Stockholm School of Economics
Department of Economics
5350 Master's thesis in economics
Academic year 2014–2015

Time-Varying Uncertainty and Durable Adjustment

Carlo Zanella (40799)

Abstract

A large share of demand fluctuations over the business cycle are due to fluctuations in consumer spending on durable goods. Households adjust their durable holdings infrequently, and adjustment is even less likely when the economy is in a recession than when it is in an expansion. This thesis investigates whether countercyclical variation in income uncertainty can explain the procyclicality of durable adjustment. Higher income uncertainty in recessions may lead households to postpone durable adjustment until the next expansion. I present a simple model of durable adjustment and show that higher uncertainty leads to an overall decline in the frequency of adjustment. The effect of time-varying uncertainty on durable adjustment is quantified with an incomplete markets model. The results suggest that a more left-skewed distribution of income growth during recessions can account for a large part of the cyclicality in durable adjustment frequencies in PSID data. Countercyclical left-skewness also performs better at explaining stylized facts than alternative hypotheses although I find that aggregate income is an important determinant of durable adjustment.

Keywords: Uncertainty, durable goods, fixed adjustment costs, business cycle
JEL: B22, D8, D91, E21, E32

Supervisors:	Winfried Koeniger[†]and Kelly Ragan[‡]
Date submitted:	May 14, 2015
Date examined:	May 26, 2015
Discussant:	Simon Costa
Examiner:	Johanna Wallenius

[†]winfried.koeniger@unisg.ch, University of St. Gallen (Swiss Institute for Empirical Economic Research)
[‡]Kelly.Ragan@hhs.se, Stockholm School of Economics

Acknowledgments

I would like to thank my supervisors Winfried Koeniger and Kelly Ragan for their invaluable input, advice and suggestions during the entire course of writing this thesis.

I also appreciate helpful comments from thesis seminar participants at the Stockholm School of Economics.

Finally, I am grateful to Adomas Balkevičius, Sigita Žvirblytė, Srinidhi Srinivasan and Yunhao Lo for reviewing and proofreading my thesis.

Contents

List of Figures

List of Tables

1. Introduction

Consumer spending on durable goods is highly cyclical and significantly more volatile over the business cycle than non-durable consumption (Stock & Watson, 1999).[1] An important determinant of consumer durable expenditures is how often consumers adjust their durable good holdings. Berger and Vavra (2015) show that households adjust durable goods significantly less frequently in recessions than in expansions. To understand demand fluctuations and implement effective stabilization policies, it is therefore of paramount importance to understand what drives the volatility of durable goods spending.

In this thesis, I show that the higher levels of uncertainty usually associated with economic disruptions can account for procyclical durable spending. Bloom (2014) shows that various measures of economic uncertainty, such as stock market volatility, sales growth rates and subjective uncertainty about GDP growth reported by professional forecasters, are substantially higher in contractionary periods. Heightened uncertainty in these periods has a direct effect on households as well: Higher job displacement rates in recessions have a long-lasting negative effect on earnings of displaced workers (Jacobson, LaLonde, & Sullivan, 1993). Storesletten, Telmer, and Yaron (2004) find that the variance of annual earnings growth is highly countercyclical. More recently, Guvenen, Ozkan, and Song (2014) find that income growth is more left-skewed in recessions than in expansions.

I study the effect of time-varying uncertainty on the frequency of durable adjustment under the assumption that households face fixed costs when adjusting their durable holdings. Fixed adjustment costs, such as transaction fees, sales commissions or search costs, can account for the low frequency of adjustment observed in microeconomic data. Inaction becomes valuable since deviations of actual durable holdings from the frictionless optimum may be corrected by random fluctuations in household income rather than by costly action. Hence, to be inactive and tolerate such deviations to some degree brings less disutility than paying transaction costs in every period. In previous studies, a rise in uncertainty leads to even more inaction, since more volatility in income implies a higher probability that a given deviation from the optimum is erased by a change in income (see, for example, Bertola and Caballero (1990); Bertola, Guiso, and Pistaferri (2005); Grossman and Laroque (1990)). Furthermore, more volatility means that benefits of adjusting are nullified faster. A household might postpone its decision to buy a more expensive car today when the probability that it would not be able to afford it tomorrow increases.

In this thesis, I aim to answer the question if and how well countercyclical uncertainty is able to explain the less frequent adjustment of durable holdings in recessions. This thesis will try to answer these questions quantitatively. To get an intuition of the mechanisms at work, however, I begin by characterizing optimal durable adjustment decisions in a

[1] In the U.S., Black and Cusbert (2010) estimate a volatility of 2.7 relative to GDP for durable goods consumption and 0.8 for non-durable consumption.

simple two-period model with fixed adjustment costs. I find that, as expected, higher uncertainty leads to less frequent adjustment overall. Surprisingly, however, downward adjustment becomes more frequent. This result is explained by the precautionary savings motive, which becomes stronger when uncertainty is higher. An increase in uncertainty pushes households who are indifferent between inaction and adjusting downwards towards adjusting because liquidating durables allows to increase savings to buffer against more uncertain future income shocks. When durable goods depreciate, more household want to adjust their durable holdings upwards than downwards. Therefore, the decline in upward adjustment outweighs the increase in downward adjustment, which leads to less frequent adjustment overall. I numerically solve an infinite-horizon version of this problem and show that downward adjustment indeed increases in more uncertain periods. Thus, in my calibration, the effect of the stronger precautionary savings motive outweighs the value of waiting for new information.

In the quantitative part of the thesis, I address the question whether countercyclical uncertainty can explain the procyclicality of durable adjustment by comparing model simulated adjustment frequencies with actual adjustment in the data. I use an incomplete markets model based on Berger and Vavra (2015) to quantify the effect of time-varying uncertainty. In this thesis, I consider two forms of income uncertainty. First, I study the case that the variance of income changes is countercyclical, which has been found by Storesletten et al. (2004). In this specification, the dispersion of income growth is much higher during recessions than it is during booms. Both large income losses as well as large gains become more likely during recessions. Second, I consider the case that in recessions, the distribution of income growth is more left-skewed than in expansions. This has recently been documented by Guvenen et al. (2014). Countercyclical left-skewness means that in recessions, large income losses become more likely, whereas large gains become less likely. In this case, only the dispersion on the negative side of the income growth distribution becomes larger in a recession, but the total dispersion of income growth does not change.

I find that countercyclical left-skewness can explain both the volatility and the timing of fluctuations in durable adjustment. However, countercyclical variance leads to variations in adjustment frequencies which are not consistent with the data.

These findings have important implications for policy makers interested in stabilizing the business cycle. If cyclical income risk intensifies business cycles, it may be more effective to stabilize fluctuations in uncertainty, rather than implementing programs to promote consumer spending. This would suggest the need for further research into the causes of the increase in income uncertainty during recessions and what measures policy makers can take to offset it.

To assess the relative importance of changes in uncertainty, this thesis investigates an alternative explanation for procyclical durable adjustment which is put forth by Berger and Vavra (2015). They show that aggregate income shocks over the business cycle can lead to less frequent adjustments during recessions because households have less wealth and lower incomes. The intuition is that as incomes decrease, fewer households want to adjust their durables stock upwards and more want to adjust them downwards. By the same argument as before, depreciation implies that the former effect outweighs the latter.

Thus, in the quantitative part of this thesis, I consider three models: countercyclical left-skewness, countercyclical variance and aggregate income shocks. The models are first calibrated to match adjustment dynamics in the Panel Study of Income Dynamics (PSID) during the period 1999–2011 using an indirect inference model proposed by Berger and Vavra (2015). To assess their explanatory power, the models are then simulated for the period 1968–1996, when the PSID survey was conducted every year and it is possible to obtain data on annual adjustment frequencies. Evaluation of the models is based on comparing the simulated adjustment frequencies with actual frequencies in the data.

Related research. This thesis is based on theoretical models of durable adjustment pioneered by Grossman and Laroque (1990) and Bertola and Caballero (1990). Inspired by the literature on investment theory, these models assume the presence of fixed adjustment costs, which leads households to adjust their durable stocks infrequently. Fixed costs imply that reversing adjustment decisions is costly and thus that inaction becomes valuable.

In these models, optimal adjustment choices are characterized by so called (S,s) rules. This means that households are inactive as long as the durable stock is in a certain interval around a target durable stock. When random income shocks push durable holdings out of the inaction region, households adjust to their target durable stock. Such models have been extensively studied and put to empirical tests by Attanasio (2000) and Bar-Ilan and Blinder (1992), and both of these papers find that automobile purchases are consistent with households behaving according to (S,s) rules.

Recently, Berger and Vavra (2015) have developed a powerful calibration method which allows them to successfully explain the dynamics of durable adjustment in an incomplete markets model. The models with time-varying uncertainty used in this thesis build upon their work. In related research, Challe and Ragot (forthcoming) study how countercyclical unemployment risk affects the dynamics of aggregate non-durable consumption. McKay (2015) studies how the countercyclical left-skewness in income growth which has been found by Guvenen et al. (2014) affects aggregate non-durable consumption dynamics.

Several studies have investigated the effects of uncertainty on durable stock adjustment in this framework. Eberly (1994) study empirically the effect of idiosyncratic uncertainty on durable adjustment in the Grossman and Laroque (1990) model, which features stock market uncertainty but no idiosyncratic labour income risk. She finds that higher uncertainty widens the consumer's inaction band as predicted by the model. Foote, Hurst, and Leahy (2000) study the effect of uncertainty in a setting with labour income risk. They find that in PSID microdata, a household's income risk is negatively correlated with the probability to adjust its durable holdings. Bertola et al. (2005) show numerically that, although the probability to adjust conditional on the deviation from the optimal durable stock is decreasing in income risk, the relationship between unconditional adjustment probability and income risk is not necessarily monotonic if there is a drift in the durable stock, such as depreciation.

A limitation common to previous studies on durable adjustment is that uncertainty is assumed to be constant over time. However, as we have seen, uncertainty is known to vary substantially over the business cycle. Unlike previous analyses which statically solve models for different uncertainty parameters (see, for example, Bertola et al. (2005)), uncertainty is explicitly modelled as a stochastic process, which allows consumers to form expectations about future uncertainty.

Furthermore, an important difference between this thesis and previous studies is that most models do not account for non-durable consumption, which Berger and Vavra (2015) find to be important to explain durable adjustment in the data. The contribution of this thesis is to explicitly model non-durable consumption and assess the effect of uncertainty on durable adjustment in this context both theoretically and quantitatively.

Knotek II and Khan (2011) use regression analysis and find that periods of higher uncertainty are associated with lower consumer spending on durable goods. However, since uncertainty is highly correlated with aggregate factors, it is hard to interpret their results causally.

The rest of this thesis is organized as follows. Section 2 presents a simple model of durable adjustment and characterizes the effect of uncertainty on optimal adjustment. The dynamic model for the quantitative analysis is presented in Section 3, and Section 4 describes the data as well as the indirect inference procedure used to calibrate the model. In Section 5, the quantitative results are presented. The thesis concludes with Section 6.

2. A Simple Model Of Durable Adjustment

This section presents and solves a simple two-period household problem of durable adjustment with fixed costs. Unlike many of the models in the previous literature, this model lets households derive utility from a non-durable consumption good in addition to a durable good.

Household problem. Consider a household that lives for two periods, $t = 1, 2$, and derives utility from durable and non-durable consumption. Consumption preferences are given by

$$u(c_1, \bar{d}) + \beta\, \mathbb{E}\, u(c_2, d),$$

with $u_c > 0$, $u_d > 0$, $u_{cc} < 0$, $u_{dd} < 0$ and $u_{cd} \leq 0$.[2] The parameter β is the discount factor, c_t denotes non-durable consumption in period t, and $\bar{d}$ and d are the stock of durable goods in period 1 and 2, respectively. The household's preferences over non-durable and durable consumption are homothetic.

In period 1, the household has an initial endowment of liquid assets a and durables $\bar{d}$. The household makes a saving decision into an asset s with gross return $1 + r$ and can choose to adjust his durable stock d, which will affect utility in the next period. If the household adjusts its durable stock, it has to pay a fixed adjustment cost $F > 0$. In period 2, the household receives a stochastic income y_2, enjoys non-durable consumption $y_2 + (1 + r)s$ and durable consumption d.

In the following we are interested in characterizing how changes in uncertainty affects the household's decision to adjust its durable holdings.

Solution. Households solve

$$\begin{aligned} V(a) &= \max_{s,d} u(c_1, \bar{d}) + \beta\, \mathbb{E}\, u(c_2, d), \\ \text{s.t.} \quad & c_1 = a - s - (d - \bar{d}) - A(d - \bar{d}) \\ & c_2 = y_2 + (1 + r)s, \end{aligned}$$

where A is adjustment costs which are zero when $d = \bar{d}$ or F otherwise. Formally,

$$A(d - \bar{d}) = \begin{cases} 0 & \text{if } d = \bar{d} \\ F > 0 & \text{otherwise.} \end{cases}$$

Considering inaction and adjustment separately, we can define the value functions for both cases as

$$\begin{aligned} V^N(a) &= \max_{s} u(a - s, \bar{d}) + \beta\, \mathbb{E}\, u(y_2 + (1 + r)s, \bar{d}), \text{ and} \\ V^A(a) &= \max_{s,d} u(a - s - (d - \bar{d}) - F, \bar{d}) + \beta\, \mathbb{E}\, u(y_2 + (1 + r)s, d), \end{aligned}$$

where the value of inaction and adjustment is denoted by $V^N(a)$ and $V^A(a)$, respectively.

[2] Subscripts to the utility function denote its partial derivatives; for instance, $u_c = \partial u / \partial c$.

The value function V is given as the upper envelope of both value functions, that is $V(a) = \max\{V^N(a), V^A(a)\}$.

The first-order conditions for interior optima are given in the non-adjustment case by

$$u_c(c_1^N, \bar{d}) = \beta(1+r)\,\mathbb{E}\,u_c(c_2^N, \bar{d}))$$

and in the case of adjustment by

$$\begin{aligned} u_c(c_1^A, \bar{d}) &= \beta(1+r)\,\mathbb{E}\,u_c(c_2^A, d)) \\ u_c(c_1^A, \bar{d}) &= \beta\,\mathbb{E}\,u_d(c_2^A, d)). \end{aligned}$$

Since the household has to pay a fixed cost F to adjust its durable stock, it does not adjust for every asset endowment. For the household to prefer to adjust, the gain from adjustment must be higher than the disutility of forgone consumption from paying the costs. For instance, if the target durable stock d is sufficiently close to $\bar{d}$, adjustment costs exceed the benefits of adjustment, and hence the household prefers inaction.

The household's optimal adjustment decision is characterized by a so called (S, s) policy. This means that there is a set of asset endowments $NA = \{a \geq 0 : V^N(a) \geq V^A(a)\}$ for which inaction is optimal. This set is called the inaction region or the non-adjustment set. Since both value functions are concave and continuous, this set is convex, closed and bounded, which allows us to write $NA = [\underline{a}, \overline{a}]$, where $\underline{a}$ and $\overline{a}$ are the bounds, or the cut-offs, of the inaction region. Note that if $\bar{d}$ is sufficiently large, then $\underline{a} > 0$.

By the continuity of the value functions, the cut-offs of the inaction set $\underline{a}$ and $\overline{a}$ satisfy the "value-matching" condition

$$V^N(a) = V^A(a), \tag{1}$$

or equivalently,

$$u(c_1^N, \bar{d}) + \beta\,\mathbb{E}\,u(c_2^N, \bar{d}) = u(c_1^A, \bar{d}) + \beta\,\mathbb{E}\,u(c_2^A, d).$$

Figure 1 shows how the value functions of adjustment and non-adjustment determine the inaction region.

Figure 1 – Value of Adjustment and Non-Adjustment

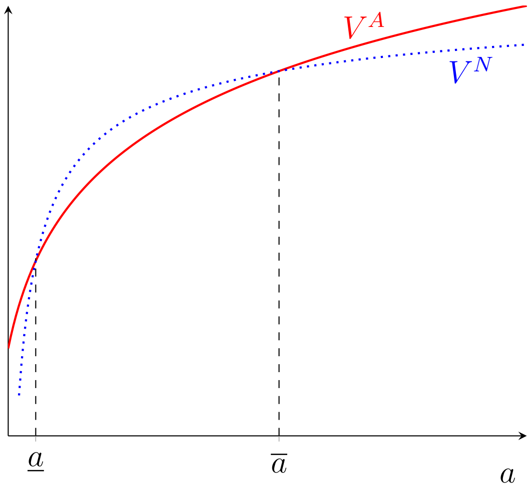

In the case when the household adjusts, the durable stock is clearly adjusted upwards at the upper cut-off and downwards at the lower cut-off.

Next, I determine how a change in uncertainty affects the inaction region. I consider the case where by a change in uncertainty we mean a mean-preserving increase in the variance σ^2 of period 2 income y_2. Appendix A shows that all results derived in this section also hold for a marginal increase in the left-skewness of the distribution.

By the implicit function theorem, the value-matching condition (1) locally defines a cut-off $\tilde{a} \in \{\underline{a}, \overline{a}\}$ as a differentiable function of σ^2 with derivative given by

$$\frac{d\tilde{a}}{d\sigma^2} = -\frac{\frac{\partial}{\partial\sigma^2}V^A(a) - \frac{\partial}{\partial\sigma^2}V^N(a)}{\frac{\partial}{\partial a}V^A(a) - \frac{\partial}{\partial a}V^N(a)} \tag{2}$$

$$= -\beta\frac{\frac{d}{d\sigma^2}\,\mathbb{E}[u(c_2^A, d) - u(c_2^N, \bar{d})]}{u_c(c_1^A, \bar{d}) - u_c(c_1^N, \bar{d})}, \tag{3}$$

where the second line follows from applying the envelope theorem.

Consider the lower cutoff $\underline{a}$. Increasing assets just a little bit makes the household prefer inaction, and thus $\partial/\partial a V^A(a) < \partial/\partial a V^N(a)$, which, by the envelope theorem, directly implies $u_c(c_1^A, \bar{d}) < u_c(c_1^N, \bar{d})$. It follows that the denominator of expression (3) is negative at the lower cut-off. Next consider the numerator. Approximate the utility function by

$$u(c_2, d) \approx u(\bar{c}_2, d) + u_c(\bar{c}_2, d)\epsilon + \frac{1}{2}u_{cc}(\bar{c}_2, d)\epsilon^2,$$

where $\bar{c}_2 = \mathbb{E}[c_2]$ and $\epsilon = c_2 - \bar{c}_2$. Taking expectations gives

$$\mathbb{E}[u(c_2, d)] \approx u(\bar{c}_2, d) + u_c(\bar{c}_2, d)\,\mathbb{E}[\epsilon] + u_{cc}(\bar{c}_2, d)E[\epsilon^2] = u(\bar{c}_2, d) + \frac{1}{2}u_{cc}(\bar{c}_2, d)\sigma^2.$$

Therefore,

$$\frac{d\tilde{a}}{d\sigma^2} \approx -\frac{\beta}{2}\frac{u_{cc}(\bar{c}_2^A, d) - u_{cc}(\bar{c}_2^N, \bar{d})}{u_c(c_1^A, \bar{d}) - u_c(c_1^N, \bar{d})}. \tag{4}$$

If preferences exhibit decreasing absolute risk aversion with respect to non-durable consumption independent from durable consumption, then $-u_{cc}(c,d)/u_c(c,d)$ is a decreasing function of c.[3] Letting $\chi(c)$ denote that function, we can thus write $u_{cc}(c,d) = -\chi(c)u_c(c,d)$ and get

$$\frac{d\tilde{a}}{d\sigma^2} \approx -\frac{\beta}{2}\frac{-\chi(\bar{c}_2^A)u_c(\bar{c}_2^A, d) + \chi(\bar{c}_2^N)u_c(\bar{c}_2^N, \bar{d})}{u_c(c_1^A, \bar{d}) - u_c(c_1^N, \bar{d})}. \tag{5}$$

At the lower cut-off, first-order conditions imply

$$\beta(1+r)\,\mathbb{E}[u_c(c_2^A, d)] = u_c(c_1^A, \bar{d}) < u_c(c_1^N, \bar{d}) = \beta(1+r)\,\mathbb{E}[u_c(c_2^N, \bar{d})]. \tag{6}$$

The opposite inequality holds at the upper cut-off. Therefore, we have that $\bar{c}_2^A > \bar{c}_2^N$ at the lower cut-off and $\bar{c}_2^A < \bar{c}_2^N$ at the upper cut-off.

I make the assumption that at the lower cut-off $\underline{a}$ it holds that $u_c(\bar{c}_2^A, d) < u_c(\bar{c}_2^N, \bar{d})$ and at the upper cut-off $\overline{a}$ that $u_c(\bar{c}_2^A, d) > u_c(\bar{c}_2^N, \bar{d})$. That is I assume that the inequalities which hold for expected marginal utilities of consumption in period 2 also hold for marginal utilities of expected consumption.[4]

From this assumption about marignal utilities of expected consumption and $\chi(\bar{c}_2^A) < \chi(\bar{c}_2^N)$, it follows that the numerator of expression (5) is positive. Hence, we have $d\tilde{a}/d\sigma^2 > 0$. The lower cutoff moves to the right as uncertainty increases. Households who are indifferent between adjusting their durable holdings downwards and inaction choose adjustment when uncertainty marginally increases.

This result is striking since previous models of durable adjustment imply that the inaction region widens at both bounds (see Grossman and Laroque (1990), Bertola et al. (2005) and Eberly (1994)). The difference between the current and previous models is due to the presence of a non-durable consumption good. The intuition behind this result is that households who are on the margin of selling their durable goods and face an increase in uncertainty want to insure against the more volatile consumption in the next period. The precautionary savings motive leads households close to the lower cut-off to save more

[3]This assumption is, for example, satisfied for any decreasing absolute risk aversion (DARA) utility function with a Cobb-Douglas aggregator. The constant relative risk aversion (CRRA) utility function, which is frequently used in the literature on durable goods, thus satisfies this condition.

[4]This assumption always holds in numerical simulations with the CRRA utility function presented in the subsequent section and when income is normally distributed or is a mixture of two normally distributed variables with close enough means. However, I neither provide a formal proof nor the conditions under which this assumption holds.

by selling off durables. It is important to emphasize that this model is not dynamic and does not consider the benefits that accrue from waiting until future income shocks correct the gap by itself. However, it sheds light on a mechanism affecting households' adjustment decisions which has not been discussed so far: More uncertainty makes previously indifferent households willing to liquidate durables to increase their buffer-stock savings. In the quantitative part of this thesis, I show that under my calibration and assumptions about consumption preferences, this result holds also in a dynamic environment.

We will see next that an increase in uncertainty also increases the upper cut-off. From $\frac{d}{da}V^A(\overline{a}) > \frac{d}{da}V^N(\overline{a})$ it becomes clear that both the numerator and the denominator in (5) flip signs at $\overline{a}$. This result implies that upward adjustment decreases in response to an increase in uncertainty. The intuition is that households on the margin of adjusting upwards want to save more to offset the higher consumption risk in the next period.

I have shown that both cut-offs increase as uncertainty increases. But which increase is of larger magnitude? The answer to this question follows readily from the characterization of the cut-off's derivative (5). Since consumption choices increase monotonically in assets, non-durable consumption is higher at the upper cut-off in the case of adjustment as well as inaction. Decreasing absolute risk aversion then implies that, in absolute terms, the numerator becomes small faster than the denominator. The derivative is thus lower at the upper cut-off than at the lower cut-off, which implies that the inaction region of the household becomes smaller as uncertainty increases — a result that is even more at odds with previous models. Though interesting, this effect is not as important as it may seem, since it is unlikely that household's positions relative to the inaction band are uniformly distributed.

We have established that higher uncertainty leads households at the lower cut-off to adjust and households near the upper cut-off to be inactive. Hence, downward adjustment increases and upward adjustment declines. Which of these effects dominate the overall effect on adjustment depends on the distribution of asset holdings a relative to the durable holdings $\bar{d}$. A simple way to understand how this distribution might look like is to add depreciation to the model. The model's implications in the presence of depreciation are very intuitive. If durable holdings depreciate over time, households are more likely to adjust their durable stock upwards than downwards. A larger number of households are near the upper cut-off, and thus the positive effect of an increase in uncertainty on downward adjustment is offset by the negative effect on adjusting upwards. Overall adjustment thus declines.

It is interesting to note that the implications of an increase in uncertainty are qualitatively the same as a decrease in aggregate income. Declining aggregate income is captured by decreasing $\mathbb{E}[y_2]$ in the model. Clearly, a decrease in aggregate income shifts the inaction region to the right by exactly that amount. Therefore, in recessions, the effects of changes in uncertainty and aggregate variables go in the same direction.

This simple model only considers two periods and does not capture the benefits of waiting that are present in a dynamic environment, where adjustment decisions can be made in each period. In a dynamic version of this model, there will also be a "wait-and-see" effect on the inaction region. For this reason and to quantify the effects, this thesis proceeds by setting up a comprehensive dynamic model which is solved numerically and evaluated on its ability to explain fluctuations in durable adjustment frequencies in the data. We will see that the implications on upward and downward adjustment are confirmed numerically in the dynamic model.

3. The Dynamic Model

3.1. Model Description

The model is an incomplete markets model based on Berger and Vavra (2015). In essence, is an infinite-horizon version of the simple model introduced in the last section.

The economy is populated by a continuum of ex-ante identical households who live forever. In each period, households earn a stochastic income and derive utility from the consumption of non-durable and durable goods. The stock of durable goods depreciates over time. Furthermore, like in the simple model presented in the previous section, adjusting durable holdings requires the household to pay a fixed cost.

Households solve

$$\begin{aligned} &\max_{\{c_t\}_{t=0}^{\infty},\{a_t\}_{t=1}^{\infty},\{d_t\}_{t=0}^{\infty}} \mathbb{E}\sum_{t=0}^{\infty}\beta^t u(c_t, d_{t+1}) \\ \text{subject to}\quad & a_{t+1}+c_t+d_{t+1} = (1+r)a_t + (1-\delta)d_t + y_t - A(d_t, d_{t+1}) \\ & \text{and } a_t \geq 0, \end{aligned}$$

where a_t, d_t, c_t and y_t denotes assets, durable stock, non-durable consumption and income, respectively, in period t. The parameter β is the quarterly discount factor and δ is the depreciation rate. The fixed adjustment cost is proportional to the stock of durables and current income. It is given by

$$A(d,d') = \begin{cases} 0 & \text{if } d' = [1-\delta(1-\chi)]d \\ F^d(1-\delta)d + F^t y & \text{otherwise,} \end{cases}$$

where F^d is the fraction of current durables and F^t is the fraction of income that the household has to pay when adjusting his durable stock. The F^t parameter can be used to model adjustment costs which also depend on the business cycle. Following Berger and

Vavra (2015), I also use "required-maintenance" parameter χ, which allows households to offset part of the depreciation without having to adjust durable stocks.[5]

For the numerical analysis, I use the constant relative risk aversion (CRRA) utility function with Cobb-Douglas aggregation:

$$u(c, d) = \frac{(c^\alpha d^{1-\alpha})^{1-\gamma} - 1}{1 - \gamma}.$$

The contribution of this thesis is to model a binary uncertainty process s_t which determines the distribution of income growth. More specifically, I let the distribution of y_{t+1} depend on (y_t, s_t). The process s_t is described later in detail. For now it is sufficient to assume that (y_t, s_t) is Markov, which will allow for a recursive formulation of the problem.

Note that prices of durable goods are normalized to unity for all periods because this thesis is not interested in modelling price effects. Surely, the rapid increase and subsequent fall of US house prices affect economic incentives to buy houses. However, since the goal of this thesis is to study fluctuations over the business cycle, long-term developments in the housing market can be safely ignored.

The model period is quarterly. At the beginning of each period, the household learns his state (a_t, d_t, y_t) and makes a consumption and saving choice (a_{t+1}, d_{t+1}, c_t). With this timing, d_{t+1} is the household's end-of-period durable stock.

The problem can be rewritten recursively as

$$\begin{aligned} V(a, d, y, s) &= \max_{c, d', a'} u(c, d') + \beta\, \mathbb{E}[V(a', d', y', s')|y, s] \\ \text{subject to} \quad a' + c + d' &= (1 + r)a + (1 - \delta)d + y - A(d, d'), \\ a &\geq 0 \\ &\text{and the law of } (y', s') \text{ given } (y, s). \end{aligned}$$

For computational feasibility, I will solve the model only in partial equilibrium, treating interest rates and wages as fixed. This is not necessarily innocuous. There is a large body of literature studying the role of equilibrium assumptions in models of lumpy firm investment. Like households when adjusting their durable holdings, firms face fixed costs when adjusting their capital stock, which leads to infrequent adjustment. Among others, Khan and Thomas (2008, 2003) find that general equilibrium dampens the aggregate implications of fixed costs. They argue that the requirement of household consumption smoothing constrains the extent to which firms can react to aggregate shocks. This leads to smooth aggregate series and dampens changes in the adjustment decisions across the business cycle.

The same argument could hold for the economy studied in this thesis. Adding general

[5] I refer to Berger and Vavra (2015) for a discussion of the parameter χ.

equilibrium conditions to the current model would impose constraints on households. However, Berger and Vavra (2015) find that when households have both illiquid durables d in addition to liquid assets a as saving instruments, this direct link between aggregate series is broken. They introduce general equilibrium into their model and show that the lumpiness in durable adjustment persists. Hence, it is reasonable to assume that the results derived in this thesis from partial equilibrium models are not undone in general equilibrium.

3.2. The Business Cycle

To keep the model conceptually simple and computationally feasible, the business cycle is modelled as a two-state Markov chain $S = \{s_t\}_{t=0}^{\infty}$ with values in $\{R, E\}$. That is, the economy is either in a recession (R) or in an expansion (E). Uncertainty jumps up immediately when transitioning from an expansion to a recession, and drops when the economy goes into a recession. Moreover, for all households the state of the economy s_t is observable at time t.

I estimate the process from NBER recession indicators using the method of maximum likelihood. The transition matrix is depicted in Table 1.

Table 1 – Transition Probabilities for Business Cycle Process S

	R	E
R	0.829	0.171
	(0.059)	(0.059)
E	0.040	0.960
	(0.015)	(0.015)

Standard errors are given in brackets.

3.3. Idiosyncratic Income

As research on earnings risk over the business cycle shows, uncertainty is extremely countercyclical (Guvenen et al., 2014; Storesletten et al., 2004). The contribution of this thesis is to model income with time-varying idiosyncratic risk and to quantify the effect of uncertainty shocks on durable adjustment.

First, let us assume that incomes of households are independent from each other and that the mean income is uncorrelated with the business cycle. In the next section, I describe how to add aggregate income to the model. The logarithm of individual earnings $\log y_t$ follows an AR(1). The innovations of the process have mean zero and their distribution is described by a cumulative distribution function $F_{s_t}(\epsilon)$, depends on the state of the business cycle s_t. That is

$$\log y_{t+1} = \rho \log y_t + \epsilon_t \text{ with } \epsilon_t \sim F_{s_t} \text{ and } \mathbb{E}[\epsilon_t] = 0.$$

This specification allows for a time-varying distribution of income growth $\log y_{t+1} - \log y_t$.

In the model with constant uncertainty, innovations are normally and identically distributed in both states of the economy. This simple AR(1) is estimated by Berger and Vavra (2015) using PSID data. They estimate a standard deviation of 0.1 and a persistence of 0.975.

For the process with countercyclical uncertainty, I will explore two distribution functions for the innovation terms ϵ_t to model the two different kinds of income risk estimated by Storesletten et al. (2004) and Guvenen et al. (2014).

The innovation of the process with countercyclical variance is a normal distribution with mean zero and variance depending on the state of the business cycle. I use the parameters estimated by Storesletten et al. (2004) to calibrate the income process. Since they estimate a yearly process, I rescale the innovations in order to match the variance of the process. This is described in appendix E.[6] I set the persistence to 0.95, the standard deviation in the expansion to $\sigma_E^{\text{Storesletten}} = 0.09$ and the standard deviation in the recession to $\sigma_R^{\text{Storesletten}} = 0.15$. Figure 2a shows the distribution of earnings growth of a household with an average income for the process with countercyclical left-skewness.

[6] To ensure comparability between the models, all income processes are further scaled such that $\mathbb{E}[y_t|s_t] = 1$ for all s_t, t.

Figure 2 – Income Growth Distribution

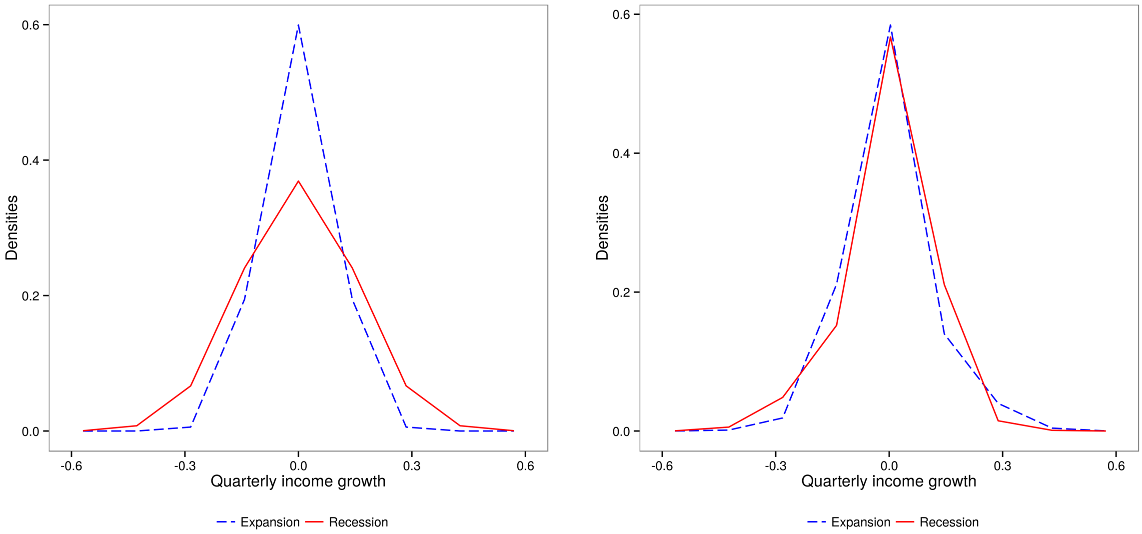

(a) Countercyclical variance (b) Countercyclical left-skewness

Source: Author's calculations

The process with countercyclical left-skewness uses a process proposed and estimated by Guvenen et al. (2014). I will explore this type of risk by using their parameterization of a mixture of two normally distributed variables with different means and variances. The means of the both components depend on the state of the business cycle. More formally, the idiosyncratic income component is modelled as

$$\log y_{t+1} = \rho \log y_t + \epsilon_t \text{ with } \epsilon_t \sim \begin{cases} N(\mu_{1s_t}, \sigma_1) & \text{with probability } p_1 \\ N(\mu_{2s_t}, \sigma_2) & \text{with probability } 1 - p_1. \end{cases}$$

The parameters are similarly rescaled and shown in Table 2.[7] Figure 2b shows the distribution of earnings growth of an average-income household for the process with countercyclical left-skewness.

[7]The discretization of the process with time-varying skewness introduces a bias in the conditional mean when using standard methods. I resort to a simulation based approximation scheme described in appendix D.

Table 2 – Parameters of the Quarterly Countercyclical Left-Skewness Process

Parameters	
ρ	0.979
μ_{1E}	0.049
μ_{2E}	-0.011
μ_{1R}	-0.042
μ_{2R}	0.038
σ_1	0.133
σ_2	0.0004
p_1	0.490

3.4. Aggregate Income

Aggregate income can be added as a further AR(1) process y_t^a independent of each household's income. Household income then consists of an aggregate and a household-specific component. Total income is thus given by

$$\log y_t = \log y_t^a + \log y_t^i,$$

with idiosyncratic income y_t^i as specified in the previous section.

Since in the model with constant uncertainty aggregate income does not need to be correlated with the state of the economy s_t, we can take aggregate income to be a simple AR(1) with normally distributed innovations. Therefore,

$$\log y_t^a = \rho^a y_{t-1}^a + e_t \text{ with } e_t \sim N(0, \sigma_e).$$

I estimate this process using hp-filtered GDP data from 1960–2013, which delivers $\rho^a = 0.876$ and $\sigma_e = 0.008$.[8]

However, this aggregate income process cannot be used for the models with time-varying uncertainty. Aggregate income would be uncorrelated with the state of the economy, which is not realistic. It should be increasing in expansions and decreasing in recessions. To model aggregate income in a more realistic way, I use a Markov switching AR(1) model with switching intercepts and persistence in the spirit of Hamilton (1989), who first showed that these models are very accurate in describing business cycles. The switching model is given by

$$\log y_t^a = c_{s_t}^a + \alpha_s \log y_{t-1}^a + e_t \text{ with } e_t \sim N(0, \sigma_e).$$

[8] All estimations are documented in Appendix E.

Unlike Hamilton (1989), I consider the states s_t to be observable, which makes it easy to estimate the process using maximum likelihood. I estimate $\mu_R = -0.008$, $\mu_E = 0.002$, $\alpha_R = 0.45$, $\alpha_E = 0.40$, and $\sigma_e = 0.009$.

4. Data and Calibration

4.1. Data

To calibrate the dynamic models, I use the 1997–2011 sample of the Panel Study of Income Dynamics (PSID) as compiled by Berger and Vavra (2015).[9] This period is selected because from 1997 on, the PSID contains detailed information on durable and non-durable consumption and assets holdings. As we will see, such detailed data are necessary for the calibration procedure to match the dynamics of durable adjustment implied by the model to the data.

From 1997, the PSID survey was conducted every two years. Following Berger and Vavra (2015), I restrict the panel to home-owning households of which the household head is at most of age 65. The key variables contained in the data set are liquid assets, durable holdings, non-durable consumption as well as a dummy variable indicating whether a household adjusts its durable stock in a given period. Data on durable holdings include houses and vehicles, which together cover almost all durable spending. The adjustment dummy is defined using a combination of survey questions and actual changes in durable holdings. A household is considered to adjust its durable stocks if it reports to have moved or sold either its home or any of its vehicles in the last three years and if the absolute change in its durable holdings exceeds a threshold of 20%.[10] This combination of self-reported adjustment and actual changes in durable holdings is used because either measure alone is not reliable to identify adjustment. There are two reasons for this. First, the survey questions on adjustment refers to the preceding three years, which might lead to counting the same adjustment twice, since the survey is conducted biennially. Moreover, these indicators would count as an adjustment when a household moves to another city for a new job, for example, although the size of the durable stock remains unchanged. Second, measurement errors in reported durable holdings may bias adjustment frequencies. Thus, defining adjustment using a combination of these indicators reduces the probability of spurious adjustments.

The quantitative performance of the models is compared with 1968–1996 PSID adjust-

[9] I use the Stata code of Berger and Vavra (2015) to generate the 1999–2011 panel.

[10] This threshold value is suggested by Berger and Vavra (2015). The median change of durable holdings is 4% if the household reports no adjustment and 40% if adjustment is reported. A threshold of 20% roughly splits this distance. I experimented with thresholds of 15% and 25% and, although the level of mean adjustment changes, the calibration as well as the numerical results remain unaffected.

ment data. During this period, the PSID survey was conducted yearly, which allows to measure the fluctuations of adjustment frequencies over the business cycle. The time series of adjustment frequencies is constructed in a similar fashion as before, but using only information on housing since information on other durable goods was not collected before 1997. The adjustment dummy is set to one if the absolute change in durable holdings exceeds 12% and the household reports to have sold its house in the preceding year.

I use seasonally-adjusted real GDP from the U.S. Bureau of Economic Analysis (BEA)[11] data to estimate aggregate income. The definition of recessionary periods is taken from the National Bureau of Economic Research (NBER). Following Guvenen et al. (2014), on whose research the calibration of the process with countercyclical left-skewness is based, I treat the 1980–1983 period as one single recession rather than two shorter ones.

4.2. Calibration Method

To decrease the dimensionality of the parameter space, I do not calibrate all the parameters of the model but use benchmark estimates from previous research. Following Berger and Vavra (2015), I set $r = 0.0125$, $\gamma = 2$ and $\delta = 0.018$. I deviate from their paper by calibrating the parameter β to the data instead of using an estimate from the literature. The reason is that different income distributions lead to differences in the level of durable stocks and asset holdings in the models. Since these levels are sensitive to the choice of β, I include it in the calibration routine. Moreover, I use a local optimization algorithm, which is more efficient than Berger and Vavra (2015)'s grid search and makes adding parameters less costly.

I use an indirect inference method, pioneered by Berger and Vavra (2015), to calibrate the parameters. This method allows to match microeconomic adjustment dynamics in the data. The idea is the following. As we have seen in the analysis of the simple adjustment model, fixed adjustment costs cause households not to adjust their durable holdings in every period. For each household, there exists some durable stock $d^\star$ that it adjusts to if the value of adjustment is higher than the value of inaction. This difference between this target durable stock and its actual durable stock is called the *gap*. More precisely, the gap x is defined as $x = \log(d^\star) - \log(d) \approx (d^\star - d)/d$.

A higher gap implies that the household is closer to one of the bounds of the inaction set. The higher the gap, the higher the opportunity cost of forgone utility and the more likely it is that the household is willing to bear the fixed adjustment cost. Therefore, for a given gap, the probability to adjust, which is called the *adjustment hazard*, is increasing in the gap. This is also one of the characteristics of the models analyzed by papers such as Grossman and Laroque (1990) or Bertola and Caballero (1990).

[11] NIPA Table 1.1.6

If the desired durable stock and thus the gap were observed in the data, this would allow us to match the adjustment dynamics of the model to the data by choosing parameters to let the distribution of gaps and the adjustment hazard in the model target their counterparts in the data. Unfortunately, this gap is not observable. However, Berger and Vavra (2015) offer a way to use this method nonetheless. They suggest to impute gaps to the data using the model. Once gaps are imputed, we can compute the distribution of gaps and the adjustment hazard in the data, which then allows to match the distribution in the model with the distribution the model implies for the data.

To impute gaps we can proceed as follows. Let z_m denote the vector of state variables in the model. The model implies a policy function $d^\star_m = g_m(z_m)$. Note that the necessary state variables consumption, assets and durable holdings are observable from PSID data. The model's policy function implies a gap for any given observation. We can use this to impute gaps to the data by applying the policy function to the observations in the data. That is we can define $\hat{d}^\star_d = g_m(z_d)$. Since adjustment decisions are known, this also allows us to compute conditional hazards.

For the more interested reader, I refer to the discussion of this method by Berger and Vavra (2015). I describe this procedure in detail in Appendix C. One of the few data sets which contains all the necessary variables on durable holdings, non-durable consumption is the PSID during 1999–2011.[12]

4.3. Calibration Results

The calibrated parameters are presented in Table 3. Column 1 reports the calibrated parameters for the model which has constant uncertainty. These parameter values are used for the model with aggregate income shocks.[13] Column 2 and 3 report the calibration results for the models with countercyclical left-skewness and countercyclical variance, respectively. The parameters are similar for constant uncertainty and for countercyclical left-skewness. The model with countercyclical variance has slightly different parameters, although it is hard to say if these differences are significant since we do not know the standard errors of the point estimates.[14]

[12]According to Berger and Vavra (2015), the Italian Survey of Household Income and Wealth (SHIW) is the only other data set besides the PSID 1999–2011 sample also containing the necessary variables.

[13]To remain comparable with Berger and Vavra (2015) and to reduce computation time, the constant uncertainty model is not calibrated using actual U.S. aggregate income shocks.

[14]It would be straightforward to obtain standard errors from bootstrapping. To economize on computation time, and since precise estimates of the parameter values are of no particular interest in this thesis, I refrain from doing so.

Table 3 – Calibrated Parameters

	Uncertainty		
	Constant	Time-Varying	
		Left-Skewness	Variance
Parameters	(1)	(2)	(3)
α (Utility flow non-dur)	0.87	0.87	0.83
χ (Required maintenance)	0.81	0.84	0.88
F_d (Fixed cost stock)	0.0596	0.0563	0.0387
F_t (Fixed cost time)	0.0024	0.0022	0.0012
β (Discount factor)	0.963	0.966	0.972
σ_ϵ (Measurement error)	0.01	0.04	0.07
Minimized objective	0.1837	0.2435	0.4113

Notes: Col. 1 reports the calibration results from the model with constant uncertainty. This model is calibrated without aggregate shocks. Cols. 2 and 3 report the calibration results from the models with time-varying left-skewness and variance, respectively. These models are calibrated using the actual sequence of binary recession indicators from NBER.

Figures 4–6 show the fit of all three models. In each figure, the upper graph shows the density of gaps in the model and in the data. The bottom part shows the adjustment hazard, that is the frequency of adjustments conditional on the gap. As predicted by the theory, the higher the gap, the more likely is the household to adjust its durable stock.

Overall, all models fit the data relatively well. The gap distribution is approximately concentrated around zero and adjustment hazard is increasing in the gap. In particular, the fits of the constant-uncertainty model and the model with countercyclical left-skewness are excellent. However, the countercyclical variance fits the data slightly worse than the other two models.

When comparing the relative fits, it is important to keep in mind that the income process in the constant-uncertainty model is explicitly targeting income in the 1999–2011 PSID data, whereas the processes estimated by Storesletten et al. (2004) and Guvenen et al. (2014) are based on an older PSID sample and an entirely different data source, respectively. Therefore, the income processes in the countercyclical-uncertainty models are not calibrated to the sample, which might put them at a relative disadvantage when comparing them with the baseline model.

The mean of the adjustment frequencies is more than two times higher in all three models than in the data. The average adjustment rates in the model with and constant uncertainty, countercyclical left-skewness and countercyclical variance are 20.7%, 20.6% and 21.7%, respectively, compared to only 9.0% in the data.[15]

[15] This is also the case in the simulations by Berger and Vavra (2015).

The worse fit of the countercyclical-variance model already foreshadows that it is probably not the right model to explain adjustment behaviour in the data, at least using the parameterization of Storesletten et al. (2004).

Figure 4 – Fit of Model with Constant Uncertainty

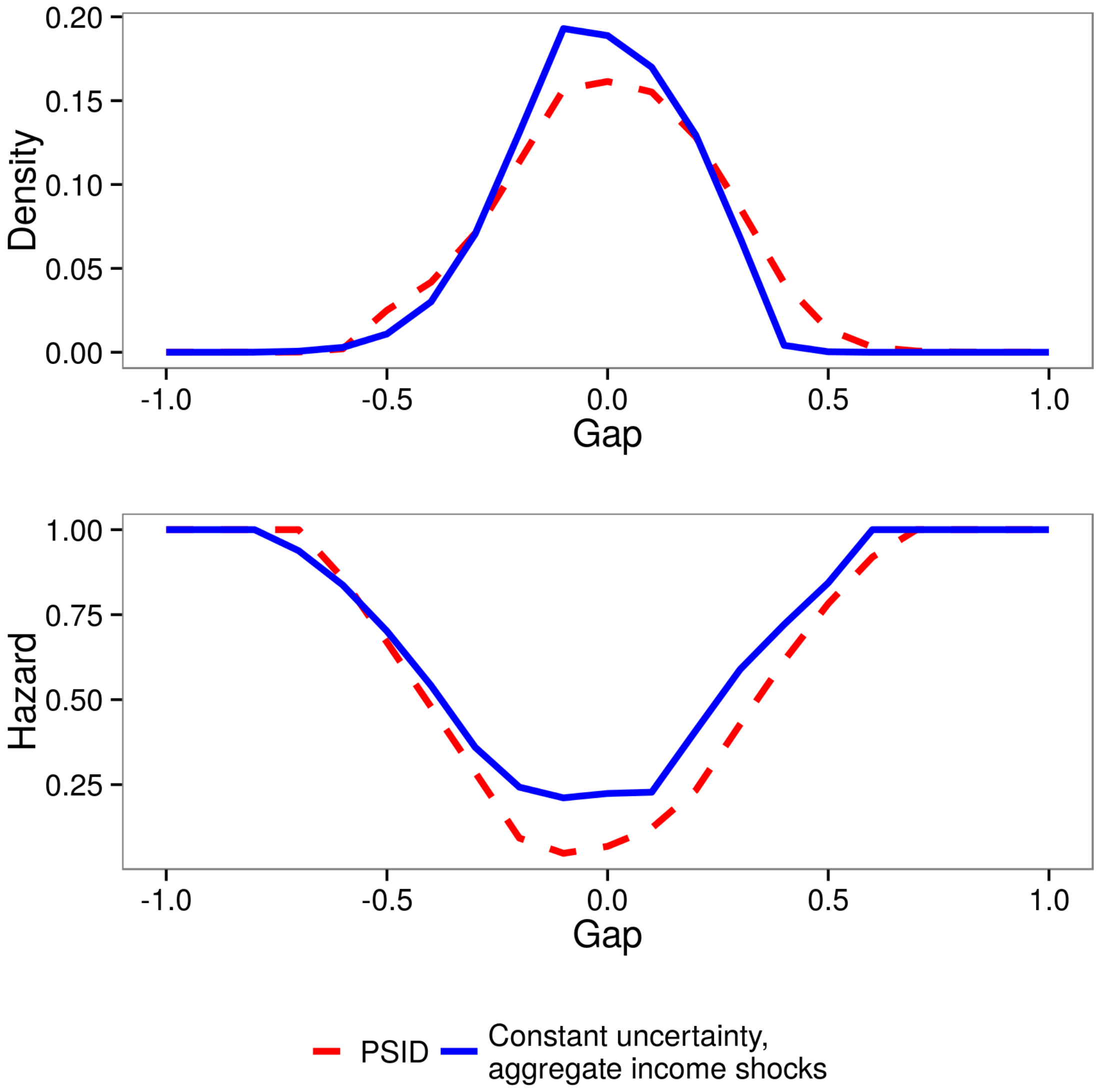

Source: Author's calculations

Figure 5 – Fit of Model with Countercyclical Left-Skewness

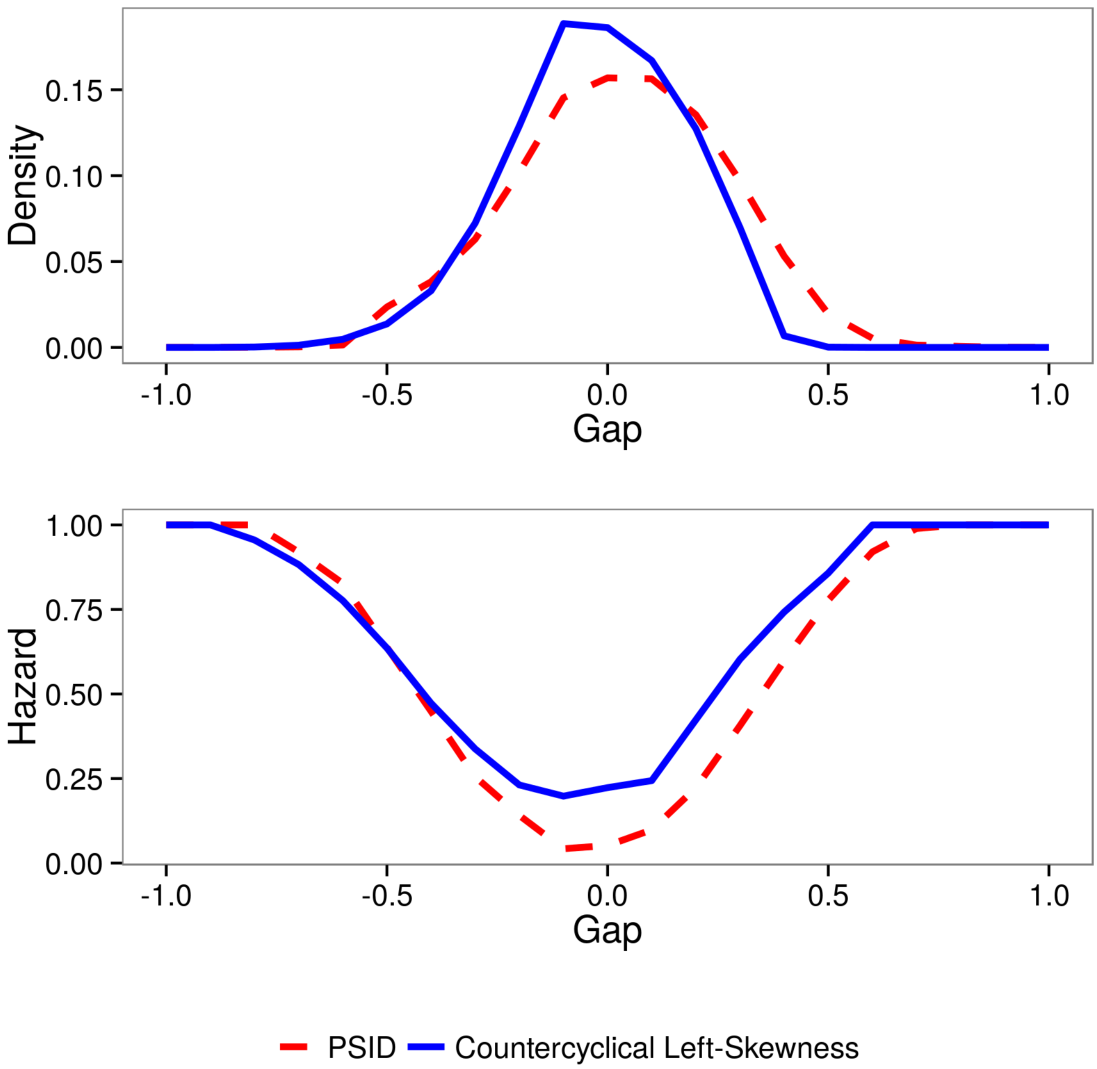

Source: Author's calculations

Figure 6 – Fit of Model with Countercyclical Variance

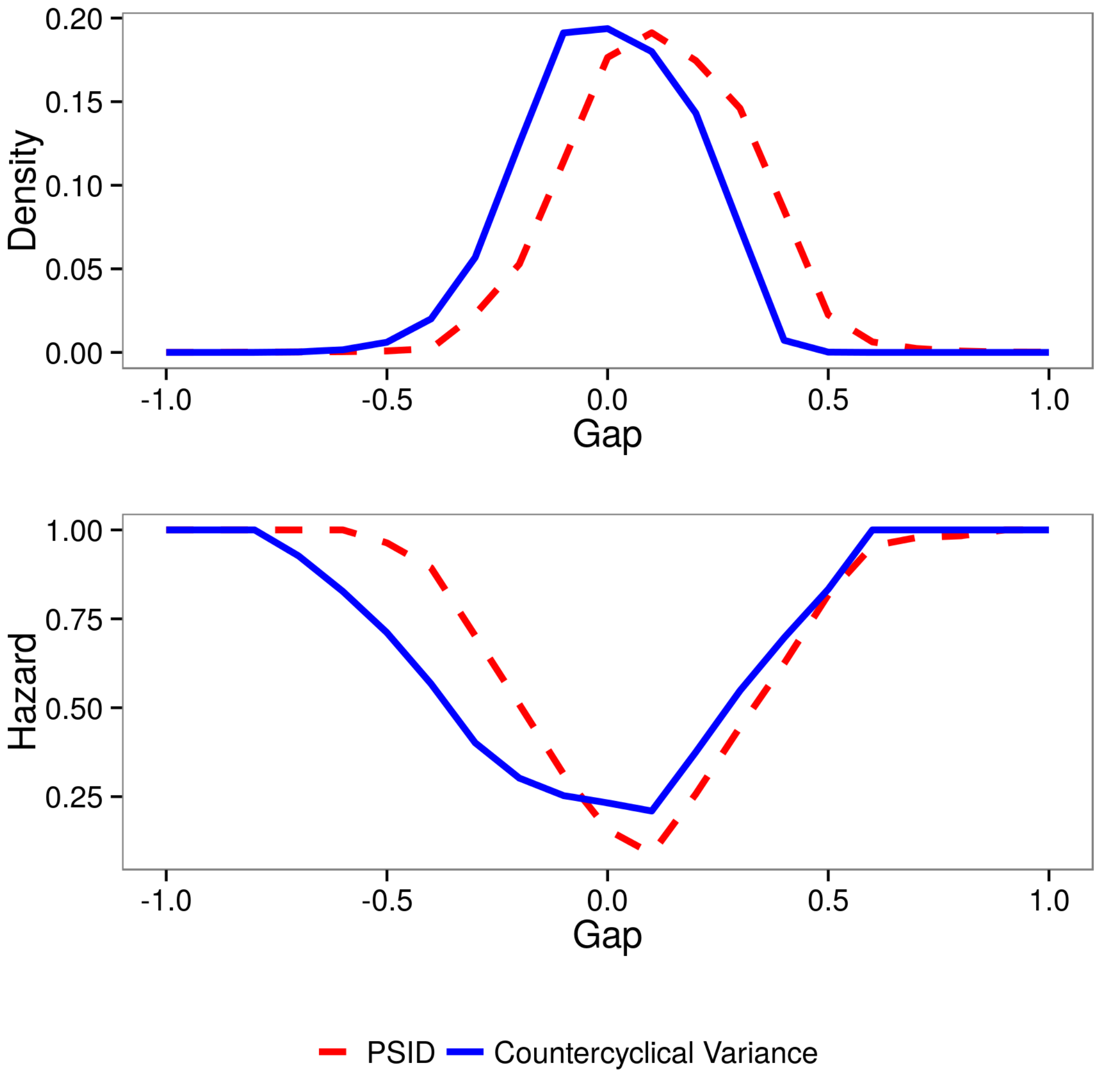

Source: Author's calculations

5. Results

The models are evaluated on their ability to predict yearly adjustment frequencies in PSID data from 1968 to 1996. To simulate adjustment frequencies, all models are solved using the calibrated parameters. Then I pick recession indicators $s_{1968q1}, \ldots, s_{1996q4}$ and aggregate income shocks $\log y^a_{1968q1}, \ldots, \log y^a_{1996q4}$ to reproduce actual U.S. business cycles and hp-filtered U.S. GDP from 1968 to 1996. Using these shocks, I simulate a panel of 200.000 households and aggregate it to annual frequency. Yearly adjustment frequencies are calculated using the same threshold criterion as for the data but without self-reported adjustment indicators (see Section 4.1).[16]

5.1. Countercyclical Left-Skewness

Figure 7 shows average adjustment frequencies for each year during the period 1968–1996 in PSID data as well as the aggregated series from the simulation of the model with countercyclical left-skewness. To get an idea of the precision of the measurements of adjustment frequencies in the data, the figure depicts the 95% confidence interval for the estimates.[17]

I begin by summarizing the stylized facts of durable adjustment in the data. Yearly adjustment is very infrequent: only 4.3% of households adjust their durable stock in a period of one year.[18] In recessionary periods, adjustment frequencies are significantly lower than otherwise. Adjustment responds very abruptly to the start of a recession or an expansion. At the onset of a recession, adjustment rates go down abruptly. Consider the 1980–1983 recession, when adjustment jumps down from its highest level to the lowest in the sample period. A similarly swift response is seen after the recession when adjustment rates shoot up rapidly.

In the model with countercyclical left-skewness, the adjustment frequencies behave as predicted by the simple model of durable adjustment presented in Section 2. Like in the data, adjustment jumps down when the economy enters a recession and shoots up when it enters an expansion.

[16] Appendix B documents in detail how the model is solved and simulated.

[17] The confidence interval is constructed by approximating the estimator's distribution with a normal distribution, that is assuming that estimates $\hat{p}_t$ are normally distributed with standard error $\sqrt{\frac{\hat{p}_t(1-\hat{p}_t)}{n_t}}$. Note that n_t is varying because the number of households satisfying the age and home-owner restriction in each survey is changing.

[18] Figure 7 and subsequent figures depicting series of adjustment frequencies do not show absolute adjustment frequencies but only deviations from the mean. The reason is that all of the models predict that households adjust more often than they do in reality. Since this thesis focuses on fluctuations of adjustment over the business cycle, the question of how well the mean is matched is not of great interest.

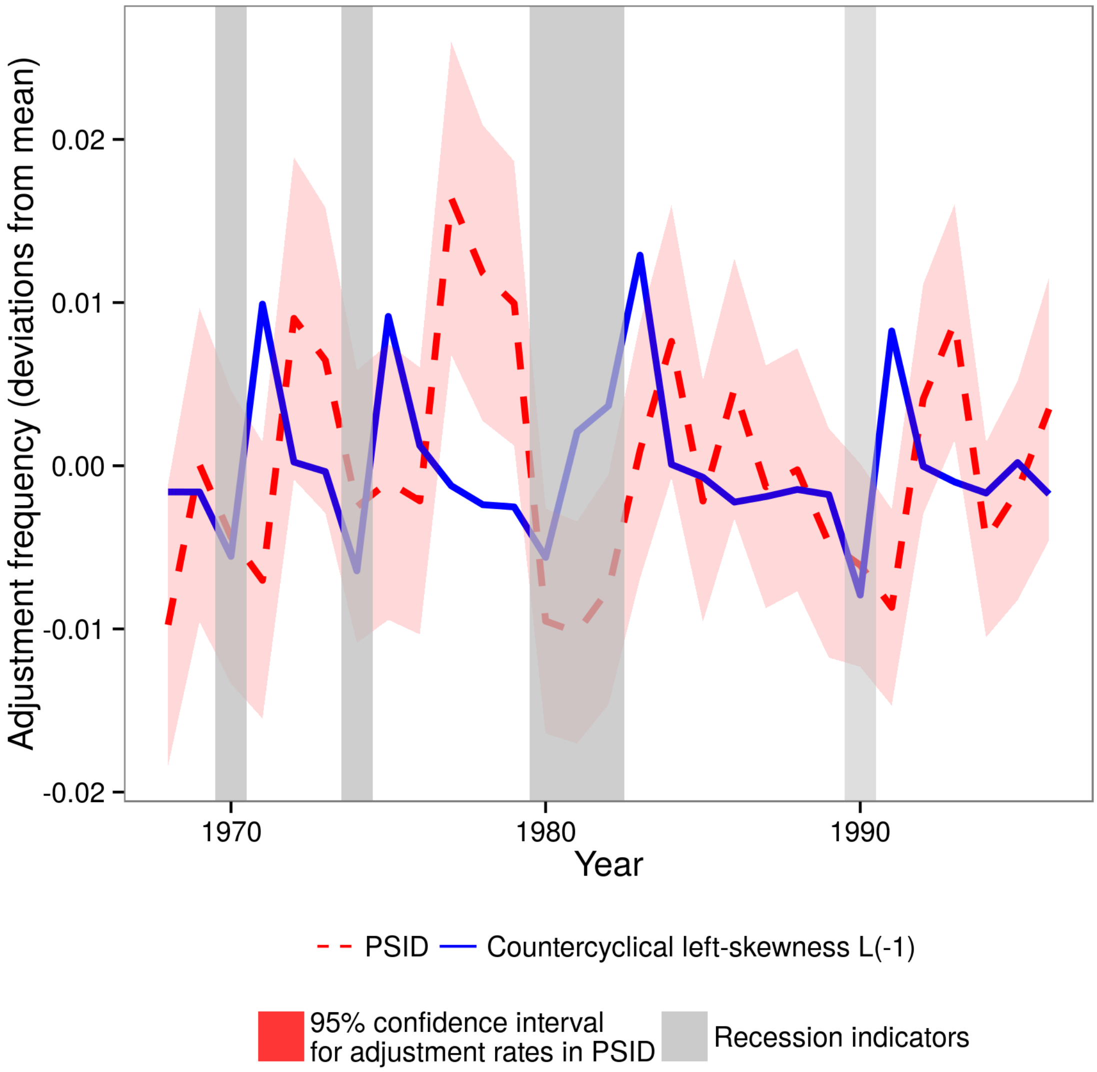

Figure 7 – Mean Adjustment Frequencies in Period 1968–1996 in PSID and Model with Countercyclical Left-Skewness

Source: Author's calculations

However, the model misses somewhat on the timing of the jumps in adjustment rates. While households in the model react immediately to changes in uncertainty, there appears to be a delay of about one year in the data. This can still be consistent with the model. Households usually do not know when the state of the economy switches into a recession until later when its effects become more apparent. It is plausible that households perceive the increase in uncertainty later than when the economy swings into a recession. Another explanation is that the increase in income uncertainty does not exactly coincide with the beginning of a recession as the National Bureau of Economic Research (NBER) dates it. I am, however, not able to provide evidence on either of these hypotheses.

There is also some discrepancy between the magnitude of the cycles in adjustment frequencies in the model in the data. In the data, fluctuations of adjustment frequencies are higher over the business cycle than in the model. This is likely due to aggregate factors which affect the decision to adjust durable holdings. In particular, as I have shown in Section 2, the implications on durable adjustment from declines in expected income and increases in uncertainty are very similar. Procyclical fluctuations in aggregate income and wealth would therefore magnify the effect of uncertainty and vice versa.

Another pattern implied by the model is that adjustment frequencies are increasing during recessions because the higher left-skewness of income shocks increases the dispersion of durable gaps. Thus, more households are forced by their changes in income to adjust. This leads to a rise in adjustment frequencies over time, which can be observed in the model during the recession from 1980 to 1983. This effect is less substantial in the data than it is in the model, although we should note that there is only one recession of this duration in the data, which is not enough to conclusively assess the dispersion effect in the data. It is also possible, that the effect might be neutralized in the data by aggregate factors.

The analytical result derived in Section 2 suggests that downward adjustment increases and upward adjustment decreases at the beginning of a recession. This result also holds in the numerical simulation of the dynamic model and is consistent with PSID data. This can be seen in Figure 13 in Appendix F, which shows the simulated frequencies of upward and downward adjustment in the model and in the data.

5.2. Countercyclical Variance

The simulated adjustment frequencies in the model with countercyclical variance are shown in Figure 8. Clearly, the fluctuation of the adjustment rates in the model are too extreme compared to the data.

Figure 8 – Mean Adjustment Frequencies in Period 1968–1996 in PSID and Model with Countercyclical Variance

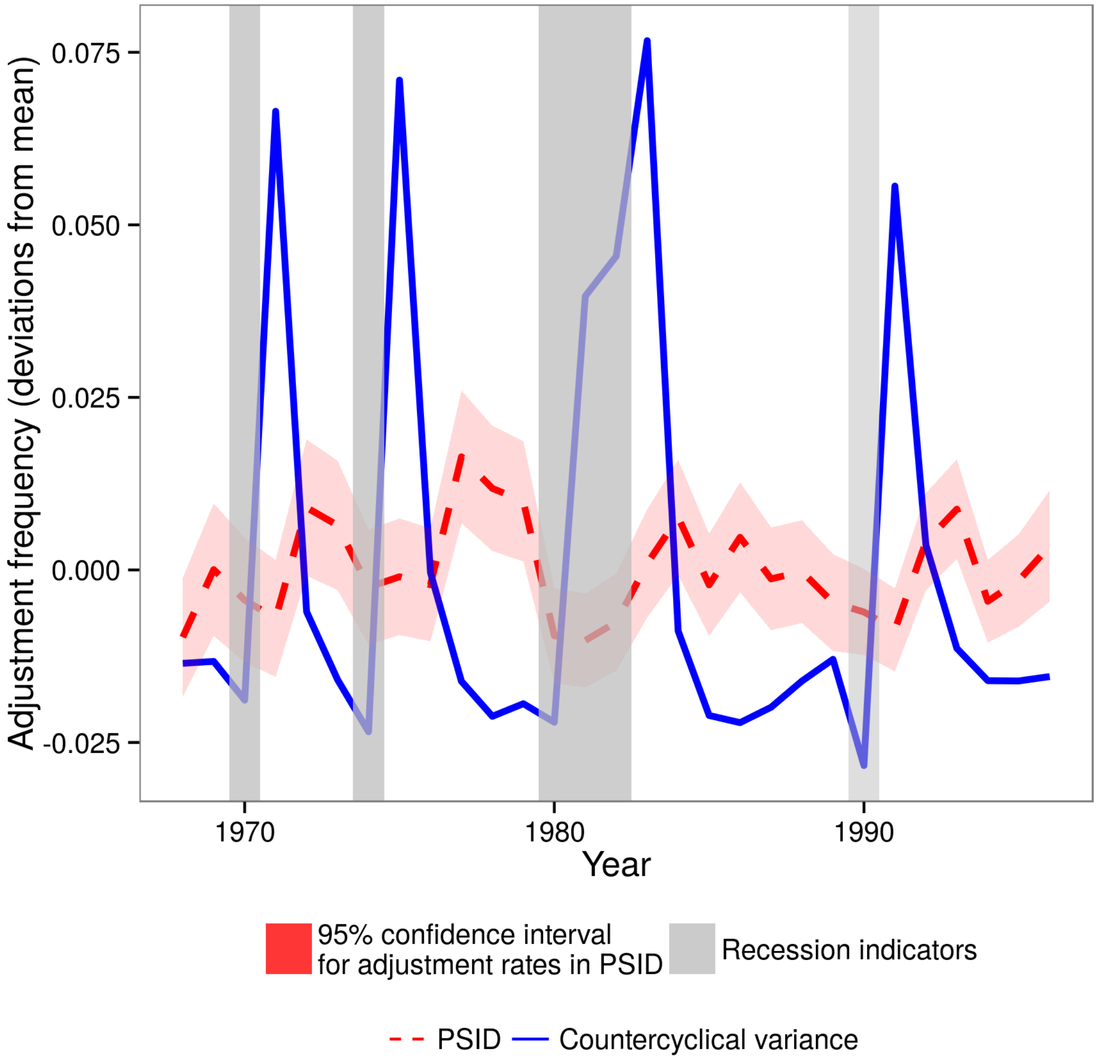

Source: Author's calculations

To understand the high volatility of adjustment rates, it is useful to separate the dispersion effect from the effect on the household's inaction region. The effect of the higher uncertainty on the inaction region is reflected in the immediate jump downwards when entering a recession. It is relatively small in the scale of this figure and its magnitude is similar to the one in the model with countercyclical left-skewness. However, the dispersion effect due to higher variance of income growth is much higher. The impact of this effect is clearly visible at the beginning of 1980s, where adjustment rates skyrocket during the recession.

The strikingly different performance of the model with time-varying variance compared to time-varying left-skewness suggests that the kind of uncertainty that households face is important. It is also possible that the estimates of Storesletten et al. (2004) are not very precise. To rule out that this result is due to the scaling of the income process to quarterly frequency, as described in Appendix E, I repeated the simulation with different scaling factors. However, the magnitude of the fluctuations does not change significantly. Hence, it is safe to conclude that the results do not give support to this specification of time-varying income risk.

5.3. Alternative Hypothesis: Aggregate Income Shocks

An alternative explanation for cyclical durable adjustment is that declines in aggregate income and wealth lead to lower desired durable holdings and thus less frequent upward adjustment. To investigate the hypothesis, I simulate the model with constant income uncertainty by picking the sequence of actual GDP and feeding them into the simulation as aggregate income shocks. Figure 9 shows the average adjustment frequencies in each year during the period 1968–1996 in PSID data and as simulated from the model.

Figure 9 – Mean Adjustment Frequencies in Period 1968–1996 in PSID and Model with Constant Uncertainty

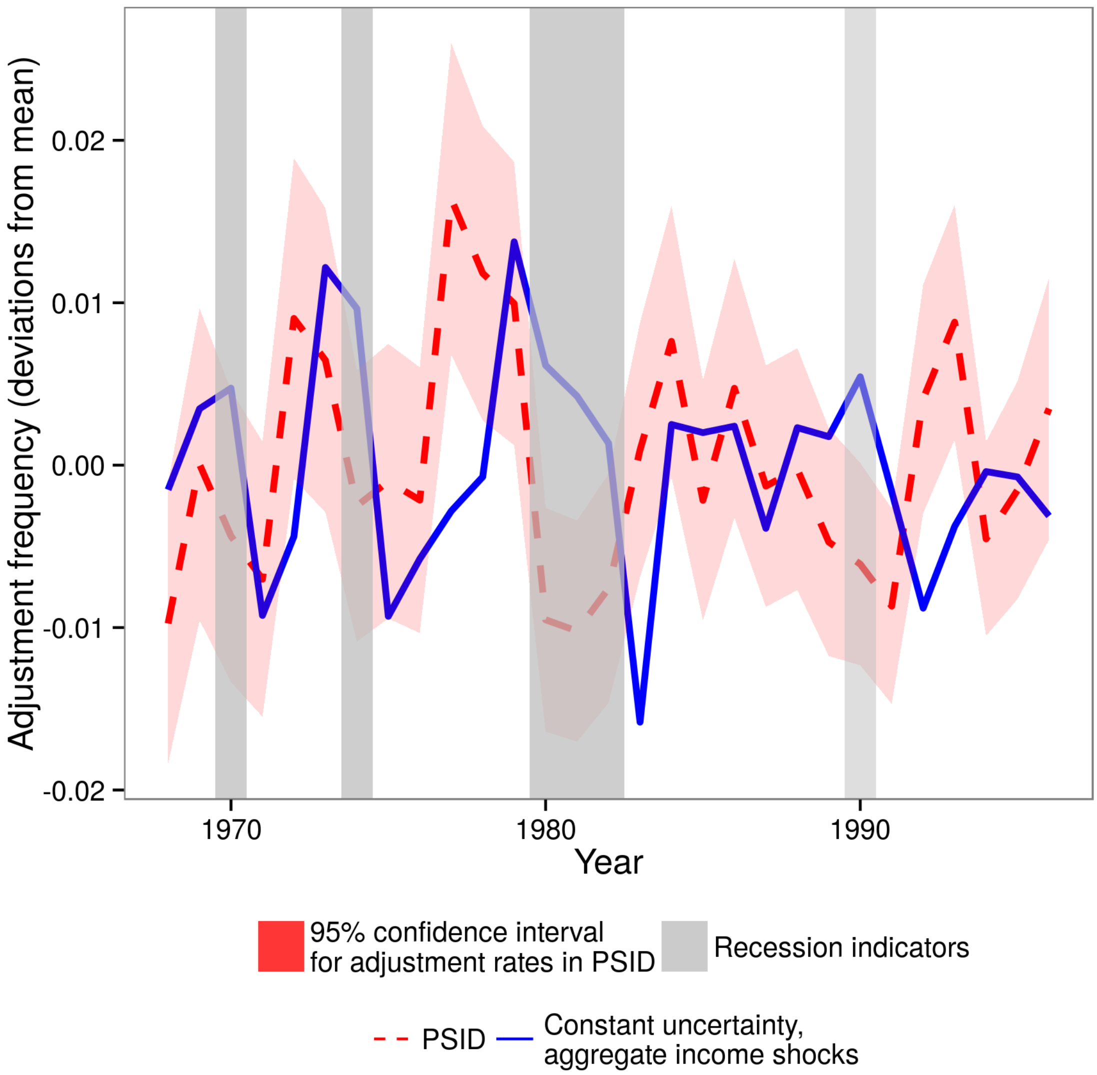

Source: Author's calculations

In general, the model does a good job matching the cyclicality of the adjustment rates. In particular, in the first half of the sample, adjustment frequencies in the model follow the data closely. The model also matches well the magnitudes of the cyclicality in adjustment rates, which seems to suggest that aggregate factors play an important role.

However, it does not manage to capture the abrupt decline of adjustment at the start of a recession and the rapid shoot-up when the economy comes out of it. Adjustment frequencies follow aggregate income and are thus at its peak when a recession starts and at its trough when the recession ends. Consider, for example, the recession from 1980 to 1983. The simulated adjustment frequencies are decreasing during the entire recession, when in the data, after the initial jump downward, there is no further decrease in adjustment.

Furthermore, in the recession in 1990, aggregate income as measured by GDP does not appear to affect adjustment frequencies in the data. On the contrary, adjustment frequencies in the model and in the data seem to move in opposite directions in the second half of the sample.

5.4. Countercyclical Left-Skewness + Aggregate Income Shocks

We have seen that both countercyclical left-skewness and aggregate income shocks are able to explain the basic pattern of the cyclicality in adjustment frequencies in the data. This suggests to explore the performance of a model with both countercyclical left-skewness and aggregate income shocks combined. The results of this model are discussed in this section. To simulate the model, I use the Markov switching AR(1) process presented and estimated in Section 3 to model aggregate income. The model is solved using the same calibration like for the countercyclical left-skewness model. When simulating this model, both NBER business cycle dates as well as actual aggregate income shocks are fed into the model.

The numerical results depicted in Figure 10 are disillusioning. Rather than complementing each other, effects of uncertainty and aggregate income appear to cancel out. An explanation for the bad performance of this combined model is that it is not calibrated to the data. Calibration is more important for this model in order to weigh the relative magnitudes of the effects of time-varying left-skewness and aggregate shocks. Moreover, the introduction of the Markov switching model for the aggregate income process might not be suitable to model household expectations about aggregate income. Thus, the bad performance of this model should not be taken as conclusive evidence against either hypothesis.

Figure 10 – Mean Adjustment Frequencies in Period 1968–1996 in PSID and Model with Countercyclical Left-Skewness and Aggregate Income Shocks

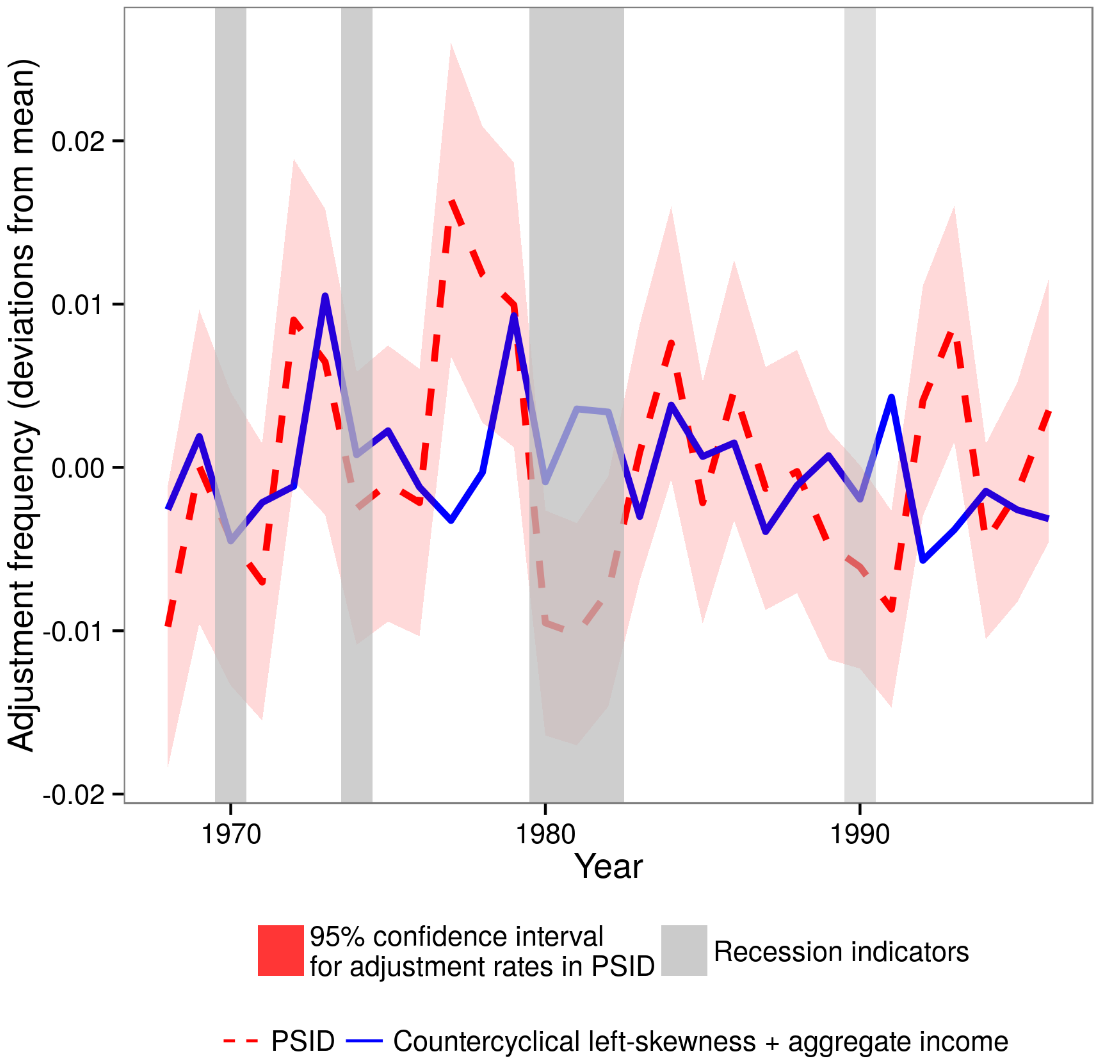

Source: Author's calculations

5.5. Discussion

The previous results suggest that the countercyclical left-skewness model and the model with aggregate income shocks can each explain part of the stylized facts of durable adjustment over the business cycle. Countercyclical variance, however, generates fluctuations in adjustment which are much more volatile compared to the data. Aggregate income explains well the magnitudes of adjustment frequencies but does not explain the abrupt decrease in adjustment. Time-varying left-skewness, on the other hand, predicts the rapid jumps that are observed in the data but does not fully account for their magnitudes. However, the effect of uncertainty is still substantial and can explain a large part of the variation in adjustment rates.

When comparing the match of these two models, it is important to take into account that the fluctuations from the model with time-varying uncertainty come from a simple binary process for uncertainty. Therefore, it is not surprising that this model cannot reproduce the richness of behaviour of the data. In reality, recessions differ substantially in severity and in the level of uncertainty. For example, Baker, Bloom, and Davis (2013) and Bloom (2014) show that economic policy uncertainty and stock market volatility, respectively, differ significantly by recessions. Similarly, uncertainty in earnings changes is not either high or low but is likely to vary continuously. This heterogeneity is captured by the model with aggregate income shocks, which uses GDP data on a much finer scale. From this perspective, the good performance of the model with countercyclical left-skewness is very promising. Further research is needed to have richer data on income uncertainty to capture the heterogeneity of recessions.

As I have pointed out before, the fluctuations in adjustment in the model with countercyclical left-skewness and in the data appear to be lagged by one year. The simulated response to uncertainty shocks precede the actual response in the data. As we know from the literature on dating business cycles, it is not easy to determine if the economy is about to enter a recession (Aastveit, Ravazzolo, & van Dijk, 2014; Askitas & Zimmermann, 2011). In this thesis, I assumed that the state of the economy is public knowledge and that households immediately learn the state and adjust their behaviour. In reality, this is often not the case, and when and how households adapt their expectations about the economy matters for aggregated adjustment decisions. If there is indeed a delay between the begin of a recession as dated by NBER and when households learn about the state of the economy, this explains the lag between the simulated series and the actual series.

6. Conclusion

This thesis offers three new results on time-varying uncertainty and durable adjustment.

First, time-varying uncertainty about future income can account for a substantial part of the empirical variation in the frequency of durable adjustment. Consistent with the data, countercyclical left-skewness implies a sharp decrease in adjustment at the beginning of a recession and an immediate and sizable increase when the economy begins to recover. This depressing effect on adjustment is further magnified by a decline in aggregate income and wealth levels.

Second, countercyclical variance as estimated by Storesletten et al. (2004) has implausible implications on durable adjustment, suggesting that the characterization of income risk over the business cycle by Guvenen et al. (2014) is more accurate. The left-skewness rather than the variance of income growth is countercyclical.

Third, I show analytically in a two-period model of durable adjustment that higher uncertainty decreases overall adjustment. However, a spike in income uncertainty increases downward adjustment because the motive for precautionary savings becomes stronger and households are willing to liquidate their durables in order to offset shocks to future consumption. Depreciation implies that the latter effect is stronger than the former, which leads to an overall decline in adjustment when uncertainty increases. The result holds numerically in an infinite-horizon version of this problem.

The evidence this thesis provides suggests that fluctuations in uncertainty may be an important determinant of the volatility of consumer durable expenditures. Thus, policy makers should focus on decreasing economic uncertainty if the goal is to stabilize the economy. However, as Bloom (2014) note, this is made difficult by the fact that research on policy implications of uncertainty is still at an early stage.

References

Aastveit, K., Ravazzolo, F., & van Dijk, H. (2014). *Nowcasting the Business Cycle in an Uncertain Enviroment* (Tech. Rep.). Norges Bank. Retrieved from `https://www.aeaweb.org/aea/2014conference/program/retrieve.php?pdfid=1239`

Askitas, N., & Zimmermann, K. F. (2011). *Nowcasting Business Cycles Using Toll Data* (Discussion Paper No. 5522). Retrieved from `http://ftp.iza.org/dp5522.pdf`

Attanasio, O. (2000). Consumer Durables and Inertial Behaviour: Estimation and Aggregation of (S, s) Rules for Automobile Purchases. *Review of Economic Studies*, *67*(4), 667-96.

Baker, S. R., Bloom, N., & Davis, S. J. (2013). *Measuring Economic Policy Uncertainty* (Working Paper No. 13-02). Chicago Booth Research Paper. Retrieved from `https://www.aeaweb.org/aea/2013conference/program/retrieve.php?pdfid=519`

Bar-Ilan, A., & Blinder, A. S. (1992). Consumer Durables: Evidence on the Optimality of Usually Doing Nothing. *Journal of Money, Credit and Banking*, *24*(2), pp. 258-272. Retrieved from `http://www.jstor.org/stable/1992740`

Berger, D., & Vavra, J. (2015). Consumption Dynamics During Recessions. *Econometrica*, *83*(1), 101-154.

Bertola, G., & Caballero, R. J. (1990). Kinked Adjustment Costs and Aggregate Dynamics. In *NBER Macroeconomics Annual, volume 5* (p. 237-296). National Bureau of Economic Research, Inc.

Bertola, G., Guiso, L., & Pistaferri, L. (2005). Uncertainty and Consumer Durables Adjustment. *Review of Economic Studies*, *72*(4), 973–1007.

Black, S., & Cusbert, T. (2010). Durable Goods and the Business Cycle. *Federal Reserve of Australia Bulletin*.

Bloom, N. (2014). Fluctuations in Uncertainty. *Journal of Economic Perspectives*, *28*(2), 153-76.

Challe, E., & Ragot, X. (forthcoming). Precautionary Saving over the Business Cycle. *The Economic Journal*. doi: 10.1111/ecoj.12189

Eberly, J. C. (1994). Adjustment of Consumers' Durables Stocks: Evidence from Automobile Purchases. *Journal of Political Economy*, *102*(3), pp. 403-436.

Foote, C., Hurst, E., & Leahy, J. (2000). Testing the (S, s) Model. *American Economic Review*, *90*(2), pp. 116-119.

Gospodinov, N., & Lkhagvasuren, D. (2014). A Moment-matching Method For Approximating Vector Autoregressive Processes By Finite-state Markov Chains. *Journal of Applied Econometrics*, *29*, 843–859.

Grossman, S. J., & Laroque, G. (1990). Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods. *Econometrica*, *58*(1), pp. 25-51.

Guvenen, F., Ozkan, S., & Song, J. (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy*, *122*(3), pp. 621-660.

Hamilton, J. D. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*, *57*(2), pp. 357-384.

Jacobson, L. S., LaLonde, R. J., & Sullivan, D. G. (1993). Earnings Losses of Displaced Workers. *American Economic Review*, *83*(4), 685 –709.

Khan, A., & Thomas, J. (2008). Idiosyncratic Shocks and the Role of Nonconvexities in

Plant and Aggregate Investment Dynamics. *Econometrica*, *76*(2), 395–436.

Khan, A., & Thomas, J. K. (2003). Nonconvex Factor Adjustments in Equilibrium Business Cycle Models: Do Nonlinearities Matter? *Journal of Monetary Economics*, *50*, 331–360.

Kimball, M. S. (1990). Precautionary Saving in the Small and in the Large. *Econometrica*, *58*(1), 53–73.

Knotek II, E. S., & Khan, A. (2011). How do Households Respond to Uncertainty Shocks? *Economic Review*, Federal Reserve Bank of Kansas City, Second Quarter.

McKay, A. (2015, April). *Time-Varying Idiosyncratic Risk and Aggregate Consumption Dynamics.* Retrieved from `http://people.bu.edu/amckay/pdfs/precsave.pdf` (Unpublished manuscript)

Rouwenhorst, K. G. (1995). Asset Pricing Implications of Equilibrium Business Cycle Models. In T. Cooley (Ed.), *Frontiers of Business Cycle Research* (pp. 294–330). Princeton, NJ: Princeton University Press.

Schmitt-Grohé, S., & Uribe, M. (2014, 11). *Finite-State Approximation Of VAR Processes: A Simulation Approach* (Note). Colombia University. Retrieved from `http://www.columbia.edu/~mu2166/tpm/tpm.pdf`

Stock, J. H., & Watson, M. W. (1999). Chapter 1 Business Cycle Fluctuations in US Macroeconomic Time Series. In J. B. Taylor & M. Woodford (Eds.), *Handbook of Macroeconomics* (Vols. 1, Part A, p. 3 - 64). Elsevier. Retrieved from `http://www.sciencedirect.com/science/article/pii/S1574004899010046` doi: http://dx.doi.org/10.1016/S1574-0048(99)01004-6

Storesletten, K., Telmer, C. I., & Yaron, A. (2004). Cyclical Dynamics in Idiosyncratic Labor Market Risk. *Journal of Political Economy*, *112*(3), pp. 695-717.

Tauchen, G. (1986). Finite State Markov-chain Approximations to Univariate and Vector Autoregressions. *Economics Letters*, *20*(2), 177 - 181.

Appendices

A. Appendix to A Simple Model Of Durable Adjustment

The results about the effect of uncertainty on adjustment presented in the text also hold for uncertainty in the form of higher left-skewness if one assumes that the household is prudent, that is if its preferences exhibit a convex marginal utility function. We can approximate expected marginal utility as

$$\mathbb{E}[u(c_2^A, d)] \approx u(\bar{c}_2^A, d) + \frac{1}{2}u_{cc}(\bar{c}_2^A, d)\sigma^2 - \frac{1}{6}u_{ccc}(\bar{c}_2^A, d)\gamma^3,$$

where $\gamma^3 = -\mathbb{E}[\epsilon^3]$ is the negative of Pearson's moment coefficient of skewness and is higher for a more left-skewed distribution. I use the negative of the skewness coefficient to have a measure that is positively related to the level of uncertainty. Therefore,

$$\frac{da}{d\gamma^3} \approx -\frac{\beta}{6}\frac{-u_{ccc}(\bar{c}_2^A, d) + u_{ccc}(\bar{c}_2^N, \bar{d})}{u_c(c_1^A, \bar{d}) - u_c(c_1^N, \bar{d})}.$$

As shown by Kimball (1990), the theory of precautionary savings is isomorphic to the Arrow-Pratt theory of risk aversion. Hence, if prudence is decreasing with consumption, then $-u_{ccc}(c,d)/u_{cc}(c,d)$ is a decreasing function of c. Let this decreasing function be denoted by $\bar{\chi}(c)$. We can write $u_{ccc}(c,d) = -\bar{\chi}(c)u_{cc}(c,d) = \bar{\chi}(c)\chi(c)u_c(c,d)$. Since the product of two positive decreasing functions is also decreasing, by the same argument as before, the numerator has the opposite sign of the denominator and thus the cut-offs shift to the right as left-skewness is higher. Similarly, the lower cut-off increases to a larger extent than the upper cut-off and the inaction region thus shrinks as left-skewness increases.

B. Solving the Dynamic Model

I describe the solution to the dynamic model when household income also includes an aggregate component. The dynamic model is solved using a Value Function Iteration (VFI) algorithm. Like in the simple model of durable adjustment presented in the text, we can consider the value of adjustment and the value of inaction separately and state the household problem:

$$
\begin{aligned}
V^A(a,d,y_i,y_a,s) &= \max_{c,d',a'} u(c,d') + \beta\,\mathbb{E}[V(a',d',y_i,y_a',s')|y_i,y_a,s] \\
V^N(a,d,y_i,y_a,s) &= \max_{c,a'} u(c,d(1-\delta(1-\chi))) + \beta\,\mathbb{E}[V(a',d(1-\delta(1-\chi)),y_i',y_a',s')|y_i,y_a,s] \\
V(a,d,y_i,y_a,s) &= \max\{V^A(a,d,y_i,y_a,s),V^N(a,d,y_i,y_a,s)\}
\end{aligned}
$$

$$
\begin{aligned}
\text{subject to} \quad & a' + c + d' = (1+r)a + (1-\delta)d + y_i + y_a - A(d,d'), \\
& a \geq 0 \\
& \text{and the law of } (y_i',y_a',s') \text{ given } (y_i,y_a,s).
\end{aligned}
$$

Moreover, we can resort to a trick used by Berger and Vavra (2015) to speed up computation by noting that the value of adjusting only depends on a household's net cash-on-hand $w = (1+r)a + d(1-\delta) - F^d(1-\delta)d - F^t(y_i + y_a)$. This allows us to eliminate one state variable and we can write

$$
\begin{aligned}
\tilde{V}^A(w,y_i,y_a,s) &= \max_{c,d',a'} u(c,d') + \beta\,\mathbb{E}[V(a',d',y_i',y_a',s')|y_i,y_a,s] \\
&\text{subject to } c = w + y_i + y_a - d' - a'.
\end{aligned}
$$

All state variables are discretized. The values for $\tilde{V}^A$ and V^N are computed separately and the associated optimal choices are saved for both. In a second step the algorithm iterates over a grid of (a,d,y_i,y_a,s) and computes cash-on-hands w. The value of adjustment $\tilde{V}^A$ is obtained by interpolating the function on the cash-on-hand grid, which is compared to V^N to get $V = \max\{\tilde{V}^A, V^N\}$.

For the calibration and the main results I use 100 grid points each for assets and durables, 85 for net cash-on-hands and 21 for idiosyncratic income. Aggregate income y_a is discretized using the method of Tauchen (1986) on a grid with 7 points. The idiosyncratic income processes are discretized using two methods. The process with time-varying variance is discretized using Tauchen (1986), whereas the process with countercyclical left-skewness is discretized using a simulation-based method described in Appendix D.

To solve for the policies, the algorithm then starts with an initial guess of the value function, computes $\tilde{V}^A$ and V^N for all states, and then updates the guess for the value function. The algorithm iterates until the maximum norm of two subsequent value functions is below 0.01.[19] The maximum of the objective function is found by an Nelder-Mead algorithm. Since the objective displays many local maxima, I run the algorithm from 4 distinct initial simplices at points in the choice space and with different volumes. When maximizing the value of adjustment, we are looking for a two-dimensional policy. The starting values are computed by using a simple heuristic which has proven to deliver

[19] I experimented with smaller convergence tolerances but results did not change significantly.

precise solutions.[20]

For each state, the algorithm computes three possible guesses of the optimal solution using a simple heuristic. These guesses are then given as the starting point for Nelder-Meade in order to ensure that the algorithm finds the global maximum. This has turned out to be a very fast and reliable solution method, and, compared to trying all possible choices to find the maximum, it has the advantage that it allows to solve the problem on a much larger grid.

The expectation $\mathbb{E}[V(a', d', y_i', y^{a\prime}, s')|y_i, y^a, s]$ is computed as

$$\mathbb{E}[V(a', d', {y_i}', {y_a}', s')|y_i, y_a, s] = \sum_{y_i', y_a', s'} V(a', d', {y_i}', {y_a}', s') \Pr({y_i}', {y_a}', s'|y_i, y_a, s),$$

where $\Pr(X|Y)$ is the conditional probability function of a random variable X given Y. By independence of ${y_i}'$, ${y_a}'$ and s', and given (y_i, y_a, s), we can write

$$\begin{aligned}
\Pr({y_i}', {y_a}', s'|y_i, y_a, s) &= \Pr({y_i}', {y_a}'|y_i, y_a, s, s') \Pr(s'|y_i, y_a, s) \\
&= \Pr({y_i}'|y_i, y_a, s, s', {y_a}') \Pr({y_a}'|y_i, y_a, s, s') \Pr(s'|y_i, y_a, s) \\
&= \Pr({y_i}'|y_i, s) \Pr({y_a}'|y_a, s) \Pr(s'|s),
\end{aligned}$$

the product of three quantities that follow readily from the discretization method and from our estimated transition matrix for the business cycle states.

Before simulating the model, the final policy functions are obtained by solving the model for one more iteration on a finer grid with each 140 grid points for assets and durables and 120 grid points for net cash-on-hands. The model is simulated for a panel of 25.000 households for the calibration and 200.000 households for the out-of-sample predictions.[21] First, the income process is simulated for 150 periods in order to start with a stationary income distribution. The algorithm uses the real shock sequence from the U.S. economy for the business cycle state and aggregate income.

The household panel is generated by starting with a uniform asset and durable distribution and then using the policy functions for each household to find its state at $t + 1$ given the state at t. The policy functions are interpolated bilinearly for (a, d) states that are not on the grid for which the model was solved. The first 200 periods are dropped in order to remove the dependence on the initial conditions.

[20] The heuristic computes the guess for d' as a fraction of net assets $w + y_a + y_i$ in the adjustment case. Similarly, the guess for a' is computed. One starting value is always (d',a')=(0,0) to capture corner solutions.

[21] I choose such a large sample size for the 1986–1996 out-of-sample predictions because for smaller sample sizes the simulated adjustment frequencies are somewhat volatile, whereas the gap distribution and adjustment hazards are precisely estimated with a smaller sample.

C. Description of the Calibration Procedure

The calibration method is based on Berger and Vavra (2015). Since the method had to be adapted to calibrate the models in this thesis, this section exactly describes the procedure by which they have been calibrated.

C.1. Measurement Errors

Since PSID data comes from a survey, it is inherently prone to measurement errors. The calibration method Berger and Vavra (2015) propose allows to deal with that in a straightforward manner. Let z be a variable of interest, such as consumption or durable holdings. Let $\hat{z}$ be a measurement of this true value z. If there are measurement errors, then $\hat{z} \neq z$ in general.

Following Berger and Vavra (2015) I assume the following relationship between variables and their measurement:

$$\hat{z} = (1 + \sigma_\epsilon \epsilon) z, \tag{7}$$

with $\epsilon \sim N(0, 1)$. The standard deviation σ_ϵ of the measurement error is a parameter and calibrated from the data.

C.2. Calibration Procedure

Let p denote the vector of parameters of the model, excluding the measurement error σ_ϵ. The objective function of the calibration is denoted by $F(p)$ and will be described in further detail below. The local optimum $p^\star = \arg\min F(p)$ is searched locally from different starting points in the parameter space using a Nelder-Meade algorithm. Since the estimates from Berger and Vavra (2015) already give a good starting guess, this speeds up the computation and increases the precision of the match. Note that Berger and Vavra (2015) solved the model for a grid of parameters.

The objective function $F(p)$ to be minimized is computed for each parameter p as follows.

1. Fix a set of shocks ζ for a panel of households with sample size $n = 25000$ using the model-specific income process. In the model with constant uncertainty, the income process is burned-in for 150 periods and then simulated for all households, which is the same as in Berger and Vavra (2015). For the models with time-varying uncertainty, I pick the NBER recession indicators $s_{1968\text{q}1}, \ldots, s_{1996\text{q}4}$ to reproduce actual U.S. business cycles and use it to simulate the idiosyncratic household incomes. This differs from Berger and Vavra (2015), who do not use actual aggregate data for the calibration part.

2. For each vector of parameters p, solve the model and simulate a panel using the shocks ζ to obtain time series for target durable stock $d^{\star}_{mi}$, actual durable stock d_{mi}, assets a_{mi} and non-durable consumption c_{mi} for each observation $i = 1, \ldots, n$.

3. Estimate the linearized policy function

$$\log(d^{\star}_{mi}) = \beta_0 + \beta_1 \log d_{mi} + \beta_2 a_{mi} + \beta_3 \log c_{mi} + \beta_4 \frac{d_{mi}}{c_{mi}}. \tag{8}$$

   Linearization is used in order to achieve an unbiased predictor in the presence of the assumed multiplicative measurement errors. Refer to Berger and Vavra (2015) for a discussion about the advantages of linear predictors and about the performance of different specifications.

4. Aggregate the model to biannual frequency. These observations are denoted by $(d^{\star a}_{mi}, d^{a}_{mi}, a^{a}_{mi}, c^{a}_{mi})$.

5. Using the estimated β vector I then impute the target durable holdings $\log^{\star}_{di}$ to the observations in the PSID data using equation (8). The gap in the data is then given by $x_{di} = \log d^{\star}_{di} - \log d_{di}$. This gives non-parametric estimates of the density function $\hat{f}_d$ and the hazard function $\hat{h}_d$.[22]

6. To obtain the model counterpart of the gap density and adjustment hazard, the algorithm proceeds as follows. Let g_{σ_ϵ} be a grid of measurement error parameters. Then fix a sequence of normally distributed shocks ϵ_{dj}, ϵ_{aj} and ϵ_{cj}, $j = 1, \ldots, n$, for durables, assets and non-durable consumption, respectively, with mean zero and standard deviation 1. For each σ_ϵ in g_{σ_ϵ}, generate new model observations $(d^{a\epsilon}_{mi}, a^{a\epsilon}_{mi}, c^{a\epsilon}_{mi})$ using identity (7). Then use (8) to impute gaps to the model and obtain $x^{a\epsilon}_{mi}$. Similarly, compute estimates of the density function $\hat{f}_{m\epsilon}$ and $\hat{h}_{m\epsilon}$.

7. The objective function is then given by

$$F(p) = \min_{\epsilon \in g_{\sigma_\epsilon}} \int (\hat{f}_d(x) - \hat{f}_{m\epsilon}(x))^2 dx + \int (\hat{h}_d(x) - \hat{h}_{m\epsilon}(x))^2 dx. \tag{9}$$

   The calibration procedure matches the gap density and the hazard in the model with their model-implied counterparts in the data.

A technical detail is that the optimal measurement parameter is found for each set of parameters. Instead of solving $\min_{(p,\epsilon)} F_\epsilon(p)$ like Berger and Vavra (2015) this procedure solves $\min_p(\min_\epsilon F_\epsilon(p))$. This brings a significant speed-up. The inner minimum is solved

[22]Following Berger and Vavra (2015) I use 21 bins to estimate the gap distribution and adjustment hazard.

very quickly because the compute-intensive step is solving and simulating the model. Recomputing gaps for different vectors of measurement errors, in contrast, is very cheap.

D. Discretization of the Process with Countercyclical Left-Skewness

The changing skewness of the income process of Guvenen et al. (2014) is critical. Unfortunately, usual methods to approximate autoregressive processes using finite-state Markov chains are not adequate for a stochastic process with non-zero skewness. In this section, I show that I cannot resort to one of the standard discretization methods such as Tauchen (1986) and instead propose another method, which I will assess numerically. Although every discretization method for processes which can take values from an unbounded set introduces skewness near the bounds of the grid, even when approximating normal distributions, this effect can be mitigated by choosing a large enough grid size. Therefore, Tauchen (1986)'s method works well for the processes in this thesis with a skewness of zero.

Alternative discretization methods such as the Rouwenhorst (1995) method focus on normal distributions as well. More recently, Gospodinov and Lkhagvasuren (2014) propose a moment-matching method based on the Rouwenhorst method which allows to match conditional moments with higher accuracy. Although in theory this method also provides a way to match the conditional skewness of the process, in the case of the process estimated by Guvenen et al. (2014), it neither matches the relevant conditional moments nor does the shape of the discretized distribution bear any resemblance to the original.

The method I use to approximate the Markov chain is the following.[23] Let $Y = \{Y_t\}_{t=0}^{\infty}$ be an X-valued stochastic process we would like to approximate on a finite grid $\bar{Y} = \{\bar{Y}_1, \dots, \bar{Y}_n\} \subset \mathbb{R}$ with $X \subset \mathbb{R}$. Let the function $d : X \mapsto \bar{Y}$ map Y_t to the nearest grid points in $\bar{Y}$, that is $d(y) = \arg\min_{Y_i \in \bar{Y}} |Y_i - Y|$. Then define the discretized process $\tilde{Y} = \{\tilde{Y}_t\}_{t=0}^{\infty}$ by $\tilde{Y}_t = d(Y_t)$. Note that $\tilde{Y}$ is not Markov in general.

The finite-state approximation Markov chain is denoted by $\hat{Y}$ and is defined as the Markov process with transition probabilities given by

$$P_{ij} = \Pr(\tilde{Y}_{t+1} = \bar{Y}_j | \tilde{Y}_t = \bar{Y}_i), \tag{10}$$

where P_{ij} is the probability of transitioning from state Y_i to Y_j. Simulating the process $\tilde{Y}$ a large number of times allows to estimate the transition matrix P.

Little is known about the approximation process $\tilde{Y}$. However, it is easy to show that conditional as well as unconditional moments converge to the moments of the original

[23]The method is also proposed in an unpublished note by Schmitt-Grohé and Uribe (2014).

process as the grid gets arbitrarily fine. Instead of providing a theoretical analysis of the model, I numerically investigate the approximation errors of the moments of the conditional distribution for both methods.

Figure 11 shows the approximation errors of the conditional moments for the simulation method and for the method of Tauchen (1986). The two regimes of the process, recession and expansion, are shown separately. In both regimes, the Tauchen method misses substantially on the conditional mean. Errors in the mean are much smaller for the method that is presented here. Approximation errors of the conditional mean play a key role in our switching model. As shown in Figure 11a and Figure 11b, the conditional mean of the approximation is biased upward in an expansion and downward in a recession. Under this approximation scheme, switching from an expansion to a recession not only changes the skewness of the process but also substantially affects the conditional mean. Changes in household behaviour between recessions and expansions are then jointly caused by changes in mean and uncertainty, and it would be impossible to distinguish the effects.

The conditional variance is well-matched by Tauchen's method, whereas the simulation method is upwards biased. However, the bias of the simulation method remains constant for between recession and expansion and does therefore not bias the results. In matching the skewness, the simulation method consistently does a better job than Tauchen's method. Especially in the recession case, when left-skewness is very high, the method has a low error, but also in the expansion the simulation method is consistently more accurate.

Although neither method can match all three conditional moments of the process, I argue that the simulation method should be preferred for two reasons. First, it better matches the skewness of the process. Since the purpose of this process is to model countercyclical left-skewness, the skewness of the discretization should be as close as possible to the exact process. Second, the simulation method more precisely matches the conditional mean and, more importantly, does not introduce a shift in the mean between recessions and expansions.

Figure 11 – Approximation errors of Tauchen's method and the simulation method

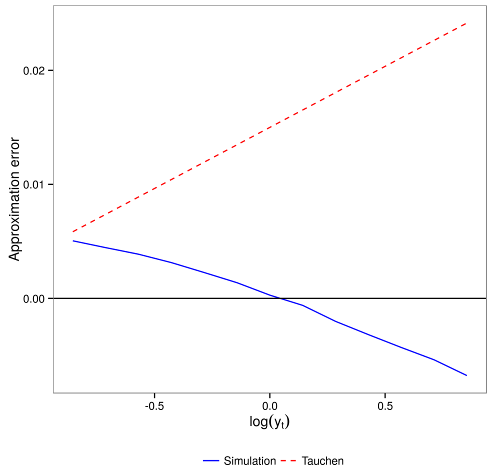

(a) Conditional mean (expansion)

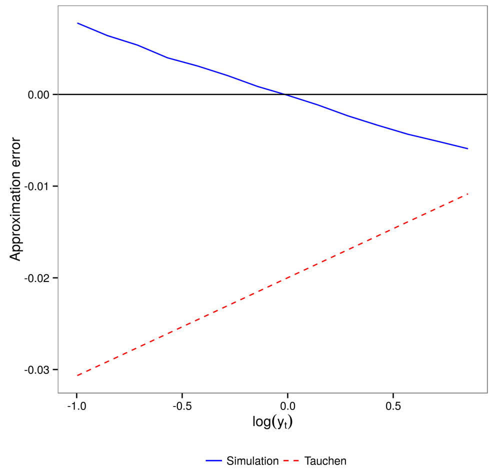

(b) Conditional mean (recession)

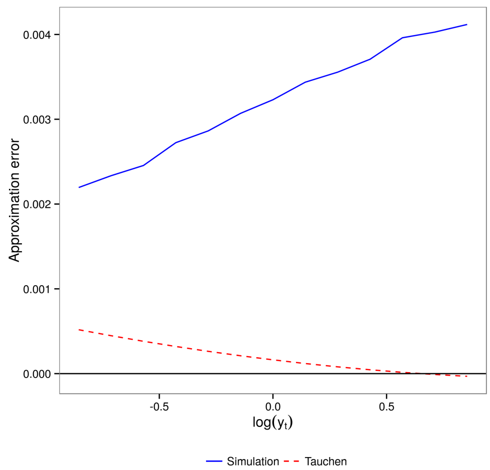

(c) Conditional variance (expansion)

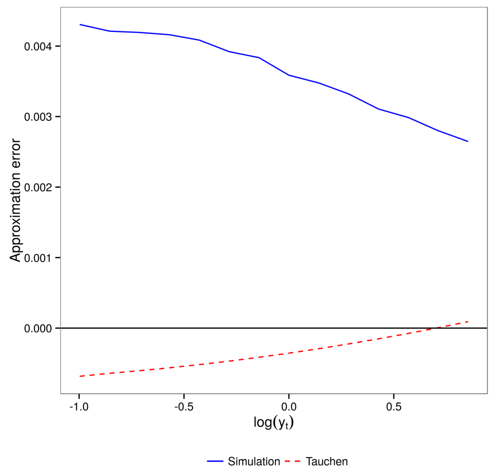

(d) Conditional variance (recession)

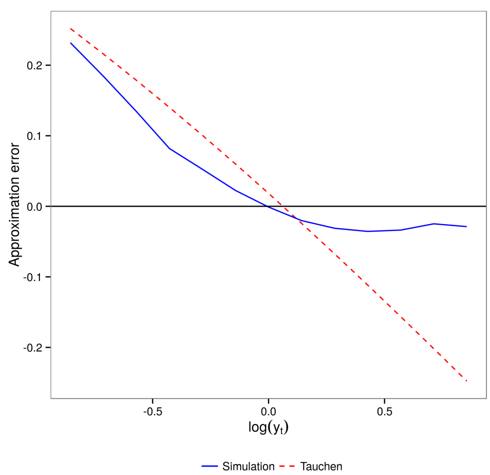

(e) Conditional skewness (expansion)

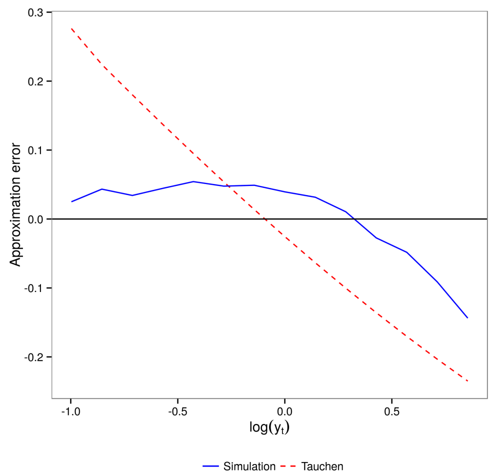

(f) Conditional skewness (recession)

Source: Author's calculations

Note: This figure shows approximation errors of the conditional moments, defined as the difference between the moment of the discretized process and the theoretical moment, for both discretization methods. The errors are depicted for the subset of the grid for income containing 95% of the values in the stationary distribution.

E. Estimation and Calibration of Stochastic Processes

E.1. Calibration of Individual Income Processes

I take the specification of the countercyclical variance income process from Storesletten et al. (2004) and the process with countercyclical left-skewness from Guvenen et al. (2014). However, both papers estimate annual processes whereas the models in this thesis are quarterly. Ideally, the income processes are re-estimated in a quarterly model. However, for computational as well as methodological reasons this thesis uses a simpler approach by rescaling the innovations such that unconditional variances of the processes coincide with estimates from PSID data.

This section briefly outlines the procedure for the process with left-skewness. In the process taken from Guvenen et al. (2014) the innovation is a mixture variable. Scaling a mixture random variable is equivalent to scaling the components of the mixture by the same factor. Therefore, to match process variances, we can compute the unconditional variance of the mixture variance $\operatorname{Var}(\epsilon)$ and set the scaling factor to $\alpha = \sqrt{\sigma^2_{\text{target}} / \operatorname{Var}(\epsilon)}$. Then the parameters of the mixture components of $\epsilon_\alpha = \alpha\epsilon$ are given by $\mu^i_\alpha = \alpha\mu^i$ and $\sigma^i_\alpha = \alpha\sigma^i$.

Let $P = (p, 1-p)$ be the stationary distribution of the business cycle regime Markov process. Then the unconditional variance is given by

$$\begin{aligned}\operatorname{Var}(\epsilon) &= \mathbb{E}[\epsilon^2] - \mathbb{E}[\epsilon]^2 \\ &= \mathbb{E}[\mathbb{E}[\epsilon^2|S]] - \mathbb{E}[\mathbb{E}[\epsilon|S]]^2,\end{aligned}$$

where S denotes the business cycle's regime. Since we assume $\mathbb{E}[\epsilon|S] = 0$ for all S, this gives

$$\operatorname{Var}(\epsilon) = p\,\mathbb{E}[\epsilon_R^2] + (1-p)\,\mathbb{E}[\epsilon_E^2],$$

where ϵ_R is the innovation in the recession case and ϵ_E is the innovation in an expansion.

To calculate the variance, the stationary distribution of the business cycle process is computed. I find $p = 0.1907$. The component parameters of the mixture are taken from Guvenen et al. (2014). However, I shift component means such that $E[\epsilon_S] = 0$ for each state S.

The parameters used are depicted in the following table.

Table 4 – Parameters of the Yearly Process, Adjusted from Guvenen et al. (2014)

Parameters	Yearly
ρ	0.979
μ_{1E}	0.119
μ_{2E}	-0.026
μ_{1R}	-0.102
μ_{2R}	0.094
σ_1	0.325
σ_2	0.001
w	0.490

These values for the parameters imply $\alpha = 0.41$, which gives the parameters in Table 2.

Similarly, the process of Storesletten et al. (2004) is rescaled to match the unconditional variance of the state-independent process. The annual estimates by Storesletten et al. (2004) are 0.21 and 0.12 in a recession and an expansion, respectively. To match the unconditional quarterly variance in the PSID, $\sigma^2_{\text{target}} = 0.1$, I find $\alpha = 0.706$ using the formula above. For quarterly income growth, this gives a standard deviation of 0.148 in a recession and 0.085 in an expansion.

E.2. Estimation of Aggregate Income Processes

E.2.1. AR(1)

Table 5 – Estimated AR(1) Process for Aggregate Income

Parameters	Estimate	Std. Error
ρ^a	0.876	0.034
σ_e	0.0076	0.0003

Maximum likelihood estimates from 1960–2013 quarterly GDP data (hp-filtered).

E.2.2. MS-AR(1)

Table 6 – Estimated Markov Switching AR(1) Process for Aggregate Income

Parameters	Estimate	Std. Error
c^a_R	-0.008	0.001
c^a_E	0.002	0.000
ρ^a_R	0.940	0.057
ρ^a_R	0.855	0.034
σ_e	0.006	0.000

Conditional maximum likelihood estimates from 1960–2013 quarterly GDP data (hp-filtered) and NBER recession indicators as described in Section 4.1.

F. Countercyclical Left-Skewness: Upward and Downward Adjustment

Figure 13 – Upward and Downward Adjustment in the Model with Countercyclical Left-Skewness and in PSID Data

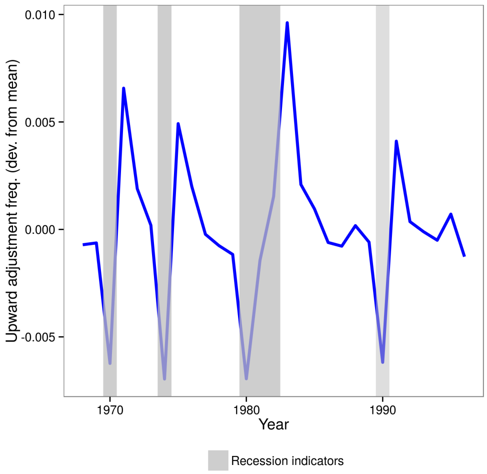

(a) Countercyclical Left-Skewness: Upward adjustment

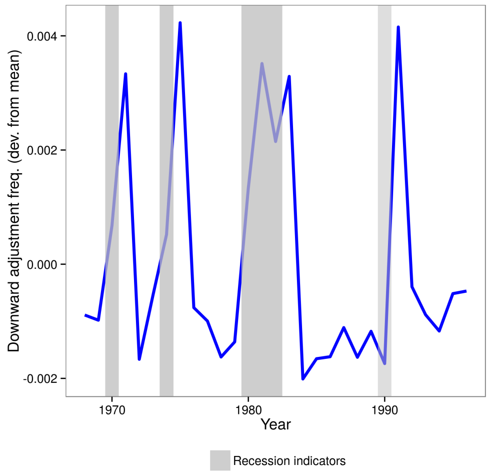

(b) Countercyclical Left-Skewness: Downward adjustment

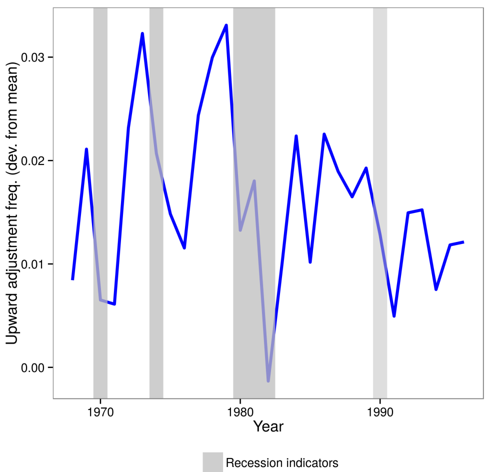

(c) PSID: Upward adjustment

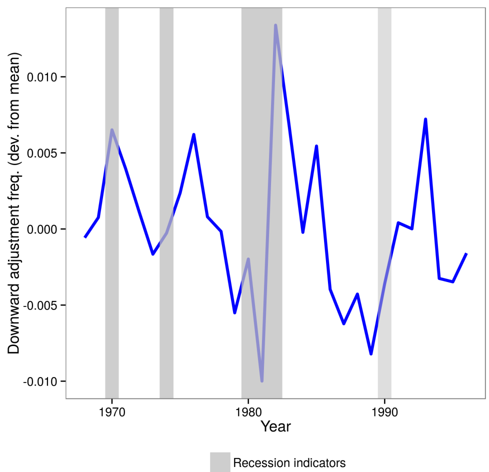

(d) PSID: Downward adjustment

Source: Author's calculations