Consumption Dynamics During Recessions

David Berger   Joseph Vavra[*]
Northwestern University   University of Chicago

10/12/2012

Abstract

When will durable expenditures respond strongly to economic stimulus? We build from micro evidence and show that in a heterogeneous agent business cycle model with fixed costs of durable adjustment, responsiveness is substantially dampened during recessions. Our model estimates imply that during the Great Recession, durable expenditures were half as responsive to additional stimulus as during the boom in the 1990s. This procyclical responsiveness is driven by changes in the distribution of households' desired durable holdings over the cycle. We directly test our model by estimating this distribution empirically using PSID micro data and find that the distribution in the data moves cyclically as predicted. In addition to this micro evidence, we also provide support for our model's procyclical responsiveness using aggregate time-series data, and we show this time-series evidence is inconsistent with simpler models featuring smooth adjustment costs
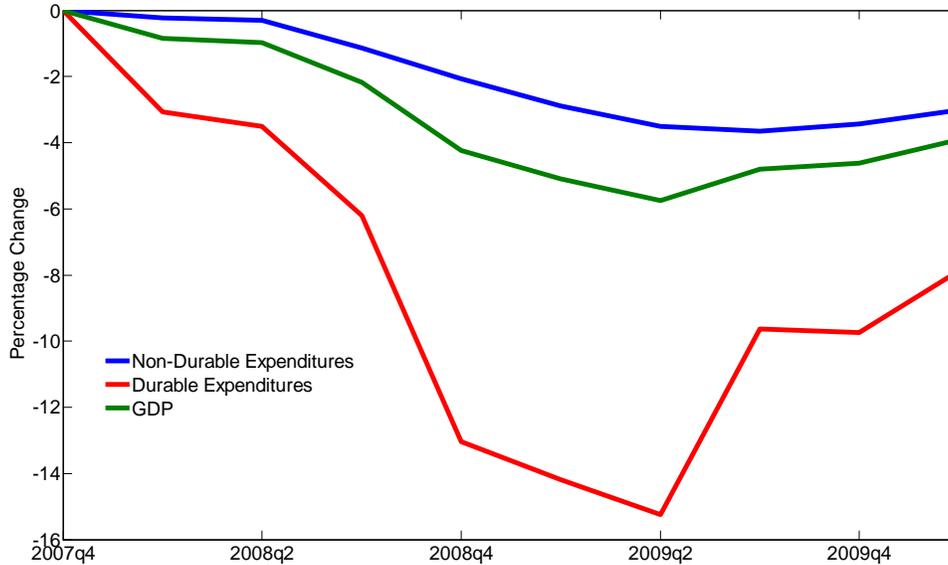
Preliminary and Incomplete

JEL Classification: E21, E32, D91
Keywords: Durables, Fixed Costs, Consumption, Non-linear Impulse Response

Figure 1: 2007 Recession

1 Introduction

Durable expenditures contribute strongly to business cycle fluctuations. Figure 1 displays the decline in durable expenditures (including residential investment), non-durable expenditures[1] and GDP[2] during the Great Recession which began in 2007.[3] Overall, the decline in durable expenditures accounted for nearly half of the decline in GDP. Thus, in a pure accounting sense, stabilizing durable expenditures would have substantially reduced the magnitude of the recession, and indeed, a number of

---

[1]We define durable expenditures as NIPA durable expenditures + residential investment. The BEA treats durable and residential investment differently, including housing services in GDP while excluding durable services. In both our model and data analysis, we define GDP as the sum of non-durable expenditures excluding housing services, consumer durable expenditures, and private domestic investment. See Appendix 1.

[2]All series are hpfiltered with smoothing parameter 1600.

[3]See Petev, Pistaferri, and Eksten [2012] for additional discussion of the behavior of consumption during the Great Recession.

policy interventions during the Great Recession were specifically designed to stimulate durable demand.[4] Despite the prevalence of such policy interventions during recessions, little is known about how the effectiveness of these policies varies with the business cycle. In this paper we argue that the demand response to such interventions is likely to be substantially dampened during recessions so that focusing on the average response can be misleading.

To show this, we start from micro evidence that household level durable purchases exhibit substantial inaction and lumpiness. We build a heterogeneous agent DSGE model with fixed costs of durable adjustment that can capture this micro behavior, and we show that these frictions matter for aggregate dynamics. Fixed costs of durable adjustment change the average response of the economy to shocks and improve the model's fit for standard business cycle moments relative to models which feature frictionless durable adjustment. More importantly, we also show that household level durable frictions have important interactions with the aggregate business cycle: non-linearities in the household durable purchase decision induced by the fixed costs lead aggregate durable expenditures to exhibit non-linear, state-dependent responses to aggregate shocks. In particular, our model implies that during recessions, durable expenditures will be substantially less responsive than during normal times to a variety of aggregate shocks and stimulus policies.

Why do fixed costs change the aggregate dynamics of the model? It is well-known (see e.g. Fisher [1997]) that simple, disaggregated RBC models exhibit a counterfactual negative comovement between investment in durables and investment in productive capital. Aggregate productivity shocks change the return to investing in durables relative to productive capital, and with no adjustment costs this leads to a strong negative correlation between the two forms of investment. In addition non-durable expenditures in the disaggregated RBC model exhibit excess smoothness with simulated volatility falling substantially short of what is observed in the data.

Fixed costs of durable adjustment help on both fronts. Unsurprisingly, adding an adjustment cost to durable purchases fixes the comovement problem.[5] In addition, fixed costs of durable adjustment amplify the volatility of non-durable expenditures.

---

[4]For example, the Cash for Clunkers and First Time Home Buyers credit.

[5]Gomme, Kydland, and Rupert [2001] Morris Davis and Heathcoate [2005] introduce features into the RBC model that essentially act like aggregate adjustment costs to solve the comovement problem. However, these models are not able to match the rich household dynamics and non-linearities that we document.

This is because we calibrate our model to match the fraction of household wealth held in durables, and when we add fixed costs of durable adjustment, this portion of wealth becomes illiquid. Increasing the fraction of wealth that is illiquid means that more households are effectively borrowing constrained. These households are similar to the "wealthy hand-to-mouth" consumers in Kaplan and Violante [2011].[6] In their model, households can hold a liquid low-return asset or a high-return illiquid asset. While they associate this illiquid asset with e.g., housing, they do not model the particulars that generate these higher returns. In contrast, since our model explicitly models illiquid durables, we generate additional predictions about the dynamic response of durables to shocks.

In particular, our model implies strong interactions between the cross-sectional distribution of household durable holdings and the response of aggregate durable expenditures to shocks. Variation in this distribution over the business cycle induces a state-dependent response to aggregate shocks and generates our most important policy implications. Furthermore, we present detailed empirical evidence using micro data to support this model implication.

Fixed costs of adjustment induce Ss dynamics into household durable purchase decisions. Let $z_{i,t} = d_{i,t}^* - d_{i,t-1}$ be the "gap" between a household's current durable stock and the value it would choose if it temporarily faced no adjustment costs. Given a fixed cost of adjustment, it is only worth adjusting if $z_{i,t}$ exceeds some Ss thresholds. Mechanically, this implies that time t aggregate durable expenditures are given by $X_t = \int z_{i,t} h_t(z_{i,t}) f(z_{i,t})$ where $h_t(z_{i,t})$ is the adjustment hazard as a function of the durable gap and $f(z_{i,t})$ is the density of households with durable gap equal to $z_{i,t}$.[7] Thus, the response of durable expenditures to aggregate shocks at any point in time depends crucially on the distribution of household durable gaps and on which households adjust to those gaps. The top panel of Figure 2 shows how this distribution and hazard move over time in our benchmark model. After a long expansion, pent-up durable demand rises, which means that many households are on the verge of adjusting their durable holdings. This makes aggregate durable expenditures more sensitive to aggregate shocks than after a recession, when fewer households are close to adjusting.[8]

---

[6]Campbell and Hercowitz [2009] has a similar mechanism.

[7]For simplicity, we are ignoring maintenance at the moment.

[8]It is worth noting that this increased responsiveness during booms does not depend on the sign of the aggregate shock. Durables respond more to aggregate shocks of either sign during expansions.

Figure 2: Distributions and Hazards Model Vs. PSID
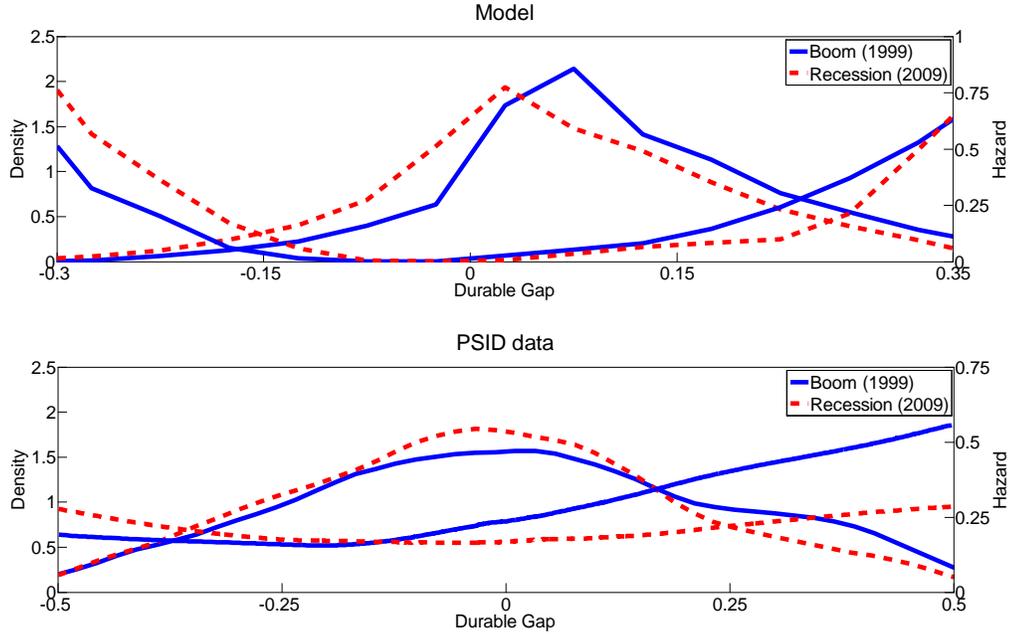
The magnitude of the time-varying responsiveness is quantitatively large. We fit our model to U.S. data and find that the impulse response (IRF) to a positive TFP shock during the Great Recession is less than half the response to the same shock if it had instead occurred during the 90s boom.[9] We also simulate a temporary subsidy to durable purchases which mirrors the "Cash-for-Clunkers" and "First-Time-Homebuyers-Credit", and we find that durable expenditures respond more strongly to the same subsidy during the boom than during the Great Recession. In addition, our model implies a non-linear relationship between the size of stimulus and the aggregate response of durable expenditures. Doubling the size of the stimulus more than doubles the aggregate response of durable expenditures. In contrast, in models without fixed costs of durable adjustment there is no state-dependent impulse response and there is a linear relationship between the size of stimulus and the aggregate durables response. Our model has important policy implications since it means that the impulse response computed from a linear VAR will substantially understate

---

[9]That our model mechanism generates procyclical responsivenss to policy shocks is important in light of the separate literature arguing that in the presence of a binding zero lower, the effectiveness of fiscal stimulus should be countercyclical .

the true response to stimulus during booms and substantially overstate the response during recessions.

Given the importance of this time-varying responsiveness, we search for additional evidence of this model mechanism using a variety of household level micro data sets as well as aggregate time series data, and we find broad empirical support relative to existing models of durable expenditures. Since the time-varying IRF that we document is driven by the interaction between household durable gaps and adjustment hazards, we attempt to directly test this implication by estimating the distribution of gaps and hazards using PSID micro data. The obvious complication with this procedure is that $d_{i,t}^*$, is unobserved in the data. However, we can exploit restrictions from the structural model to generate a mapping between variables that we do observe in the data and $d_{i,t}^*$. This allows us to use PSID data to estimate the distribution of gaps as well as the empirical adjustment hazard as a function of these gaps to directly test our model. The bottom panel of Figure 2 shows that the resulting estimates using PSID data line up remarkably well with our model.

Finally, we also use a parsimonious time-series model to document that aggregate durable expenditures exhibit conditional heteroscedasticity, with aggregate durable expenditures exhibiting substantially greater volatility during expansions than during recessions. The procyclical responsiveness in our model with fixed costs naturally generates such conditional heteroscedasticity. In contrast, existing mechanisms that improve the business cycle performance of disaggregated RBC models do not generate conditional heteroscedasticity. Thus, in addition to providing a better fit for average business cycle statistics than models with frictionless durable adjustment, our model with fixed costs provides a better fit for household level micro data as well as additional aggregate consumption dynamics.

There is a long line of literature studying models with durable consumption. Baxter [1996] investigates the implications of durables for business cycles in a two-sector representative agent framework. Mankiw [1982], Bernanke [1985] and Caballero [1990] study the implications of durables for test of the permanent income hypothesis. Bertola and Caballero [1990], Caballero [1993] and Eberly [1994] investigate stylized heterogeneous agent durables models with fixed cost of adjustment and argue that fixed costs can help explain aggregate dynamics. Diaz and Luengo-prado [2010], Luengo-prado [2006] and Flavin [2011] study similar questions quantitatively but the former paper has no aggregate shocks and the latter papers abstract from general

equilibrium. Leahy and Zeira [2005] and Browning and Crossley [2009] argue that the timing of durable purchases can insulate non-durable purchases from shocks.[10] However, these papers must make strong assumptions to obtain analytical results or to simplify the analysis, and are thus less useful for quantitative analysis. Our estimation of time-variation in household durable holdings builds on Eberly [1994] as well as Attanasio [2000] who estimate the distribution of households' desired vehicle holdings in more stylized (S,s) models. However, we have a more realistic, general equilibrium model, and the restrictions necessary to estimate these earlier papers do not hold in our model. Furthermore, we show that our model is a substantially better fit for the micro data. Caballero, Engel, and Haltiwanger [1995] performs a closely related exercise for business investment and also argues that fixed adjustment costs are important for aggregate dynamics. In addition, a large literature including Dunn [1998], Luengo-prado [2006], and Krueger and Fernandez-Villaverde [2010] has studied durable consumption in quantitative life-cycle models.

Perhaps most related to our model is a recent working paper, Iacoviello and Pavan [2009]. They build a similar incomplete markets model with fixed costs of housing adjustment and aggregate shocks, however they focus on different 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. They instead focus on explaining secular changes in aggregate volatility. In addition, Bajari, Chan, Krueger, and Miller [2010] estimates a micro model of housing demand and explores similar questions to our paper but in partial equilibrium.

The remainder of the paper proceeds as follows: Section 2 describes our benchmark models and discusses their fit along standard business cycle dimensions. Section 3 discusses the model's implications for time-varying impulse response functions. Section 4 tests our model using household level micro data. Section 5 provides time-series evidence for a time-varying impulse response function, and Section 6 concludes.

---

[10]While this insulation margin is active in our model, we find that it is quantitatively small relative to the illiquidity effect.

2 Heterogeneous Households with Fixed Costs

It is well-known that household level durable purchases are infrequent and lumpy. In PSID microdata only 2.2% of prime-age[11] homeowners sell houses[12] each year. Households purchase automobiles more frequently than housing, but even this broader notion of durables is only adjusted on average every five years. Given this pervasive lumpiness, we then ask whether a model with fixed costs of durable adjustment can jointly explain micro behavior together with aggregate consumption dynamics.

2.1 Model Setup

Our model is similar to Krusell and Smith [1998] with the addition of household durable consumption subject to fixed costs of adjustment. Households maximize expected utility of a consumption aggregate, and they are subject to idiosyncratic earnings shocks as well as borrowing constraints. Households solve

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

$$s.t.$$

$$c_t^i = w_t h \eta_t^i + (1 + r_t) a_{t-1}^i + d_{t-1}^i (1 - \delta_d) - d_t^i - a_t^i - A(d_t^i, d_{t-1}^i)$$

$$a_t^i \geq 0; \ d_t^i \geq 0$$

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

where c_t^i, d_t^i and a_t^i are household i's non-durable consumption, durable stock, and assets, respectively. η_t^i represents shocks to idiosyncratic labor earnings, h is a household's fixed[13] hours of work while w_t and r_t are the aggregate wage and interest rate. Finally, $A(d_t^i, d_{t-1}^i)$ is the fixed proportional adjustment cost that households

---

[11] 18-65 years old.

[12] This is the one half the average value of the PSID "sold home" variable: [In the last two years], did you (or anyone in your family living there) sell any home you were using as your main dwelling?

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

face when adjusting their durable stock. We assume that A takes the form

$$A(d, d_{-1}) = \begin{cases} 0 \text{ if } d \in ((1 - \delta_d) d_{-1}, d_{-1}) \\ f (1 - \delta_d) d_{-1} \text{ else} \end{cases}$$

This specification implies that households can maintain their durable stock or let part of it depreciate without paying an adjustment cost, but if they want to adjust their durable stock by larger amounts then they must pay a fixed adjustment cost. Thus, we allow households to engage in routine maintenance that does not incur a fixed cost but assume that they must pay a cost when actually buying or selling their current durable stock. We associate these fixed costs with explicit transaction costs[14] together with time costs.

This adjustment cost specification lies between two extremes that are common in the literature. Under one extreme, households must pay the adjustment cost if $d \neq d_{-1}$. This specification implies that households cannot let their houses depreciate without paying a fixed cost. Alternatively, it is also common to assume that the adjustment cost occurs if $d \neq (1 - \delta_d) d_{-1}$. Under this specification, households cannot maintain their durable stock without paying an adjustment cost. While these alternatives are somewhat simpler,[15] we think that our specification is likely to better capture the realities of the costs associated with durable adjustment. Nevertheless, we have experimented with these alternative specifications and it did not materially affect our results.

A representative firm rents capital and labor and its first order conditions pin down these prices:

$$w_t = (1 - \alpha) Z_t K_t^{\alpha} H^{1-\alpha}$$
$$r_t = \alpha Z_t K_t^{\alpha-1} H^{1-\alpha} - \delta_k$$

where as before, aggregate productivity evolves as an AR process

$$\log Z_t = \rho_Z \log Z_{t-1} + \xi_t.$$

---

[14] Brokers fees, titleing fees, etc.

[15] In these extreme cases, when households choose to not buy or sell durables, the consumption decision becomes one dimensional. In our specification, households must still choose how much to let the durable stock depreciate.

As usual, equilibrium requires that the aggregate resource constraint

$$C_t + D_t + K_{t+1} + A_t = Z_t K_t^\alpha H^{1-\alpha} + (1 - \delta_k) K_t + (1 - \delta_d) D_{t-1}$$

be satisfied, where

$$
\begin{aligned}
K_t &= \int a_{t-1}^i \\
D_t &= \int d_t^i \\
C_t &= \int c_t^i \\
A_t &= \int A(d^i, d_{-1}^i) \\
H &= \int h \eta_t^i.
\end{aligned}
$$

Solving the household problem requires forecasting aggregate prices and thus the aggregate capital stock, which is determined by the continuous distribution of household states, so as usual solving the model requires making computational assumptions. Following Krusell and Smith [1998], we conjecture that after conditioning on aggregate productivity, aggregate capital is a linear function of current aggregate capital:[16]

$$K_{t+1} = \gamma_0 (Z) + \gamma_1 (Z) K_t.$$

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

---

[16]The forecasting rule might also depend on the previous durable stock. An earlier version of this paper found that this added little explanatory power and had substantial computational cost.

$$V\left(a_{-1}, d_{-1}, \eta; Z, K\right) = \max\left[V^{adjust}\left(a_{-1}, d_{-1}, \eta; Z, K\right), V^{noadjust}\right]$$

with

$$V^{adjust}\left(a_{-1}, d_{-1}, \eta; Z, K\right) = \max_{c,d,a} \frac{\left[c^v d^{1-v}\right]^{1-\theta}}{1-\theta} + \beta E_{\varepsilon,\xi} V\left(a, d, \eta'; Z', K'\right)$$

$$s.t.$$

$$c = wh\eta + (1+r)a_{-1} + d_{-1}\left(1-\delta_d\right) - d - a - f\left(1-\delta_d\right)d_{-1}$$

$$a > 0; \quad \text{equilibrium conditions and prod. processes}$$

$$V^{noadjust}\left(a_{-1}, d_{-1}, \eta; Z, K\right) = \max_{\substack{c,a, \\ d\in\left((1-\delta_d)d_{-1}, d_{-1}\right)}} \frac{\left[c^v d^{1-v}\right]^{1-\theta}}{1-\theta} + \beta E_{\varepsilon,\xi} V\left(a, d, \eta'; Z', K'\right)$$

$$s.t.$$

$$c = wh\eta + (1+r)a_{-1} + d_{-1}\left(1-\delta_d\right) - d - a$$

$$a > 0; \quad \text{equilibrium conditions and prod. processes}$$

We solve the model by conjecturing an aggregate law of motion, approximating the value function by linearly interpolating[17] 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 2 for additional details on the solution method.

2.2 Business Cycle Results

We now assess our model's business cycle performance relative to a simpler incomplete markets model with no adjustment costs as well as a frictionless RBC model with durables. Table 1 shows our benchmark calibration. Our calibration strategy uses a broad measure of the durable stock including both housing and consumer durables. See Appendix 1 for discussion of our empirical moments.

The discount rate is picked to generate a quarterly interest rate of 1%, and we assume an intertemporal elasticity of substitution[18] θ of 2. The depreciation rate of

---

[17]Linear interpolation gives speed advantages relative to cubic spline or other interpolation methods. While linear interopolation 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 problematic.

[18]Raising θ makes both non-durable expenditures and durable expenditures less volatile. With no adjustment costs, we find that non-durable expenditures are too smooth while durable expenditures are too volatile, altering θ cannot simultaneously improve the fit along both dimensions.

capital $\delta_k = 0.022$ is set to match the long-run average investment to capital ratio. The average depreciation rate of consumer durables is moderately higher than that of productive capital while the depreciation rate of residential capital is somewhat lower, so in our benchmark results we impose an intermediate value and set $\delta_d = \delta_k$. If anything, this is an over estimate of the actual weighted depreciation rate, and using a higher depreciation rate for the durable stock improves the fit of the models with no adjustment costs, so we view this as a conservative calibration.

Table 1

Model Parameters

Parameter	With Fixed Cost	No Fixed Cost	RBC
β	0.986	0.987	0.999
ρ	0.95	0.95	0.95
σ_z	0.008	0.088	0.008
α	0.3	0.3	0.3
δ_k	0.022	0.022	0.022
δ_d	0.022	0.022	0.022
θ	2	2	2
ν	0.73	0.74	0.73
h	0.33	0.33	0.33
ρ_η	0.975	0.975	N/A
σ_η	0.1	0.1	N/A
f	0.025	0	N/A

The weight on non-durable consumption, v, is set to match an average ratio of non-durable to durable expenditures[19] of 4.0. We set the fixed cost of adjustment at 2.5%, so that households lose 2.5% of their pre-adjustment durable stock when adjusting.[20] We calibrate the idiosyncratic productivity process to match the persis-

[19]A value of 4 may seem low relative to standard numbers from NIPA, but our measure of non-durable expenditures excludes housing services while our measure of durable expenditures includes residential investment. Increasing the target value did not affect any of our qualitative conclusions.

[20]This is less than the the value used by Diaz and Luengo-prado [2010], who report that the typical fee charged by real estate brokers in the U.S. is around 6%. Since costs of consumer durable adjustment are smaller, we pick a lower value for our benchmark. In our model, this implies a quarterly adjustment frequency of just under 3%. While this is higher than the empirical frequency of adjustment of housing, it is lower than the frequency of consumer durable adjustment. Our conclusions were not sensitive to decreasing adjustment costs to 1% or raising adjustment costs to 10% as long as ν is recalibrated to match the same ratio of non-durable to durable expenditures.

tence 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.

Given these parameter choices, Table 2 shows our business cycle results, and Appendix 2 discusses the sensitivity of our results to these parameter choices as well as to relaxing the Cobb-Douglas specification for the utility function and shows that these results are very robust to a range of reasonable parameter choices.

Table 2

Business Cycle Volatility

	Data	W/ Fixed Costs	No Fixed Costs	RBC
	Standard Deviation Relative to Y	Standard Deviation Relative to Y	Standard Deviation Relative to Y	Standard Deviation Relative to Y
Durable	3.04	2.81	14.22	13.71
Non-Durable	0.57	0.51	0.40	0.36
Investment	2.24	5.60	35.18	21.23

Why is are durable expenditures and investment substantially too volatile in the models with no adjustment costs? It is because these models feature the comovement problem identified in Greenwood and Hercowitz [1991] and further explored in Fisher [1997]. Aggregate productivity shocks change the return to investing in durables relative to productive capital. Increases in productivity make it more valuable to save in productive capital, and the additional output produced can later be used to finance durable consumption. This generates a strong negative correlation between durable expenditures and investment in productive capital in models with no adjustment costs and increases the volatility of both variables. Table 3 shows the comovement between these two forms of investment in the data and models:

Table 3

Business Cycle Comovement

	Data	W/ Fixed Costs	No Fixed Costs	RBC
Correlation(I_K, I_D)	0.36	0.24	-0.90	-0.85

Clearly, the model with fixed costs generates a dramatic improvement in the correlation between the two forms of investment as well as in the relative volatility of these variables. This is because adjustment costs break the incentive to rapidly adjust between saving in the two forms of capital. More interestingly, we also find that the model with fixed costs is a substantially better fit for the volatility of non-durable consumption. This is because the presence of adjustment costs on durables means that a fraction of household wealth is effectively illiquid. Since we target the same level of total wealth in all models, a greater fraction of households are temporarily liquidity constrained in the model with fixed costs since much of their total wealth is illiquid. These households are similar to the wealthy hand-to-mouth households emphasized in Kaplan and Violante [2011]. While households may have a large amount of total wealth, households rationally choose to avoid paying fixed costs associated with using illiquid wealth to smooth non-durable consumption. In their model, wealth is illiquid because households put some wealth into an illiquid higher return investment while in our model, wealth is illiquid because households want to consume durables, which are subject to adjustment costs. While the general mechanism is similar, we have additional panel data on income, durable and non-durable expenditures and wealth, which allows us to test directly for this mechanism.[21]

We conclude this section by noting that fixed costs of adjustment are not the only mechanism that can solve the comovement problem and improve the business cycle volatility of consumption.[22] Gomme, Kydland, and Rupert [2001] add time-to-build to a disaggregated model while Davis and Heathcote [2005] introduce a fixed factor of production. These model features end up acting like aggregate adjustment costs and so they help to improve the comovement of disaggregated investment components. Indeed, we find that by introducing quadratic adjustment costs into the RBC model, we can also better match the average response to shocks.

However, we think fixed costs of adjustment are a particularly attractive mecha-

---

[21]Previous versions of this paper explored the implications of this illiquid wealth for household level insurance using Italian Survey of Household Income and Wealth Data. In ongoing work, we explore these implications using PSID data.

[22]Another force that could dampen durable volatilty and amplify non-durable volatility would be movements in the relative price of these two forms of consumption. If the relative price of durables is procyclical, then this would mute durable increases during booms and dampen them during recessions. However, for the period from 1960-2012, the relative price of durables is actually mildly countercyclical, so while this mechanism could improve the real business cycle fit, it would require counterfactual relative price movements. Matching empirical price movements over the business cycle would exacerbate the model fit.

nism for two separate reasons: First, we will bring a wealth of micro data to bear in testing our model that is unavailable for representative agent models: there is substantial evidence for lumpy durable adjustment at the micro level which will not be replicated by convex adjustment costs. In addition, we will show that models with fixed costs of adjustment induce non-linearities that generate conditional heteroscedasticity for aggregate durable expenditures: the residual variance of the aggregate durable investment rate is substantially larger during booms than during recessions. We will show shortly that this prediction is supported by U.S. time-series data. In contrast, models with smooth aggregate adjustment costs do not generate such heteroscedasticity. Thus, in addition to being a good fit for the average response to shocks, our model will be consistent with the empirical time-variation around that average. Nevertheless, we view our analysis as complementary to the existing literature rather than in direct competition with it. Nothing precludes both smooth and non-convex adjustment costs from being simultaneously present, but we explore how far we can get with non-convex adjustment costs alone.

3   Non-Linear Dynamics

We now show that in addition to better matching business cycle moments, fixed costs of durable adjustment also induce impulse responses that vary with the state of the economy. Durable expenditures are significantly less responsive to changes in aggregate conditions during recessions than during booms. Since fiscal stimulus does not happen at random, and it is typically timed during recessions, assessing the response of durable expenditures to aggregate shocks during recessions is of particular interest.

How do we define a boom and a recession in our model? In order to replicate U.S. time-series data in general, and the Great Recession in particular, we hit the economy with two aggregate shocks. We first pick the aggregate TFP shocks in the model to match observed Solow Residuals in aggregate data. We then hit the simulated economy with an additional unanticipated 4 percent decline in the capital stock in the fourth quarter of 2008.[23] We choose these particular shocks because we believe they

---

[23]We have also experimented with shocks to the durable stock, but we believe these map less naturally into the recession. The decline in housing value was largely a decline in the price of housing rather than a decline in the real stock of housing. Furthermore, large declines in the housing stock lead to counterfactual housing booms.

Figure 3: Recession Simulation with Fixed Costs Vs. Data
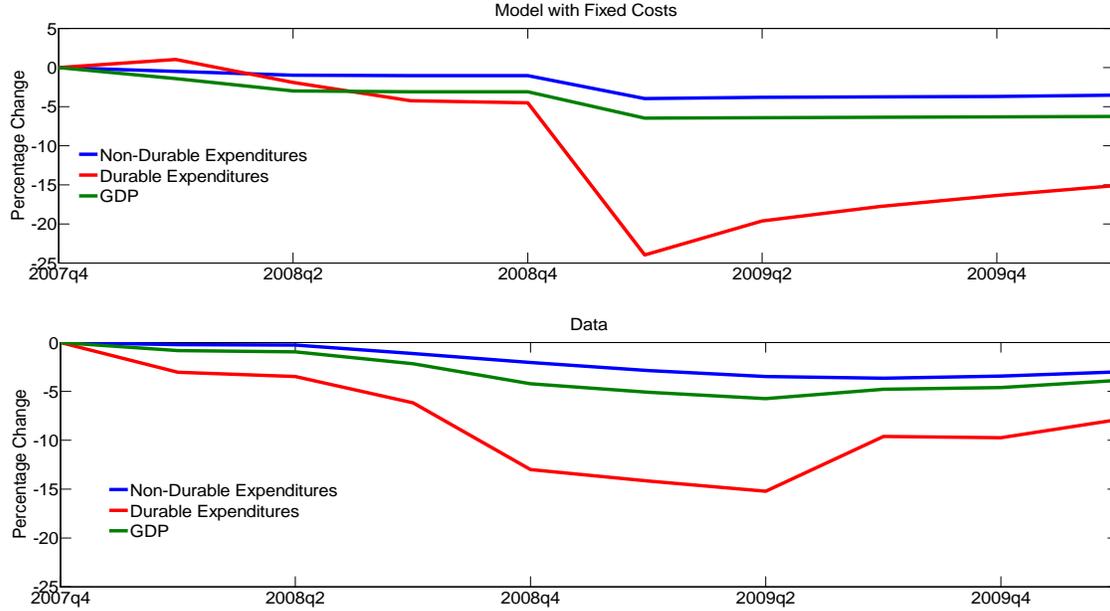
map typical descriptions of the recession into the mechanisms of our model. The one time decline in the capital stock corresponds to a decline in household wealth while the decline in TFP that is fed into the model leads to a decline in household earnings. We pick the magnitude of the capital shock to roughly match the declines in capital actually observed in the most recent recession.

How well do our simulated models do at matching the dynamics of consumption during the Great Recession? Figure 3 shows that the model with fixed costs closely matches the behavior of consumption aggregates over the Great Recession. In contrast, Figure 4 shows that when the fixed adjustment costs are turned off, the model now does an extremely poor job of matching consumption during the Great Recession. In particular, durable expenditures decline too much relative to the data while non-durables decline too little. Again, the model with no adjustment costs does a bad job of matching the response of consumption to aggregate shocks.

Since the model with fixed costs of adjustment does a good job of replicating time-series behavior, we can then ask whether it implies an IRF that varies with the state of the economy. Figure 5 shows the IRF to a one standard deviation increase in TFP

Figure 4: Recession Simulation Incomplete Markets Model (No Fixed Costs)
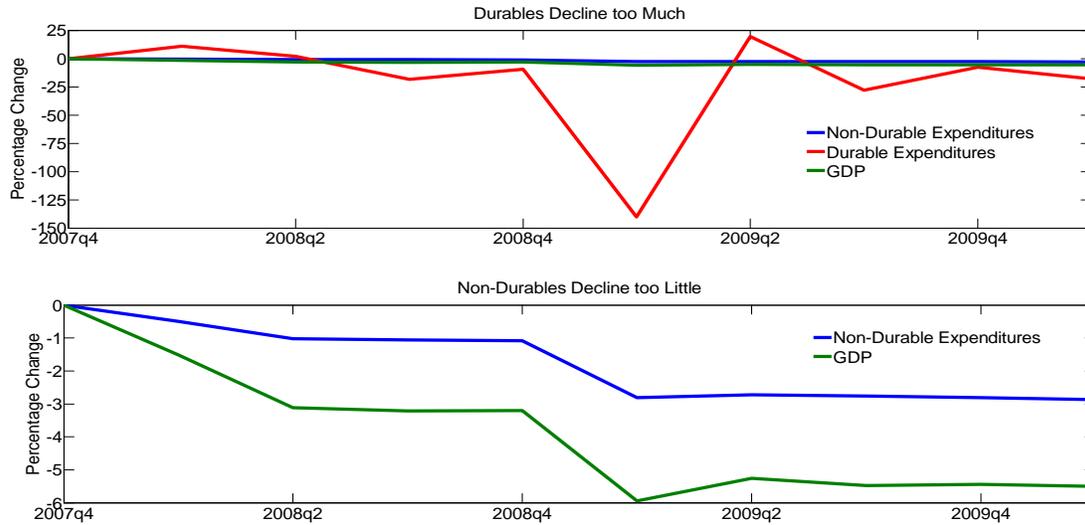
in 1999 and compares it to the IRF calculated in a 2009.[24] The impulse response is much larger during the 90s than during the Great Recession. The cumulative impulse response in 1999 is roughly twice that in 2009.

It is interesting to look at impulse responses to TFP shocks since these are the driving shocks in our model. However, while one can imagine government policies that might manifest themselves as increases in TFP, it is also interesting to model more realistic stimulus policies that were actually implemented during the Great Recession. Towards that end, we next investigate the response of the economy to a one-time subsidy to durable purchases modeled after the "Cash-for-Clunkers" and "First-Time-Home-Buyers" credit. In particular, we assume that in the period of the stimulus, households face a one-time subsidy so that their new budget constraint is given by

---

[24] The hump-shaped IRFs eventually return to zero. Interestingly the hump-shape is consistent with the impulse response of aggregate durable expenditures to observed changes in TFP in the data. The hump-shape arises due to equilibrium movements in the interest rate across time. On impact, as TFP increases, interest rates rise, which increases financial income as well as the return to saving in liquid assets. So initially, most of the increase in savings goes into liquid assets. However, as additional capital is accumulated, the return to saving in liquid assets falls and households begin to accumulate more in durable assets so that the response of durable expenditures grows with time. As the TFP shock dies out, this process reverses itself and the economy returns to steady-state.

Figure 5: State-Dependent Impulse Response of Aggregate Durable Expenditures
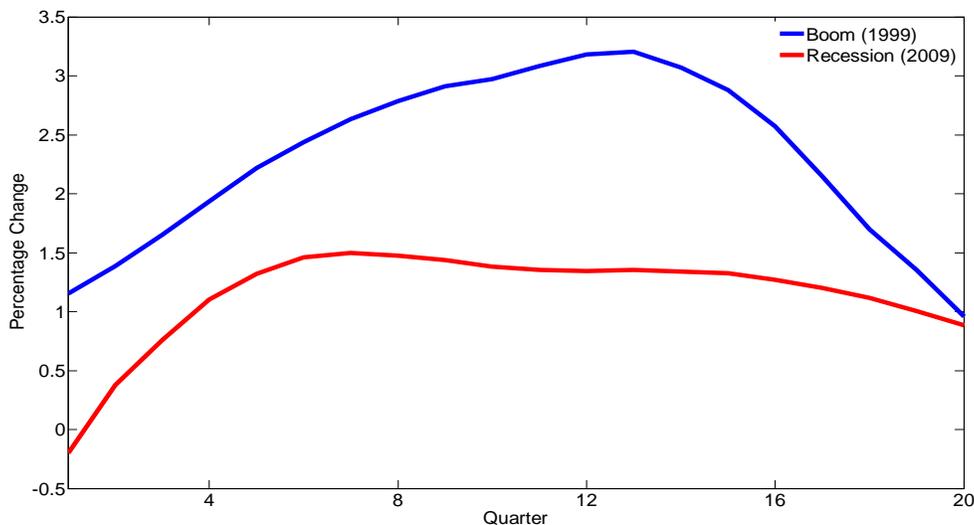
$$c_t^i = w_t h \eta_t^i + (1 + r_t) a_{t-1}^i + d_{t-1}^i (1 - \delta_d) - d_t^i - a_t^i - A(d_t^i, d_{t-1}^i) + s \left(d_t^i - d_{t-1}^i \right),$$

where the last term reflects a subsidy to adjusting durable holdings upwards.[25] After this one-time shock, we assume that the economy returns to the ergodic steady-state and that households then use their original value functions to determine optimal behavior. This is a strong assumption, but it has large computational advantages relative to computing a transition path to the ergodic distribution, and it is likely to hold approximately for small values of the subsidy. In our benchmark results, we use a subsidy value of 1%, which is in line with the actual size of the stimulus programs after accounting for eligibility and phase-outs.[26] In addition, we evaluate the validity of our Krusell-Smith forecasting rule after the one-time shock and find

---

[25]Modeling the government sector and the source of these funds in detail is beyond the scope of this paper, but this subsidy could be financed by a one-time reduction in unmodeled steady-state government expenditures.

[26]We view a 1% shock as a reasonable approximate to the actual policy since the First-Time-Home-Buyers credit since the subsidy as a fraction of the mean home price in 2008 was \$8000/\$292,000 = 2.7%. However, only households that had not purchased a primary residence in the last three years were eligible.

Figure 6: State-Dependent Impulse Response to Durable Subsidy
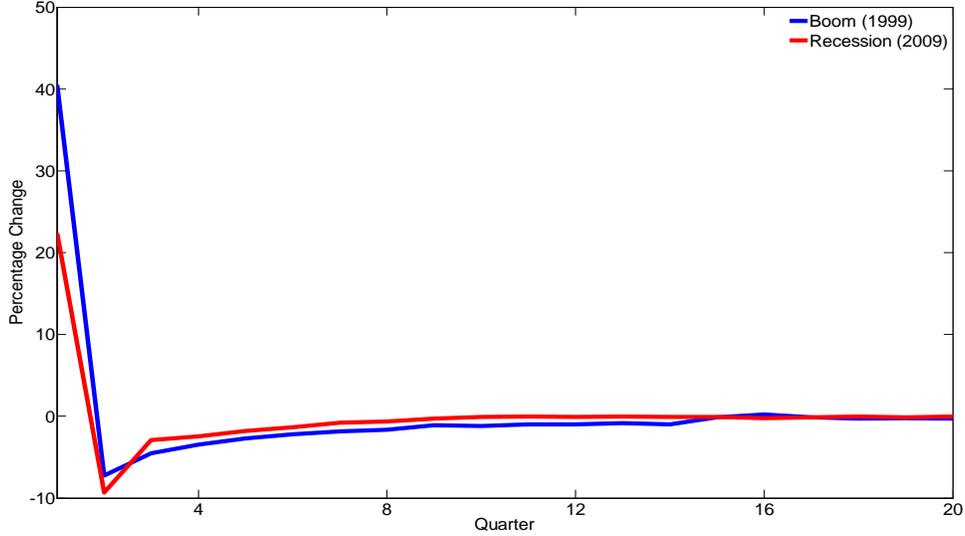
that its accuracy is only mildly reduced.

Figure 6 shows that durable expenditures respond much more strongly to the durable subsidy if it occurs during a boom than during a recession. Interestingly, we also find that much of the effect of the subsidy is undone in the ensuing periods, as household purchases are pulled forward from the near future by the stimulus. This is similar to the empirical result found in Mian and Sufi [2012]'s evaluation of the Cash-for-Clunkers Program. Nevertheless, we still find that both the effect of the subsidy on impact as well as the cumulative response are very procyclical.

Why does our model generate a procyclical IRF? Fixed costs of adjustment imply that an individual household's response to aggregate shocks is highly non-linear. Since households are continuously hit with both aggregate and idiosyncratic shocks, in the presence of fixed costs of durable adjustment there is a gap between a household's current durable holdings and the durable holdings it would choose if it temporarily faced no adjustment costs. When this gap is small, it is not worth paying the fixed cost of adjustment. As this gap becomes large, it becomes optimal to pay the fixed cost. Thus, fixed costs of adjustment induce lumpiness into adjustment, and changes across time in the fraction of households choosing to adjust can induce

significant time-variation in aggregate durable expenditures. Let $z_{i,t} = d_{i,t}^* - d_{i,t-1}$ be the "gap" between a household's current durable stock and the value it would choose if it temporarily faced no adjustment costs. Mechanically, time t aggregate durable expenditures are given by $X_t = \int z_{i,t} h_t(z_{i,t}) f(z_{i,t})$ where $h_t(z_{i,t})$ is the adjustment hazard as a function of the durable gap and $f(z_{i,t})$ is the density of households with durable gap equal to $z_{i,t}$.[27]

What determines the response of X_t to aggregate shocks that increase households' desired during holdings by Δd^*? The total response of X_t can be decomposed into two components. The first component is the intensive margin: conditional on adjusting all households will now choose durable holdings that are Δd^* larger than before the aggregate shock. The second component is the extensive margin: some households close to increasing their durables will be pushed into action by a positive shock, and some households who previously would have sold durables instead choose inaction after the shock.

When will these margins be more important? The intensive margin response increases with the frequency of adjustment. The more households that are adjusting before the aggregate shock, the greater the response to that shock along the intensive margin. The extensive margin response to a shock will be more important when more households' adjustment decisions are changed by that shock. This will be true when more households are in the upward sloping region of the hazard function $h_t(z_{i,t})$.

During a boom, both of these margins become more important, and so aggregate durable expenditures become more responsive to aggregate shocks. Figure 7 plots the distribution of durable gaps and adjustment hazard in a boom and in a recession, for the model with fixed costs of durable adjustment. The distribution of durable gaps is skewed right because depreciation means that more households want to increase than to decrease durable holdings. During the boom, households' desired durable holdings rise so that the distribution of durable gaps shifts to the right, and more households are pushed into the steep part of the hazard function. Furthermore, households are more likely to adjust up and less likely to adjust down for a given durable gap (as can be seen by the shifting hazard[28] in Figure 7), and the net effect

---

[27]Where the gaps and hazards are defined relative to durable holdings after maintenance so that X_t is durable investment excluding maintenance.

[28]The shift in the hazard is a combination of two forces. During recessions, fixed costs represent a larger fraction of household wealth, so adjustment to both positive and negative gaps is reduced. However, households are also more likely to hit the borrowing constraint, which increases the prob-

Figure 7: Boom and Recession: Distribution and Adjustment Hazard
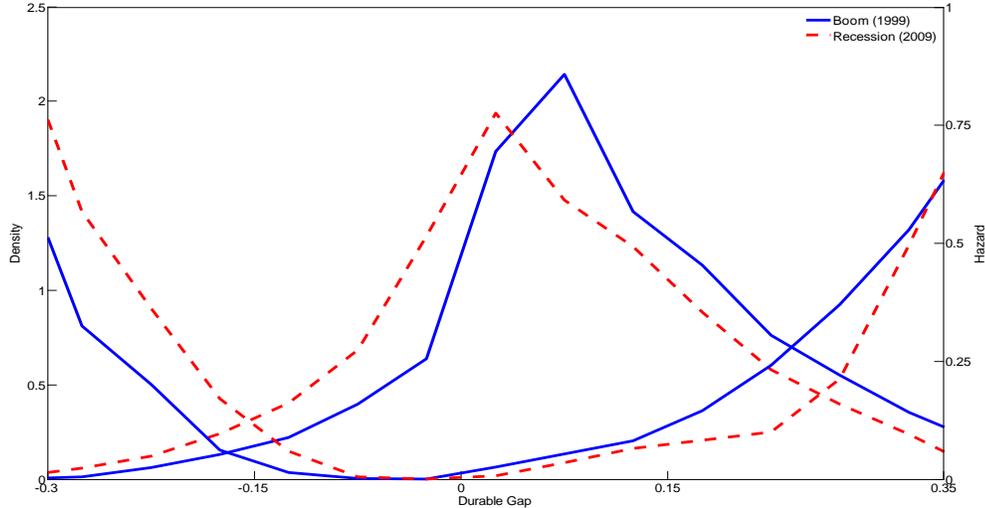
is an increase in the probability of adjustment.[29] Together, this increase in the mass of households adjusting as well as in the mass of households close to the margin of adjustment leads to an increase in both the intensive and extensive margin response to aggregate shocks so that total responsiveness rises.

This time-variation in distributions and hazards induced by fixed costs of adjustment leads to a procyclical IRF, but there are other features of the model that could also lead to IRFs that vary across time. We now argue that it is the fixed costs that are quantitatively important for our aggregate dynamics rather than these alternative features. In particular, household borrowing constraints also introduce non-linearities, and the particular sequence of shocks might introduce GE effects that matter for aggregate dynamics.[30] However, we can show that neither the sequence of

---

ability of adjustment to negative gaps. The net effect is a decrease in the probability of increasing and an increase in the probability of decreasing.

[29]It is worth noting that aggregate durable expenditures become more responsive to both positive and negative aggregate shocks during a boom. During a boom (after a sequence of positive aggregate shocks), more households are in the region with a steep hazard of adjustment, and these households become more responsive to both positive and negative aggregate shocks. Our model implies a state-dependent IRF, not an asymmetric IRF.

[30]In particular, the economy may respond differently to shocks in a recession, when aggregate income is low and more households are close to hitting their borrowing constraint than when aggregate income is high and most households are unconstrained.

Figure 8: No State-Dependent IRF without Fixed Costs
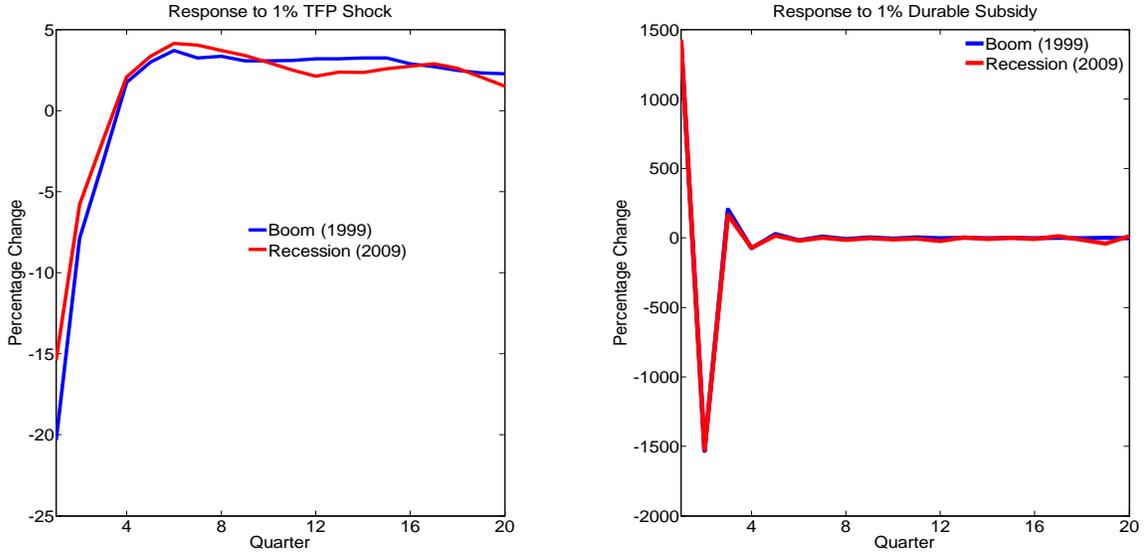
shocks nor the presence of borrowing constraints are driving our procyclical IRF by recomputing impulse response functions for an otherwise identical model with borrowing constraints but with no costs of durable adjustment. Figure 8 shows that there is essentially no difference between the IRF at different points in time in the model with no adjustment costs. This implies that it must be the fixed costs of adjustment that drive our procyclical IRF.

Thus, our model with fixed costs of adjustment delivers aggregate dynamics that differ sharply from models with frictionless durable adjustment, even with relatively small fixed costs of adjustment. This stands in stark contrast to the debate in the investment literature between Khan and Thomas [2008] and Bachmann, Caballero, and Engel [2010] that finds that large adjustment costs are necessary to generate aggregate non-linearities. What differentiates our model from this previous work? In both Khan and Thomas [2008] and Bachmann, Caballero, and Engel [2010], the representative household can only save in one asset, productive capital. If fixed costs on the firm side of the model generate results for investment that differ dramatically from the frictionless model, this necessarily implies that households must also face consumption that differs dramatically from the frictionless environment. In equilibrium, household consumption smoothing has large quantitative effects on the interest

rate that dampen the effects of fixed costs.

In contrast, in our model, households can save in two assets: productive capital and durables. Non-linearities in durable expenditures can then be absorbed by movements in capital rather than movements in consumption, so even small fixed costs of durable adjustment can have big effects on aggregate dynamics without inducing huge effects on consumption dynamics.[31]

Thus our model suggests that fixed costs may lead policies designed to stimulate durable expenditures during recessions to be less effective than suggested by linear VAR evidence that assumes there is no time-variation in the response to shocks. This has important implications for policy makers. For example, both "Cash for Clunkers" and "First-Time Home Buyers Tax Credit" were enacted to stimulate durable expenditures during the most recent recession. However, since pent up durable demand was relatively low, this likely dampened the effectiveness of the programs. Conversely, if each program led to an increase in durable demand and pushed more households towards an active region of adjustment, then the total effect of the two programs may have been greater than the sum of the individual programs in isolation. We find support for such non-linearities by computing the IRF to shocks of varying sizes: doubling the size of the TFP shock or durable stimulus roughly triples the cumulative impulse response.

4   Household Micro Data

Given the important policy implications of these non-linearities, we now look for direct evidence that the cross-sectional distribution of household durable demand moves in the manner predicted by our model. The obvious difficulty with constructing an empirical counterpart to Figure 7 is that "durable gaps" are not observed in the microdata. However, our structural model imposes strong restrictions on the relationship between variables which are observed in microdata and durable gaps, which are observed in the model but not in the data.[32] We use the model to estimate

---

[31] As noted previously, we do find that fixed costs amplify non-durable volatility, but the magnitude of this effect is much smaller than in a model with fixed costs of investment and no GE effects.

[32] Our empirical strategy in general is similar to that in Eberly [1994], but we have a more complicated empirical model that does not require us to exclude borrowing constrained households. Furthermore, we test the predictive power of our estimated durable gaps for the actual probabilities of adjustment, and find an upward sloping adjustment hazard as predicted by our model. In contrast, Eberly [1994] imposes and estimates a single adjustment threshold.

Figure 9: Estimated Gaps and Hazard, PSID
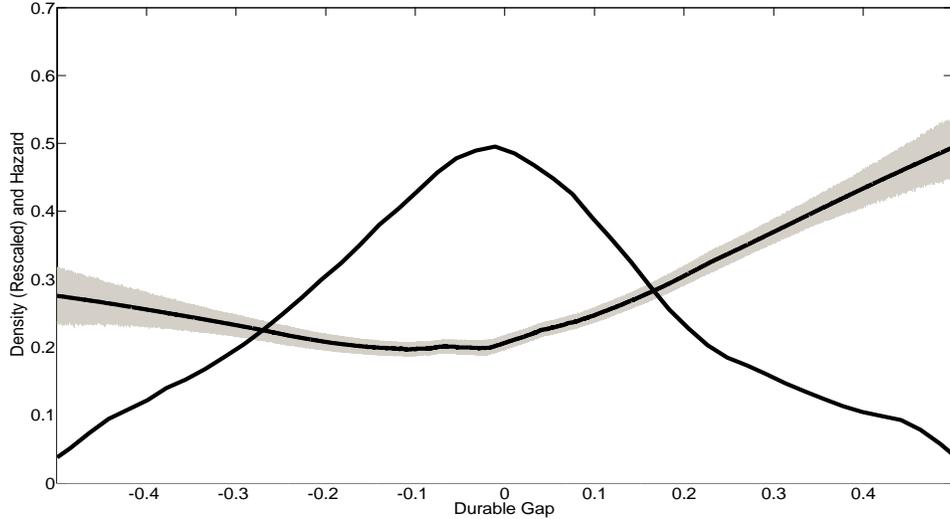
a flexible functional form that relates d_{-1}, c, a to the durable gap, and find that within the model these variables are sufficient to explain more than 99% of the variation in durable gaps. Once we have estimated the relationship in the model, we then apply the same relationship to measures of d_{-1}, c, a in PSID microdata, which allows us to estimate an empirical cross-section of durable gaps for each of the sample years in the PSID from 1999-2009. Once we have estimated this distribution, we can test our empirical fit by estimating the probability that a household actually adjusts its durable stock, conditional on our imputed durable gap. We first test the overall predictive power of our model and then discuss implications for variation over the business cycle as in Caballero, Engel, and Haltiwanger [1997]. See Appendix 3 for more detailed discussion.

Figure 9 shows estimated gaps and hazards for the PSID data, pooling across all years. Given that our model is abstracting from life-cycle considerations and many other idiosyncratic taste shocks that will affect durable adjustment, the fit between the model and the data is remarkably strong. Our estimated durable gaps have strong predictive power for the actual probability of adjustment. The empirical probability of adjustment rises from approximately 20% for a durable gap of zero to

24

more than 50% for a durable gap of 0.5.[33]   (While the adjustment probability of 20% at zero may appear to be a failure, it is worth noting that our measure of adjustment is calculated over a two-year period while our gap measure is instantaneous, so a positive probability despite a gap of zero may reflect a positive gap at some other point during the two year period which led to adjustment).   In contrast, if the model we used was uninformative for actual household behavior, we would find no relationship between estimated durable gaps and the empirical adjustment hazard.

Figure 10: Estimation Based on Maintaining a Constant Ratio of Liquid and Illiquid Wealth
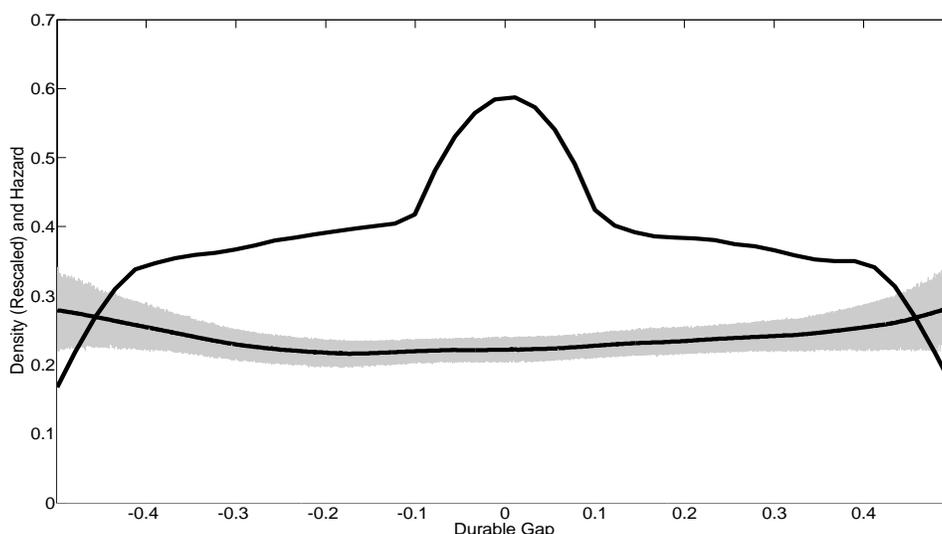
Indeed, in addition to our benchmark model, we can also redo this estimation procedure using alternative models of durable consumption.   In particular, Grossman and Laroque [1990] build a model of optimal durable consumption with fixed costs that implies that all households maintain a constant ratio of d/a when adjusting their durable stocks.   We repeat our estimation procedure using a generalized version of this model where we allow for household level heterogeneity in the optimal ratio of liquid and illiquid wealth.   Figure 10 shows that, not only does this model generate strange distributions of implied durable gaps, these gaps also have essentially no predictive power for actual household durable adjustment.   The point estimate for the hazard

---

[33]Bootstrapped 90 Percent Confidence Intervals in Gray.

rises from only 24% to 28% over the range of durable gaps, and the 90% confidence intervals show that a flat hazard cannot be rejected. Thus, the close match between our benchmark model and our PSID estimates is not a result of using our model restrictions to impute durable gaps and is indeed a strong test of the model. Figure 10 shows that this procedure need not yield self-confirming results.

In addition to these results using PSID data, we have also repeated the analysis using household panel data from the Italian Survey of Household Income and Wealth. While our model is estimated to match aggregate U.S. data, similar forces are likely to apply in other countries as well. Again, Figure 11 shows that our model generates durable gaps with strong predictive power for actual adjustment, in contrast to the simpler alternative model.

Figure 11: Estimates Using SHIW Data
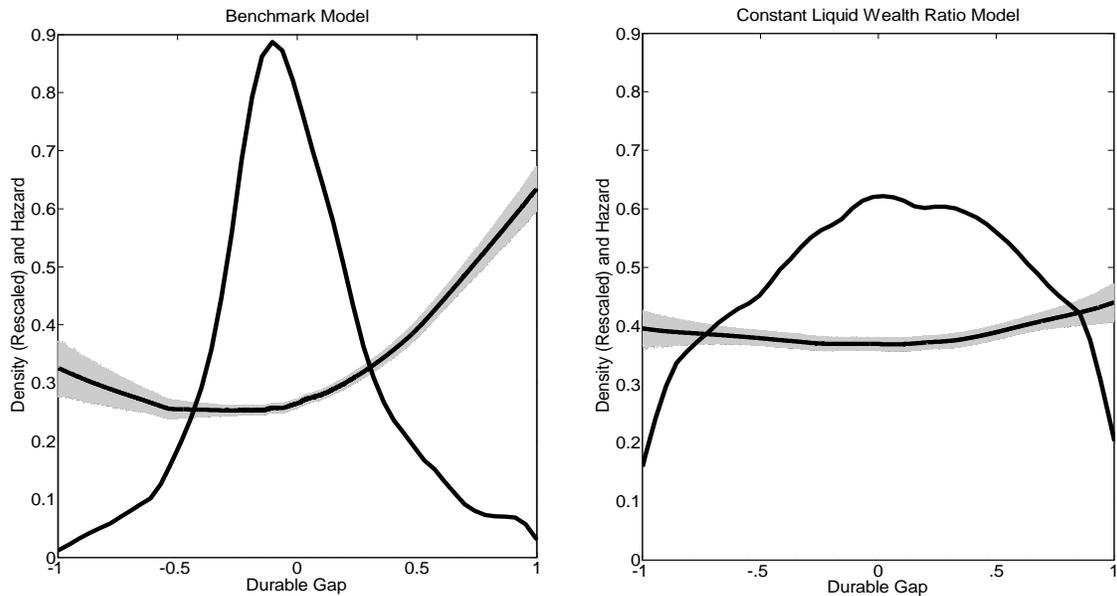
Now that we have shown our model provides a good empirical fit for household micro data on average, we test whether the estimated distributions and hazards move over the business cycle in the way predicted by our model. Figure 2 compares the estimates from the PSID data to the predictions from our model. We find that both the empirical distribution of durable gaps as well as the corresponding adjustment hazards shift between 1999 and 2009 in a way that is remarkably consistent with

26

our model.[34]   Thus, there is strong empirical evidence that both the distribution of households' durable gaps and adjustment probabilities move in the manner necessary to generate a procyclical IRF.[35]

5   Aggregate Time-Series Evidence

In addition to the test of our model using household panel data, we also look for additional evidence of a time-varying IRF that does not rely on a particular structural model. We now provide simple time-series evidence that the response of durable expenditures to aggregate shocks is systematically weaker during recessions than during booms. While one could imagine many estimation procedures to explore this question, we proceed in a parsimonious manner that is agnostic about the underlying structural model.

In particular, we follow Bachmann, Caballero, and Engel [2010] and estimate a two-stage time-series model. In the first stage, we estimate an AR process for durable expenditures. Then, in the second stage, we regress the absolute value of the residuals from the first stage on the average of lagged durable expenditures to assess whether residual variance is different during booms than it is during recessions. (See Appendix 1 for details). We find clear evidence that aggregate durable expenditures exhibit conditional heteroscedasticity, with durable expenditures exhibiting much larger (absolute) residuals following expansions. Figure 12 shows that our estimates of residual variance rise dramatically with previous durable expenditures.[36]   At the peak of the nineties boom (when durable investment rates were very high), the residual variance was 50% larger than at the trough of the 2007 recession (when durable investment rates were very low).

Why does greater variance during expansions imply that durable expenditures are more responsive to shocks during expansions? In principle, conditional heteroscedasticity could arise for two reasons: (1) The economy is subject to shocks of constant variance but durable expenditures respond more to these shocks during booms than

---

[34]Appendix 3 shows that the differences across time are statistically significant.

[35]While Figure 5 shows that the IRF on impact is only a partial summary of the differences in the cumulative IRFs, we can nevertheless compute the IRF on impact implied by our empirical distribution of gaps and hazards (ignoring GE effects) and we find that the implied IRF on impact is nearly 30 percent larger in 1999 than in 2009.

[36]Bootstrapped 90 percent confidence interval in gray.

Figure 12: Conditional Heteroscedasticity
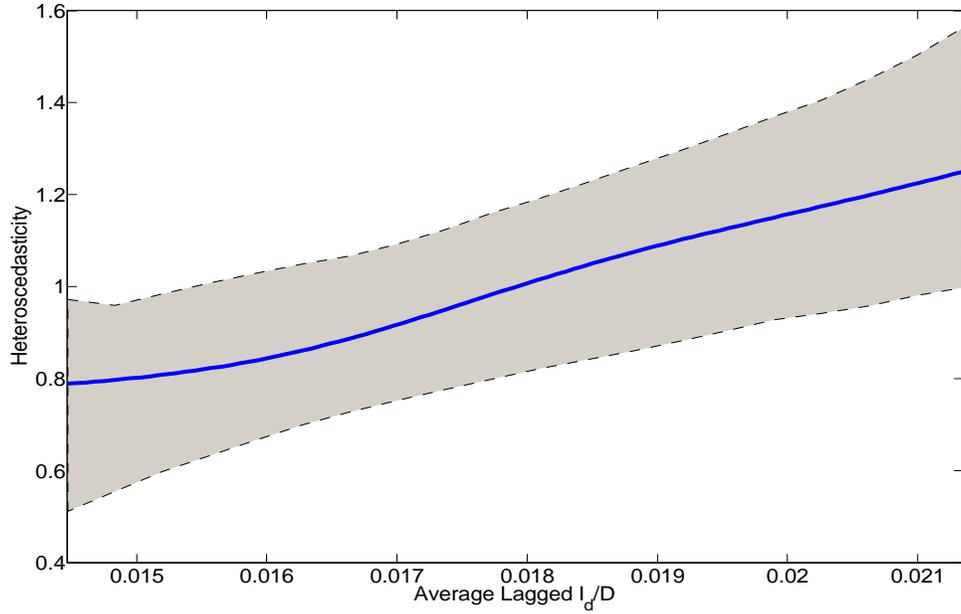
during recessions. (2) The size of shocks during booms is greater than the size of shocks during recessions.

Appendix 1 shows that while durable expenditures exhibit conditional heteroscedasticity, there is no evidence for heteroscedasticity for either productivity or monetary shocks. Thus, there is no evidence for heteroscedasticity among the most common shocks used in business cycle models. Furthermore, aggregate GDP also does not exhibit conditional heteroscedasticity. Since the aggregate shocks affecting durable expenditures and total GDP should be similar, our evidence supports the first explanation: durable expenditures become more responsive to aggregate shocks during business cycle booms.

We find similar effects in our model with fixed costs. Appendix 1 presents estimates of conditional heteroscedasticity for the model with fixed costs of durable adjustment as well as for models with no adjustment costs and for an RBC model with convex adjustment costs. The model with fixed costs of durable adjustment implies conditional heteroscedasticity that is in line with the empirical estimates, in contrast to the alternative models without fixed costs. Conditional heteroscedasticity

arises in the model with fixed costs despite the fact that the shocks are of constant variance precisely because fixed costs of adjustment lead to a procyclical IRF.

6 Conclusion

It is widely recognized that there are substantial fixed costs of durable adjustment. Nevertheless, business cycle models with fixed costs of durable adjustment have received limited attention due to their computational complexity. In this paper, we argue that while the introduction of fixed costs of adjustment is computationally challenging, it is not intractable, and it provides a significantly better fit for both macro and micro dynamics. An incomplete markets model with fixed costs and aggregate shocks is able to match the decline in durable expenditures in the Great Recession as well as business cycle volatilities more generally. In addition, since durable holdings are also a form of household wealth, fixed costs of adjustment make a fraction of wealth illiquid and costly to adjust. Since some households have large durable holdings with small levels of liquid wealth, this amplifies the response of non-durable consumption to shocks and implies that the model is able to match the relatively large decline in non-durable expenditures during the recession.

More importantly, introducing realistic adjustment costs is important for evaluating the efficacy of fiscal stimulus. Fixed costs of durable adjustment induce significant non-linear impulse responses of aggregate durable expenditures to shocks. Our estimates imply that during the Great Recession, durable expenditures were half as responsive to shocks to household income and wealth as during the late 90s. This time-varying IRF induces conditional heteroscedasticity in the time-series behavior of aggregate durable expenditures: the residuals from an AR regression on durable expenditures are greater during booms than during recessions. We show that such conditional heteroscedasticity is present in actual U.S. time-series data.

The procyclical IRF in the model is driven by changes in the distribution of households' desired durable holdings over the cycle. During a boom, households have more pent-up durable demand and are more likely to adjust their durable holdings. We test this implications of our model by estimating the empirical distribution of household level desired during holdings using household level panel data, and we find that the empirical distribution moves as predicted by our model. Thus, the micro data is consistent with a large time-varying IRF, as implied by our model.

29

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 (2010): "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model," *Cowles Foundation Discussion Paper.*

BAJARI, P., P. CHAN, D. KRUEGER, AND D. MILLER (2010): "A Dynamic Model of Housing Demand: Estimation and Policy Implications," *Mimeo.*

BAXTER, M. (1996): "Are Consumer Durables Important for Business Cycles?," *The Review of Economics and Statistics*, 78(1).

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).

CAMPBELL, J., AND Z. HERCOWITZ (2009): "Liquidity Constraints of the Middle Class," *Federal Reserve Bank of Chicago Working Paper 2009-20*.

DAVIS, M., AND J. HEATHCOTE (2005): "Housing and the Business Cycle," *International Economic Review*, 46(3).

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).

FISHER, J. (1997): "Relative Prices, Complementarities, and Co-movement Among Components of Aggregate Expenditures," *Journal of Monetary Economics*, 101(8).

FLAVIN, M. (2011): "Housing, Adjustment Costs, and Macro Dynamics," *CESinfo Economic Studies*.

GOMME, P., F. KYDLAND, AND P. RUPERT (2001): "Home Production Meets Time to Build," *Journal of Political Economy*, 109(5).

GREENWOOD, J., AND Z. HERCOWITZ (1991): "The Allocation of Capital and Time over the Business Cycle," *Journal of Political Economy*, 99(6).

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 (2009): "Housing and Debt Over the Life Cycle and Over the Business Cycle," *Mimeo*.

KAPLAN, G., AND G. L. VIOLANTE (2011): "A Model of the Consumption Response to Fiscal Stimulus Payments," *NBER Working Paper 17338*.

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-VILLAVERDE (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.

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).

MORRIS DAVIS, AND J. HEATHCOATE (2005): "Housing and the Business Cycle," *International Economic Review*, 46(3).

PETEV, I., L. PISTAFERRI, AND I. S. EKSTEN (2012): "Consumption and the Great Recession," *Analyses of the Great Recession*.

TAUCHEN, G. (1986): "Finite state markov-chain approximations to univariate and vector autoregressions," *Economics Letters*, 20(2).

7 Appendix 1: Empirical Results

7.1 Data Adjustments and Construction

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..

We have also experimented with using analogous constructions with chained GDP rather than constructing real GDP with the GDP deflator. However, the necessary chained GDP components are only available going back to 1995. Nevertheless, the results for the most recent recession are similar. In addition, it should be noted that due to the construction of the price deflators, the real series constructed by deflating the nominal series individually do not add up to the aggregate real series. This problem is more problematic the further from the base year we move. Since the current NIPA data uses 2005 as the base year, this is likely to introduce only small errors into our results for the most recent recession, but it makes the older series somewhat less reliable. The business cycle volatilities we estimate are sensitive to the exact procedure used to construct the individual consumption components (e.g. using non-durable consumption + services - housing services to get our measure of non-durable consumption does not obtain the exact same series as using PCE - housing services - durable expenditures even though these two measures contain the same components). These series are nearly identical in the 2000s, but diverge as we move further back in time.

Constructing durable expenditure rates requires constructing quarterly measures of the durable stock. Following Bachmann, Caballero, and Engel [2010], we construct measures of real quarterly durable expenditures 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 [2010] for the more detailed procedure.

7.2 Conditional Heteroscedasticity

In this section we show that durable expenditures exhibit conditional heteroscedasticity, rising in booms and falling in recessions. As in Bachmann, Caballero, and Engel [2010], we assume that our series of interest can be described by an AR process:

$$x_t = \sum_{j=1}^{p} \phi_j x_{t-j} + \sigma_t e_t,$$

where $x_t \equiv \frac{I_D}{D}$ is durable expenditures divided by the durable stock,[37] $e_t \sim$i.i.d. with zero mean and unit variance, and

$$\sigma_t = \alpha + \eta \overline{x}_{t-1}$$
$$\overline{x}_{t-1} = \frac{1}{k} \sum_{j=1}^{k} x_{t-j}.$$

That is, we allow the variance of the residuals in the AR process for durable expenditures to vary with past durable expenditures. This specification implies that the impulse response of x to e on impact at time t is given by $\alpha + \eta \overline{x}_{t-1}$. If $\eta = 0$ then the impulse response of x to e does not vary with past durable expenditures while $\eta > 0$ implies that the IRF rises with lagged expenditures.

We estimate the time-series model using quarterly data on $\frac{I_D}{D}$ from 1960-2010. The estimation follows a 2-stage procedure. In the first stage, we estimate the AR process via OLS to obtain residuals ε_t. The second stage then estimates by OLS η using

$$|\varepsilon_t| = \left(\frac{2}{\pi}\right)^{1/2} (\alpha + \eta \overline{x}_{t-1}) + \text{error.}$$

---

[37] The ratio of durable expenditures to the stock is stationary while durable expenditures are not.

We repeat the estimation for all combinations of $p, k \leq 12$ and choose the best fit, p^*, k^* using AIC. For more details on the time-series model, see Bachmann, Caballero, and Engel [2010][38]. Table 6 contains the time-series estimates. Both total durable expenditures as well as residential investment exhibit strongly significant[39] $\eta > 0$. The estimated $\eta > 0$ for consumer durables, but it is only marginally significant. While $\eta > 0$ implies that there is a statistically significant increase in the IRF with lagged expenditures, it does not imply that the increase is economically significant. In the 4th and 5th rows of Table 6, we report statistics that show that there is quantitatively large variation in the IRF across time. The maximum IRF is 2.82 times larger than the minimum IRF while the 95th percentile is 1.82 times higher than the 5th percentile. Thus, the estimated heteroscedasticity is both statistically and economically significant.

Table 6

Conditional Heteroscedasticity Data							
Series:	Total Dur. Exp.	Resid.	Cons. Dur.	Non Dur.	GDP	TFP	FF
η :	0.05	0.03	0.04	-0.003	0.002	-0.11	0.03
t-η :	2.63	2.04	1.52	-1.33	0.84	-1.40	0.83
Bootstrap p-value ($\eta > 0$) :	0.007	0.03	0.04	0.83	0.25	0.86	0.56
$\pm (\sigma_{\max}/\sigma_{\min})$:	2.82	2.52	1.67	1.85	1.25	1.61	3.89
$\pm (\sigma_{95}/\sigma_5)$:	1.82	1.65	1.59	1.39	1.20	1.35	1.12
no. obs. :	192	192	192	192	192	192	630

Table 7

Conditional Heteroscedasticity Models				
Series:	RBC	RBC Adj Costs	No Fixed Costs	W/ Fixed Costs
η :	-0.02	-0.003	0.02	0.08
t-η :	-0.34	-0.14	0.42	2.13
$\pm (\sigma_{\max}/\sigma_{\min})$:	1.25	1.24	1.17	1.83
$\pm (\sigma_{95}/\sigma_5)$:	1.15	1.14	1.11	1.49

It is important to note that while we have interpreted conditional heteroscedastic-

---

[38]In addition to the model we presented, their paper presents an alternative time-series model. We obtained similar results for this model, so for brevity we did not report these results.

[39]The bootstrap p-value row constructs bootstrapped p-values for $\eta > 0$, accounting for the fact that errors in the first stage estimation increase the standard errors in the second stage.

ity as a time-varying impulse response to aggregate shocks with a constant variance, an alternative interpretation is that aggregate shocks themselves are larger during booms than during recessions. We test for this by performing the same regressions on Baxter-King bandpass filtered GDP.[40] Table 6 shows that, as in the estimates in Bachmann, Caballero, and Engel [2010], there is no evidence of conditional heteroscedasticity for GDP. Presumably the aggregate shocks hitting Y should be similar to the aggregate shocks hitting D, so we interpret our results as evidence that it is not the shocks that drive heteroscedasticity, and is rather the mechanism that translates those shocks into durable expenditures that drives our estimates. As further evidence for this point, we also estimate the time-series model for changes in TFP[41] as well as the Federal Funds rate.[42]

To construct Figure 12, we allowed for a more flexible non-parametric second stage. We also bootstrapped the procedure to construct corresponding confidence intervals. To do this, we redraw residuals from the first stage to create one-thousand bootstrap samples when estimating the non-parametric second-stage. We tried a variety of bandwidths for the second-stage kernel estimator, and it did not change the qualitative conclusions.

In addition to these empirical estimates, we also compute heteroscedasticity estimates for our simulated models. Since we know that the true shocks in the model follow an AR(1) process, we report benchmark results restricted $k = 1$, $p = 1$, but results are not sensitive to this restriction. Table 7 shows that the frictionless model does not generate procyclical IRFs. If anything, the model without fixed costs implies $\eta < 0$. In contrast, the model with fixed costs exhibits conditional heteroscedasticity that is in line with the empirical estimates. The estimated $\eta > 0$, and the time-variation in the impulse response on impact is similar to that in the data.

8 Appendix 2: Model Solutions

8.1 RBC Model

---

[40]Unlike expenditure rates, GDP is non-stationary and so must be filtered. Using alternative filters did not substantively change the results.

[41]Available at http://www.frbsf.org/economics/economists/jfernald/quarterly_tfp.xls

[42]To increase the sample size we use monthly FF rates. While we could also use FF residuals or surprises, it is likely that the actual rate is more relevant for durable purchases as households should respond to both the anticipated and unanticipated component.

The RBC model (for the Cobb-Douglas case) has the following first order conditions:

$$
\begin{aligned}
C_t &: \quad vC_t^{v-1}D_t^{1-v}\left[C_t^v D_t^{1-v}\right]^{-\theta} = \lambda_t \\
K_{t+1} &: \quad \lambda_t = \beta E\left[\lambda_{t+1}\left(\alpha\frac{Y_{t+1}}{K_{t+1}} + (1-\delta_k)\right)\right] \\
D_t &: \quad \lambda_t = (1-v)D_t^{-v}C_t^v\left[C_t^v d_t^{1-v}\right]^{-\theta} + \beta E\lambda_{t+1}(1-\delta_d)
\end{aligned}
$$

The steady-state equations (fixing $Z = 1$) are then given by

$$
\begin{aligned}
Y &= K^\alpha \\
\lambda &= vC^{v-1}D^{1-v}\left[C^v D^{1-v}\right]^{-\theta} \\
1 &= \beta\left[\left(\alpha\frac{Y}{K} + (1-\delta_k)\right)\right] \\
\lambda &= (1-v)D^{-v}C^v\left[C^v D^{1-v}\right]^{-\theta} + \beta\lambda(1-\delta_d) \\
C &= ZK^\alpha - \delta_k K - \delta_d D
\end{aligned}
$$

Solving for the steady-state gives:

$$
\frac{C}{D} = \frac{v(1-\beta(1-\delta_d))}{1-v} \tag{1}
$$

$$
K = \left[\frac{1}{\alpha}\left[\frac{1}{\beta} - (1-\delta_k)\right]\right]^{\frac{1}{\alpha-1}} \tag{2}
$$

$$
Y = K^\alpha \tag{3}
$$

$$
D = \frac{Y - \delta_k K}{\frac{v(1-\beta(1-\delta_d))}{1-v} + \delta_d} \tag{4}
$$

We set $\delta_k = .022$ to match the long-run ratio of investment to capital, and we set $\alpha = .3$ to match the long-run labor share. We then pick β so that the capital stock in 2 implies an interest rate (in a decentralized economy) of 1%. We then pick v so that 1 implies $\frac{C}{\delta_d D} = 4$. Given the steady-state solution, we then compute the dynamic solution to the model by taking a second order perturbation around the steady-state solution.[43]

To investigate the sensitivity of our conclusions to our parameter values and cali-

---

[43] Additional details available on request.

Figure 13: Robustness of RBC results
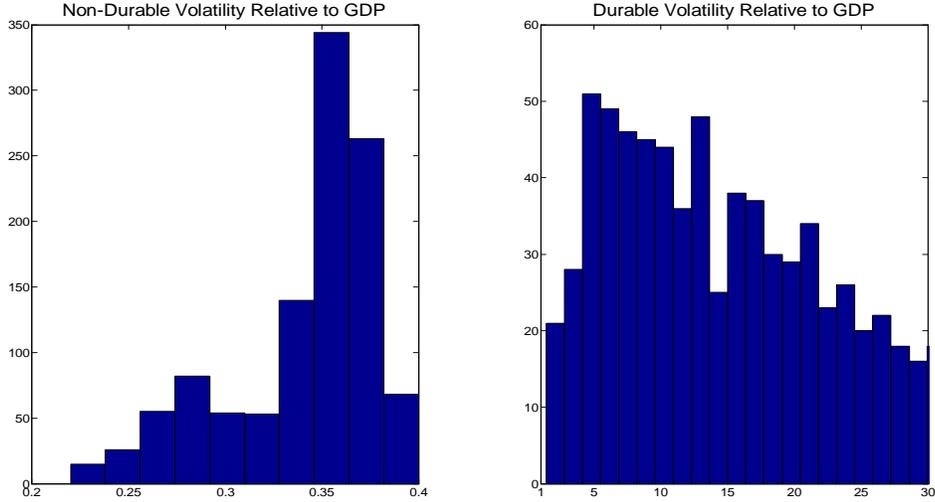
bration targets, we have also resolved the model using a variety of different specifications. We varied δ_d, choosing 12 values evenly spaced between .011 and .16, 7 values for θ evenly spaced between 0.5 and 3, three values for the $\frac{C}{\delta_d D}$ (4, 5 and 6), and we also allowed for a CES utility specification with elasticities ranging from 0.7 to 1.7.[44] Figure 13 shows the resulting standard deviations of durable and non-durable expenditures relative to GDP. None of the parameter specifications (including allowing for CES between durable and non-durable consumption) is able to match the empirical volatility of non-durable expenditures, and only a tiny substet of parameters delivers durable expenditures that are not substantially more volatile than the data.

8.2 Incomplete Markets Model

We describe the solution for the model with fixed costs. The solution for the model without fixed costs is similar. Reproducing the recursive formulation from the text,

---

[44]This includes the range typically estimated in the literature. (See Krueger and Fernandez-Villaverde [2010]). Expanding this range further did not matter for the general conclusions. For high enough elasticity, we can match the volatility of non-durable expenditures relative to output, but for these parameter values, durable expenditures are more than 40 times too volatile.

households solve

$$V\left(a_{-1}, d_{-1}, \eta; Z, K\right) = \max\left[V^{adjust}\left(a_{-1}, d_{-1}, \eta; Z, K\right), V^{noadjust}\right]$$

with

$$V^{adjust}\left(a_{-1}, d_{-1}, \eta; Z, K\right) = \max_{c,d,a} \frac{\left[c^v d^{1-v}\right]^{1-\theta}}{1-\theta} + \beta E_{\varepsilon,\xi} V\left(a, d, \eta'; Z', K'\right)$$

$$s.t.$$

$$c = wh\eta + (1+r)a_{-1} + d_{-1}\left(1-\delta_d\right) - d - a - f\left(1-\delta_d\right)d_{-1}$$

equilibrium conditions and prod. process

$$V^{noadjust}\left(a_{-1}, d_{-1}, \eta; Z, K\right) = \max_{\substack{c,a, \\ d\in\left((1-\delta_d)d_{-1}, d_{-1}\right)}} \frac{\left[c^v d^{1-v}\right]^{1-\theta}}{1-\theta} + \beta E_{\varepsilon,\xi} V\left(a, d, \eta'; Z', K'\right)$$

$$s.t.$$

$$c = wh\eta + (1+r)a_{-1} + d_{-1}\left(1-\delta_d\right) - d - a$$

equilibrium conditions and prod. process

We begin by substituting the budget constraint into the utility function to eliminate non-durable consumption as a choice-variable. We discretize η and Z using the algorithm of Tauchen [1986] and approximate $V^{adjust}\left(:,:,\eta; Z,:\right)$ and $V^{noadjust}\left(:,:,\eta; Z,:\right)$ as multilinear[45] functions in the two continuous idiosyncratic states and one continuous aggregate state. Initializing the grid for aggregate capital requires knowledge of the steady-state level of capital, so before solving the model with aggregate shocks, we solve for the steady-state of the model. The solution method is similar and simpler than the solution with aggregate shocks, so we only describe the latter:

Given an initial guess for the value functions and transition function, we solve for the optimal two-dimensional policy functions using a Nelder-Meade algorithm initialized from 3 different starting values to reduce the problems of finding local maxima in the policy function. The values of adjusting and not adjusting are compared, to generate the overall policy function and to update the overall value function. We iterate until the separate value functions change by less[46] than 0.001. Once the value functions have converged, we then solve for the optimal policy function an additional time on a finer grid, to use for simulation.

---

[45] We have experimented with cubic spline interpolation and have found that the speed advantages of linear interpolation appear to be worth potential decreases in accuracy (especially since fixed costs imply that the value functions may not be well approximated by cubic splines).

[46] Finer converge values didn't appear to affect the results.

We then simulate a panel of households and compute the evolution of the aggregate capital stock to update the aggregate transition rule $K' = \gamma_0 (Z) + \gamma_1 (Z) K$. We then repeat the above procedure until the coefficients in the value function change by less than 1%. Once the transition rule has converged, aggregate forecasts are highly accurate, with $R^2 > 0.999$. We have experimented with including the aggregate durable stock in the transition rule and found that it did little to improve forecasts, at considerable additional computational cost.

For the benchmark results, we use 25 grid points each for interpolating a_{-1} and d_{-1}, 3 grid points for interpolating K, and we discretize aggregate and idiosyncratic productivity using 5 and 7 points, respectively. (We recompute the model using 9 points for aggregate productivity when computing impulse response functions). We construct a finer grid with 90 points for a_{-1} and d_{-1} to compute the final policy function used for simulation. Thus, our fine policy function must be solved for approximately 2 million grid points and their associated expectations. Our simulation uses 25,000 households for 5,000 periods with an initial burnin of 1,000 periods.

In order to simulate the recession shock, we first feed Solow residuals computed for the U.S. economy into the model and then hit the model with a simulated 4% reduction in capital in the period corresponding to 2008q4. While our Krusell-Smith approximation is highly accurate on average, it may be a poor approximation following this one-time shock. We address this by comparing forecast errors following the shock to forecast errors in general. Indeed we find that forecast errors are substantially worse immediately following the one-time shock, but they nevertheless remain highly accurate. Computing across various simulations, the maximum forecast error following the one-time shock is 0.1% of the capital shock, which translates into less than a 0.1% error in the forecast interest rate. This is substantially larger than the maximum .01% forecast errors arising without the shock, but is still extremely small. Exogenously raising and lower the interest rate by 0.1% only has tiny effects on simulated outcomes, so we view these errors as economically inconsequential.

9 Appendix 3: Estimating Durable Gaps

This appendix discusses the estimation of the empirical adjustment gaps and hazards using PSID microdata. The procedure begins by first estimating the relationship between variables in the structural model that have empirical counterparts and unob-

served quantities of desired durable holdings. In the benchmark results, we estimate a linear regression of $c, d_{-1}, a, \frac{c}{d_{-1}}, \frac{a}{d_{-1}+a}$ on d^*_{adjust} and $d^*_{noadjust}$ (where d^*_{adjust} is the quantity of durables a household would choose today if its optimal policy is to pay the fixed cost and $d^*_{noadjust}$ is the optimal policy if it chooses to not pay the fixed cost). We find that this specification has R^2 of over 0.99 so that these observables in the model are sufficient to explain desired durable holdings.[47]

To estimate this procedure empirically, we must construct measures of these variables using PSID data. 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. Missing values are pulled through between observations. After constructing measures of c, d_{-1}, a per household member (with household head < age 65), we then deflate these nominal values using NIPA price indices (we use state-level price indices to adjust housing values), adjust for household age and remove a household level fixed effect.

This then gives us an empirical counterpart to $c, d_{-1}, a, \frac{c}{d_{-1}}, \frac{a}{d_{-1}+a}$ to which we apply the same structural relationship estimated from our model to generate empirical measures of $d^*_{adjust,estimate}$ and $d^*_{noadjust,estimate}$. The empirical gap in Figure 2, is defined as

$$\log\left(d^*_{adjust,estimate}/d^*_{noadjust,estimate}\right)$$

. Given estimates of the durable gap, we can then calculate the empirical adjustment hazard as a function of the estimated durable gap. In our benchmark results, we define durable adjustment as a self-reported house or vehicle sale together with a 25% 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

---

[47]Allowing for a more flexible functional form did not substantially enhance the predictive power. Similar results are also obtained by directly estimating the durable gap.

Figure 14: 90% Confidence Intervals for PSID Estimates
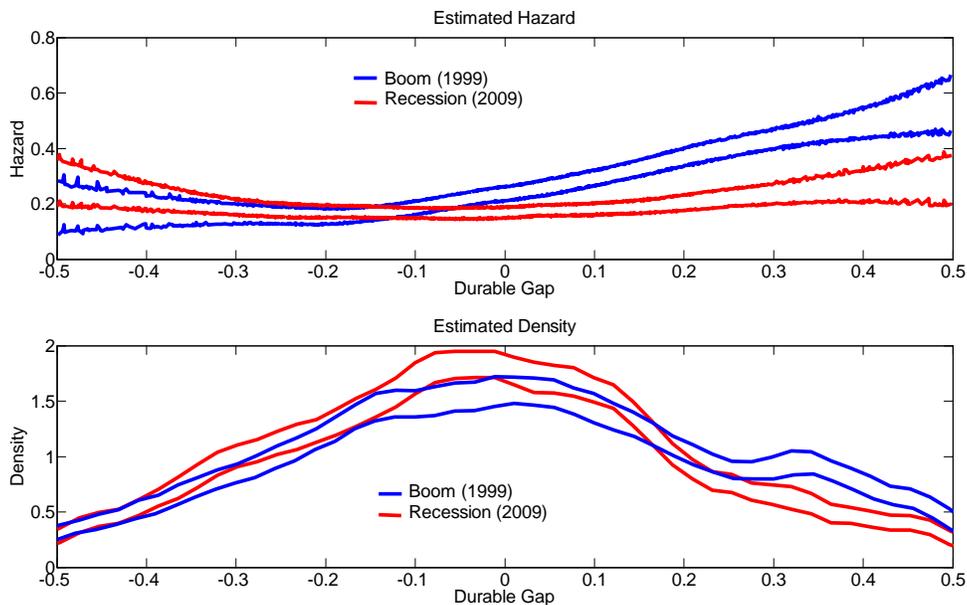
which may not lead to any substantial adjustment in the size of the stock. 3) Finally, and most importantly, self-reported adjustment indicators ask about adjustment over the previous three years while the sample is conducted every two years. This implies that the same adjustment may be counted twice. 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 13%, so a 25% threshold roughly splits this distance. This adjustment definition generates an overall two-year adjustment probability of roughly 30%.

In addition to the results shown in Figure 2, we have also estimated bootstrapped confidence intervals for the average density and hazard to assess the statistical significance of the business cycle shifts that we estimate. Figure 14 shows estimated confidence intervals,[48] which show that the time-variation we estimate is statistically significant.

While we believe our benchmark empirical specification is reasonable, we also as-

---

[48]With a bootstrap sample of 1000.

sess the robustness of our results to alternative choices. The main empirical object of interest is the slope of the empirical adjustment hazard as a function of the estimated durable gap. A positive relationship between the absolute value of the durable gap and the actual probability of adjustment is an indication that our estimated durable gap has predictive power. 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 empirical specifications.

$$adjust_{i,t} = \alpha + \beta abs(\log \left(d^*_{adjust,estimate,i,t}/d^*_{noadjust,estimate,i,t} \right)$$

In all cases, we continue to find a significant positive relationship between the durable gap and the probability of adjustment. The empirical specification is weakest when using only self-reported adjustment and not requiring any minimum threshold for a change in the size of the stock. For the reasons reported above, we believe this specification introduces substantial measurement error, but we nevertheless still find that it contains significant predictive power.

<div align="center">

Table 8

Specification	β	t-stat
Benchmark	0.26	18.1
Adjust Threshold of 0.1 instead of 0.25	0.19	12.4
Adjust Threshold of 0.4 instead of 0.25	0.29	21.4
Adjust Threshold of 0.01 instead of 0.25	0.06	4.7
No Adjust Threshold	0.04	2.9
Control for Year Fixed Effects in HH estimation	0.26	17.9
No adjustment for household size	0.28	19.9
Exclude business value from a	0.26	18.8
Restrict analysis to housing	0.09	8.7
Restrict analysis to vehicles	0.19	12.5
0.25 Threshold, Ignore self-reported adj	0.41	33.2
Do not pull-through missing obs	0.16	10.4
Don't use deflators to adjust values	0.24	16.8
Use $d^*_{noadjust,estimate,i,t} = d_{t-1}$	0.28	19.34
Do not control for HH age	0.09	6.4

</div>

In addition to PSID data, we have also performed our analysis using an alternative data set. The Bank of Italy Survey and Household Income and Wealth (SHIW) collects detailed information on demographics, households' consumption and assets. 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, and 2008 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 durable flows for means of transport, furniture, and jewelry. Net assets are defined as the sum of all deposits, CDs, securities, businesses and valuables minus the value of all liabilities to banks, corporations and other households.

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. We define durable adjustment as times when the household either had non-zero expenditure in a period on means of transport or a 100% 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 close to 35%, whereas a 100% threshold implies an annual frequency of adjustment closer to 10%, which we view as a more empirically reasonable value.