state of the business cycle, it is not evidence of an asymmetric response to positive and negative shocks. Negative income shocks in booms also have bigger effects on durable spending than negative income shocks in recessions.

Figure 6: How Responsive Are Durable Expenditures to Income Shocks?



On average, the IRF on impact in recessions is only 54% of that in expansions, indicating an economically significant amount of state-dependence. Table 2 shows that the 95th percentile of the IRF on impact is 174% larger than the 5th percentile and that the 95th percentile of the cumulative IRF is 46% larger than the 5th percentile.

Table 2

Aggregate Shock	$\frac{IRF_{impact}^{95}}{IRF_{impact}^5}$	$\frac{IRF_{cum}^{95}}{IRF_{cum}^5}$
Income	2.74	1.46
Wealth	6.17	4.72
Interest Rate	2.29	2.17
Tax	1.60	1.52
Durable Purchase Subsidy	1.85	1.91

95 is the 95th percentile across time. 5 is 5th percentile across time. Impact computes the first element of the IRF and cum is the total area under the IRF

The third way we examine the extent of state-dependence is by plotting the entire impulse response function for particular dates. The years 1999 and 2009 are boom and recession years which also overlap with dates in our PSID data, so we focus on the average IRF in

these years:⁴³

Figure 7: Durable Expenditure Impulse Responses to 1% Aggregate Income Shock

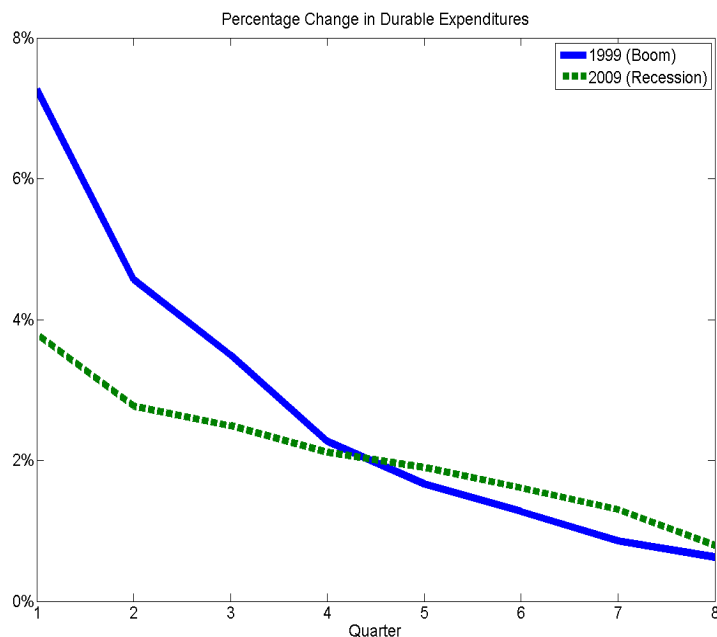


Figure 7 shows that the IRF on impact in 1999 is estimated to be almost twice as large as that in 2009. While the IRF on impact is most relevant for assessing the short-run impact of economic shocks during recessions, these differences persist for several quarters: the cumulative IRF in 1999 is more than 30% larger than that in 2009.

Before turning to an explanation for this procyclical durable spending IRF, we show that the same result holds for a variety of other aggregate shocks. Beyond just providing a simple robustness check, this is important because we want to argue that the aggregate implications of lumpy micro adjustment apply to a wide-class of aggregate shocks. Essentially all shocks that are commonly used to explain business cycles yield similar implications.

3.2 Aggregate Wealth Shocks

While we consider income shocks to be the most natural proxy for U.S. business cycles in a partial equilibrium model, we next show that wealth shocks deliver similar results. We think of these shocks as proxying for declines in stock market value or other asset holdings during recessions which will affect households' consumption decisions. We assume that households'

⁴³Other boom and recession years yield similar results.

liquid wealth is subject to aggregate shocks which follow some AR process in logs. That is, $a'_{actual} = a'_{choice} * w'$ with $\log w' = \rho_w \log w + \varepsilon_w$. We calibrate these shocks to match the persistence and standard deviation of the hpfitered quarterly U.S. capital stock, which we construct using a perpetual inventory method as in Bachmann, Caballero, and Engel [2013]. This yields a quarterly persistence of 0.95 and a standard deviation of 0.003 so that aggregate wealth shocks are small but highly persistent.

Figure 8: How Responsive Are Durables to Wealth Shocks?

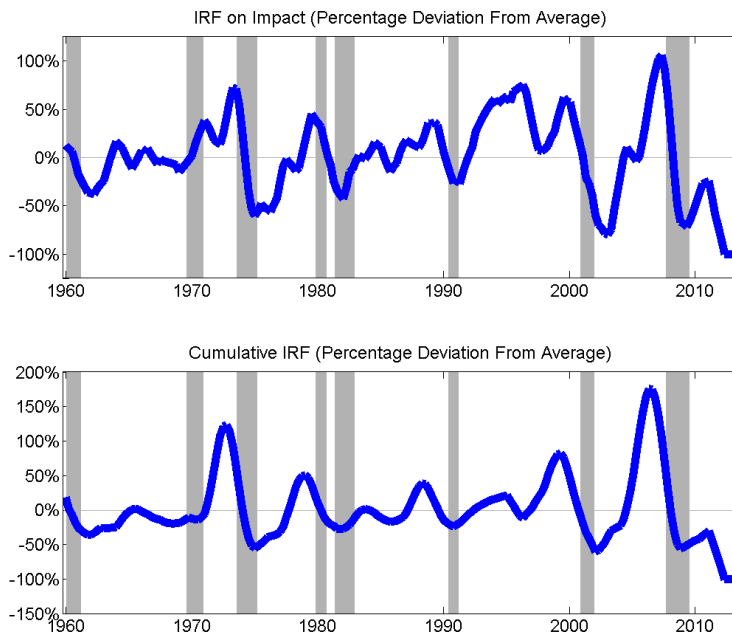


Figure 8 shows that the durable response to wealth shocks is even more procyclical than the response to income shocks.⁴⁴ Table II shows that the 95th percentile of the IRF on impact is more than 6 times as large as the 5th percentile. The 95th percentile of the cumulative IRF is almost 5 times larger than the 5th percentile.⁴⁵

Our baseline wealth shock is a proportionate equal decline in all households wealth so that all households face the same shock. However, rich households have a greater proportion of their total wealth in liquid assets and so are more affected by these shocks. Nevertheless, this proportional shock may still understate distributional effects of wealth shocks if rich

⁴⁴Note that the model business cycles in the two versions of the model are not identical, since in one aggregate income exactly matches U.S. GDP, in the other aggregate wealth exactly matches U.S. capital. Solving a model with both shocks simultaneously would be much more computationally difficult.

⁴⁵Wealth shocks induce greater time-series variation in IRFs because they are more persistent than income shocks and lead to larger movements in households desired durable holdings’.

households hold assets which are riskier and lose more value during recessions. To assess the importance of this channel, we have resolved a version of the model with wealth shocks that only affect wealthy households as well as with wealth shocks that are increasing in the level of household wealth. For brevity we do not plot the results but note that all of our conclusions are strengthened under these alternative specifications: if wealth shocks mainly affect the rich then the IRF becomes even more procyclical.⁴⁶

3.3 Policy Shocks

In addition, we can compute impulse responses to shocks that roughly correspond to various policy experiments. Since we do not think the business cycle is primarily driven by any of these policy experiments, we now perform a slightly different experiment. Rather than assuming that there are stochastic shocks to policy and picking these shocks to match the behavior of GDP, we introduce one-time unanticipated policy shocks on top of our previous model with aggregate income shocks. That is, we assume that households are subject to aggregate income shocks which, as before, are picked to match the behavior of U.S. GDP. We then compute the optimal response of households to a one-time unanticipated policy experiment at different points in the business cycle (as defined by aggregate income).

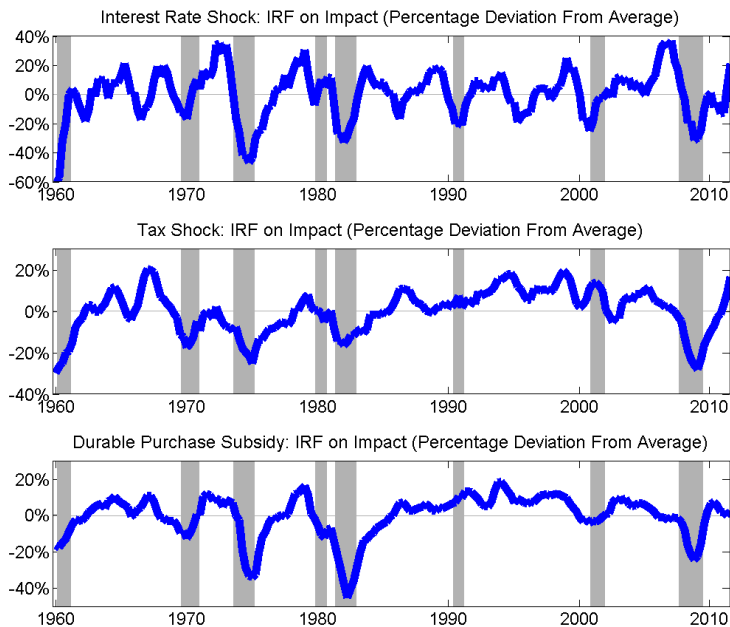
While it is not computationally feasible to simultaneously introduce stochastic policy shocks together with stochastic aggregate income shocks, we can compute the durable response to changes in policy that are either completely temporary or are fully permanent. For brevity we only report results for permanent policy shocks, but temporary shocks deliver similar time variation. We compute the impulse response to three policy shocks: a permanent decline in the interest rate, a permanent decline in the payroll tax, and a subsidy to durable adjustment (which is financed by an increase in taxes).⁴⁷ We view these policy experiments as rough approximations to the various stimulus policies such as reductions in payroll taxes and "Cash-for-Clunkers" that were implemented during the recession of 2007-2009. While we believe a more detailed quantitative study of these particular policies is an important subject for future research, in this paper we want to focus on the broad fact that micro-level household behavior has important implications for a broad variety of aggregate shocks, which necessitates abstracting from some of the institutional details important to each of these policies.

Figure 9 plots the impulse response to each of these three policy shocks. Again, the IRF

⁴⁶This is because as we show shortly, in a model with liquidity constraints but no illiquid wealth, the IRF is mildly countercyclical. Since this countercyclical effect of liquidity constraints is entirely driven by households close to the liquidity constraint, if these constrained households do not face wealth shocks then this effect is shut down and the IRF becomes more procyclical.

⁴⁷In the modeling appendix we describe each of these experiments in more detail.

Figure 9: Impulse Response to Policy Shocks



is procyclical. Table II shows that across time, the 95th percentile of the IRF on impact is 60-129% larger than the 5th percentile for the interest rate, tax and durable subsidy shock.

3.4 Robustness Results

The previous sub-sections show that in response to a variety of aggregate shocks, durable expenditures exhibit strongly procyclical impulse response functions. This conclusion is highly robust to a number of model extensions. As previously discussed, our benchmark analysis focuses on the broadest interpretation of durables with fixed adjustment costs. Nevertheless, this forces us to abstract from features that may make housing respond differently to shocks than automobiles or other consumer durables. The use of partial equilibrium simplifies the computation of the model so that it is feasible to explore some of these questions. In Appendix 2, we introduce rental markets and collateralized borrowing into our model and show that the model continues to deliver a quantitatively significant procyclical IRF.

In addition, Section 5 introduces general equilibrium into our benchmark model and shows that results continue to go through.

4 Understanding Procyclical IRFs: Fixed Costs and Cross-Sectional Implications

4.1 Importance of Fixed Costs

Why is the IRF of durable expenditures to aggregate shocks procyclical? These aggregate patterns arise because of the household-level non-linearities induced by fixed costs of durable adjustment. We first show this by documenting that the procyclical impulse response disappears when durable adjustment is frictionless. We then discuss the microeconomic mechanism that drives our result and provide additional evidence for this channel by further exploiting our PSID data.

Figure 10: Impulse Response to 1% Income Shocks (Frictionless Durable Adjustment)




Figure 10 shows the impulse response to income shocks in a model which is otherwise identical to our benchmark model but with $F^d \equiv F^t \equiv 0$. Clearly there is much less variation in impulse responses across time than in the model with fixed costs. Furthermore, what variation there is now countercyclical instead of procyclical. The reason the IRF becomes countercyclical when there are no fixed costs of adjustment is that during recessions, more households are close to the borrowing constraint, which increases the response of their durable expenditures to income shocks. This is just a manifestation of the classic result that

marginal propensities to consume out of income shocks are larger for liquidity constrained households.

This experiment with no fixed costs of adjustment is important because it shows that our results are driven by fixed costs rather than just by the sequence of aggregate shocks. In a model with incomplete markets, state-dependent IRFs could arise even without fixed costs of adjustment as the business cycle interacts with borrowing constraints. Indeed, we find evidence of this effect, but it works in the opposite direction of our headline result and is relatively mild.

4.2 The Role of the Cross-Section

Thus, in the model with no fixed costs of adjustment, which is inconsistent with micro data, the IRF is mildly countercyclical. In contrast, in our benchmark model with fixed costs, that matches micro data, there is an extremely procyclical IRF. Why do fixed costs of adjustment induce a procyclical IRF? We can see this by returning to the expression for aggregate durable investment: $ID = \int xh_t(x) f_t(x) dx$. The more households that choose to adjust their durable holdings and the larger the size of the gaps, the more responsive will be aggregate durable investment. Caballero and Engel [2007] show that this formula can be used to calculate the response of the economy on impact to aggregate shocks. In particular, if there is a positive shock Δd^* to households' desired durable holdings then the IRF on impact is given by:

$$IRF_t^{impact} = \lim_{\Delta d^* \rightarrow 0} \frac{\Delta ID}{\Delta d^*} = \int h_t(x) f_t(x) dx + \int xh'_t(x) f_t(x) dx. \quad (1)$$


The more households that are adjusting ($\int h_t(x) f_t(x) dx$) or that are close to the margin of adjustment ($\int xh'_t(x) f_t(x) dx$), the greater will be the aggregate response of durable expenditures.

Figure 11 plots the distribution of durable gaps and adjustment hazard in a boom and in a recession, for the model with aggregate income shocks.⁴⁸ On average the distribution has negative skewness because depreciation means that more households want to increase than to decrease durable holdings. This becomes more pronounced during the boom, as households' desired durable holdings rise and the distribution of durable gaps shifts to the right.⁴⁹ As more households are now further from their desired level of durables, they move

⁴⁸Note that here we are plotting the true model hazards and gaps (with no measurement error) while Figure 5 plots the distribution and hazard for model data with measurement error. While the true hazard is zero when the durable gap is equal to zero, measurement error leads the measured hazard to be strictly positive at all points.

⁴⁹The model with business cycles driven by aggregate wealth shocks delivers stronger movements in the

Figure 11: Model Gap Distribution and Hazard: Boom Vs. Bust



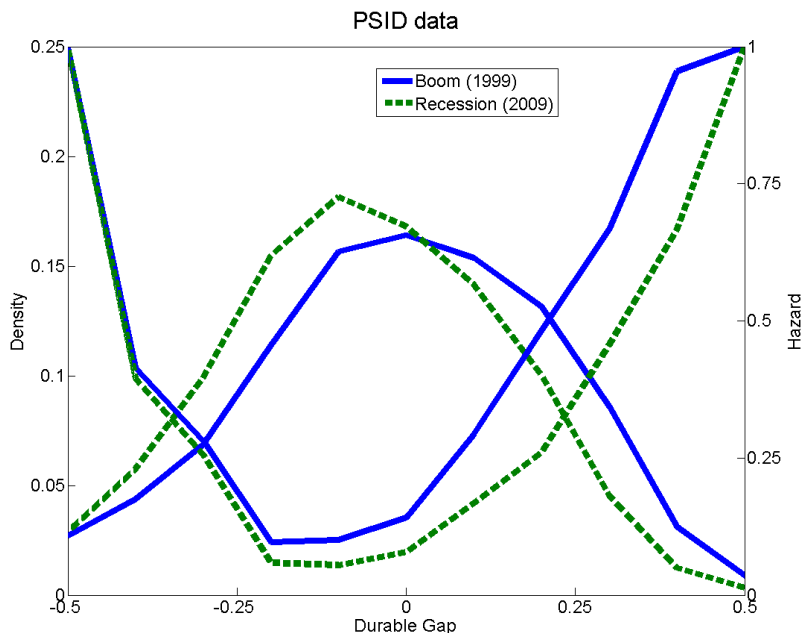
into the region with a higher probability of adjustment, and since all households that adjust will respond to aggregate shocks, aggregate durable expenditures become more responsive to these shocks. This is amplified by the increase in the probability of adjustment during a boom. Households are more likely to adjust to a given durable gap during a boom than during a recession as the fixed costs of durable adjustment represent a smaller fraction of household income.

Note that this increase in responsiveness is symmetric in the sign of the aggregate shock. During booms, a shock that increases households' desired durable holdings will raise aggregate durable expenditures by more than if this same shock occurs in a recession. But it is also true that during booms, a shock that lowers households' desired durable holdings will lower aggregate durable expenditures by more than if the same shock occurs in a recession. Our model implies an IRF that depends on the state of the business cycle; it does not imply an asymmetric IRF. Together the rightward shift of $f(x)$ and the vertical shift in $h(x)$ greatly amplify the response of aggregate durable expenditures to any shock that changes households' desired durable stocks.

Given the importance of shifts in the distribution and hazard for explaining our pro-cyclical IRF, it is important to provide additional support for this theoretical mechanism.

distribution of gaps since wealth shocks are more persistent.

Figure 12: Estimated Gap Distribution and Hazard in PSID: Boom Vs. Recession



Since our estimation procedure delivers values for $h^d(\hat{x}^d)$ and $f^d(\hat{x}^d)$ for each PSID sample year between 1999 and 2011, it is straightforward to test whether empirical hazards and gap distributions move across time as predicted by the model.⁵⁰ Furthermore, we can use (1) to calculate a reduced form responsiveness index IRF_t^{impact} implied by the PSID data and compare it to the model. Figure 12 shows that exactly as predicted by the model, the distribution of households' desired durable holdings shifts to the right and that the hazard of durable adjustment shifts up during booms. If anything, the variation in the data is even stronger than that predicted by our model which suggests that the simultaneous presence of wealth, income and other shocks over the business cycle all push households' decisions in the same direction.


Given that our estimation targeted only the average distribution and hazard in the PSID data and exploited no time-series variation in these distributions, this serves as another strong support for our model. Matching the average distribution and hazard in the data provides no guarantee that the time-series variation in the data will conform to the predictions of our theoretical model.

Figure 13 shows the PSID estimates of the IRF on impact computed using Formula (1) from 1999-2011. Comparing IRF_t^{impact} in Figure 13 to that implied by the model in the

⁵⁰That is, for PSID observables in a particular year z_t^d , we can compute $f_t^d(G^m(z_t^d))$ and then compute the empirical adjustment hazard as the actual probability of adjustment given imputed gaps in that year.

first panel of Figure 6 shows that the PSID micro data implies procyclical responsiveness that is both qualitatively and quantitatively similar to our structural model.⁵¹

Figure 13: Impulse Response Implied by PSID Gap Distribution and Hazard



In Appendix 3, we again explore the robustness of our empirical results to the inclusion of rental markets and collateralized borrowing. Since changing the structural model changes both our parameter estimates as well as the imputed gaps in the data, our estimates of $h^d(\hat{x}^d)$ and $f^d(\hat{x}^d)$ are slightly different in these alternative specifications. Nevertheless, we show that we again find shifts in the distribution and hazards, as well as time-variation in the implied impulse response on impact, that conform with our theoretical predictions.

5 Robustness to General Equilibrium

There is a large and important literature studying the role of general equilibrium in models of lumpy firm investment. In an extremely influential paper Khan and Thomas [2008] show that general equilibrium can eliminate the aggregate effects of micro lumpiness that had been found in earlier partial equilibrium work such as Caballero, Engel, and Haltiwanger [1995].

⁵¹Unfortunately, as we note in the discussion of the data for our estimation, prior to 1999, the PSID does not collect the necessary data to estimate gaps and hazards so this responsiveness index cannot be calculated further back in time.

Given that our evidence thus far is purely partial equilibrium, it is important to explore whether a similar effect arises in our model. Does the inclusion of general equilibrium price movements eliminate the time-varying IRF that we find in partial equilibrium? In this section we provide evidence that it does not, and we provide intuition for why general equilibrium is less important for lumpy household durable adjustment than is often found for lumpy firm investment.

Our general equilibrium model is identical to our benchmark partial equilibrium model, but we now endogenize the aggregate wage and interest rate. To ease comparison of our model’s aggregate dynamics with those in the existing literature, we focus on an RBC version of the model with aggregate TFP shocks Z_t . The setup is extremely similar to the partial equilibrium model, so we leave the details to Appendix 4. The main difference is that the previously exogenous interest rate and wage must now satisfy the first order conditions of a representative firm. We assume that firms forecast these prices using the methods in Krusell and Smith [1998]. In addition, aggregate state variables must be consistent with individual household decision rules.

Where possible, we choose all parameters in the general equilibrium model to be identical to those in our benchmark estimation, but there are several new parameters and restrictions imposed by general equilibrium. Since the interest rate is endogenous, we now choose β to target the steady-state interest rate used in partial equilibrium: $r = 0.0125$. We pick the depreciation rate of capital $\delta_k = 0.022$ to match the average ratio of investment to capital.⁵² We choose a capital share of $\alpha = 0.3$, and we pick $\rho_Z = 0.95$ and $\sigma_Z = 0.008$ to match the behavior of U.S. TFP.

We solve the model by conjecturing an aggregate law of motion, approximating the value function by linearly interpolating⁵³ between continuous grid points, solving the contraction, simulating the household problem and updating the aggregate law of motion until convergence is obtained. In equilibrium, the aggregate law of motion is highly accurate. See Appendix 4 for additional details on the solution method.

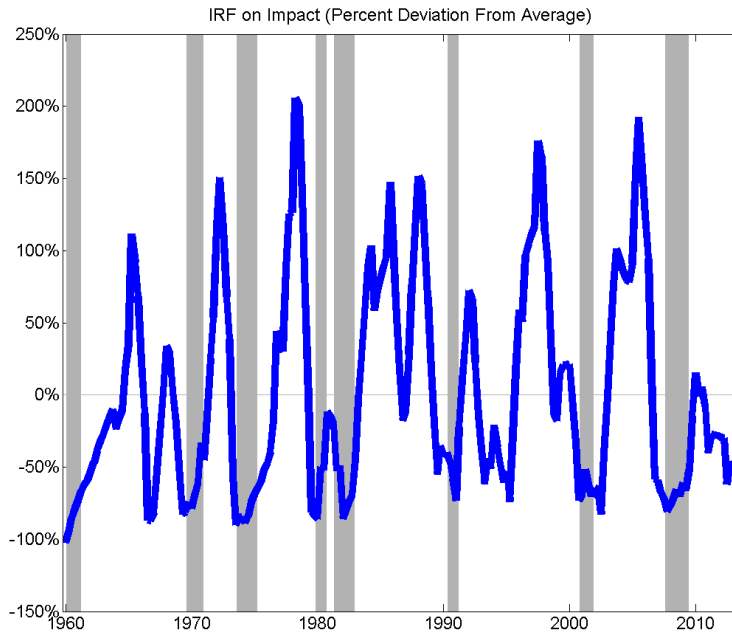
As is typical in general equilibrium models, there are now fewer degrees of freedom along which we can add shocks to the model, so the experiments we can perform are simpler in nature. Since income and wealth are now endogenous, we can no longer directly introduce aggregate shocks to these variables. Instead, we focus on the response of durable expenditures to the exogenous TFP shocks in our model. We do this not because we want to take

⁵²Changing δ_k to higher or lower values does not affect our conclusions.

⁵³We have found that linear interpolation gives speed advantages that make it attractive relative to cubic spline or other interpolation methods. While linear interpolation will introduce kinks into the value function, we do not rely on derivative based methods for solving the household problem, so this does not prove particularly problematic.

a firm stand on TFP shocks as the most important driver of U.S. business cycles but rather for illustrative simplicity. In the partial equilibrium section of the paper we showed that our results apply to a large variety of aggregate shocks and we simply want to argue that general equilibrium does not undo our basic conclusions.

Figure 14: How Responsive are Durable Expenditures to TFP Shocks in General Equilibrium?



Just as in partial equilibrium, we find a quantitatively large procyclical IRF. Figure 14 shows that there is large and procyclical variation in the IRF across time. It is worth noting that movements in the IRF in this version of the model do not line up quite as sharply with recessions as in Figure 6. However, recall that we are feeding very different aggregate shocks into these two models. In Figure 6, we were hitting the economy with aggregate income shocks that exactly correspond to actual U.S. GDP, while in Figure 14 we are hitting the economy with TFP shocks that correspond to U.S. Solow Residuals. Since TFP in the data does not perfectly comove with GDP, it is not surprising that the resulting IRF would line up less sharply with observed recessions. Nevertheless, our general conclusion remains: after sequences of TFP shocks that increase household income and wealth, durable responsiveness rises.

Why does the addition of general equilibrium have little effect for aggregate dynamics in our environment while it has large effects in Khan and Thomas [2008]? The main reason is

because in our model, households have two sources of savings: households can save in liquid assets a or illiquid assets d . In contrast, in Khan and Thomas [2008] households only have access to savings through a . Khan and Thomas [2008] argue that the main reason general equilibrium is so important in their model is because of household consumption smoothing motives. If lumpy investment at the firm level causes aggregate investment in capital to move differently than in a frictionless RBC model, the representative household would face more consumption volatility. That is because in their model, $Y = C + I_k$, so a large change in I_k necessitates a large change in C . Since households have a strong consumption smoothing motive, there are then large price movements in general equilibrium that undo the partial equilibrium effects of lumpy investment.

In contrast, in our model $Y = C + I_k + I_d$. In this environment, if lumpy durable adjustment induces aggregate dynamics for I_d that depart from the frictionless model, these changes can be absorbed by I_k without implying a more volatile consumption process. That is, with multiple sources of savings, large changes in the behavior of some component of savings do not necessarily imply that households must violate consumption smoothing. This is similar to the intuition in Bachmann and Ma [2013] who argue that the presence of inventories in a model with lumpy investment reduces the importance of general equilibrium effects.

6 Geographical Evidence

We now exploit geographic variation to provide additional reduced form evidence for our theoretical model with fixed costs of adjustment. The basic punchline of our structural model is that the elasticity of durable spending to economic shocks should be lower during recessions than during booms. We now exploit cross-sectional geographic variation to show that this is indeed the case.

In particular, we show that MSA-level auto spending responds much more to housing wealth shocks in locations experiencing local booms than in locations experiencing local recessions. Our basic empirical strategy identifies local wealth shocks using the empirical methodology of Mian and Sufi [2013], Mian and Sufi [2014b] and Mian and Sufi [2014a]. In particular, we instrument for local house price changes using the geography based measure of housing supply elasticity constructed by Saiz [2010]. Mian and Sufi [2013] show that this instrument has a strong first stage in explaining house price movements. Furthermore, they provide various evidence that this instrument plausibly satisfies the exclusion restriction by showing that it is uncorrelated with various other local economic variables, such as the change in wages over the housing boom and bust.

Our empirical strategy augments that in Mian and Sufi by showing that the response of MSA-level auto spending to housing wealth shocks strongly interacts with the state of the local business cycle. In particular, our benchmark empirical results use the following two stage least squares specification:

$$\begin{aligned}\Delta \log \text{AutoSales}_{i,t} &= \alpha^{IV} + \beta_1^{IV} \Delta \log \widehat{HP}_{i,t} \times \Delta U_{i,t} + \beta_2^{IV} \Delta \log \widehat{HP}_{i,t} + \beta_3^{IV} \Delta U_{i,t} + \Lambda X_i + \varepsilon_{i,t} \\ \Delta \log \widehat{HP}_{i,t} &= \omega + \eta_1 \text{Elasticity}_i \times \Delta U_{i,t} + \eta_2 \text{Elasticity}_i + \eta_3 \Delta U_{i,t} + \Psi X_i + \epsilon_{i,t} \\ \Delta \log \widehat{HP}_{i,t} \times \Delta U_{i,t} &= \psi + \lambda_1 \text{Elasticity}_i \times \Delta U_{i,t} + \lambda_2 \text{Elasticity}_i + \lambda_3 \Delta U_{i,t} + \Pi X_i + \xi_{i,t},\end{aligned}$$

where X_i is a vector of MSA-level controls including local employment shares and income.⁵⁴ We estimate this specification using annual data from $t = 2002, \dots, 2012$. See Appendix 1 for discussion of the underlying data construction.

The main coefficient of interest in the above specification is β_1^{IV} . Our structural model implies that this coefficient should be negative: auto spending in recession MSAs should respond less to a given wealth shock than in boom MSAs.

Table 3 shows that β_1^{IV} is indeed highly significant and negative.⁵⁵ Our benchmark specification in column (2) shows that a one-percentage point increase in an MSA's unemployment rate lowers the elasticity of its auto-spending to housing wealth shocks by 1.3 (relative to a median elasticity across MSAs of 3.2).⁵⁶

Columns (3)-(6) show a number of robustness checks. In column (3) we include year fixed effects in our regressions. Since these year fixed effects absorb the aggregate house price movements over the housing boom and bust, all identification then comes off of cross-MSA variation.⁵⁷ The interaction of these purely local housing wealth shocks with local unemployment changes is nearly unchanged. In column (4), we use the alternative instrument of Charles, Hurst, and Notowidigdo [2014] to identify exogenous housing price movements. This instrument uses only the "bubble" component of house price movements, under the identifying assumption that fundamental factors should move smoothly so that structural breaks in house prices should capture movements uncorrelated with fundamentals. Again, the importance of the interaction is unchanged.

⁵⁴These control for sector-specific shocks or local income shocks that may be correlated with our instrument for house prices. See Mian and Sufi [2013] for additional discussion. Removing these controls or adding additional controls such as local wages did not affect our results.

⁵⁵Regressions are weighted by MSA population and standard errors are clustered by state.

⁵⁶The median change in unemployment is 0, so that the elasticity equals β_2^{IV} for the median MSA.

⁵⁷We also find similar results when splitting the sample separately into the housing boom and bust.

Table 3: Response of Automobile Spending to Wealth Shocks

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	Saiz IV	Saiz IV + Year FE	Bubble IV	Saiz IV + Bartik IV	Saiz IV
Δ House Price X Δ U	-0.114* (0.063)	-1.314*** (0.373)	-1.174*** (0.353)	-1.286*** (0.382)	-1.429*** (0.414)	
Δ House Price	1.522*** (0.119)	3.206*** (0.769)	3.634*** (0.906)	3.184*** (0.617)	2.339 (2.367)	9.391*** (2.209)
Δ U	-0.020 (0.012)	-0.056*** (0.012)	0.024 (0.023)	-0.054*** (0.018)	-0.085* (0.051)	
Δ House Price X U						-1.034*** (0.360)
Unemployment Rate						-0.038* (0.020)
Construction Share 2002	-0.140 (0.782)	-0.835 (0.544)	-0.776 (0.545)	-0.820* (0.440)	-0.883 (0.602)	-0.786* (0.474)
Non Tradeable Share 2002	0.380** (0.187)	0.644*** (0.177)	0.683*** (0.179)	0.639*** (0.207)	0.602*** (0.209)	0.494** (0.214)
Ln (AGI / Capita) 2002	0.036** (0.017)	0.003 (0.022)	0.004 (0.022)	0.004 (0.016)	0.001 (0.022)	0.039** (0.016)
N	2237	2237	2237	2237	2237	2237

Standard errors in parentheses, clustered at state-level. Results weighted by population. Sample period 2002-2012.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Our benchmark specification treats unemployment movements as exogenous. This is clearly a strong assumption, and there are reasons to believe that both local unemployment and local house prices should be subject to endogeneity concerns. To address this, in Column (5), we instrument for local house price movements using the Saiz instrument and instrument for local unemployment movements using the popular Bartik [1991] instrument. This IV strategy interacts the preexisting local composition of manufacturing with nationwide manufacturing employment changes to construct exogenous movements in unemployment, under the assumption that the local composition of the manufacturing sector is predetermined at the time of the employment shock and that aggregate employment shocks are exogenous to individual MSAs. Column (5) shows that auto sales in locations with high predicted increases in unemployment again respond less to wealth shocks.

Finally, Column (6) uses an alternative measure of local business cycles. The procyclical responsiveness in our structural model is driven by temporary business cycle shocks that

move the distribution of households' durable gaps away from steady-state. Our model does not predict that permanently moving to a more prosperous steady-state should lead to a permanent increase in the responsiveness of durable spending to shocks. While differences in unemployment across MSAs are not permanent, they are highly persistent. For this reason, our benchmark empirical specification uses the change in unemployment as our measure of local economic conditions. Nevertheless, Column (6) repeats our empirical exercise using the level of MSA unemployment rather than the change in unemployment. We again find a significant interaction effect, although as predicted, it is slightly less strong than under our benchmark specification.

In addition to this cross-sectional evidence, Appendix 5 uses time-series data on durable spending to provide further support for our theoretical mechanism. In particular, we show that the model with fixed costs of durable adjustment better fits business cycle moments than either frictionless models or models with convex adjustment costs. The model with fixed costs of adjustment implies that aggregate durable spending should exhibit greater volatility during booms, when responsiveness is high, than during recessions. Alternative models make no such prediction. We show that time-series data strongly confirms the prediction of the model with fixed costs of adjustment. In addition, Berger and Vavra [2014] provide additional time-series evidence for procyclical durable spending responsiveness. In that paper, we use an STVAR to estimate the durable spending multiplier in response to identified government spending changes and show that it rises dramatically during expansions.

7 Conclusion

In this paper, we argue that household-level durable adjustment frictions matter for aggregate dynamics. We use a novel indirect inference procedure to estimate an incomplete markets model with fixed costs of durable adjustment and show that it does a very good job of explaining various microeconomic consumption patterns. More importantly, this model implies that as household wealth and income falls during a recession, fewer households adjust their durable holdings or are on the margin of doing so. This means that the elasticity of aggregate durable expenditures to shocks which affect durable demand falls substantially.

We provide support for this mechanism in various ways. In addition to showing that the average frequency of durable adjustment indeed falls in recessions, we show that cross-sectional distributions in the PSID data move as predicted by our model and that MSA-level auto spending is less responsive to wealth shocks that occur during local recessions.

Our results have implications for estimating the efficacy of durable stimulus. The response of durable expenditures to changes in policy is highly dependent on the aggregate

state of the economy, which means that using estimates from linear VARs to estimate the effects of any such program is likely to be misleading. While a growing body of research argues that the government spending multiplier should be countercyclical,⁵⁸ the forces we identify push in the opposite direction. In Berger and Vavra [2014] we provide STVAR evidence that while the overall government spending multiplier is indeed countercyclical, the "durable spending" multiplier is instead procyclical, just as predicted by our model.

In this paper, we emphasized the general mechanism that causes micro adjustment frictions to lead to procyclical IRFs. As such, we considered a variety of aggregate shocks and explored a very broad definition of durables that should include all durable goods subject to transaction costs of adjustment. In future work, we plan to explore our model implications for particular policies such as the "Cash-for-Clunkers" program or the "First-Time-Homebuyer" credit. Realistic policy analysis will require enriching our model to include various institutional details that are beyond the scope of this paper. Nevertheless, our modeling insights should continue to apply and we hope to quantify their importance for specific policies. More generally, understanding durable spending patterns requires understanding the level and distribution of wealth in an economy and how this distribution moves across time.

⁵⁸Due to, e.g., excess capacity or to the ZLB.

References

- ATTANASIO, O. (2000): “Consumer Durables and Inertial Behavior: Estimation and Aggregation of (S,s) Rules for Automobile Purchases,” *Review of Economic Studies*.
- BACHMANN, R., R. J. CABALLERO, AND E. M. ENGEL (2013): “Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model,” *American Economic Journal: Macroeconomics*.
- BACHMANN, R., AND L. MA (2013): “Lumpy Investment, Lumpy Inventories,” *Mimeo*.
- BAJARI, P., P. CHAN, D. KRUEGER, AND D. MILLER (2013): “A Dynamic Model of Housing Demand: Estimation and Policy Implications,” *International Economic Review*.
- BAR-ILAN, A., AND A. BLINDER (1992): “Consumer Durables: Evidence on the Optimality of Usually Doing Nothing,” *Journal of Money, Credit and Banking*, 24(2).
- BARTIK, T. (1991): “Who Benefits From State and Local Economic Development Policies?,” *Upjohn Press*.
- BERGER, D., AND J. VAVRA (2014): “Measuring How Fiscal Shocks Affect Durable Spending in Recessions and Expansions,” *American Economic Review: Papers and Proceedings*.
- BERNANKE, B. (1985): “Adjustment Costs, Durables and Aggregate Consumption,” *Journal of Monetary Economics*, 15(1).
- BERTOLA, G., AND R. J. CABALLERO (1990): “Kinked Adjustment Costs and Aggregate Dynamics,” *NBER Macroeconomics Annual*, 5.
- BERTOLA, G., L. GUISO, AND L. PISTAFERRI (2005): “Uncertainty and Consumer Durables Adjustment,” *Review of Economic Studies*, 72(4).
- BROWNING, M., AND T. CROSSLEY (2009): “Shocks, Stocks, and Socks: Smoothing Consumption Over a Temporary Income Loss,” *Journal of the European Economic Association*, 7(6).
- CABALLERO, R., E. ENGEL, AND J. HALTIWANGER (1995): “Plant-Level Adjustment and Aggregate Investment Dynamics,” *Brookings Papers on Economic Activity*, 1995(2).
- (1997): “Aggregate Employment Dynamics: Building from Microeconomics,” *American Economic Review*, 87(1).

- CABALLERO, R. J. (1990): “Expenditure on Durable Goods: A Case for Slow Adjustment,” *The Quarterly Journal of Economics*, 105(3).
- (1993): “Durable Goods: An Explanation for Their Slow Adjustment,” *The Journal of Political Economy*, 101(2).
- CABALLERO, R. J., AND E. M. ENGEL (2007): “Price stickiness in Ss models: New interpretations of old results,” *Journal of Monetary Economics*, 54(Supp 1).
- CHARLES, K., E. HURST, AND M. NOTOWIDIGDO (2014): “Masking and Aggregate Employment Changes: Housing Booms and Manufacturing Decline During the 2000s,” *Mimeo*.
- DIAZ, A., AND M. LUENGO-PRADO (2010): “The Wealth Distribution with Durable Goods,” *International Economic Review*, 51(1).
- DUNN, W. (1998): “Unemployment Risk, Precautionary Saving, and Durable Goods Purchase Decisions,” .
- EBERLY, J. C. (1994): “Adjustment of Consumers’ Durables Stocks: Evidence from Automobile Purchases,” *Journal of Political Economy*, 102(3).
- GROSSMAN, S., AND G. LAROQUE (1990): “Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods,” *Econometrica*, 58(1).
- IACOVIELLO, M., AND M. PAVAN (2013): “Housing and Debt Over the Life Cycle and Over the Business Cycle,” *Journal of Monetary Economics*.
- KAPLAN, G., AND G. L. VIOLANTE (2014): “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*.
- KHAN, A., AND J. THOMAS (2008): “Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics,” *Econometrica*, 76(2).
- KRUEGER, D., AND J. FERNANDEZ-VILLAYERDE (2010): “Consumption and Saving over the Life Cycle: How Important are Consumer Durables?,” *Macroeconomic Dynamics*.
- KRUSELL, P., AND A. A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *The Journal of Political Economy*, 106(5).
- LEAHY, J. V., AND J. ZEIRA (2005): “The Timing of Purchases and Aggregate Fluctuations,” *Review of Economic Studies*, 72.

- LEAMER, E. E. (2007): “Housing IS the Business Cycle,” *Proceedings, Federal Reserve Bank of Kansas City*.
- LUENGO-PRADO, M. (2006): “Durables, Nondurables, Down Payments and Consumption Excesses,” *Journal of Monetary Economics*, 53(1).
- MANKIW, N. G. (1982): “Hall’s Consumption Hypothesis and Durable Goods,” *Journal of Monetary Economics*, 10(3).
- MIAN, A., AND A. SUFI (2012): “The Effects of Fiscal Stimulus: Evidence from the 2009 ‘Cash for Clunkers’ Program,” *Quarterly Journal of Economics*, 127(3).
- (2013): “Household Balance Sheets, Consumption, and the Economic Slump,” *Quarterly Journal of Economics*.
- (2014a): “House Price Gains and U.S. Household Spending from 2002 to 2006,” *Mimeo*.
- (2014b): “What Explains the 2007-2009 Drop in Employment?,” *Econometrica*.
- PIAZZESI, M., AND M. SCHNEIDER (2007): “Asset Prices and Asset Quantities,” *Journal of the European Economic Association*.
- SAIZ, A. (2010): “The Geographic Determinants of Housing Supply,” *Quarterly Journal of Economics*.

8 Appendix 1: Data Definitions and Cleaning

8.1 PSID Data

This appendix discusses provides additional discussion of our PSID data analysis. We restrict our analysis to 1999-2011 because prior to 1999, the PSID did not collect the data necessary for our analysis. Beginning in 1999, the PSID contains detailed information on non-durable consumption, the value of housing and vehicles as well as various wealth holdings. Although more detailed non-durable consumption data is available beginning in 2003, for comparability we use only variables that are available beginning in 1999. The value for non-durable expenditures is the sum of all components of food consumption, utilities, transportation expenses, schooling expenses and health services. Our measure of d_{-1} is the sum of last periods housing value and vehicle values. Assets are the sum of business value, stocks, iras, cash, bonds, minus the value of outstanding debt.

Our benchmark analysis is restricted to home-owners with household head $<$ age 65. After constructing each of our variables, we deflate these nominal values using NIPA price indices and remove a household fixed effect. To define durable adjustment, we combine several questions in the PSID. In our benchmark results, we define durable adjustment as a self-reported house or vehicle sale together with a 20% change in the reported value of the durable stock. We use a combination of self-reported adjustment and a minimum threshold for several reasons. 1) Combining these indicators is likely to reduce spurious adjustments due to measurement error. 2) Some house sales are likely to be the results of idiosyncratic moves across location which may not lead to any substantial adjustment in the size of the stock. 3) Finally, and most importantly, self-reported adjustment indicators ask about adjustment over the previous three years while the sample is conducted every two years. This implies that the same adjustment may be counted twice. Requiring a simultaneous change in value and self-reported adjustment reduces this concern. We chose a 25% threshold because the median change in the reported durable stock conditional on self-reported adjustment is 40% while the median change conditional on no adjustment is 4%, so a 20% threshold roughly splits this distance. This adjustment definition generates an adjustment probability of roughly 10%.

Figure 1 applies this adjustment definition to our broad measure of durables that includes housing and vehicles as well as to a more narrow definition that focuses just on housing after removing deterministic age effects. We split the sample in 1999 for the housing series because the sampling frequency and thus questions change slightly and there appears to be a trend break in the series. Table A1 reports results a for a panel logit for the probability

of adjustment on recession indicators. Overall, the probability of broad durable adjustment falls by around 20% during recessions while the probability of buying/selling a house falls by around 15%.

Table A1

Outcome	Sample Period	Odds Ratio	Std. Err.	#obs	#households	Age Controls
Sold (Broad d)	1999-2011	0.78***	0.074	5316	1460	NO
	1999-2011	0.84**	0.078	5316	1460	YES
Sold (House)	1969-1999	0.88***	0.035	76851	8954	NO
	1969-1999	0.85***	0.033	76851	8954	YES

Since Figure 1 shows that for the 1999-2011 there is an overall downward trend in the frequency of durable adjustment, there is some concern that the broad durable panel regressions are capturing time-trends rather than something about recessions. To argue that this is not the case, we perform a similar exercise using cross-state variation in unemployment rates. Running a logit of broad durable adjustment on local unemployment rates shows that a two standard deviation in unemployment lowers the odds of durable adjustment by 30-40%. This statistically significant decline in the frequency of durable adjustment is robust to a variety of location, time and age controls.

Table A2

Odds Ratio (1-Std Effect)	Std. Err.	Year FE	State FE	Age Controls
0.85***	0.05	YES	YES	YES
0.79***	0.03	YES	NO	YES
0.85***	0.02	NO	YES	YES
0.84***	0.02	NO	NO	YES
0.85***	0.05	YES	YES	NO
0.79***	0.03	YES	NO	NO
0.81***	0.02	NO	YES	NO
0.80***	0.02	NO	NO	NO

8.2 Robustness of PSID Cross-Sectional Results to Alternative Data Cleaning Procedures

While we believe that our benchmark empirical specification is reasonable, we also assess the robustness of our results to alternative choices. To explore this, we apply the model point

estimates $x_{benchmark}^d \widehat{=} G^m(z_{benchmark}^d \widehat{=})$ to observables computed under various alternative assumptions: $x_{alternative}^d \widehat{=} G^m(z_{alternative}^d \widehat{=})$. That is, we use the model point estimates computed from our benchmark data definitions and apply them to alternative data definitions. While it would be desirable to reestimate the entire model under different data assumptions, this is numerically infeasible. The main empirical object of interest is whether the slope of the empirical adjustment hazard as a function of (absolute) imputed durable gaps is upward sloping. Towards that end, Table 8 displays the results of a regression of the probability of adjustment on the absolute value of the durable gap for a range of alternative empirical specifications.

$$adjust_{i,t} = \alpha + \beta \left| x_{alternative,i,t}^d \widehat{=} \right|.$$

Table A3 reports results for a number of robustness checks:

Specification	β	$t - stat$
Benchmark	0.67	57.7
(1) Adj Threshold of 0.1 instead of 0.2	1.14	69.5
(2) Adj Threshold of 0.3 instead of 0.2	0.19	24.2
(3) No Adj Threshold	0.76	25.0
(4) 0.2 Threshold, Ignore self-reported adj	0.40	49.6
(5) Control for Year Fixed Effects in HH estimation	0.70	62.1
(6) No adjustment for HH size	0.60	51.6
(7) Exclude business value from a	0.67	57.3
(8) Keep only ages 25-55	0.74	45.9
(9) Do not use price deflators	0.71	60.2
(10) Do not adjust for age effects	0.65	54.5

8.3 Aggregate Data

The top panel of figure 2 is constructed using proprietary data from the CNW auto market research firm. They collect data on both new and used auto sales across time. To construct turnover rates, we merge this sales data with data on total registered vehicles from the DOT ([http://www.rita.dot.gov/bts/sites/rita.dot.gov.bts/files/publications/](http://www.rita.dot.gov/bts/sites/rita.dot.gov.bts/files/publications/national_transportation_statistics/html/table_01_11.html)

[national_transportation_statistics/html/table_01_11.html](http://www.rita.dot.gov/bts/sites/rita.dot.gov.bts/files/publications/national_transportation_statistics/html/table_01_11.html)). This measure of the vehicle stock is available annually beginning in 1990 and is available every five years before 1990. To construct annual measures of vehicle registration before 1990 we use a perpetual inventory method. For example, we observe the stock of vehicles in 1985 and 1990 and we observe total purchases in each year from 1985-1990. We assume that there is a constant

depreciation rate over each five year period and we pick this depreciation rate so that the beginning and ending stock is consistent with annual purchases.

The bottom panel of Figure , we merge data from HUD and the census. HUD reports data on existing home sales from 1969-2008:

(http://www.huduser.org/periodicals/ushmc/fall09/hist_data.pdf), which we merge with data from the national association of realtors for recent years:

(<https://research.stlouisfed.org/fred2/series/EXHOSLUSA495S/downloaddata>), and the census reports data on total housing stocks:

(<http://www.census.gov/housing/hvs/data/histtab7.xls>) as well as on new house purchases: (<https://research.stlouisfed.org/fred2/series/HSN1F>).

In our time-series analysis in Appendix 5, we define durable expenditures as real consumer durable expenditures + real residential investment where real consumer durables are NIPA Table 1.1.5 line 4 divided by NIPA Table 1.1.9 line 4 and real residential investment is NIPA table 1.1.5 line 12 divided by NIPA Table 1.1.9 line 12. Non-durable consumption is defined as non-durable goods (NIPA Table 1.1.5 line 5 divided by NIPA Table 1.1.9 line 5) + services (Table 1.1.5 line 6 divided by Table 1.1.9 line 6) - housing services (Table 2.3.5 line 14 divided by Table 2.4.4 line 14). Our measure of GDP is then the sum of non-durable consumption, durable expenditures and private non-residential investment..

Constructing durable investment rates requires quarterly measures of the durable stock. Following Bachmann, Caballero, and Engel [2013], we construct measures of real annual durable stocks using nominal data from BEA Domestic Product and Income Tables 1.1.5 and price deflators from Table 1.1.9. We then next construct quarterly depreciation estimates using annual nominal measures of depreciation from BEA Fixed Asset Table 1.1 together with the price deflators from Table 1.1.9. Since the BEA publishes annual measures of the stock of durables and housing in Fixed Asset Table 1.1, we just need to construct quarterly measures in between these annual observations. To do this, we combine the annual observations with the quarterly expenditure and depreciation measures together with a standard stock accumulation expression to construct quarterly stock measures. See Bachmann, Caballero, and Engel [2013] for the more detailed procedure.

8.4 Geographic Data

We describe here the data used in Section 6. Our auto sales data is produced by R.L. Polk. These zip-code level monthly auto sales numbers are proprietary and cannot be shared without data provider approval. Please contact robert_sacka@polk.com for access to the underlying data. We aggregate these zip-code level auto sales to MSAs and years,

excluding the zip codes 74153 and 74117. All auto registrations by one large national auto rental company occur in these two zip codes, so they are both large outliers and poor measures of local auto demand. Following Mian and Sufi [2012], we compute auto sales relative to steady-state sales, as measured by average sales over 1990-1999.

Our house price data is constructed by CoreLogic and is also proprietary. Please contact rmeyers@corelogic.com. We use two instruments to measure exogenous house price movements in our empirical specification. The first instrument is the geography based instrument of Saiz [2010]. This instrument measures the difficulty of constructing new housing in an MSA due to the presence of geographic features such as mountains and oceans. This data is available at <http://real.wharton.upenn.edu/~saiz/>. In addition, we use the bubble IV instrument from Charles, Hurst, and Notowidigdo [2014]. This instrument is constructed by searching for the presence of structural breaks in house price growth from 2002-2006. We reproduce their procedure exactly, so see their paper for additional details.⁵⁹

We construct the Bartik [1991] IV using data from County Business Patterns. For each MSA i , we compute the share of employment ω_{jt} in each 4-digit manufacturing sector j and year t . Let Δe_{jt} be aggregate U.S. manufacturing employment growth. Predicted manufacturing employment in a given MSA is then $\widehat{\Delta e_{it}} = \sum_j \omega_{jt-1} \Delta e_{jt}$, and we use this instrument to predict an exogenous component of MSA unemployment.

Our data for local employment share controls also comes from County Business Patterns. We classify industries as construction, and non-tradeable using the definitions in Mian and Sufi [2014b]. Local income data comes from the IRS Statistics of Income.

⁵⁹We thank Matt Notowidigdo for providing us the code to construct this instrument.

9 Appendix 2: Estimation and Identification

9.1 Estimation Algorithm and Construction of Standard Errors

In this section we provide additional details on our estimation algorithm, the particular functional form we choose for G^m and the construction of standard errors for point estimates and model targets.

A key requirement for our estimation algorithm is the construction of a function and set of observables that solves: $x^m = G^m(z)$. First, note that as long as $d_{-1} \in z$ then predicting x^m is equivalent to predicting d^* since $x = \log d^* - \log(d_{-1})$. Since we will assume that d_{-1} is in the econometrician's information set, we thus change notation and search for a function $d^* = G^m(z)$. The most straightforward way to construct $G^m(z)$ is to make z equal to the agents' (empirically observable) state-variables and then G^m will be equal to the policy function that solves $d^*(a_{-1}, d_{-1}, \eta) = \arg \max V^{adjust}(a_{-1}, d_{-1}, \eta)$. While this function would exactly map observable states to choices in a world with perfect data, it is problematic in a world with measurement error. Since this function is non-linear, even if measurement error is on average zero it need not produce estimates which are on average correct. That is, $E[d^*(a_{-1}, d_{-1}, \eta)] \neq d^*(Ea_{-1}, Ed_{-1}, E\eta)$: if state-variables are measured with noise and used as inputs to a highly non-linear function, then the resulting estimate for d^* need not be an unbiased estimate of the truth. Given this concern, we instead approximate G by a linear function of various observable variables.

In particular, we assume that $G^m(z) = \beta_0 + \beta_1 a_{-1} + \beta_2 d_{-1} + \beta_3 c + \beta_4 \frac{d_{-1}}{c}$. In practice, this functional form is highly accurate along various different metrics. First, we can ask how well this linear function does at predicting actual model d^* when there is no measurement error. That is, given the true values for z , how well do we predict the actual d^* in the model? Running our regression of d^* on z delivers an $R^2 = 0.98$. Thus, we do not perfectly match model gaps in an environment with no measurement error, but this function does an extremely good job. (Note that using the true policy function would by construction deliver an $R^2 = 1$). We can next ask how well we do at matching the true d^* implied by z if we instead use a noisy measure \hat{z} as the input to our function. That is, how well does our model do when regressing $\hat{d}^* = \beta_0 + \beta_1 \hat{a}_{-1} + \beta_2 \hat{d}_{-1} + \beta_3 \hat{c} + \beta_4 \frac{\hat{d}_{-1}}{c}$ on actual d^* . Overall we find an $R^2 = .85$, so even with noisily measured inputs, we are able to well predict actual durable gaps in the model. In addition we find that $E\hat{d}^* = d^*$. In contrast, if we apply the model's true policy function to noisily measured state-variables then we find an $R^2 = 0.60$ and we also find that $E\hat{d}^* \neq d^*$. That is, imposing a simple linear relationship between inputs with measurement error and outputs does a more successful job of producing the true model d^*

then does imposing actual model policy functions on these mismeasured inputs.⁶⁰

Thus, the G^m that we choose performs quite well. Nevertheless, we have explored various alternative functional forms including adding $\frac{d_{-1}}{a_{-1}+d_{-1}}$, y , and $\frac{y}{d}$ as additional predictors. However, none of these alternatives provide much additional predictive power in the environment with no measurement error and they perform less well in an environments where variables are measured with noise. Since many households hold no liquid assets, introducing $\frac{d_{-1}}{a_{-1}+d_{-1}}$ as a predictor introduces collinearity issues when identifying the effect of d_{-1} . We use information on consumption rather than earnings in our baseline specification because earnings is more frequently missing as households have spells of unemployment. Simulations in the model suggest that we lose essentially no accuracy by using c instead of y when there is no measurement error, and we gain substantial additional predictive power when there is random missing data on y . Using ratios of $\frac{y}{d}$ is again problematic for households with no earnings. Nevertheless, while these alternative functions seem to do slightly less well in simulation at predicting true gaps in the model, we have re-run our benchmark model using various alternative specifications and arrived at quite similar results implications.

Armed with our function $G^m(z) = \beta_0 + \beta_1 a_{-1} + \beta_2 d_{-1} + \beta_3 c + \beta_4 \frac{d_{-1}}{c}$, we now describe in additional detail our estimation algorithm and construction of standard errors.

1) For a given set of parameters p , solve the model and regress $d^{*m} = \beta_0 + \beta_1 (a_{m-1}) + \beta_2 (d_{-1}^m) + \beta_3 c^m + \beta_4 \left(\frac{d_{-1}^m}{c}\right)$.

2) Given a measurement error parameter, simulate the model using sample sizes equal to PSID, with measurement error and aggregate this simulated data to biannual frequencies. Then compute estimates of gaps in the model: $\widehat{d^{*m}} = \beta_0 + \beta_1 (\widehat{a_{-1}^m}) + \beta_2 (\widehat{d_{-1}^m}) + \beta_3 \widehat{c^m} + \beta_4 \left(\frac{\widehat{d_{-1}^m}}{c}\right)$. In this step (and for identification of measurement error) it is important to note that we estimate the vector β using the true model relationship with no measurement error and then apply that function to observables with simulated measurement error. This implies that as we change the degree of measurement error, we do not change β and change only the relationship between z and \widehat{z} .

3) Compute $\widehat{d^{*d}}$ in PSID data: $\widehat{d^{*d}} = \beta_0 + \beta_1 (\widehat{a_{-1}^d}) + \beta_2 (\widehat{d_{-1}^d}) + \beta_3 \widehat{c^d} + \beta_4 \left(\frac{\widehat{d_{-1}^d}}{c}\right)$.

4) Convert estimates of d^* to measures of gaps: $\widehat{x^m} = \log \widehat{d^{*m}} - \log \widehat{d_{-1}^m}$ and $\widehat{x^d} = \log \widehat{d^{*d}} - \log \widehat{d_{-1}^d}$.

5) Compute the density of gaps in the model and data f_p , and calculate the probability of adjustment as a function of imputed gap h_p^m (using the threshold for adjustment defined

⁶⁰We can also assess the accuracy directly in data by looking at the relationship between actual durable decisions when adjusting and those predicted by G . Again, the simple functional form is substantially more accurate.

in Appendix 1).

6) Compute $L_p = \int \left[\left(f_p^m(\widehat{x}^m) - f^d(\widehat{x}^d) \right)^2 + \left(h_p^m(\widehat{x}^m) - h^d(\widehat{x}^d) \right)^2 \right] dx$

7) Repeat 1-6 over parameters to minimize L_p .

Steps 1-7 describe the procedure for constructing point estimates. To construct bootstrapped standard errors we do the following:

8) Given our parameter point estimates, simulate data as in step 2) above.

9) Replace the PSID data in the previous estimation with the "fake data" simulated from the model best fit point estimates.

10) Re-estimate the model to find point estimates which best fit the new "fake" PSID data.

11) Record point estimates and implied hazards for this bootstrap replication. Also record the implied distribution of gaps and hazards in the actual PSID data under this new point estimate.

12) Repeat 8-12 1000 times to construct a distribution of bootstrapped standard errors that accounts for sampling error.

9.2 Identification

How are parameters identified in our model? As usual in numerical models, we have no proof of global identification, but in practice we parameterize our hazard and density using 21 bins for each. This means that we have 42 targeted moments and only five parameters so that the model appears to be overidentified. In addition, starting the search for the best fit parameters at various starting values converges to the same best-fit results, which suggests that the model is globally identified.

Furthermore, we can argue more strongly for local identification by varying individual parameters holding others fixed at their best fit estimates. Each of our parameters induces independent variation on the model and data densities and hazards. In the following figures, we change one parameter at a time and explore its implications for model and data densities and hazards. In each figure, the blue lines correspond to model objects and green lines to PSID objects.

Figure 15: Changing v (Relative Utility of Non-Durables)

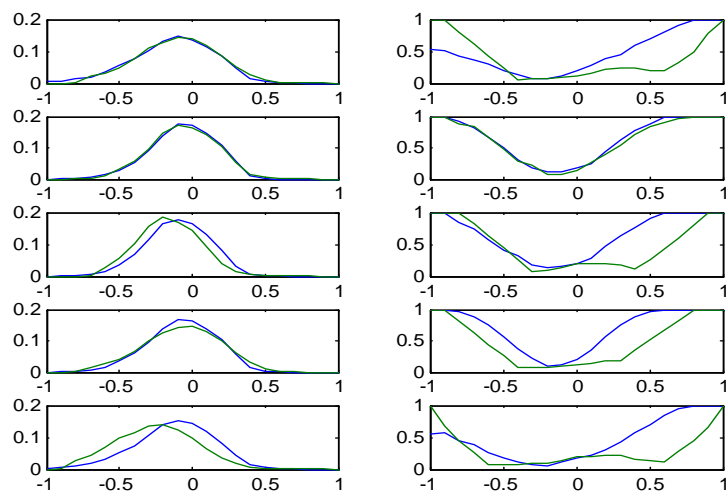


Figure 16: Changing χ (Maintenance)




Figure 17: Changing F_t (Time cost of adjustment)




Figure 18: Changing F_d (Fixed Cost Proportional to Stock)


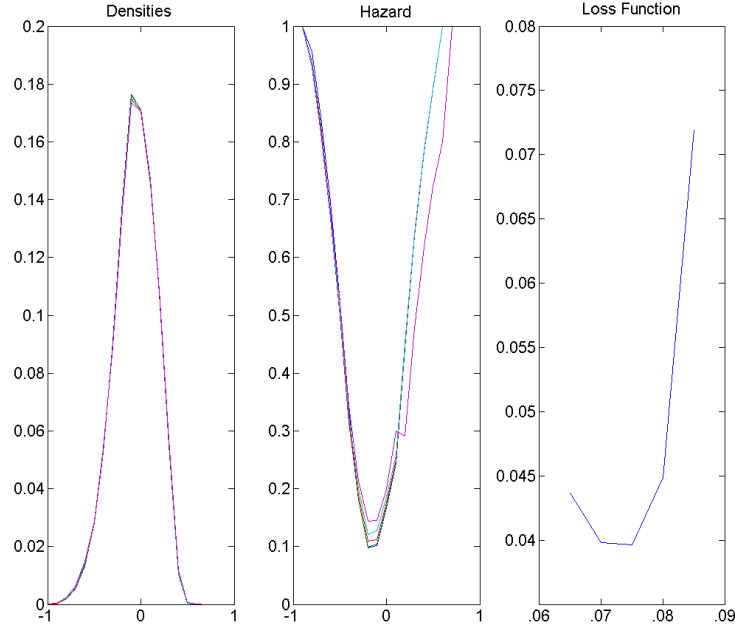


Figure 19: Changing Proportional Measurement Error



Changes in the first four parameters induce changes in both model and empirical densities. This is because as we change these parameters we alter the function G^m that maps observables to gaps in the model. This in turn induces variation in PSID gaps: $\hat{x}^d = G^m(\hat{z}^d)$. In contrast, introducing measurement error affects $\hat{z}^m = (1 + \hat{\epsilon})z^m$ but it does not affect the mapping between true model observables (with no measurement error) and outcomes. That is, measurement error does not affect G^m . As such, changing the degree of measurement error has no effect on the gaps and hazard imputed in PSID and only affects the gaps and hazards imputed for model simulated data (with measurement error). Since changing the measurement error parameter has no effect on PSID we plot its affect only on model densities and hazards. In addition, we show how the loss function varies as we change measurement error to demonstrate that there is a clear minimum.

However, it is worth noting that the variation in densities and hazards induced by measurement error is substantially less than that induced by the other parameters of our model, so if any parameter is not particularly well-identified it is probably the amount of measurement error. However, this is not a huge concern for our results as we are not particularly interested in assessing the amount of measurement error in PSID variables. Furthermore, measurement error does not really affect any of our conclusions aside from their affect on the overall fit and thus on the other parameters of their model. True impulse responses

in the model are calculated without measurement error, so these IRFs are not affected by changes in measurement error. Since the measurement error parameter does not affect the actual PSID data, changing its importance has no effect on the conclusions for the implied IRF given by the PSID cross-section.

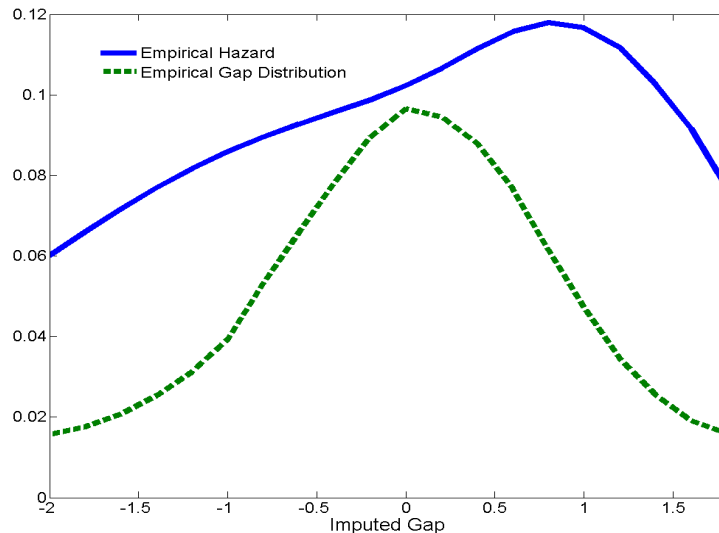
Finally, we can also construct surface plots for the loss-function as two different parameters are varied simultaneously. For brevity we do not report these plots, but again the model appears to be globally identified.

9.3 Applications to Alternative Models and Data

In this section, we explore two additional aspects of our estimation. 1) Can our estimation procedure identify misspecified models? 2) Does our model deliver useful out-of-sample predictions?

9.3.1 Grossman and Laroque 1990 Model

Figure 20: Empirical Results for Grossman Laroque (1990) Model



First, to what extent is the empirical hazard actually a test for misspecification in our model? Since we target the density and hazard directly, perhaps it is just mechanical that we find a good fit for these variables. To assess this, we apply our gap imputation procedure to PSID data using the model from Grossman and Laroque [1990]. In this model, households target a constant fraction of liquid wealth when adjusting, so it is straightforward to impute

durable gaps. To what extent do these durable gaps provide predictive power for actual adjustment patterns? Figure 20 shows that the answer is: not at all. Gaps imputed using the structural model of 20 generate a nearly flat hazard (varying from 0.06 to 0.12) but more importantly, the hazard is not upward sloping in the absolute gap. Households which are predicted to adjust by the model are actually less likely to adjust.

Clearly the model of 20 is highly stylized since it ignores liquidity constraints and has no non-durable consumption, and clearly these things matter empirically. The point of this section is not that 20 is a bad model but is instead to illustrate that using an (S,s) model with a *theoretical* upward sloping hazard to impute gaps does not imply that the resulting *empirical* hazard need be upward sloping. In this sense, the behavior of the empirical hazard implied by the structural model is indeed a good test for misspecification of the structural model.

9.3.2 SHIW Data

Our model was estimated to match PSID data, and we showed that we can find parameters so that the model is a good fit to the behavior of households in this sample. In this sense, our extremely strong fit is constructed "in-sample". If we apply our same model estimates to data on which our model is not directly estimated, do we continue to find strong predictive power? In this section we argue that the answer is yes.

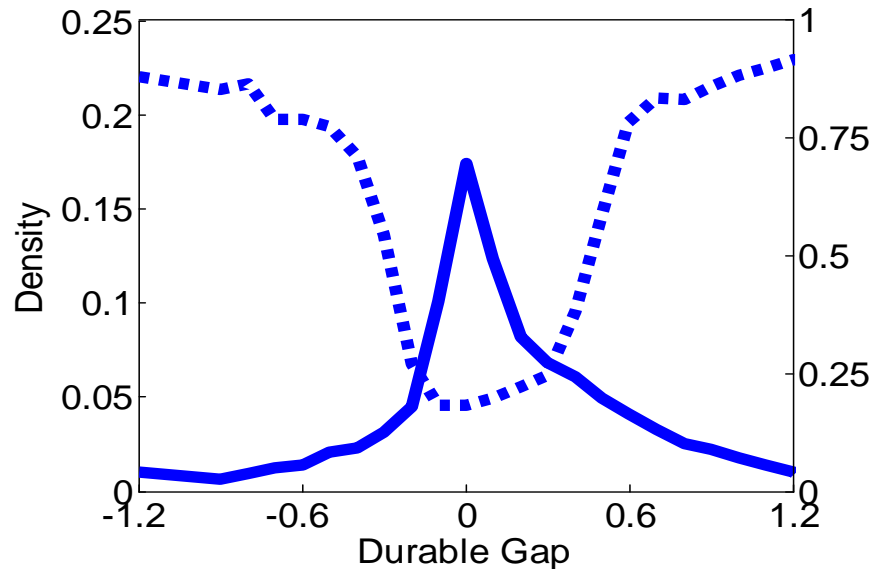
To our knowledge the Bank of Italy Survey of Household Income and Wealth (SHIW) data is the only other data set that exists with the necessary variables to apply our model estimates.⁶¹ The SHIW collects detailed information on demographics, households consumption and assets.⁶² Following Bertola, Guiso, and Pistaferri [2005], we only use the waves after 1987 as the survey methodology has remained roughly constant over this time period. In particular, we use the 1989, 1991, 1993, 1995, 1998, 2000, 2002, 2004, 2006, 2008, 2010 and 2012 waves in our analysis. Each wave surveys a representative sample of 8000 Italian households. We focus our analysis on head of households. The value of non-durable consumption in the data is defined as the sum of expenditure on apparel, schooling, entertainment, food, medical expenses, housing repairs and additions and imputed rents. Our preferred measure of the durable stock is the sum of end-of-period value of means of transport (includes autos, motorcycles, caravans, boats and bicycles) and the value of real estate (housing and land). Our results are robust to including other measures of durable adjustment including the value of end-of-period stocks for furniture and jewelry. The SHIW also includes information on

⁶¹An alternative out-of-sample test would be to estimate the model on half of the PSID data and test it on the other half of the PSID data. This produces similar results.

⁶²The dataset can be downloaded here: https://www.bancaditalia.it/statistiche/indcamp/bilfait;internal&action=__setlanguage.action?LANGUAGE=en

durable flows for means of transport, furniture, and jewelry. Net assets are the sum of all deposits, CDs, securities, businesses and valuables minus the value of all liabilities to banks, corporations and other households.

Figure 21: Imputed Gap and Hazard in SHIW data



Next, we impose the same structural relationship from on the model on this data (as we did in the PSID) to generate empirical measures of the empirical durable gap. Given estimates of these gap, we can then calculate the probability of adjustment as a function of the durable gap. Unfortunately, unlike the PSID, there are fewer variables explicitly asking households about their durable adjustment, so we must define adjustment purely in terms of some adjustment threshold. We define durable adjustment as times when the household either had non-zero expenditure in a period on means of transport or a 40% change in the reported value of real estate. The results are qualitatively robust (the hazard rate is increasing in the durable gap) to using different minimum thresholds including the 25% threshold we used in our benchmark specification in the PSID. The main difference is that a 25% threshold implies that the annual frequency of adjustment is approximately 30%, whereas a 40% threshold implies an annual frequency of adjustment closer to 15%.

Figure 21 shows the distribution of gaps and hazards that arise in SHIW data when applying our model which is estimated on PSID data. Clearly, we continue to find an extremely strong upward sloping hazard. In addition, the gap distribution continues to have negative skewness.

In addition to computing the average distribution and hazard, we can also redo the

exercise in Section 4.2 using SHIW data. Figure 22 is the counterpart to Figure 13 in the SHIW. Since there is no comparable business cycle dating committee for Italy, we plot the implied impulse response on impact from SHIW against the Italian unemployment rate. Just as in PSID data, the implied IRF is strongly procyclical. As unemployment falls, the IRF increases substantially.

