Title
Essays in Quantitative Macroeconomics

Author
Neumuller, Seth

Publication Date
2013

Peer reviewed|Thesis/dissertation

University Of California
Los Angeles

Essays in Quantitative Macroeconomics

A dissertation submitted in partial satisfaction
of the requirements for the degree of
Doctor of Philosophy in Economics

by

Seth Neumuller

2013

Abstract of the Dissertation

Essays in Quantitative Macroeconomics

by

Seth Neumuller
Doctor of Philosophy in Economics
University of California, Los Angeles, 2013
Professor Lee Ohanian, Chair

These essays contribute to the study of labor economics and consumer finance, and fall within the broader category of quantitative macroeconomics. Chapter 1 investigates the trade-off between wage volatility (*risk*) and wage differentials (*return*) across industries through the lens of a general equilibrium, incomplete markets, life cycle model which allows for inter-industry mobility. While standard economic reasoning tells us that risk averse workers will demand a premium for exposure to wage volatility, for plausible calibrations of the model, I find that precisely the opposite is true – industries which expose workers to relatively low (high) wage volatility pay relatively high (low) wages. This chapter argues that inter-industry mobility is a quantitatively important insurance channel against labor market risk which is responsible for this counter-intuitive result. Chapters 2 and 3, which are both co-authored by Matthew Nelson Luzzetti, address issues in consumer finance. In Chapter 2, we introduce statistical learning and aggregate uncertainty into an otherwise standard model of consumer default. We show that learning by households and creditors endogenously generates a credit boom during a prolonged economic expansion like the Great Moderation and a severe and protracted credit crunch in response to an economic contraction like the recent financial crisis. This chapter illustrates that learning by households and creditors is an important driver of

aggregate debt dynamics. Chapter 3 develops an equilibrium model of consumer default with both long-term collateralized mortgages and short-term unsecured debt. We use this framework to evaluate whether the Bankruptcy Abuse Prevention and Consumer Protection Act of 2005 contributed to the severity of the housing crisis by inducing homeowners to default on their mortgage who would otherwise have declared bankruptcy and remained in their home. We find that although this reform significantly increased mortgage default rates upon implementation, it likely had only a minor impact on the severity of the subsequent housing crash if lenders rationally adjusted their mortgage interest rates to account for its impact on the incentives of households to repay their debt.

The dissertation of Seth Neumuller is approved.

Andrew Atkeson

Hanno Lustig

Pierre-Olivier Weill

Lee Ohanian, Committee Chair

University of California, Los Angeles

2013

To my wife, Jessica.

Thank you for your unwavering love,

encouragement, and support.

Table of Contents

List of Tables

List of Figures

Acknowledgments

I am particularly indebted to my dissertation chair, Lee Ohanian, for his guidance and support throughout this process. I would also like to thank Andrew Atkeson and Pierre-Olivier Weill for their invaluable feedback and suggestions. Chapters 2 and 3 of this dissertation are co-authored by my good friend and colleague, Matthew Nelson Luzzetti, without whom this dissertation would not have been possible. These two chapters are reprinted here with his permission. Generous financial support from the Department of Economics, the Graduate Division, and the Welton Family is also gratefully acknowledged.

Vita

2000	Diploma, Thousand Oaks High School, Thousand Oaks, CA
2005	B.S. (Chemical Engineering), Rice University, Houston, TX
2005–2007	Production Engineer, Lyondell Chemical Company, La Porte, TX
2009	M.A. (Economics), University of California, Los Angeles
2010	C.Phil. (Economics), University of California, Los Angeles

1 Wage Volatility and the Option Value of Mobility: A Macro-Labor Analysis of the Trade-off between Risk and Return across Industries

1.1 Introduction

While virtually all occupations require workers to invest in specific skills that tie them to that occupation, industry specific skills are typically acquired through work experience.[1] If occupation specific skills are transferable across industries, then inter-industry mobility may be a low cost way to insure against labor market risk, particularly for individuals with only a few years of labor market experience. This paper investigates both the empirical and theoretical relationships between the volatility of persistent idiosyncratic shocks to wages (*risk*) and inter-industry wage differentials (*return*). While standard economic reasoning tells us risk averse workers will demand a premium for exposure to wage volatility, a preliminary analysis of the data suggests that precisely the opposite is true. Figure 1.1 (a) plots the average log wage versus the volatility of log wages for working age males in the 1968 – 2009 survey years of the Panel Study of Income Dynamics (PSID), where each data point represents the trade-off between risk and return offered by one of twenty-four distinct industries.[2] Contrary to standard economic reasoning, the relationship between risk and return implied by the raw data is decidedly negative – for every 1 percent increase in wage volatility, the average wage falls by approximately 1.5 percent.[3] In other words, industries which expose workers to relatively low (high) wage volatility pay relatively high (low) average wages.

While this simple empirical exercise does not control for observable heterogeneity or

[1] See, for example, Neal (1995) or Parent (2000).

[2] Log wage changes are computed on a biennial basis for individuals who do not switch industries. A detailed discussion of the PSID and my sample selection procedures is included in Section 1.3.

[3] A similar relationship exists between average log annual income, defined as total labor income plus transfers, and the volatility of log annual income as depicted in Figure 1.1 (b). Log wages and log annual income are also highly correlated, as are the volatility of log wages and the volatility of log annual income, as shown in Figures 1.1 (c) and (d).

selection, I demonstrate that a similar inverse relationship between risk and return across industries arises naturally within the context of a general equilibrium, incomplete markets, life cycle model which allows for inter-industry mobility. In the model, risk averse agents are each endowed with a vector of industry-specific idiosyncratic productivities. Industries differ in their equilibrium wage rates and the volatility of the idiosyncratic uninsurable productivity shocks that workers face. Each period, agents optimally decide in which industry to work subject to an exogenous resource cost of switching industries. Conditional on their industry choice and realized productivity shocks, agents make a labor supply and consumption-savings decision. The equilibrium relationship between inter-industry wage differentials and the volatility of persistent idiosyncratic shocks to wages generated by the model is the outcome of a horse race between two opposing forces. On one hand, since markets are incomplete, agents cannot perfectly insure against labor market risk. As a result, risk averse agents demand a premium for exposure to wage volatility (*risk premium effect*). But since agents can always exercise their option to switch industries, they also value wage volatility as it creates the potential for faster wage growth (*option value effect*).

I take the theory to the data in order to determine the relative strength of these opposing forces and the resulting equilibrium trade-off between risk and return across industries. The key parameters to be calibrated are the within-industry volatility of idiosyncratic productivity shocks and the exogenous resource cost of switching industries, the latter of which is the cost of exercising this option. Using individual-level data for working age males from the PSID, I employ a combination of indirect inference and simulated method of moments in order to calibrate these critical parameters. For plausible levels of risk aversion, the option value effect dominates, leading to an inverse relationship between inter-industry wage differentials and the volatility of persistent idiosyncratic shocks to wages – for every 1 percent increase in annual wage volatility that an agent accepts, their expected wage falls by 1.5 percent. Thus, the model predicts that industries which expose workers to the least (most) idiosyncratic wage volatility pay the highest (lowest) wages. A series of counter-factual exercises reveals that inter-industry mobility is a quantitatively important insurance channel against labor market risk which is responsible for generating this negative relationship.

The theory has sharp predictions for patterns of inter-industry mobility over the life cycle and industry-specific expected wage profiles which can be compared directly with the data. In particular, young workers in the model value wage volatility more than their older counterparts since they have more time before retirement to either realize large productivity gains or to exercise their embedded option to switch industries. The option value of inter-industry mobility is also increasing in the volatility of the productivity shocks that workers face. Therefore, young agents in the model disproportionately seek employment in industries with higher than average wage volatility. As they accumulate labor market experience, agents who receive sufficiently positive productivity shocks remain, while those who don't find it optimal to seek employment in another industry in which they are more productive. Although wages are initially lower for agents who choose to work in an industry with higher than average wage volatility, in expectation they will grow faster and eventually exceed the wages earned by their counterparts. Critically, I demonstrate that these predictions of the model are fully consistent with the patterns of inter-industry mobility over the life cycle and expected wage profiles conditional on initial industry of employment derived from the data.

The model also highlights an important selection effect – agents who have success in their current industry are likely to remain, while those who receive adverse productivity shocks have an incentive to switch. Thus, the model predicts that industry tenure and industry-specific productivity are positively correlated. Since the latter is unobservable to the econometrician, if the theoretical model that I develop here is the true data generating process, inter-industry wage differentials estimated using standard econometric techniques will be biased. In addition, the volatility of idiosyncratic shocks to wages can only be estimated using the observed wage changes of individuals who do not switch industries between two adjacent points in time. Since the theory predicts that those individuals who receive sufficiently adverse shocks will exercise their option to switch industries, estimates of wage volatility measured directly from the data will also suffer from selection bias. By employing a structural equilibrium model which explicitly controls for selection effects and by using individual-level data to carefully calibrate the model's key parameters, I am able to successfully deal with these inherent empirical complications and arrive at an unbiased estimate of the trade-off between risk and return across industries.

1.1.1 Related Literature

This paper contributes to the literature along a number of dimensions. Most importantly, it demonstrates that inter-industry mobility is a quantitatively important insurance channel against labor market risk, which is responsible for generating the negative relationship between risk and return across industries. Switching industries in my model plays a similar role to moving in and out of the parental home in Kaplan (2012), which allows young adults to pursue riskier jobs that offer faster wage growth. This idea is also closely related to Vereshchagina and Hopenhayn (2009) in which the ability of entrepreneurs to close down their business and enter the workforce aids in explaining why entrepreneurs are willing to take on a disproportionate amount of risk in starting their own business in exchange for a relatively low expected return.[4] In each case, the presence of an outside option reduces the downside risk of pursuing a career which offers larger, but more uncertain, potential gains.

Previous related studies which estimate idiosyncratic earnings risk and assess its macroeconomic implications within the context of fully optimizing general equilibrium framework include Storesletten, Telmer, and Yaron (2004a), Heathcote, Storesletten, and Violante (2008), and Low, Meghir, and Pistaferri (2010). This paper contributes to the existing literature by estimating the within-industry volatility of persistent and transitory shocks to wages. It also assesses their impact on the pattern of inter-industry wage differentials and the allocation of workers across sectors. In contrast to the existing literature, this paper demonstrates that when agents have the option to switch industries, wage volatility is valued as it leads to

[4] Abbring and Campbell (2005) find that the largest component of a new firm's value is actually embedded in the option to exit.

higher expected wage growth rates by limiting downside risk.

This paper also demonstrates that, although the theoretical relationship between wage differentials and wage volatility is ambiguous, for reasonable levels of risk aversion it is decidedly negative. Interestingly, this finding stands in stark contrast to Cubas and Silos (2012) who analyze data from the Survey of Income and Program Participants (SIPP) and report a positive correlation between wage volatility and average wages across industries.[5] The authors then construct a general equilibrium model of sorting in the labor market in which agents are ex-ante heterogeneous in terms of their comparative advantage to work in each sector in order to decompose how much of the variation in average earnings across industries can be attributed to compensation for earnings risk as opposed to sorting on unobserved ability. While their model abstracts from inter-industry mobility, Kambourov and Manovskii (2008) argue that workers change their industry affiliation frequently and at a rate that has been steadily increasing over time. In light of this evidence, I explicitly allow for inter-industry mobility and demonstrate that this feature of the model has important quantitative implications for the resulting equilibrium trade-off between risk and return across industries.

The link between the variability of labor income and average earnings has been studied extensively in the literature, dating back to the seminal work of Friedman and Kuznets (1954) and, more recently, contributions by Abowd and Ashenfelter (1981), Feinberg (1981), Leigh (1983), and Carroll and Samwick (1997). In a related study, Dillon (2012) structurally estimates a model of occupational choice using data from the PSID and concludes that compensation for earnings risk is a key factor in explaining variations in expected lifetime earnings across careers. If the cost of switching occupations is sufficiently high, then the

[5] One potential source of the divergence in our findings is due to our differing sample selection procedures. While I restrict my analysis to males, both married and single, Cubas and Silos (2012) utilize data for both males and females, but require them to be married. Another key difference is that I use the volatility of biennial shocks to wages as my proxy for risk, whereas they use the volatility of quarterly fluctuations in total labor income. The prevalence of imputed values for labor income in the SIPP relative to the PSID could also be responsible for our conflicting results. According to Dahl, DeLeire, and Schwabish (2011), in 1998, 54 percent of households in the SIPP had imputed earnings; in 2002, 60 percent; and in 2005, 46 percent. Dahl et al. (2011) go on to argue that although imputed earnings can improve estimates of cross-sectional means and variances, the use of imputed data can be problematic for estimating income volatility.

risk premium effect will dominate the option value effect, leading to a positive relationship between risk and return across occupations.[6] Thus, the theoretical framework that I develop here is capable of rationalizing the inverse relationship between risk and return across industries as well the positive relationship between risk and return across occupations.

Murphy and Topel (1987) analyze two-year panels of individuals in the 1977 – 1984 Current Population Survey (CPS) and document wide differences in the volatility of hours worked and annual earnings across industries. The authors conclude, though, that the lack of a valid instrument calls into question their estimates of compensating wage differentials based on variations in job characteristics across industries. My approach improves upon this study in two ways. First, the PSID offers a longer panel structure than the CPS which allows me to disentangle transitory fluctuations in wages from long-term, persistent shocks. Second, by employing a structural general equilibrium model to estimate the relationship between risk and return across industries, I avoid the need to identify a valid instrument.

There is a large body of literature which aims to explain the existence of inter-industry wage differentials.[7] Identifying the underlying source of these differentials, however, has proven illusive. One strand of the existing literature, beginning with the seminal work of Krueger and Summers (1988), argues such differentials are only compatible with non-competitive theories of wage determination. Another strand of the literature, building on the work of Gibbons and Katz (1992), contends measured inter-industry wage differentials result from differences in workers' productive abilities not captured by individual-level data sets.[8] This paper adds to the literature by providing an alternative and complementary explanation for the emergence of inter-industry wage differentials within the context of a

[6] In the limit, the cost of switching is prohibitively high and agents are locked into their initial occupation of choice until retirement.

[7] Slichter (1950) was one of the first to provide empirical evidence supporting the existence of inter-industry wage differentials. Dickens and Katz (1987) offer a comprehensive review of the early empirical literature on this subject. More recently, Caju, Ktay, Lamo, Nicolitsas, and Poelhekke (2010) provide an updated literature review with an emphasis on studies of labor markets outside of the U.S.

[8] Workers with higher unobserved abilities will earn higher wages. Industries which employ proportionately more of these workers will then appear to pay higher average wages to observationally equivalent workers.

competitive model of the labor market. In particular, I find that variation across industries in the volatility of persistent idiosyncratic shocks to wages can explain roughly one-fifth of inter-industry wage differentials at the one-digit level.

Finally, the theoretical framework that I develop here builds upon the multi-armed bandit problem analyzed by Miller (1984). I extend his framework along a number of dimensions. In particular, I allow for endogenous labor supply and consumption-savings decisions, as well as the equilibrium determination of wages. While the former is important for quantifying the relative importance of inter-industry mobility as an insurance channel against labor market risk, the latter is essential for both addressing the puzzle of inter-industry wage differentials and estimating the trade-off between risk and return across industries. While I do not explicitly model learning dynamics, the persistent idiosyncratic productivity shocks in my model can be interpreted as learning about an agent's innate industry-specific ability.

The remainder of the paper proceeds as follows. Section 2 develops a stylized theoretical model which permits an analytical characterization of the risk premium and option value effects. Section 3 extends this framework into a general equilibrium overlapping generations model which can be taken to the data. Section 4 describes the PSID and outlines my calibration strategy. Section 5 presents the main quantitative results. Section 6 concludes.

1.2 Theoretical Model

In this section, I develop a stylized partial equilibrium model for which I can sharply characterize the risk premium and option value effects, as well as their roles in determining the equilibrium relationship between inter-industry wage differentials and wage volatility.

1.2.1 Environment

Consider the decision problem of a worker who lives for two periods $t = 1, 2$ and is endowed with one unit of time each period. The worker's period utility from consumption is given by, $u(c) = c^{1-\gamma}/1 - \gamma$, where $\gamma > 1$ is the coefficient of relative risk aversion and the discount factor $\beta \in (0,1)$. There are J islands on which the worker can seek employment, each representing a different industry. Prior to $t = 1$, nature draws a vector of island-specific log efficiency units of labor for the worker, $\mathbf{z} \in \mathbb{R}^J$, where $z_j \sim N\left(-\frac{1}{2}\sigma_j^2, \sigma_j^2\right)$ is i.i.d. and revealed to the worker only upon arrival on island j. There is no storage technology, and therefore the worker consumes her entire labor income, $w_j e^{z_j}$, each period, where w_j is the wage rate on island j in units of the consumption good per efficiency unit of labor.[9]

1.2.2 Analytical Results

For simplicity, suppose there are only two islands, $j \in \{1,2\}$. Furthermore, assume the worker's productivity on island 1 is known, $\sigma_1 = 0$, and on island 2 is uncertain, $\sigma_2 > 0$. First, consider the case in which the worker selects an island on which to work at $t = 1$ and must remain on that same island for both periods. In this case, the value of selecting island 1 at $t = 1$ is deterministic and given by

$$V_1 = \frac{w_1^{1-\gamma}}{1-\gamma} + \beta \frac{w_1^{1-\gamma}}{1-\gamma}. \tag{1.1}$$

The value of selecting island 2 at $t = 1$, on the other hand, is given by

$$V_2 = \mathbb{E}_{z_2}\left[\frac{\left(w_2 e^{z_2}\right)^{1-\gamma}}{1-\gamma} + \beta \frac{\left(w_2 e^{z_2}\right)^{1-\gamma}}{1-\gamma}\right]. \tag{1.2}$$

The following theorem states that for the worker to be indifferent between selecting islands 1 and 2 at $t = 1$, the worker will demand a premium to work on the island for which the payoff is uncertain.

Theorem 1.1 *Suppose the worker selects an island on which to work at $t = 1$ and must remain on that same island for both periods. If $V_1 = V_2$, then $w_1 < w_2$.*

[9] In order to focus the reader on inter-industry mobility as an insurance channel against labor market risk, I have ruled out precautionary savings, borrowing in credit market, and adjusting labor supply as insurance mechanisms in this illustrative example. Here I also abstract from the resource cost of switching industries. These assumptions will be relaxed in the full quantitative model which I develop in the following section.

Proof The log wage differential, $\ln(w_1/w_2)$, which makes the worker indifferent between selecting islands 1 and 2 is given by

$$\ln\left(\frac{w_1}{w_2}\right) = -\gamma\frac{\sigma_2^2}{2}, \tag{1.3}$$

which is strictly less than zero for all $\gamma > 1$ and $\sigma_2^2 > 0$. It follows that $w_1 < w_2$.

In what follows, I refer to the risk premium effect as the right hand side of equation (3).

Now consider the case in which the worker selects an island on which to work at $t = 1$, and then is given the option to switch islands prior to the start of $t = 2$. In this case, the value of selecting island 1 at $t = 1$ is given by

$$V_1 = \frac{w_1^{1-\gamma}}{1-\gamma} + \beta \max\left\{\frac{w_1^{1-\gamma}}{1-\gamma}, \mathbb{E}_{z_2}\left[\frac{(w_2 e^{z_2})^{1-\gamma}}{1-\gamma}\right]\right\}, \tag{1.4}$$

while the value of selecting island 2 at $t = 1$ is given by

$$V_2 = \mathbb{E}_{z_2}\left[\frac{(w_2 e^{z_2})^{1-\gamma}}{1-\gamma} + \beta \max\left\{\frac{w_1^{1-\gamma}}{1-\gamma}, \frac{(w_2 e^{z_2})^{1-\gamma}}{1-\gamma}\right\}\right], \tag{1.5}$$

where, in each value function, the max operator represents the worker's option to switch islands prior to the start of $t = 2$.

The following lemma states that if the worker selects island 1 at $t = 1$, then she will never find it optimal to switch islands prior to the start of $t = 2$.

Lemma 1.2 *Suppose the worker selects an island on which to work at $t = 1$, and then is given the option to switch islands prior to the start of $t = 2$. If $V_1 \geq V_2$, then*

$$\frac{w_1^{1-\gamma}}{1-\gamma} \geq \mathbb{E}_{z_2}\left[\frac{(w_2 e^{z_2})^{1-\gamma}}{1-\gamma}\right]. \tag{1.6}$$

Proof This result follows directly from linearity of the expectation operator.

It follows that the value of selecting island 1 at $t = 1$ is given by (1).

The next lemma states that if the worker selects island 2 at $t = 1$, then she will exercise her option to switch islands prior to the start of $t = 2$ if her realized productivity z_2 is less than a threshold value z_2^*. Otherwise, she will remain on island 2 for $t = 2$.

Lemma 1.3 *Suppose the worker chooses to work on island 2 at* $t = 1$ *and has the option to switch islands prior to the start of* $t = 2$*. Then* $\exists z_2^* \in \mathbb{R}$ *such that if* $z_2 < z_2^*$*, then she will exercise her option to switch islands. Moreover,* $z_2^* = \ln(w_1/w_2)$*.*

Proof Consider a worker who selected island 2 at $t = 1$. The value of remaining on island 2 for $t = 2$ is given by

$$V_{2,2} = \frac{\left(w_2 e^{z_2}\right)^{1-\gamma}}{1-\gamma}, \tag{1.7}$$

which is strictly increasing in z_2 (i.e. $\partial V_{2,2}/\partial z_2 > 0 \; \forall z_2 \in \mathbb{R}$). The value of switching to island 1 prior to the start of $t = 2$, on the other hand, is given by

$$V_{2,1} = \frac{w_1^{1-\gamma}}{1-\gamma}, \tag{1.8}$$

which is independent of z_2. Thus, there exists a threshold $z_2^* = \ln(w_1/w_2)$ such that $V_{2,2} = V_{2,1}$. Moreover, if $z_2 < z_2^*$, then $V_{2,2} < V_{2,1}$ and the worker will find it optimal to switch to island 1 prior to the start of $t = 2$. Otherwise, $V_{2,2} \geq V_{2,1}$ and the worker will find it optimal to remain on island 2 for $t = 2$.

It follows that the value of selecting island 2 at $t = 1$ is given by

$$V_2 = \mathbb{E}_{z_2}\left[\frac{\left(w_2 e^{z_2}\right)^{1-\gamma}}{1-\gamma}\left(1 + \beta \mathbb{I}\left[z_2 \geq \ln\left(\frac{w_1}{w_2}\right)\right]\right) + \beta \mathbb{I}\left[z_2 < \ln\left(\frac{w_1}{w_2}\right)\right]\frac{w_1^{1-\gamma}}{1-\gamma}\right], \tag{1.9}$$

where $\mathbb{I}[X]$ is an indicator function set equal to a value of one if the logical expression X is true and zero otherwise.

The final theorem states that for the worker to be indifferent between choosing islands 1 and 2 at $t = 1$, the log wage differential between islands is the solution to a non-linear fixed-point problem.

Theorem 1.4 *Suppose the worker selects an island on which to work at $t = 1$, and then is given the option to switch islands prior to the start of $t = 2$. If $V_1 = V_2$, then* $\ln(w_1/w_2)$ *solves the following fixed-point problem:*

$$\ln\left(\frac{w_1}{w_2}\right) = -\gamma\frac{\sigma_2^2}{2} + G\left(\beta, \gamma, \sigma_2, \ln\left(\frac{w_1}{w_2}\right)\right). \tag{1.10}$$

Proof See Appendix A for a derivation of this result and a proof of existence.

The first term on the right hand side of (10) is the risk premium effect, while the second represents the option value effect. The relationship between w_1 and w_2 is then determined by the relative strength of these opposing forces.

Since an analytical solution to (10) does not exist, I proceed numerically. Figure 1.2 (a) plots the log wage differential which solves (10) as a function of σ_2 for the case in which $\gamma = 2$. As σ_2 increases, the option value effect increases at an approximately constant rate while the risk premium effect decreases at a rate that is increasing in σ_2. For this particular parameterization of the model, the option value effect dominates the risk premium effect for $\sigma_2 \in [0, 0.16]$, leading to a positive inter-island wage differential, which implies that $w_1 > w_2$. In this case, the worker is willing to accept a lower wage to work on the island for which their productivity is uncertain in exchange for the possibility of realizing a high positive productivity draw. As σ_2 increases beyond 0.16, the risk premium effect dominates and the worker demands a premium to work on the island for which their payoff is uncertain.

Figure 1.2 (b) plots the log wage differential which solves (10) as a function of γ for the case in which $\sigma_2 = 0.10$. As γ increases, the option value effect remains nearly constant while the risk premium effect decreases at a constant rate. For this particular parameterization of the model, the option value effect dominates the risk premium effect for $\gamma \in (1.0, 3.2]$, leading to a positive inter-island wage differential. Again, in this case, the worker is willing to accept a lower wage to work on the island for which their productivity is uncertain in exchange for the possibility of realizing a high positive productivity draw. As γ increases beyond 3.2, the risk premium effect dominates and the worker demands a premium to work on the island for which their payoff is uncertain.

The key feature of this stylized model which generates this inverse relationship between wages and wage volatility is the option to switch islands. Inter-industry mobility in this context acts like a put option by providing the worker insurance against the possibility of a low island-specific productivity draw and thus serving as an insurance channel against labor market risk. Using the full quantitative model which I develop in the following section, I will demonstrate that for empirically plausible parameterizations, the option value effect dominates the risk premium effect, leading to a similar negative relationship between risk and return across industries. Moreover, I will show that inter-industry mobility is the key feature of the model which is responsible for generating this result.

1.3 Quantitative Model

In this section, I outline the general equilibrium model which I later take to the data in order to estimate the relationship between inter-industry wage differentials and the volatility of persistent idiosyncratic shocks to wages. This model builds upon the stylized theoretical model developed in the previous section by relaxing the assumption that agents cannot save or borrow, allowing for a labor supply decision, and introducing a nonzero resource cost of switching islands. I also introduce an overlapping generations structure and explicitly model the production technologies of the economy.

1.3.1 Environment

The economy consists of overlapping generations of ex ante identical agents who live for $N + R$ periods and have time separable preferences over consumption and leisure. There are J islands, each representing an industry and containing a representative firm which hires labor to produce a differentiated intermediate good. A representative firm combines intermediate goods and capital to produce a homogeneous final consumption good. There is also a financial intermediary which rents capital to the final goods producer and borrows and lends in credit markets using one-period uncontingent bonds. I assume the lender faces a proportional transaction cost τ of issuing debt contracts.

1.3.2 Problem of an Agent

Each agent is endowed with a vector of island-specific idiosyncratic productivities $\mathbf{z} \in \mathbb{R}^J$, where z_j is the log efficiency units of labor the agent has available to supply to the representative firm on island j. While working on island j, z_j follows a random walk, while $z_{k \neq j}$ remain fixed.[10] Islands differ in the volatility of the idiosyncratic productivity shocks that workers face. Agents decide on which island to work optimally each period during their first N periods of life, subject to a resource cost $\chi(a) > 0$ of switching islands which may vary with an agent's age a. During their final R periods of life, agents are restricted from working and therefore simply enjoy leisure and consume their savings.

Each period, the problem of an agent is divided into two stages. In the first stage, the agent selects an island j on which to work conditional on their age a, bond holdings b, productivity vector $\mathbf{z}_{-1}$, and previous island of employment m:

$$V_m^a(b, \mathbf{z}_{-1}) = \max_{j \in \{1,\ldots,J\}} \mathbb{E}_{\zeta,e} \left[W_{j,m}^a(b, \mathbf{z}, e) \right] \tag{1.11}$$

subject to

$$z_j = z_{j,-1} + \zeta \tag{1.12}$$

and

$$z_k = z_{k,-1} \text{ for all } k \neq j, \tag{1.13}$$

[10] A worker's productivity is industry-specific and evolves only with industry-specific experience. These assumptions are consistent with Jovanovic (1979) who develops a model of learning about job-specific ability to explain the negative empirical relationship between job turnover and job tenure. Neal (1995) analyzes wage data from Displaced Worker Surveys and concludes that industry-specific human capital is an important factor in explaining the observed relationship between wages and seniority. Parent (2000) finds similar patterns in data from the PSID and the National Longitudinal Survey of Youth (NLSY79) and concludes that what matters most for the wage profile is the accumulation of industry-specific human capital.

where $W^{a}_{j,m}(b,\mathbf{z},e)$ is the value of working on island j conditional on the realizations of an idiosyncratic island-specific permanent productivity shock $\zeta \sim N\left(-\frac{1}{2}\sigma^2_{\zeta,j},\sigma^2_{\zeta,j}\right)$ and an idiosyncratic transitory productivity shock $e \sim N\left(-\frac{1}{2}\sigma^2_{e,j},\sigma^2_{e,j}\right)$.[11] Importantly, the volatility of both productivity shocks is allowed to vary across islands.[12]

In the second stage, the agent first draws ζ and e, and then optimally selects the bond holdings to carry into the following period b' and the fraction of time to spend working in the current period h, subject to their time endowment $\overline{h}$:

$$W^{a}_{j,m}(b,\mathbf{z},e) = \max_{h\in\left[0,\overline{h}\right],b'} \frac{\left[c^{\nu}(1-h)^{1-\nu}\right]^{1-\gamma}}{1-\gamma} + \beta V^{a+1}_{j}(b',\mathbf{z}) \tag{1.14}$$

subject to

$$c + b' + \chi(a)\mathbb{I}[j \neq m] = \left(1 + r + \tau\mathbb{I}[b \leq 0]\right) b + w_j \phi_j h, \tag{1.15}$$

and

$$\ln \phi_j = f(a) + z_j + e, \tag{1.16}$$

where γ is the coefficient of relative risk aversion, ν governs time spent enjoying leisure relative to consumption, r is the interest rate, w_j is the wage on island j in units of consumption per efficiency unit of labor, and ϕ_j is the efficiency units of labor the agent has available to supply to the representative firm on island j, which itself is a function of the agent's age a, island-specific productivity z_j, and transitory shock e. For agents in their first N periods of life, $\overline{h} = 1$, while for agents in their final R periods of life, $\overline{h} = 0$. The interest rate r will

[11] I also assume that both productivity shocks have bounded support, $\zeta \in [\underline{\zeta},\overline{\zeta}]$ and $e \in [\underline{e},\overline{e}]$, which, when combined with my assumption that agents cannot work during retirement, leads to an endogenous borrowing constraint that varies with age a, current industry of employment j, and productivity vector $\mathbf{z}$.

[12] Industry-specific productivity shocks are intended to capture a multitude of factors which affect the wages earned by an individual in their current industry of employment such as changes in the relative demand for the individual's skill set (occupation, experience, education, etc.) within the industry in which they are currently employed, learning about industry-specific ability (think of the consulting industry, for example, in which many enter but only a select few have both the talent and desire to become partner), or varying degrees of insurance against aggregate, industry, and firm-level shocks provided by employers. Lagakos and Ordonez (2011), for example, use a model of limited commitment to demonstrate that workers who are employed in industries with high displacement costs get more insurance from their employer, and, as a result, wages are smooth for workers in these industries even when firm productivity is volatile.

be determined in equilibrium by the market clearing condition for capital, while the wages $\{w_j\}_{j=1}^J$ will be pinned down by the market clearing condition for labor on each island.

1.3.3 Intermediate Goods Producers

The representative firm on island j produces a differentiated intermediate good according to a technology that is linear in efficiency units of labor L_j. The firm chooses the quantity of labor that maximizes their static profits π_j taking the wage on their island w_j and the price of their output p_j as given:

$$\pi_j = \max_{L_j} p_j x_j - w_j L_j \tag{1.17}$$

subject to the production function

$$x_j = L_j. \tag{1.18}$$

The first order condition for profit maximization implies that the equilibrium wage on island j is equal to the price of the intermediate good produced on island j:

$$w_j = p_j. \tag{1.19}$$

1.3.4 Final Goods Producer

A representative firm produces a homogeneous final consumption good using capital and a CES aggregation of the differentiated intermediate goods produced on islands $j \in \{1, ..., J\}$. Each period, the firm maximizes its static profits Π taking the rental rate of capital r and the prices of intermediate goods $\{p_j\}_{j=1}^J$ as given:

$$\Pi = \max_{K, \{X_j\}_{j=1}^J} Y - (r + \delta)K - \sum_{j=1}^{J} p_j X_j \tag{1.20}$$

subject to the production function

$$Y = K^{\alpha} \left(\sum_{j=1}^{J} X_j^{\rho} \right)^{(1-\alpha)/\rho}, \tag{1.21}$$

where δ is the depreciation rate of capital and $1/(1-\rho)$ is the elasticity of substitution between intermediate goods. The first order conditions for profit maximization pin down the equilibrium interest rate

$$r = \alpha K^{\alpha-1} \left(\sum_{j=1}^{J} X_j^{\rho} \right)^{(1-\alpha)/\rho} - \delta \tag{1.22}$$

and the equilibrium price of each intermediate good $j \in \{1, ..., J\}$

$$p_j = (1-\alpha) X_j^{\rho-1} K^{\alpha} \left(\sum_{j=1}^{J} X_j^{\rho} \right)^{(1-\alpha-\rho)/\rho}. \tag{1.23}$$

1.3.5 Market Clearing

Let $\Gamma(b, \mathbf{z}_{-1}; a, m)$ be the distribution of agents over states. There are markets for labor on each island, capital, intermediate goods, and the final consumption good, all of which must clear in equilibrium.

The market clearing condition for labor on each island $j \in \{1, ..., J\}$ is given by:

$$\int d_j(b, \mathbf{z}_{-1}; a, m) \int \phi_j(b, \mathbf{z}, e; a, m) h(b, \mathbf{z}, e; a, m) dF(\zeta, e) d\Gamma(b, \mathbf{z}_{-1}; a, m) = L_j, \tag{1.24}$$

where $d_j(\cdot)$ is equal to one if the agent is employed on island j and zero otherwise. This condition states that the total efficiency units of labor supplied by agents to the representative firm on island j must be equal to the total efficiency units of labor demanded by the representative firm on island j. The relative wages $\{w_j\}_{j=1}^{J}$ adjust to clear the market for labor on each island.

Assuming that risk-free bonds are in zero net supply, the market clearing condition for the capital market is given by:

$$\int\int b'(b,\mathbf{z},e;a,m)\ dF(\zeta,e)d\Gamma(b,\mathbf{z}_{-1};a,m) = K. \tag{1.25}$$

The interest rate r adjusts to clear the market for capital.

The market for each intermediate good $j \in \{1,...,J\}$ clears when the quantity produced by the representative firm on island j is equal to the quantity demanded by the final goods producer:

$$X_j = x_j. \tag{1.26}$$

The price of each intermediate good p_j adjusts to clear the market for the intermediate good produced on island j.

Finally, the market for the final consumption good clears when the aggregate quantity demanded by agents equals the total quantity produced by the final goods producer:

$$\int\int c(b,\mathbf{z},e;a,m)\ dF(\zeta,e)d\Gamma(b,\mathbf{z}_{-1};a,m) = Y. \tag{1.27}$$

Note that by Walras' Law, if the markets for labor, capital, and intermediate goods clear, then the market for the final consumption good also clears.

1.3.6 Defining a Stationary Equilibrium

A stationary equilibrium is a set of decision rules for agents $\{d_j(b,\mathbf{z}_{-1};a,m), b'(b,\mathbf{z},e;a,m), l(b,\mathbf{z},e;a,m)\}$, demand for labor by intermediate goods producers $\{L_j\}_{j=1}^J$, demand for capital and intermediate goods by the final goods producer $\{K,\{X_j\}_{j=1}^J\}$, wages and output prices on each island $\{w_j,p_j\}_{j=1}^J$, an interest rate r, and a distribution of agents over states $\Gamma(b,\mathbf{z}_{-1};a,m)$ such that:

- Taking $\{w_j\}_{j=1}^J$ and r as given, $\{d_j(b,\mathbf{z}_{-1};a,m), b'(b,\mathbf{z},e;a,m), h(b,\mathbf{z},e;a,m)\}$ solve each agent's optimization problem.

- Taking $\{w_j, p_j\}_{j=1}^J$ as given, $\{L_j\}_{j=1}^J$ maximizes the static profits π_j earned by the intermediate goods producer on each island $j = \{1, ..., J\}$.

- Taking r and $\{p_j\}_{j=1}^J$ as given, $\{K, \{X_j\}_{j=1}^J\}$ maximizes the static profits Π earned by the final goods producer.

- Given $\Gamma(b, \mathbf{z}_{-1}; a, m)$ and the decision rules for agents and firms, the markets for labor, capital, and intermediate goods clear.

1.4 Calibration

I start by selecting those parameters which can be reasonably calibrated outside of the model. I assume agents are born at age 18 with zero assets and productivity vector $\mathbf{z} = (0, ..., 0)$, retire at age 58, and die with probability one at age 78. Each period in the model represents one year, and therefore I take $N = 40$ and $R = 20$. As is standard in the literature, I set the coefficient of relative risk aversion γ equal to 2 and the discount factor β equal to 0.96.[13] I set the depreciation rate of capital δ equal to 0.06 and follow Hsieh, Hurst, Jones, and Klenow (2012) in setting ρ, which determines the final goods producer's elasticity of substitution between intermediates, equal to 0.75.

The remaining model parameters are the age-earnings profile $f(a)$, the industry-specific shock volatilities $\{\sigma_{\zeta,j}, \sigma_{e,j}\}_{j=1}^J$, the resource cost of switching industries $\chi(a)$, the preference parameter governing time spent enjoying leisure relative to consumption ν, and the proportional transaction cost of borrowing in credit markets τ. I now turn to micro-level data on the wages and inter-industry mobility patterns of working age males in order to identify a useful set of target moments in the data. I then employ a combination of indirect inference and simulated method of moments in order to calibrate these remaining parameters.

[13] In the following section, I explore the sensitivity of my quantitative results to alternative, empirically plausible calibrations of the coefficient of relative risk aversion γ.

1.4.1 The Data

The University of Michigan's Panel Study of Income Dynamics (PSID) is a longitudinal household survey that began in 1968 with a nationally representative sample of over 18,000 individuals living in 5,000 families in the United States. Of these families, about 3,000 were representative of the US population as a whole (Survey Research Center Sample) while the remaining 2,000 were an over-sampling of low-income families (Survey of Economic Opportunity). Information on individuals and their descendants was collected annually through 1997 and biennially thereafter. At the time of the most recent survey, the PSID has grown to include more than 22,000 individuals living in 9,000 families.

The PSID survey is comprehensive and includes questions pertaining to employment, income, wealth, education, and health, as well as numerous other topics. I use data from the merged family– and individual–level files on age, years of education, labor force participation, self-employment status, labor income, hours worked, union membership status, and job tenure.[14] The PSID also reports each individual's occupation and industry affiliation at the three-digit level. Between 1968 and 1980, this information is based on the Retrospective Occupation-Industry Files, while from 1981 onward it is based on the main survey data.[15] Due to sample size restrictions and reporting practices, I choose to classify individuals using the 2000 Standard Occupational Classification System (SOC) and the 1970 Census One-digit Industry Codes. Descriptions of each classification system are included in Tables 1.1 and 1.2.[16]

[14]Labor income is defined as the sum of wages, bonuses, commissions, and overtime pay.

[15]Between 1968 and 1980, the PSID recorded occupations and industries using various combinations of one- and two-digit codes. The 1968–1980 Retrospective Occupation-Industry Files provide 1970 Census Three-digit Codes for the occupation and industry of each individual's main job for all sample years prior to 1981 based on a recoding of handwritten job descriptions.

[16]For the purpose of studying wage risk, the PSID has many advantages and relatively few disadvantages compared with other longitudinal panels such as the SIPP, NLSY79 and NLSY97, or the CPS. For example, although the CPS is much larger than the PSID and offers a two year panel structure, the lack of a third consecutive observation on each individual prevents the disentangling of persistent wage shocks from transitory fluctuations. While the NLSY79 and NLSY97 meet this requirement, the PSID offers data on a wider range of cohorts within each sample year. The main disadvantage of the PSID is its relatively small sample size which prevents a finely disaggregated analysis. Moreover, the PSID records the total labor earnings for each individual within each sample year, prohibiting the study of high frequency wage changes related

My sample selection procedure is as follows. I keep only heads of household who were in, or are descendants of, the nationally representative Survey Research Center sample. I also restrict my sample to include only those individuals aged 18 to 55 who were in the labor force, not self-employed, and worked at least 520 hours in the sample year.[17] I keep only males, and therefore abstract from the large increase in female labor force participation rates during my sample period. I compute the average real wage for each individual-year observation by first deflating total labor income using the Consumer Price Index reported by the Bureau of Economic Analysis with 1982 taken as the base year and then dividing by total hours worked in the sample year. In order to reduce errors in either the reporting or recording of labor income and hours worked, I drop those individual-year observations with real hourly wages less than \$2.50 or greater than \$250.[18] These criteria are similar to the ones used in previous related studies by Abowd and Card (1989), Guvenen (2009), and Heathcote, Storesletten, and Violante (2010), among others. There are a total of 51,059 individual-year observations for which I have non-missing data on age, years of education, labor income, hours worked, union membership status, firm tenure, and occupation and industry affiliations. Due to the relatively small number of individual-year observations in either the personal services or entertainment and recreation services industries, I choose to combine these two industries to form a merged industry group which hereafter I refer to as "Other Services." This procedure ensures a balance between sufficient disaggregation of industry classifications and the need for enough observations within each industry group to generate meaningful statistical inference.

to movements into and out of unemployment, which I abstract from in both my theoretical and empirical analysis. While the SIPP offers information on income and hours worked at a quarterly frequency, the PSID allows for the analysis of wage changes over multiple years which is critical for identifying persistent shocks to wages.

[17]This restriction corresponds to 13 weeks of full-time employment. While this step eliminates all individuals who were unemployed for more than three-quarters of the given sample year, this fraction of individuals is quite small, and hence its adverse effect on my estimates is minimal. The main purpose of this restriction is to ensure that my sample consists mainly of individuals who have a strong attachment to the labor force.

[18]The minimum wage in 1982 was \$3.36. This criteria helps avoids issues related to changes in the PSID structure in 1992 which, according to Dynan, Elmendorf, and Sichel (2012), resulted in a dramatic increase in the number of respondents who simultaneously reported positive hours worked and zero labor income.

Key summary statistics for my base sample are reported in Tables 1.3 and 1.4. Table 1.3 lists the number of observations and the average age, years of education, union membership rate, hours worked, and real hourly wage for each sample year.[19] The number of observations increases from just over 1,000 to an average of around 2,000 by the end of my sample period.[20] The average age in all sample years is between 34 and 39, while the average years of education increases gradually from about 12 in 1967 to nearly 14 by 2008, the latter of which is broadly consistent with the known increase in college enrollment. Interestingly, the union membership rate declined by over 50% during my sample period. Table 1.4 reports the distribution of individual-year observations across the eleven industries under consideration in each sample year. Notably, the manufacturing industry experienced a sharp decline in its employment share between 1967 and 2008, while most service industries, as well as the construction industry, realized gains in their share of employment.

1.4.2 Auxiliary Model of Wage Dynamics

Allowing for inter-industry mobility in my quantitative model leads to an important selection effect – agents who have success in their current industry are likely to remain, while those who receive adverse productivity shocks have an incentive to switch. Thus, the model predicts that industry tenure and industry-specific productivity should be positively correlated. Since the latter is unobservable to the econometrician, if the theoretical model that I develop here is the true data generating process, inter-industry wage differentials estimated using standard econometric techniques will be biased. In addition, the volatility of idiosyncratic shocks to wages can only be estimated using the observed wage changes of individuals who do not switch industries between two adjacent points in time. Since the theory predicts that those individuals who receive sufficiently adverse shocks will exercise their option to switch industries, estimates of wage volatility measured directly from the data will also be biased.

Controlling for selection using a purely reduced form approach would require identifying

[19] There are no observations for 1977, 1978, or 1979 due to a lack of data on firm tenure.
[20] The size of the PSID grows as descendants of individuals in the original SRC sample enter the panel.

a valid instrument. Instead, I use my quantitative model to control for selection and employ a calibration strategy based on the method of indirect inference first introduced by Smith (1993) and later extended by Gourieroux, Monfort, and Renault (1993).[21] I proceed by defining the following auxiliary model which I use as a device for estimating a useful set of target moments in the actual data:

$$\ln \tilde{w}_{i,j,t} = \xi_t + \eta_{j,t} + x_{i,j,t}\psi + f(a_{i,t}) + \omega_j + z_{i,j,t} + e^b_{i,j,t}, \tag{1.28}$$

where

$$z_{i,j,t} = z_{i,j,t-2} + \zeta^b_{i,j,t}, \tag{1.29}$$

$\tilde{w}_{i,j,t}$ is the real wage for individual i employed in industry j at time t, ξ_t denotes a year fixed-effect that is common across all industries, $\eta_{j,t}$ is an industry-specific year fixed-effect, $x_{i,j,t}$ is a vector of individual-level controls including years of education, union membership status, occupation, firm tenure, occupational tenure, and industry tenure, $f(a_{i,t})$ is a quartic in age, ω_j is a fixed-effect of employment in industry j, $\zeta^b_{i,j,t} \sim N(0, \sigma^b_{\zeta,j})$ is a biennial permanent wage shock, and $e^b_{i,j,t} \sim N(0, \sigma^b_{e,j})$ is a biennial transitory wage shock.[22] I assume $\mathbb{E}[z_{i,j,t}|t, j, x_{i,j,t}, a_{i,t}] = 0$, which is a necessary condition for identification.[23]

Critically, my auxiliary model is chosen to allow for a simple decomposition of wage risk into its persistent and transitory components following the methods of Moffitt and Gottschalk (2002), Meghir and Pistaferri (2004), and Low et al. (2010). I start by taking first differences of equation (28), after substituting in for $z_{i,j,t}$ using equation (29), for individuals who do not

[21]See Smith (2008) for an overview of indirect inference and more recent applications.

[22]Specifically, $f(a_{i,t}) = \theta_1\ a_{i,t} + \theta_2\ a^2_{i,t}/10^2 + \theta_3\ a^3_{i,t}/10^4 + \theta_4\ a^4_{i,t}/10^6$. Given that the PSID transitioned from an annual to biennial survey after 1997, I focus on two-year wage changes in order to maintain consistency throughout my entire sample period. Dynan et al. (2012) use a similar two-year panel structure to estimate changes in the volatility of household income in data from the PSID over the same sample period.

[23]If my theoretical model is the true data generating process, however, this condition will certainly not hold due to endogenous selection effects. Misspecification, however, is not a problem for indirect inference since the auxiliary model is merely a lens through which to view both the actual data and the data generated by my quantitative model. My quantitative model, which explicitly controls for selection, will be used to uncover the true trade-off between risk and return across industries.

switch industries between dates t and $t-2$ to obtain the following wage growth equation:

$$\Delta \ln \tilde{w}_{i,j,t} = \Delta\left(\hat{\xi}_t + \hat{\eta}_{j,t} + x_{i,j,t}\hat{\psi} + \hat{f}(a_{i,t})\right) + \left(\zeta^b_{i,j,t} + \Delta e^b_{i,j,t}\right), \tag{1.30}$$

where $\Delta y_{i,j,t} \equiv y_{i,j,t} - y_{i,j,t-2}$. The left hand side of equation (30) is observed wage growth, while the first term on the right hand side is that which is predicted by the auxiliary model based solely on observable factors. The second term on the right hand side is the unpredicted component of wage growth, or the cumulative shock $g_{i,j,t} \equiv \zeta^b_{i,j,t} + (e^b_{i,j,t} - e^b_{i,j,t-2})$. Computing the variance and first-order autocovariance of $g_{i,j,t}$ for individuals employed in industry j at dates t and $t-2$ leads to a system of equations which can be solved for the volatility of biennial permanent wage shocks for industry j:

$$\sigma^b_{\zeta,j} = \sqrt{\mathbb{E}[g^2_{i,j,t}] + 2\mathbb{E}[g_{i,j,t}g_{i,j,t-2}]}, \tag{1.31}$$

and the volatility of biennial transitory wage shocks for industry j:

$$\sigma^b_{e,j} = \sqrt{-\mathbb{E}[g_{i,j,t}g_{i,j,t-2}]}. \tag{1.32}$$

1.4.3 Estimation of the Auxiliary Model using Actual Data

A subset of the results from an OLS estimation of equation (28) using the base sample of data from the PSID are presented in column (1) of Table 1.5. The first four estimates represent the coefficients of the quartic in age $\hat{f}(a)$. Also reported are the estimated effects of an additional year of education and union membership. While not listed in the table, the vast majority of both the industry and industry-year fixed effects are statistically significant, which confirms that industry affiliation is an important factor in explaining variations in the real wage for observationally equivalent workers employed in different industries.

Estimates of inter-industry wage differentials are presented in column (1) of Table 1.6, computed as the difference between the cumulative fixed effect of affiliation with each industry j, Ω_j, and the employment weighted average of these cumulative fixed effects across all industries. The magnitude and pattern of these estimated wage differentials are similar to those previously reported in the literature. In particular, workers employed in the mining, manufacturing, and transportation, communications, and utilities industries earn higher wages on average than observationally equivalent workers employed in the agriculture, forestry, and fisheries, wholesale and retail trade, and other services industries. A typical worker employed in the mining industry, for example, earns a wage that is, on average, 27 percent higher than that of an observationally equivalent worker employed in the agriculture, forestry, and fishing industry. Moreover, the weighted standard deviation of these differentials is 0.061, which is roughly equivalent to the effect on wages of 1–2 additional years of education. Hence, wage differentials, as viewed through the lens of the auxiliary model, are large and vary quite substantially across industries.

After employing equation (30) to compute $g_{i,j,t}$ for industry-stayers, equations (31) and (32) were used to obtain estimates of the industry-specific biennial shock volatilities $\hat{\sigma}^b_{\zeta,j}$ and $\hat{\sigma}^b_{e,j}$ which are reported in column (1) of Table 1.7.[24] The employment weighted average and standard deviation of $\hat{\sigma}^b_{\zeta,j}$ are 0.139 and 0.021, respectively. Thus, through the lens of the auxiliary model, there is also substantial variation in the volatility of biennial permanent shocks to wages across industries. The transportation, communications, and utilities industry offers workers the lowest permanent shock volatility of 0.101, which is nearly than 32 percent less than that offered to workers by the agriculture, forestry, and fishing industry. The employment weighted average and standard deviation of $\hat{\sigma}^b_{e,j}$, on the other hand, are 0.152 and 0.008, respectively, which also reflect measurement error in either the reporting or recording of total labor income and hours worked.[25] Assuming measurement error is

[24] In order to reduce the impact of measurement error, I drop observations of $\Delta\tilde{w}_{i,j,t}$, $g_{i,j,t}$, and $g_{i,j,t}g_{i,j,t-2}$ that fall in either the 1% or 99% quantiles of their respective distributions across all industries.

[25] Given that my empirical strategy does not allow for the disentangling of true transitory wage shocks from measurement error, I rely on external estimates of the magnitude of measurement error in the data. In particular, Bound, Brown, Duncan, and Rodgers (1994) conduct a validation study of the PSID data on earnings and conclude that measurement error explains 22 percent of the overall variance in the rate of earnings growth in the PSID.

independent of industry affiliation, the data suggest that there is far less variation in the volatility of biennial transitory shocks to wages across industries.

The auxiliary model thus provides estimates of both the volatility of persistent idiosyncratic shocks to wages (*risk*) and inter-industry wage differentials (*return*) which can be used to measure the empirical trade-off between risk and return across industries. Figure 1.3 (a) plots the estimated inter-industry wage differentials versus the estimated volatility of permanent biennial shocks to wages, where each data point in the figure represents the risk-return trade-off offered by one of the eleven industries under consideration. The dotted green line depicts the weighted OLS regression line which best fits the data, the slope of which is equal to -1.34 with a standard error of 0.19. Thus, the data suggest that the trade-off between risk and return across industries is negative and significant at the $p = 0.0001$ level.[26] While this result confirms the preliminary result presented in the introduction is robust to controlling for observable heterogeneity, it fails to address the issue of selection. In the following section, I will use the calibrated quantitative model to explicitly control for selection and arrive at an unbiased estimate of this trade-off.

Given that agents can self-insure against transitory income fluctuations quite well using only one-period uncontingent bonds (as is true in the quantitative model), I choose to focus on the volatility of permanent shocks to wages as the measure of risk for my analysis. For comparison purposes, however, Figure 1.3 (b) depicts the relationship between the estimated inter-industry wage differentials and the estimated volatility of transitory biennial shocks to wages. In contrast to the volatility of permanent shocks to wages, there is little variation across industries in this measure of wage risk.

In order to verify whether or not the results presented here are robust to sample selection, I repeat my analysis of the auxiliary model for the following sub-groups of individuals in the

[26] The outlier in the northeast portion of Figure 1.3 (a) represents the risk-return trade-off offered by the mining industry. The difference between the actual wage differential and that predicted given the volatility of persistent idiosyncratic shocks to wages is likely due, at least in part, to compensating differentials for health risk. Thus, while variations in wage volatility can explain a large portion of inter-industry wage differentials, there is still considerable room for existing theories.

PSID: non-union (those not affiliated with a labor union), unskilled (those with at most a high school diploma), skilled (those with at least a 4 year college degree), 1967–1984 (individual-year observations in the first half of my sample period), and 1985–2008 (individual-year observations in the second half of my sample period). Results for each of these sub-groups are presented in columns (2) – (6) of Tables 1.5 – 1.8. While there is some variation in my estimates of the volatility of the permanent shocks across sub-groups, the rank correlation of each subgroup with the base sample is always positive and significant at the $p = 0.05$ level. There is substantially more variation in my estimates of the volatility of the transitory shocks across sub-groups, but the rank correlation of each subgroup with the base sample is still always positive. Finally, the slope of the weighted OLS regression line indicating the measured trade-off between risk and return through the lens of the auxiliary model is always negative and significant at the $p = 0.02$ level.

1.4.4 Indirect Inference and Simulated Method of Moments

The remaining model parameters to be chosen are the resource cost of switching industries $\chi(a)$, the age-earnings profile $f(a)$, the industry-specific shock volatilities $\{\sigma_{\zeta,j}, \sigma_{e,j}\}_{j=1}^{J}$, the preference parameter governing time spent enjoying leisure relative to consumption ν, and the proportional transaction cost of borrowing in credit markets τ. I now describe how these parameters are jointly calibrated using a combination of indirect indifference and simulated method of moments.

Before proceeding further, I set the number of islands in the economy J equal to 2 in order to reduce the computational burden involved in repeatedly simulating my general equilibrium model. Without loss of generality, let island 1 be the low volatility island and island 2 be the high volatility island. Mapping this simplifying assumption to the data, I divide the eleven industries into two mutually exclusive groups based on the estimated value of $\hat{\sigma}_{\zeta,j}^{b}$ for each industry j relative to the employment weighted average across all industries. Following this procedure, the low volatility group accounts for 48.5 percent of the individual-year observations in my base sample and includes the following three industries: (1) manufacturing, (2) transportation, communications, and utilities, and (3) public administration. The remaining eight industries then make up the high volatility group which accounts for the remaining 51.5 percent of the individual-year observations in my base sample.

Next, I use the auxiliary model to identify a natural set of target moments for use in selecting the age earnings profile and industry-specific shock volatilities. In particular, I choose $f(a)$ and $\{\sigma_{\zeta,j}, \sigma_{e,j}\}_{j=1}^{2}$ within the model such that the estimates $\hat{f}(a)$ and $\{\hat{\sigma}^b_{\zeta,j}, \hat{\sigma}^b_{e,j}\}_{j=1}^{2}$ obtained using the auxiliary model on model-generated data matches the corresponding estimates obtained using the auxiliary model on actual data.

The resource cost of switching industries, $\chi(a)$, can be interpreted as the cost of exercising the embedded option to switch industries, and therefore directly affects the option value of inter-industry mobility which, in turn, affects the equilibrium allocation of workers across industries. This cost is intended to capture the retraining or relocation required by industry switchers, the magnitudes of which may vary over the life cycle.[27] I assume $\chi(a)$ takes on one value for each of the following age groups: 18–24, 25–34, 35–44, and 45–58. In the base sample, 68.5 percent of individuals started their working life in one of the nine high volatility industries. I select $\chi(18 \leq a < 25)$ such that an equivalent fraction of 18 year olds in the model choose to work on the high volatility island. It is well know that the annual rate of inter-industry mobility declines with age. Interestingly, a similar relationship holds for movement between the high and low volatility groups described above. Since $\chi(a)$ determines the ease with which workers can move between islands, I choose $\chi(25 \leq a < 35)$, $\chi(35 \leq a < 45)$, and $\chi(45 \leq a < 58)$ to match the inter-group mobility rate for each age group relative to that for the 18–24 age group observed in the data.

The preference parameter ν governs the average fraction of time agents spend working relative to leisure. Individuals in my base sample worked an average of 2,223 hours per year. Assuming a time endowment of 16 hours per day, individuals worked an average of 38% of

[27] Young workers, for example, may face lower relocation costs than middle age workers since they are less likely to be married or own their home.

their free time. I choose ν to match this moment in the data. The final remaining model parameter is the proportional transaction cost of borrowing in credit markets τ, which I select to match the average unsecured debt-to-income ratio in 1985 of 0.058.[28]

I search for the set of parameters $\left\{f(a), \{\sigma_{\zeta,j}, \sigma_{e,j}\}_{j=1}^{2}, \chi(a), \nu, \tau\right\}$ that minimizes the weighted sum of the square deviations of each moment implied by the model-generated data from its target in the actual data as described above. The results of this procedure, which combines elements of both indirect inference and simulated method of moments, are summarized in Table 1.9. Of particular note are the annual shock volatilities and the resource cost of switching islands. Specifically, I find that the annual shock volatilities required to match the biennial shock volatilities observed in the data are $\{\sigma_{\zeta,1}, \sigma_{\zeta,2}\} = \{0.080, 0.103\}$ and $\{\sigma_{e,1}, \sigma_{e,2}\} = \{0.121, 0.125\}$, respectively. Importantly, these estimates control for both observable heterogeneity and selection effects. The resource costs of switching islands required to match the degree of risk taking in the labor force by young workers and the relative mobility rates of experienced workers range from 0.5 times median annual income for the 45–58 year old age group to as high as 1.9 times median annual income for the 25–34 year old age group. These values are similar to those reported by Artuç and McLaren (2012) who find that industry switching costs of between 1.0 and 1.5 times average annual income are needed to explain the patterns of inter-industry mobility observed in CPS data. Table 1.10 compares the calibration targets generated by the quantitative model with their counterparts in the actual data. The model is able to match most moments in the data quite well, however, it does generate slightly higher relative inter-industry mobility rates.

1.5 Quantitative Results

In this section, I first explore the calibrated quantitative model's predictions for the trade-off between risk and return across industries. After performing a rigorous sensitivity analysis, I use the model to quantify the importance of inter-industry mobility as an insurance channel

[28]See Livshits, MacGee, and Tertilt (2010).

against labor market risk. I then validate the model by comparing its predictions for the patterns of inter-industry mobility and expected wage profiles directly with the data.

1.5.1 Estimating the Risk-Return Trade-off across Industries

The trade-off between risk and return across industries estimated using the calibrated quantitative model, expressed as the log wage differential per unit of excess annual permanent wage volatility, is given by:

$$\frac{\ln(w_1) - \ln(w_2)}{\sigma_{\zeta,1} - \sigma_{\zeta,2}} = -1.54, \tag{1.33}$$

where $\{\sigma_{\zeta,1}, \sigma_{\zeta,2}\}$ are calibrated model parameters and $\{w_1, w_2\}$ are equilibrium objects. Given that $\sigma_{\zeta,1} > \sigma_{\zeta,2}$, the negative sign above indicates that the wage per efficiency unit of labor on the low volatility island exceeds that on the high volatility island. Hence, the calibrated quantitative model, which controls for observable heterogeneity and selection, predicts that industries which expose workers to higher (lower) than average wage volatility, pay relatively low (high) wages – for every 1 percent increase in the volatility of annual permanent shocks to wages that a worker accepts, their expected wage falls by 1.54 percent.

As an alternative, one can view data generated by the calibrated quantitative model through the lens of the auxiliary model. This approach facilitates a direct comparison between the model-implied trade-off between risk and return and that which is present in the actual PSID data. Re-estimating the auxiliary model on data generated by the calibrated quantitative model yields a trade-off between risk and return across industries, expressed as the log wage differential per unit of excess biennial permanent wage volatility, of:

$$\frac{\hat{\Omega}_1^m - \hat{\Omega}_2^m}{\hat{\sigma}_{\zeta,1}^b - \hat{\sigma}_{\zeta,2}^b} = -2.20. \tag{1.34}$$

Importantly, while this estimate controls for observable heterogeneity, it abstracts from selection. Thus, the difference between (33) and (34), when both are expressed on an annual basis, provides an estimate of the selection bias. Using actual data from the PSID, the comparable risk-return trade-off is given by:

$$\frac{\hat{\Omega}_1 - \hat{\Omega}_2}{\hat{\sigma}^b_{\zeta,1} - \hat{\sigma}^b_{\zeta,2}} = -2.44. \tag{1.35}$$

where $\hat{\Omega}_1$ and $\hat{\Omega}_2$ are the employment weighted averages of $\hat{\Omega}_j$ for industries in the low and high volatility groups, respectively.

Figure 1.3 (a) depicts the trade-off between risk and return estimated using the auxiliary model on both model-generated data (dashed red line) and actual data from the PSID (solid blue line). Both convey a quantitatively similar negative relationship between wage differentials and the volatility of persistent shocks to wages. Moreover, the R^2 for both lines is 0.21, implying that variations in the volatility of idiosyncratic shocks to wages across industries explains roughly 21% of inter-industry wage differentials.

In order to verify whether or not the results presented here are robust, I re-estimate the auxiliary model on the following sub-groups of individual-year observations in the PSID: non-union (those not affiliated with a labor union), unskilled (those with at most a high school diploma), skilled (those with at least a 4 year college degree), 1967–1984 (individual-year observations in the first half of my sample period), and 1985–2008 (individual-year observations in the second half of my sample period). Results for each of these sub-groups, along with the base sample, are presented in Table 1.11. While there is modest variation across sub-groups, the measured wage differential is always positive and the risk-return trade-off is always negative. Thus, the qualitative nature of my estimates from the auxiliary model using the base sample of individual-year observations from the PSID is robust to sample selection. Table 1.12 reports the trade-off between risk and return implied by the calibrated quantitative model for a range of plausible values of the coefficient of relative risk aversion.[29] While the risk premium effect is certainly increasing in γ, in all cases the option

[29] Cogley (2002) formulates an asset-pricing model with idiosyncratic risk and uses data on consumption growth from the Survey of Consumer Expenditures (CEX) to determine that a risk aversion coefficient of 8 is required to account for the empirical equity premium. In a related study, Storesletten, Telmer, and Yaron (2007) consider a model with idiosyncratic time varying income risk and conclude that a similar level of risk

value effect still dominates, leading to a negative relationship between risk and return, both within the economic model and when viewed through the lens of the auxiliary model.

1.5.2 The Quantitative Significance of Inter-industry Mobility

Individuals in the quantitative model have three ways to insure themselves against labor market risk, in addition to precautionary savings: (1) borrowing in credit markets, (2) adjusting their labor supply, and (3) switching islands of employment. I compute the stationary equilibrium in which all three of these insurance channels are eliminated and find that the resulting trade-off between risk and return implied by the model is 0.43, meaning that for every 1 percent increase in wage volatility that a worker accepts, their expected wage rises by 0.43 percent. Thus, when the only available insurance channel against labor market risk is precautionary saving, risk-averse agents demand a premium for exposure to wage volatility. I decompose the total change in the model-implied trade-off between risk and return (from -1.54 to 0.43) by sequentially eliminating these insurance channels one-by-one and computing the resulting stationary equilibrium. Since the order in which each is removed may influence the magnitude of its incremental effect, I consider all possible permutations in order to obtain upper and lower bounds on their relative importance. The results of this exercise are reported in Table 1.13, where each case represents one of the six unique orderings in which these three insurance channels can be removed from the model.

First, consider the importance of borrowing in credit markets. Young agents who receive positive permanent productivity shocks respond by borrowing heavily against the increase in their expected future income, an effect which is increasing in the volatility of the permanent productivity shocks that agents face. Conversely, agents who receive negative permanent productivity shocks tend to save in order to offset the resource cost of switching islands. Any additional resources required to switch islands are then financed by borrowing in credit markets. The attractiveness of seeking employment on the high volatility island when young,

aversion is required to match the empirical equity premium.

thus, depends critically on an agent's ability to borrow in credit markets. When agents have the option to switch islands, removing their ability to borrow in credit markets accounts for between 23.7 and 35.8 percent of the total increase in the trade-off between risk and return (Cases 1–3). When the option to switch islands is eliminated, however, borrowing in credit markets has a markedly diminished effect, accounting for merely 0.7 to 17.6 percent of the total (Cases 4–5). Therefore, the relative importance of borrowing in credit markets as an insurance channel against labor market risk is influenced heavily by whether or not agents also have the option to switch islands.

Next, consider labor supply as insurance against labor market risk. Agents low productivity may find it optimal to enjoy leisure rather than spend their time working. Those employed on the high volatility island will therefore value the ability to adjust their labor supply more than their counterparts since they are more likely to find themselves in a position where leisure is more valuable than work. Eliminating the labor supply decision thus reduces the relative attractiveness of employment on the high volatility island. This insurance channel, however, has a relatively minor quantitative effect on the trade-off between risk and return across islands, as it accounts for only 1.1 to 18.0 percent of the total change.

While the relative importance of borrowing in credit markets and adjusting labor supply are non-trivial, the option to switch islands consistently accounts for the majority of the total change in the trade-off between risk and return across islands, representing between 62.5 and 97.5 percent of the total change. Moreover, when this insurance channel is eliminated, the risk-return trade-off is always positive. In the absence of inter-industry mobility, agents are restricted to their initial industry of choice. In this case, agents demand a premium in exchange for wage volatility since they are fully exposed to the risk of realizing adverse permanent productivity shocks. This effect is particularly acute for agents employed on the high volatility island. The net result is a large increase in the risk-return trade-off across islands when the option to switch is eliminated. Inter-industry mobility is thus a quantitatively important insurance channel against labor market risk which is critical for understanding the negative relationship between inter-industry wage differentials and wage volatility observed in the data and predicted by the quantitative model.

Inter-industry mobility also has important implications for aggregate productivity. When agents are exposed to persistent idiosyncratic industry-specific productivity shocks, the ability to switch industries allows them to seek employment in industries in which they are most productive. When inter-industry mobility is eliminated from the model, measured productivity is 8.4 percent lower than in the baseline economy.[30] Thus, inter-industry mobility is not only an important insurance channel against labor market risk, it is also a productivity enhancing device which improves the allocation of workers across industries.

1.5.3 Patterns of Inter-industry Mobility

The frequency with which workers move between industries at the one-digit level is substantial. Moreover, the majority of these movements represent switches between the low and high volatility groups defined above. In fact, between 1967 and 1996, the average annual inter-industry mobility rate at the one-digit level for working age males was 11.1 percent, nearly 60 percent of which is accounted for by the movement of workers between the high and low volatility industries.[31] If the movement of workers between industries was random, then roughly one half of all industry switches at the one-digit level should have occurred between these volatility groups since each contains about one half of the individual-year observations in my sample. Given that the average annual inter-island mobility rate in the calibrated quantitative model is 1.8%, we can conclude that differences in wage volatility across industries is responsible for roughly one quarter of all inter-group mobility.[32]

[30] Measured productivity, or the Solow residual, is computed assuming the aggregate production function is Cobb-Douglas in capital and labor. Specifically, measured productivity is given by the ratio $Y/(K^{\alpha}L^{1-\alpha})$, where Y is output of the final consumption good, K is capital rented by the final goods producer, and L is aggregate hours worked.

[31] In order to identify genuine industry switches in the PSID data, I count only those changes in industry affiliation that are accompanied by a concurrent change of employer. This procedure, when combined with using the 1968–1980 Retrospective Occupation-Industry Files, has been shown by Kambourov and Manovskii (2012) to greatly reduce the impact of industry coding errors prevalent in the PSID data.

[32] Since productivity shocks are only one of many reasons for which individuals might choose to switch industries, it would be quite troubling if the calibrated quantitative model generated the same level of

In addition, the net flow of workers between these volatility groups over the life cycle is far from random. Figure 1.4 depicts the average fraction of agents working in a high volatility industry as a function of age, both in the model and in the actual PSID data. In the data, an average of 68.5 percent of young workers enter the workforce in a high volatility industry, while only 45.0 percent are employed in a high volatility industry at age 50.[33]

The pattern of inter-island mobility over the life cycle generated by the model matches the data quite well. In the model, young workers value wage volatility more than their older counterparts since they have more time before retirement to either realize large productivity gains or to exercise their embedded option to switch industries. The option value of inter-industry mobility is also increasing in the volatility of the productivity shocks that workers face. Therefore, young agents in the model disproportionately seek employment in industries with higher than average wage volatility. As they accumulate labor market experience, agents who receive sufficiently positive productivity shocks remain, while those who don't find it optimal to seek employment in another industry in which they are more productive.

The delay in inter-industry mobility predicted by the model for agents with less than 3 years of labor market experience results from the nonzero resource cost of switching islands – given that agents enter the model with zero assets, those who receive negative productivity shocks face a trade-off between borrowing to finance switching islands and biding their time to see if their island-specific productivity improves. Since the resource cost of switching islands for agents between 18 and 24 is equivalent to 1.08 times median annual income, young agents in the model tend to delay exercising their option to switch islands rather than borrowing in credit markets against their future income.

mobility as observed in the data.

[33]This pattern of inter-industry mobility over the life cycle is reminiscent of Miller (1984) who argues that because young workers are less experienced than their older counterparts, they are more willing to try out jobs for which success is rare. A similar notion of "young and foolish" was also put forth by Adam Smith who wrote "The contempt of risk and the presumptuous hope of success, are in no period of life more active than at the age at which young people choose their professions." (Smith (2009))

1.5.4 Expected Wage Profiles

Young agents will only be willing to expose themselves to highly volatile and relatively low average wages if they can credibly expect their wages to grow faster than they otherwise would if employed in an industry with less volatile, but relatively high average wages. Using data generated by the calibrated quantitative model, I regress log wages on a constant and a quartic in age separately for each of two groups of agents: those who initially chose to work on the low volatility island, and those who initially chose to work on the high volatility island. Figure 1.5 (a) then plots the results of this exercise as a function of age, where each curve in the figure represents the expected wage profile conditional on initial industry of employment. While the expected wage of agents who initially chose the high volatility island is lower than that of their counterparts, it grows faster and eventually the expected wage profiles cross. Young agents who initially chose to work on the high volatility island, thus, do so knowing that exposure to higher wage volatility offers the potential for faster wage growth over the life cycle. This feature of the model results from an important selection effect – since agents always have the option to switch to another island in which they are more productive, those who remain on the high volatility island can expect to earn higher wages on average than agents of the same age who were initially employed on the low volatility island.

To determine if this implication of the model is corroborated by the data, I regress log wages from the PSID separately for similarly constructed groups on a constant and a quartic in age, controlling for year fixed-effects, union membership status, and years of education. Figures 1.5 (b) and (c) illustrate the results of this exercise for unskilled and skilled workers, respectively.[34] Both figures depict a similar single-crossing pattern in which the expected wage of individuals who initially chose to work in a high volatility industry begin lower but eventually exceed those of their counterparts. Models which abstract from mobility will fail to generate these same systematic patterns, and thus are firmly rejected by the data.

[34] Unskilled workers are classified based on their first industry of employment before age 20, while skilled workers are classified based on their first industry of employment between ages 23 and 25. I allow for a range of ages so that individuals who enter the PSID at age 19, for example, can still be included in my analysis.

1.6 Conclusion

This paper investigates the trade-off between risk and return across industries using a macro–labor approach. While standard economic reasoning tells us that risk averse workers should demand a premium for exposure to wage volatility, the data suggests that precisely the opposite is true – industries which expose workers to relatively low (high) wage volatility pay high (low) average wages. I demonstrate this counter-intuitive feature of the data can be quantitatively rationalized by a general equilibrium, incomplete markets, life cycle model in which risk-averse workers are allowed to optimally select their industry of employment each period. Inter-industry wage differentials arise in equilibrium as young agents are willing to accept a relatively low initial wage in exchange for the possibility of rapid wage growth in the future. The resulting wage differentials depend on the relative sizes of the option value and risk premium effects, which, in turn, are determined largely by the resource cost of switching industries and the volatility of productivity shocks that workers face. My quantitative results suggest that inter-industry mobility is an important insurance channel against labor market risk, which is responsible for generating the negative trade-off between risk and return across industries. The model is consistent with the empirical trade-off between risk and return, patterns of inter-industry mobility, and expected wage profiles.

A natural next step in this line of research is to determine whether or not the findings described here are compatible with the positive relationship between risk and return across occupations documented by Dillon (2012). A significant challenge will be to simultaneously explain the fact that rates of occupational mobility far exceed rates of inter-industrial mobility.[35] Another promising area for future research is to understand why idiosyncratic wage volatility varies so widely across industries in the first place. Differences across industries in the amount of risk sharing between firms and workers, as documented by Lagakos and Ordonez (2011), may prove to be a useful starting point for this analysis.

[35]See Kambourov and Manovskii (2008) for occupational and inter-industry mobility rates in the U.S.

1.7 Appendix

In this section, I provide formal derivations of the results presented in Section 2. To start, consider the case in which the agent is restricted to their initial island of choice. In this case, the value of choosing island 1 can be simplified as follows:

$$V_1 = \frac{w_1^{1-\gamma}}{1-\gamma} + \beta \frac{w_1^{1-\gamma}}{1-\gamma} = \frac{w_1^{1-\gamma}}{1-\gamma}(1+\beta),$$

while the value of selecting island 2 can also be simplified as follows:

$$V_2 = \mathbb{E}_{z_2}\left[\frac{(w_2 e^{z_2})^{1-\gamma}}{1-\gamma} + \beta \frac{(w_2 e^{z_2})^{1-\gamma}}{1-\gamma}\right] = \frac{w_2^{1-\gamma}}{1-\gamma}(1+\beta)\mathbb{E}_{z_2}\left[e^{z_2(1-\gamma)}\right]$$

Equating V_1 and V_2 yields:

$$\frac{w_1^{1-\gamma}}{1-\gamma} = \frac{w_2^{1-\gamma}}{1-\gamma}\mathbb{E}_{z_2}\left[e^{z(1-\gamma)}\right]$$

$$\left(\frac{w_1}{w_2}\right)^{1-\gamma} = \mathbb{E}_{z_2}\left[e^{z_2(1-\gamma)}\right]$$

Taking logs, and recognizing that $e^{z_2(1-\gamma)}$ is log normally distributed with mean $-\frac{\sigma_2^2}{2}(1-\gamma)$ and variance $\sigma_2^2(1-\gamma)^2$, yields the following expression the log wage differential:

$$\ln\left(\frac{w_1}{w_2}\right) = \frac{1}{1-\gamma}\ln\left(\mathbb{E}_{z_2}\left[e^{z_2(1-\gamma)}\right]\right) = \frac{1}{1-\gamma}\ln e^{-\frac{\sigma_2^2}{2}(1-\gamma)+\frac{\sigma_2^2}{2}(1-\gamma)^2} = -\gamma\frac{\sigma_2^2}{2}$$

which is strictly negative for any $\gamma > 1$ and $\sigma_2 > 0$.

Assuming the agent can switch islands prior to the second period, the value of choosing island 1 is given by:

$$V_1 = \frac{w_1^{1-\gamma}}{1-\gamma} + \beta \max\left\{\frac{w_1^{1-\gamma}}{1-\gamma}, \mathbb{E}_{z_2}\left[\frac{(w_2 e^{z_2})^{1-\gamma}}{1-\gamma}\right]\right\}$$

while the value of selecting island 2 is given by:

$$V_2 = \mathbb{E}_{z_2}\left[\frac{(w_2 e^{z_2})^{1-\gamma}}{1-\gamma} + \beta \max\left\{\frac{w_1^{1-\gamma}}{1-\gamma}, \frac{(w_2 e^{z_2})^{1-\gamma}}{1-\gamma}\right\}\right]$$

First, note that if $V_1 \geq V_2$, linearity of the expectations operator implies that

$$\frac{w_1^{1-\gamma}}{1-\gamma} \geq \mathbb{E}_{z_2}\left[\frac{(w_2 e^{z_2})^{1-\gamma}}{1-\gamma}\right]$$

and therefore any worker who chooses island 1 in period 1 will never find it optimal to switch to island 2 in period 2. Hence, the value of choosing island 1 can be simplified as follows:

$$V_1 = \frac{w_1^{1-\gamma}}{1-\gamma} + \beta\frac{w_1^{1-\gamma}}{1-\gamma} = \frac{w_1^{1-\gamma}}{1-\gamma}(1+\beta),$$

Second, for every $\{w_1, w_2\}$ there exists a threshold productivity shock z_2^* such that workers who choose island 2 in period 1 decide to stay on island 2 in period 2 if $z_2 \geq z_2^*$ and decided to switch to island 1 in period 2 otherwise. Note that this threshold value is precisely that for which $w_1 = w_2 e^{z_2^*}$, and hence $z_2^* = \ln(w_1/w_2)$. Equating V_1 and V_2 and substituting in for w_1 inside of the max operator yields:

$$\frac{w_1^{1-\gamma}}{1-\gamma}(1+\beta) = \mathbb{E}_{z_2}\left[\frac{(w_2 e^{z_2})^{1-\gamma}}{1-\gamma} + \beta \max\left\{\frac{(w_2 e^{z_2^*})^{1-\gamma}}{1-\gamma}, \frac{(w_2 e^{z_2})^{1-\gamma}}{1-\gamma}\right\}\right]$$

$$\frac{w_1^{1-\gamma}}{1-\gamma}(1+\beta) = \frac{w_2^{1-\gamma}}{1-\gamma}\mathbb{E}_{z_2}\left[e^{z_2(1-\gamma)} + \beta \max\left\{e^{z_2^*(1-\gamma)}, e^{z_2(1-\gamma)}\right\}\right]$$

$$\left(\frac{w_1}{w_2}\right)^{1-\gamma}(1+\beta) = \mathbb{E}_{z_2}\left[e^{z_2(1-\gamma)}\right] + \beta\mathbb{E}_{z_2}\left[\max\left\{e^{z_2^*(1-\gamma)}, e^{z_2(1-\gamma)}\right\}\right]$$

$$= e^{-\frac{\sigma_2^2}{2}(1-\gamma)+\frac{\sigma_2^2}{2}(1-\gamma)^2} + \beta\mathbb{E}\left[e^{z_2(1-\gamma)} | e^{z_2(1-\gamma)} > e^{z_2^*(1-\gamma)}\right]\Pr\left[e^{z_2(1-\gamma)} > e^{z_2^*(1-\gamma)}\right]$$

$$+\beta e^{z_2^*(1-\gamma)}\Pr\left[e^{z_2(1-\gamma)} \leq e^{z_2^*(1-\gamma)}\right]$$

$$= e^{-\frac{\sigma_2^2}{2}(1-\gamma)+\frac{\sigma_2^2}{2}(1-\gamma)^2} + \beta e^{-\frac{\sigma_2^2}{2}(1-\gamma)+\frac{\sigma_2^2}{2}(1-\gamma)^2}\Phi\left[\left(\frac{1}{2}-\gamma\right)\sigma_2 - \frac{z_2^*}{\sigma_2}\right] + \beta e^{z_2^*(1-\gamma)}\Phi\left(\frac{z_2^*}{\sigma_2} + \frac{\sigma_2}{2}\right)$$

$$\vdots$$

$$\ln\left(\frac{w_1}{w_2}\right) = \underbrace{-\gamma\frac{\sigma_2^2}{2}}_{\text{Risk Premium}} + \cdots$$

$$\cdots + \underbrace{\frac{1}{1-\gamma}\ln\left\{\frac{1}{1+\beta}\left[1+\beta\Phi\left[\left(\frac{1}{2}-\gamma\right)\sigma_2-\frac{z_2^*}{\sigma_2}\right]+\beta e^{(1-\gamma)\left(z_2^*+\gamma\frac{\sigma_2^2}{2}\right)}\Phi\left(\frac{z_2^*}{\sigma_2}+\frac{\sigma_2}{2}\right)\right]\right\}}_{\text{Option Value of Mobility}}$$

The first term on the right hand side is the risk premium effect while the second term on the right hand side is the option value of mobility. This final equation can be expressed as the following fixed point problem by subtracting $z_2^* = \ln(w_1/w_2)$ from both sides: $f(z_2^*) = 0$. Suppose $\gamma > 1$. When $z_2^* \to +\infty$, $f(z_2^*) \to -\infty$. Alternatively, when $z_2^* \to -\infty$, $f(z_2^*) \to +\infty$. Given that $f(z_2^*)$ is continuous, there must exist a threshold $z_2^* \in \mathbb{R}$ such that $f(z_2^*) = 0$.

Table 1.1: Industry Classifications

Code	Industry Description
01	Agriculture, forestry, and fisheries
02	Mining
03	Construction
04	Manufacturing
05	Transportation, communications, and utilities
06	Wholesale and retail trade
07	Finance, insurance, and real estate
08	Business and repair services
09	Personal services
10	Entertainment and recreation services
11	Professional and related services
12	Public administration

Table 1.2: Occupation Classifications

Code	Occupation Description
11	Management
13	Business and financial operations
15	Computer and mathematical
17	Architecture and engineering
19	Life, physical, and social science
21	Community and social services
23	Legal
25	Education, training, and library
27	Arts, design, entertainment, sports, and media
29	Healthcare practitioners and technical
31	Healthcare support
33	Protective services
35	Food preparation and serving related
37	Building and grounds cleaning and maintenance
39	Personal care and services
41	Sales and related
43	Office and administrative support
45	Farming, forestry, and fishing
47	Construction and extraction
49	Installation, repair, and maintenance
51	Production
53	Transportation and material moving

Table 1.3: Summary Statistics

Year	Observations	Age	Years of Education	Union Members	Hours Worked	Real Wage*
1967	1,092	37.3	12.0	0.349	2,312	$11.07
1968	1,125	37.5	12.0	0.358	2,286	$11.37
1969	1,184	36.9	12.2	0.354	2,260	$11.67
1970	1,231	36.7	12.2	0.348	2,210	$11.84
1971	1,307	36.1	12.3	0.334	2,226	$11.61
1972	1,374	35.2	12.4	0.322	2,244	$11.93
1973	1,442	34.9	12.5	0.316	2,224	$12.12
1974	1,501	34.6	12.6	0.322	2,189	$11.70
1975	1,472	34.5	12.7	0.314	2,175	$11.37
1976	1,484	34.4	12.7	0.309	2,210	$11.86
1980	1,618	34.5	12.9	0.297	2,171	$11.59
1981	1,596	34.6	13.0	0.274	2,142	$11.58
1982	1,565	34.7	13.2	0.267	2,132	$11.40
1983	1,609	34.6	13.2	0.241	2,142	$11.46
1984	1,623	34.8	13.4	0.243	2,221	$11.44
1985	1,657	34.9	13.5	0.225	2,207	$11.60
1986	1,665	35.3	13.5	0.234	2,219	$11.96
1987	1,681	35.4	13.5	0.220	2,233	$12.02
1988	1,711	35.7	13.6	0.208	2,247	$11.93
1989	1,724	36.0	13.6	0.211	2,272	$11.81
1990	1,715	36.3	13.6	0.209	2,269	$11.89
1991	1,709	36.5	13.6	0.211	2,236	$11.96
1992	1,734	36.8	13.6	0.201	2,223	$12.29
1993	1,578	37.5	13.7	0.215	2,203	$13.02
1994	1,435	37.2	13.6	0.194	2,235	$12.31
1995	1,193	37.9	13.7	0.200	2,235	$12.45
1996	1,430	38.1	13.7	0.202	2,227	$12.78
1998	1,318	38.7	13.7	0.187	2,235	$13.69
2000	1,348	39.0	13.6	0.197	2,217	$14.25
2002	1,999	38.8	13.7	0.190	2,270	$12.70
2004	1,986	38.3	13.8	0.178	2,285	$12.46
2006	2,013	38.2	13.8	0.171	2,251	$12.59
2008	1,905	38.1	13.8	0.160	2,176	$13.29

*The agerage real wage is expressed in 1982 dollars.

Table 1.4: Distribution of Observations across Industries

Year	Agriculture, forestry, and fisheries	Business and repair services	Const-ruction	Finance, insurance, and real estate	Mining	Manufact-uring	Professional and related services	Public admin-istration	Other services	Transport-ation, comm-unications, and utilities	Wholesale and retail trade
1967	0.024	0.038	0.087	0.036	0.011	0.381	0.114	0.078	0.009	0.096	0.125
1968	0.022	0.033	0.087	0.035	0.008	0.396	0.118	0.075	0.011	0.092	0.122
1969	0.016	0.035	0.085	0.043	0.014	0.386	0.120	0.079	0.011	0.090	0.122
1970	0.015	0.028	0.089	0.040	0.012	0.367	0.119	0.087	0.011	0.092	0.140
1971	0.015	0.029	0.090	0.050	0.011	0.351	0.120	0.091	0.015	0.094	0.135
1972	0.023	0.028	0.089	0.047	0.014	0.352	0.116	0.086	0.017	0.088	0.142
1973	0.020	0.030	0.091	0.048	0.011	0.343	0.112	0.092	0.018	0.096	0.140
1974	0.025	0.023	0.087	0.043	0.016	0.341	0.116	0.099	0.021	0.093	0.136
1975	0.020	0.028	0.079	0.047	0.016	0.334	0.115	0.106	0.023	0.094	0.139
1976	0.021	0.028	0.078	0.042	0.016	0.337	0.114	0.102	0.022	0.096	0.145
1980	0.017	0.040	0.085	0.042	0.015	0.331	0.109	0.095	0.015	0.110	0.142
1981	0.016	0.039	0.078	0.035	0.019	0.335	0.111	0.088	0.019	0.114	0.145
1982	0.024	0.040	0.076	0.037	0.012	0.300	0.117	0.104	0.013	0.126	0.150
1983	0.019	0.040	0.084	0.036	0.014	0.306	0.112	0.095	0.017	0.127	0.149
1984	0.021	0.044	0.074	0.033	0.015	0.305	0.112	0.099	0.019	0.126	0.153
1985	0.019	0.048	0.091	0.034	0.012	0.291	0.116	0.100	0.018	0.116	0.155
1986	0.015	0.050	0.082	0.038	0.011	0.292	0.113	0.100	0.020	0.118	0.161
1987	0.020	0.046	0.079	0.040	0.011	0.296	0.121	0.096	0.020	0.105	0.165
1988	0.017	0.051	0.075	0.043	0.009	0.285	0.122	0.098	0.023	0.119	0.157
1989	0.016	0.055	0.084	0.046	0.009	0.289	0.121	0.096	0.020	0.110	0.155
1990	0.019	0.055	0.087	0.038	0.008	0.290	0.128	0.093	0.014	0.107	0.159
1991	0.016	0.054	0.094	0.040	0.011	0.281	0.126	0.096	0.019	0.111	0.153
1992	0.017	0.061	0.085	0.038	0.010	0.264	0.145	0.100	0.017	0.103	0.161
1993	0.015	0.064	0.090	0.044	0.006	0.252	0.135	0.100	0.014	0.110	0.170
1994	0.022	0.062	0.086	0.045	0.008	0.253	0.129	0.097	0.016	0.107	0.176
1995	0.018	0.065	0.076	0.045	0.008	0.260	0.138	0.095	0.012	0.118	0.164
1996	0.018	0.077	0.091	0.049	0.008	0.243	0.133	0.101	0.021	0.113	0.146
1998	0.020	0.076	0.093	0.049	0.007	0.266	0.127	0.097	0.021	0.107	0.138
2000	0.019	0.077	0.085	0.043	0.010	0.243	0.134	0.110	0.016	0.103	0.159
2002	0.015	0.051	0.101	0.052	0.011	0.210	0.157	0.083	0.093	0.086	0.143
2004	0.017	0.053	0.106	0.054	0.006	0.209	0.166	0.073	0.088	0.077	0.151
2006	0.015	0.049	0.114	0.055	0.010	0.195	0.174	0.077	0.090	0.077	0.143
2008	0.013	0.047	0.115	0.061	0.012	0.175	0.182	0.083	0.100	0.072	0.139

Table 1.5: Auxiliary Model Estimation Results

VARIABLES	(1) Base Sample log(w)	(2) Non-union log(w)	(3) Unskilled log(w)	(4) Skilled log(w)	(5) 1967-1984 log(w)	(6) 1985-2008 log(w)
age	0.172***	0.113**	0.388***	0.200	0.302***	0.134*
	(0.0475)	(0.0562)	(0.0563)	(0.153)	(0.0631)	(0.0707)
$age^2/10^2$	-0.453**	-0.209	-1.419***	-0.287	-1.103***	-0.221
	(0.202)	(0.240)	(0.242)	(0.625)	(0.270)	(0.298)
$age^3/10^4$	0.536	0.110	2.317***	-0.0846	1.882***	0.000626
	(0.371)	(0.443)	(0.448)	(1.104)	(0.498)	(0.545)
$age^4/10^6$	-0.250	0.0180	-1.416***	0.290	-1.229***	0.166
	(0.249)	(0.298)	(0.302)	(0.716)	(0.335)	(0.364)
years of education	0.0466***	0.0447***	0.0470***	0.000157	0.0565***	0.0416***
	(0.00450)	(0.00533)	(0.00168)	(0.00206)	(0.00132)	(0.00661)
union membership	0.198***		0.230***	0.110***	0.170***	0.215***
	(0.00447)		(0.00550)	(0.0118)	(0.00621)	(0.00648)
Constant	-1.031**	-0.522	-2.653***	-1.586	-2.022***	-1.044*
	(0.402)	(0.472)	(0.477)	(1.381)	(0.536)	(0.615)
Observations	51,024	38,573	25,707	14,201	21,223	29,801
R-squared	0.428	0.452	0.387	0.387	0.445	0.428

Additional regressors include industry-specific tenure, occupation-specific tenure, occupation dummies, industry dummies, year dummies, and industry-year dummies. Standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 1.6: Inter-industry Wage Differentials

	(1) Base Sample	(2) Non-union	(3) Unskilled	(4) Skilled	(5) 1967-1984	(6) 1985-2008
Industry	Ω_j	Ω_j	Ω_j	Ω_j	Ω_j	Ω_j
Agriculture, forestry, and fisheries	-0.1349	-0.1048	-0.1428	-0.0568	-0.1533	-0.1212
	(0.0274)	(0.0302)	(0.0320)	(0.0663)	(0.0417)	(0.0366)
Business and repair services	0.0256	0.0352	-0.0114	0.0761	0.0166	0.0309
	(0.0166)	(0.0187)	(0.0228)	(0.0325)	(0.0265)	(0.0218)
Construction	0.0378	0.0105	0.0344	-0.0029	0.0258	0.0317
	(0.0145)	(0.0175)	(0.0190)	(0.0402)	(0.0223)	(0.0205)
Finance, insurance, and real estate	-0.0303	-0.0285	0.0294	-0.0517	-0.0371	-0.0125
	(0.0188)	(0.0209)	(0.0322)	(0.0281)	(0.0274)	(0.0263)
Manufacturing	0.0542	0.0772	0.0310	0.0668	0.0542	0.0557
	(0.0122)	(0.0150)	(0.0169)	(0.0217)	(0.0173)	(0.0179)
Mining	0.1394	0.1528	0.0713	0.2544	0.1247	0.1214
	(0.0302)	(0.0395)	(0.0416)	(0.0692)	(0.0407)	(0.0466)
Other services	-0.1086	-0.1087	-0.0413	-0.1532	-0.1842	-0.0681
	(0.0213)	(0.0245)	(0.0358)	(0.0416)	(0.0382)	(0.0266)
Professional and related services	-0.0376	-0.0218	-0.0432	-0.0186	-0.0578	-0.0261
	(0.0146)	(0.0173)	(0.0229)	(0.0223)	(0.0210)	(0.0206)
Public Administration	-0.0024	-0.0071	-0.0316	-0.0067	-0.0414	0.0110
	(0.0142)	(0.0175)	(0.0210)	(0.0265)	(0.0204)	(0.0205)
Transportation, communications, and utilities	0.0445	0.0441	0.0128	0.1336	0.0307	0.0572
	(0.0138)	(0.0175)	(0.0191)	(0.0274)	(0.0204)	(0.0197)
Wholesale and retail trade	-0.0968	-0.0996	-0.0535	-0.1342	-0.0576	-0.1159
	(0.0133)	(0.0156)	(0.0180)	(0.0265)	(0.0198)	(0.0189)
Employment Weighted Standard Deviation	0.0610	0.0673	0.0424	0.0765	0.0589	0.0634
Rank Correlation with Base Sample	1.0000	0.9720	0.9091	0.9371	0.9790	0.9860
p-value	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Standard errors in parentheses.

Table 1.7: Shock Volatilities

	(1) Base Sample		(2) Non-union		(3) Unskilled		(4) Skilled		(5) 1967-1984		(6) 1985-2008	
Industry	$\sigma_{\zeta,j}$	$\sigma_{e,j}$	$\sigma_{\zeta,j}$	$\sigma_{e,j}$	$\sigma_{\zeta,j}$	$\sigma_{e,j}$	$\sigma_{\zeta,j}$	$\sigma_{e,j}$	$\sigma_{\zeta,j}$	$\sigma_{e,j}$	$\sigma_{\zeta,j}$	$\sigma_{e,j}$
Agriculture, forestry, and fisheries	0.1803	0.1593	0.1816	0.1601	0.1787	0.1380	0.1621	0.1684	0.1583	0.1552	0.2047	0.1460
	(0.0350)	(0.0160)	(0.0380)	(0.0170)	(0.0370)	(0.0195)	(0.0981)	(0.0396)	(0.0693)	(0.0279)	(0.0491)	(0.0282)
Business and repair services	0.1648	0.1555	0.1709	0.1563	0.1675	0.1415	0.1684	0.1390	0.1603	0.1470	0.1818	0.1548
	(0.0241)	(0.0107)	(0.0248)	(0.0113)	(0.0328)	(0.0156)	(0.0368)	(0.0171)	(0.0466)	(0.0196)	(0.0293)	(0.0140)
Construction	0.1543	0.1462	0.1566	0.1439	0.1501	0.1382	0.1407	0.1715	0.1614	0.1388	0.1680	0.1412
	(0.0140)	(0.0062)	(0.0179)	(0.0084)	(0.0157)	(0.0072)	(0.0655)	(0.0188)	(0.0232)	(0.0114)	(0.0186)	(0.0095)
Finance, insurance, and real estate	0.1821	0.1535	0.1736	0.1576	0.1330	0.1761	0.1457	0.1738	0.1439	0.1512	0.1827	0.1700
	(0.0186)	(0.0094)	(0.0208)	(0.0097)	(0.0737)	(0.0217)	(0.0372)	(0.0120)	(0.0390)	(0.0140)	(0.0305)	(0.0136)
Manufacturing	0.1291	0.1435	0.1350	0.1427	0.1286	0.1393	0.1249	0.1328	0.1175	0.1386	0.1288	0.1558
	(0.0075)	(0.0029)	(0.0094)	(0.0038)	(0.0093)	(0.0037)	(0.0158)	(0.0062)	(0.0117)	(0.0043)	(0.0133)	(0.0047)
Mining	0.1765	0.1328	0.1912	0.1301	0.1690	0.1309	0.1445	0.1728	0.0911	0.1591	0.1531	0.1376
	(0.0340)	(0.0175)	(0.0521)	(0.0318)	(0.0418)	(0.0212)	(0.1171)	(0.0500)	(0.0846)	(0.0245)	(0.0717)	(0.0306)
Other services	0.1733	0.1726	0.1767	0.1744	0.1896	0.1507	0.2071	0.1722	0.1807	0.1556	0.1579	0.1838
	(0.0347)	(0.0140)	(0.0416)	(0.0169)	(0.0587)	(0.0307)	(0.0633)	(0.0295)	(0.0694)	(0.0297)	(0.0542)	(0.0175)
Professional and related services	0.1425	0.1547	0.1504	0.1504	0.1303	0.1439	0.1397	0.1607	0.1488	0.1338	0.1372	0.1709
	(0.0123)	(0.0049)	(0.0137)	(0.0058)	(0.0322)	(0.0112)	(0.0164)	(0.0060)	(0.0185)	(0.0089)	(0.0198)	(0.0066)
Public Administration	0.1231	0.1543	0.1368	0.1516	0.1296	0.1520	0.1115	0.1521	0.1088	0.1489	0.1260	0.1601
	(0.0154)	(0.0051)	(0.0186)	(0.0070)	(0.0211)	(0.0072)	(0.0363)	(0.0096)	(0.0321)	(0.0088)	(0.0231)	(0.0073)
Transportation, communications, and utilities	0.1005	0.1612	0.1078	0.1631	0.0795	0.1573	0.1519	0.1596	0.0885	0.1532	0.1128	0.1719
	(0.0193)	(0.0048)	(0.0277)	(0.0069)	(0.0344)	(0.0061)	(0.0358)	(0.0129)	(0.0391)	(0.0083)	(0.0286)	(0.0072)
Wholesale and retail trade	0.1462	0.1586	0.1515	0.1604	0.1305	0.1535	0.1604	0.1603	0.1287	0.1517	0.1419	0.1729
	(0.0125)	(0.0049)	(0.0133)	(0.0054)	(0.0168)	(0.0060)	(0.0324)	(0.0130)	(0.0242)	(0.0085)	(0.0209)	(0.0072)
Employment Weighted Average	0.139	0.152	0.147	0.152	0.131	0.146	0.141	0.154	0.127	0.144	0.142	0.163
Employment Weighted Standard Deviation	0.021	0.008	0.018	0.009	0.024	0.009	0.019	0.014	0.022	0.007	0.022	0.011
Rank Correlation with Base Sample	1.000	1.000	0.9441	0.9441	0.8671	0.5385	0.6084	0.1678	0.5874	0.5000	0.9301	0.7762
p-value	0.000	0.000	0.0000	0.0000	0.0002	0.0683	0.0338	0.6010	0.0424	0.0951	0.0000	0.0025

Bootstrap Standard Errors in Parentheses.

Table 1.8: Empirical Risk-Return Trade-off

	(1) Base Sample	(2) Non-union	(3) Unskilled	(4) Skilled	(5) 1967-1984	(6) 1985-2008
Slope of weighted OLS regression line for plot of Ω_j versus $\sigma_{\zeta,j}$ [b]	-1.340	-1.723	-0.398	-1.600	-1.272	-0.842
Standard error	0.19	0.23	0.13	0.27	0.17	0.21
t-statistic	-7.01	-7.37	-3.09	-5.87	-7.49	-4.07
p-value	0.00004	0.00002	0.01154	0.00016	0.00002	0.00226

Table 1.9: Calibration Summary

Parameter	Description	Value	Source / Target
γ	Coefficiency of relative risk aversion	2	n/a
β	Discount factor	0.96	n/a
δ	Depreciation rate of capital	0.06	n/a
ρ	Elasticity parameter	0.75	Hsieh et al. (2012)
θ_0	Age earnings profile coefficient	0.209	θ_1 = 0.172 (Auxiliary Model using Actual Data)
θ_1	Age earnings profile coefficient	-0.509	θ_2 = -0.453 (Auxiliary Model using Actual Data)
θ_2	Age earnings profile coefficient	0.515	θ_3 = 0.536 (Auxiliary Model using Actual Data)
θ_3	Age earnings profile coefficient	-0.165	θ_4 = -0.250 (Auxiliary Model using Actual Data)
$\{\sigma_{\zeta,1},\sigma_{\zeta,2}\}$	Annual permanent shock volatilities	{0.080,0.103}	$\{\sigma_{\zeta,1}^b,\sigma_{\zeta,2}^b\}$ = {0.122,0.155} (Auxiliary Model using Actual Data)
$\{\sigma_{e,1},\sigma_{e,2}\}$	Annual transitory shock volatilities	{0.121,0.125}	$\{\sigma_{e,1}^b,\sigma_{e,2}^b\}$ = {0.116,0.121} (Auxiliary Model using Actual Data)
$\chi(18 \leq a < 25)$	Resource cost of switching islands for $18 \leq a < 25$ relative to median income	1.082	Fraction entering workforce in high volatility industry = 0.685 (PSID)
$\chi(25 \leq a < 35)$	Resource cost of switching islands for $25 \leq a < 35$ relative to median income	1.890	Mobility rate of 25-34 year olds relative to 18-24 year olds = 0.769 (PSID)
$\chi(35 \leq a < 45)$	Resource cost of switching islands for $35 \leq a < 45$ relative to median income	0.651	Mobility rate of 35-44 year olds relative to 18-24 year olds = 0.664 (PSID)
$\chi(45 \leq a < 58)$	Resource cost of switching islands for $45 \leq a < 58$ relative to median income	0.501	Mobility rate of 45-58 year olds relative to 18-24 year olds = 0.575 (PSID)
ν	Preference parameter	0.347	Average fraction of time spent working = 0.38 (PSID)
τ	Proportional transaction cost of borrowing in credit markets	0.072	Average debt-to-income ratio = 0.058 (Livshits et al. 2010)

The age earnings profile is defined as the following quartic in age: $f(a) = \theta_1 a + \theta_2 a^2/10^2 + \theta_3 a^3/10^4 + \theta_4 a^4/10^6$. The first four parameters listed in Table 8 are calibrated outside of the model, while the remaining parameters are jointly chosen using simulated method of moments.

Table 1.10: Calibration Results

Statistic	**Economic Model**	**Actual Data**
Age earnings profile coefficient implied by auxiliary model, θ_1	0.170	0.172
Age earnings profile coefficient implied by auxiliary model, θ_2	-0.436	-0.453
Age earnings profile coefficient implied by auxiliary model, θ_3	0.524	0.536
Age earnings profile coefficient implied by auxiliary model, θ_4	-0.263	-0.250
Biennial permanent shock volatilities implied by auxiliary model, $\{\sigma_{\zeta,1}^{b},\sigma_{\zeta,2}^{b}\}$	{0.122,0.155}	{0.122,0.155}
Biennial transitory shock volatilities implied by auxiliary model, $\{\sigma_{e,1}^{b},\sigma_{e,2}^{b}\}$	{0.116,0.121}	{0.116,0.121}
Fraction entering workforce in high volatility industry	0.685	0.685
Inter-industry mobility rate of 25-34 year olds relative to 18-24 year olds	0.794	0.769
Inter-industry mobility rate of 35-44 year olds relative to 18-24 year olds	0.690	0.664
Inter-industry mobility rate of 45-58 year olds relative to 18-24 year olds	0.613	0.575
Average fraction of time spent working	0.380	0.380
Average debt-to-income ratio	0.059	0.058*

The age earnings profile is defined as the following quartic in age: $f(a) = \theta_1 a + \theta_2 a^2/10^2 + \theta_3 a^3/10^4 + \theta_4 a^4/10^6$

*See Livshits et al. (2010).

Table 1.11: Model versus Empirical Risk-Return Trade-off

	(1) Model	(2) Base Sample	(3) Non-union	(4) Unskilled	(5) Skilled	(6) 1967-1984	(7) 1985-2008
Log wage differential implied by auxiliary model, $\Omega_1 - \Omega_2$	0.073	0.080	0.095	0.040	0.096	0.072	0.085
Risk-return trade-off implied by auxiliary model, $(\Omega_1 - \Omega_2) / (\sigma_{\zeta,1}^b - \sigma_{\zeta,2}^b)$	-2.20	-2.44	-3.27	-1.46	-3.95	-2.06	-2.72

Table 1.12: Model Statistics and Sensitivity Analysis

Statistic	Data	Quantitative Model $\gamma = 2$	$\gamma = 3$	$\gamma = 4$	$\gamma = 5$	$\gamma = 6$	$\gamma = 7$	$\gamma = 8$
Risk-return trade-off implied by economic model, $[\ln(w_1) - \ln(w_2)] / (\sigma_{\zeta,1} - \sigma_{\zeta,2})$	n/a	-1.54	-1.41	-1.36	-1.31	-1.24	-1.14	-1.10
Risk-return trade-off implied by auxiliary model, $(\Omega_1 - \Omega_2) / (\sigma_{\zeta,1}^{b} - \sigma_{\zeta,2}^{b})$	-2.44	-2.20	-1.98	-1.91	-1.80	-1.77	-1.62	-1.29
Average inter-island mobility rate	0.067	0.018	0.018	0.019	0.019	0.020	0.019	0.020
Fraction of workers employed on high volatility island	0.52	0.52	0.51	0.51	0.51	0.50	0.50	0.49
Capital-to-output ratio	3.50*	3.24	3.16	3.12	3.32	3.20	3.20	3.20

*See Birkland and Prescott (2007).

Table 1.13: Quantitative Significance of Insurance Channels

Insurance Channel and Order Removed	$\frac{\ln(w_1/w_2)}{(\sigma_1 - \sigma_2)}$	**Percentage of Total Change**
Case 1		
(i) Borrowing	-1.07	23.7%
(ii) Inter-industry Mobility	**0.40**	**75.1%**
(iii) Labor Supply	0.43	1.1%
Case 2		
(i) Borrowing	-1.07	23.7%
(ii) Labor Supply	-0.80	13.8%
(iii) Inter-industry Mobility	**0.43**	**62.5%**
Case 3		
(i) Labor Supply	-1.50	1.8%
(ii) Borrowing	-0.80	35.8%
(iii) Inter-industry Mobility	**0.43**	**62.5%**
Case 4		
(i) Labor Supply	-1.50	1.8%
(ii) Inter-industry Mobility	**0.41**	**97.5%**
(iii) Borrowing	0.43	0.7%
Case 5		
(i) Inter-industry Mobility	**0.06**	**81.3%**
(ii) Labor Supply	0.41	18.0%
(iii) Borrowing	0.43	0.7%
Case 6		
(i) Inter-industry Mobility	**0.06**	**81.3%**
(ii) Borrowing	0.40	17.6%
(iii) Labor Supply	0.43	1.1%

Table 1.14: Expected Wage Profiles

	(1)	(2)	(3)	(4)	(5)	(6)
	Model		PSID: Unskilled		PSID: Skilled	
	Low Volatility	High Volatility	Low Volatility	High Volatility	Low Volatility	High Volatility
VARIABLES	log(w)	log(w)	log(w)	log(w)	log(w)	log(w)
Constant	-2.452	-3.059	-0.539	0.240	-13.98**	-12.18***
			(1.218)	(1.272)	(5.689)	(4.528)
age	0.172	0.231	0.157	0.0498	1.675**	1.409***
			(0.155)	(0.163)	(0.667)	(0.535)
$age^2/10^2$	-0.365	-0.567	-0.212	0.337	-6.497**	-5.091**
			(0.718)	(0.763)	(2.876)	(2.318)
$age^3/10^4$	0.329	0.638	-0.183	-1.381	11.44**	8.421*
			(1.444)	(1.547)	(5.408)	(4.384)
$age^4/10^6$	-0.093	-0.270	0.409	1.358	-7.593**	-5.319*
			(1.062)	(1.146)	(3.748)	(3.053)
union membership			0.288***	0.277***	-0.00690	-0.115***
			(0.00948)	(0.0124)	(0.0264)	(0.0237)
Observations			8,146	7,507	2,359	3,884
R-squared			0.229	0.200	0.324	0.292

Additional regressors include year dummies. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

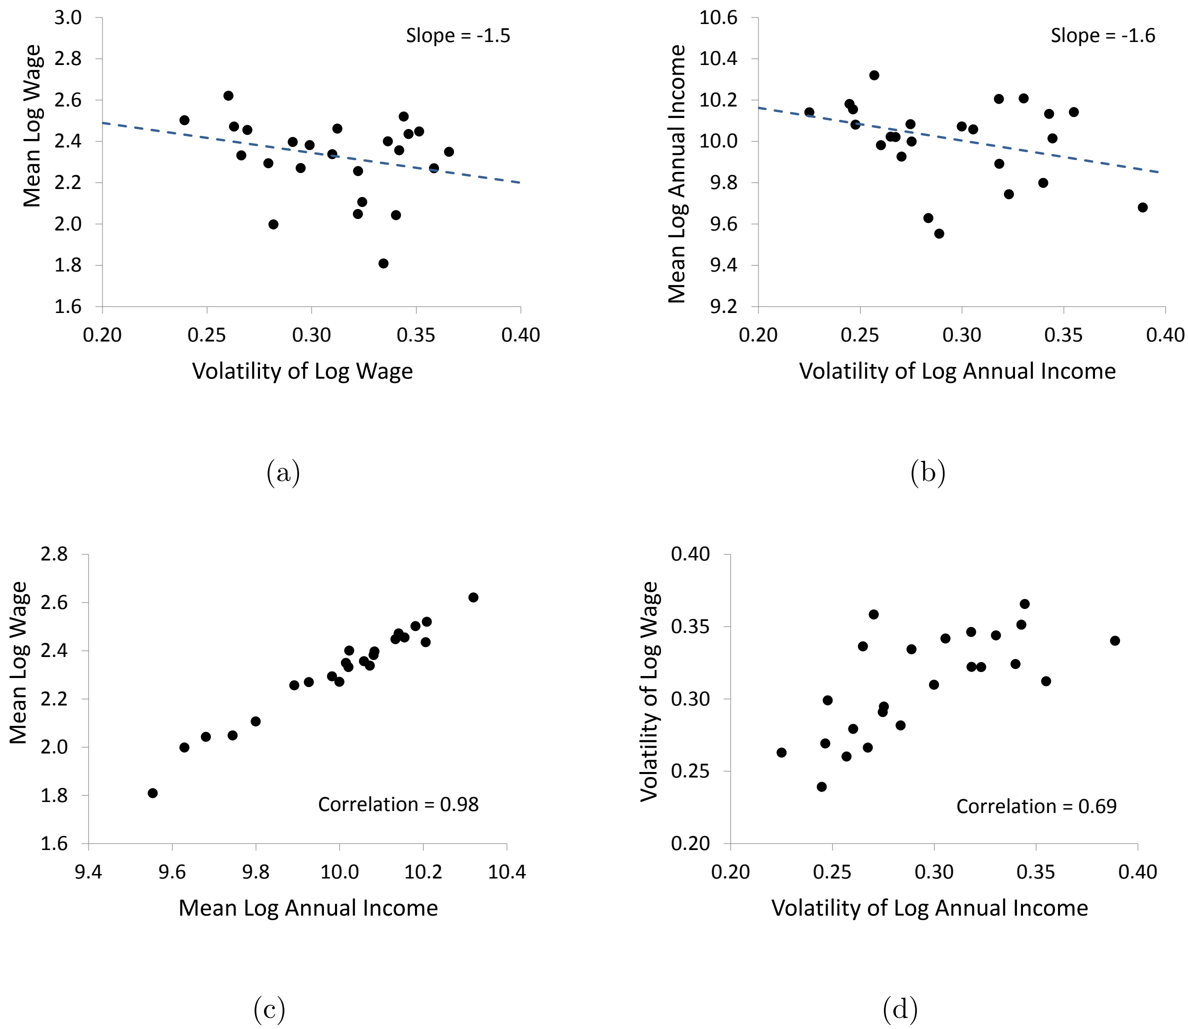

Figure 1.1: Empirical Relationship between Risk and Return across Industries

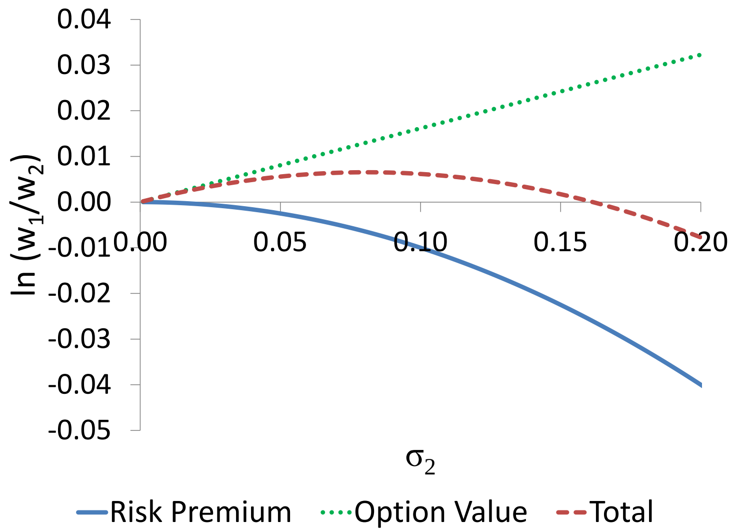

(a) $\gamma = 2$

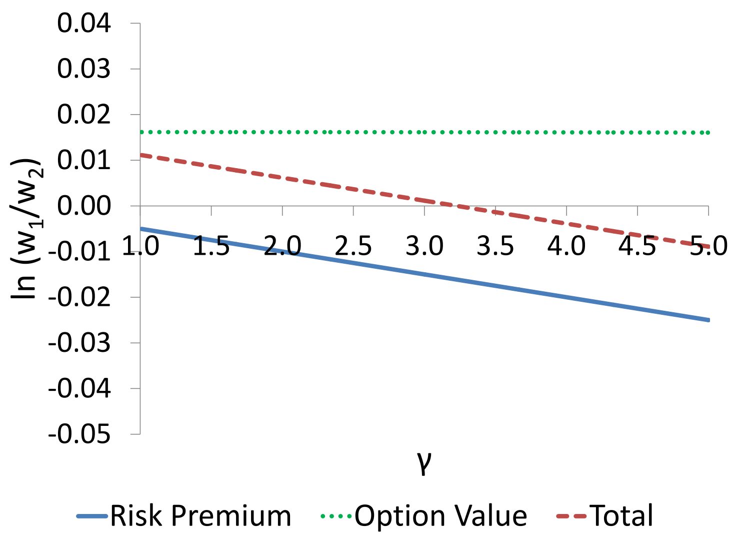

(b) $\sigma_2 = 0.10$

Figure 1.2: Numerical Results for Theoretical Model

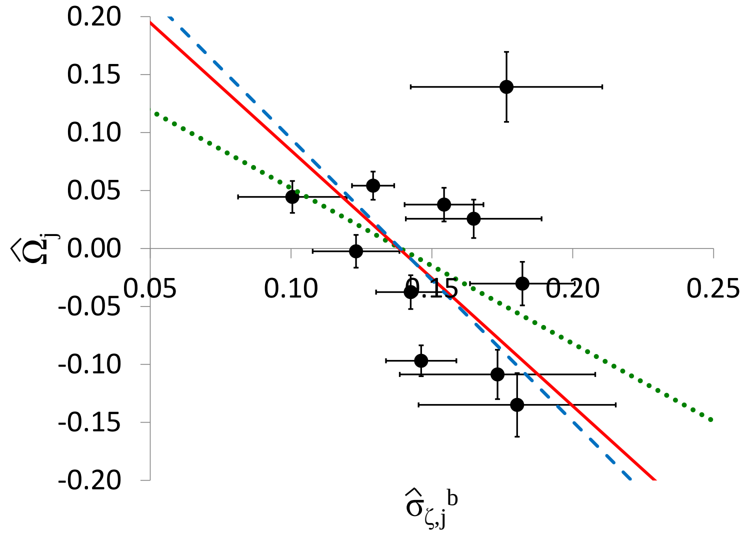

(a) Permanent Shock

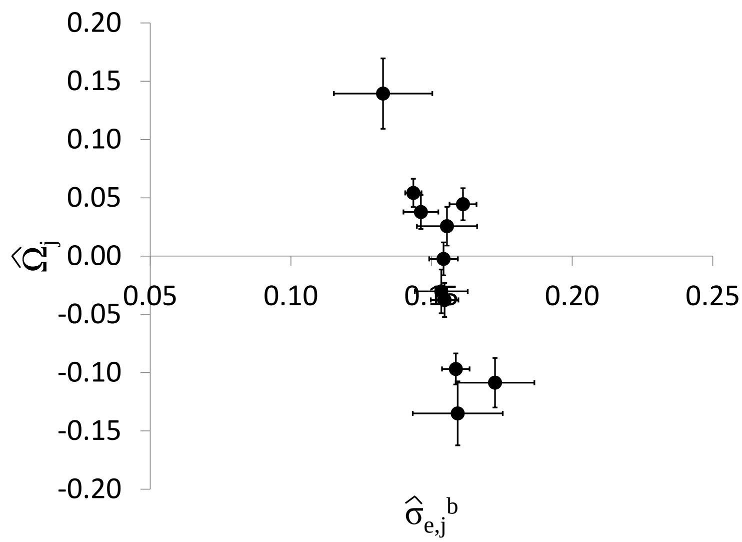

(b) Transitory Shock

Figure 1.3: Trade-off between risk and return implied by auxiliary model.

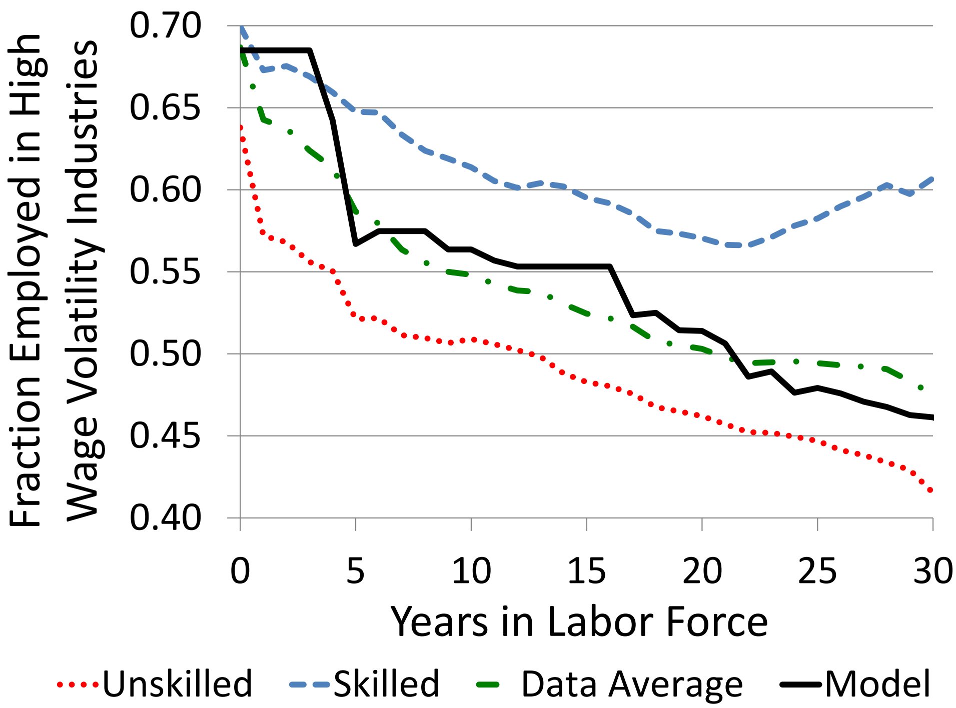

Figure 1.4: Fraction of Individuals Employed in a High Volatility Industry by Age

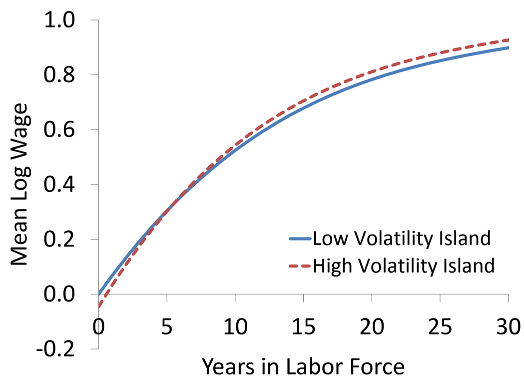

(a) Model

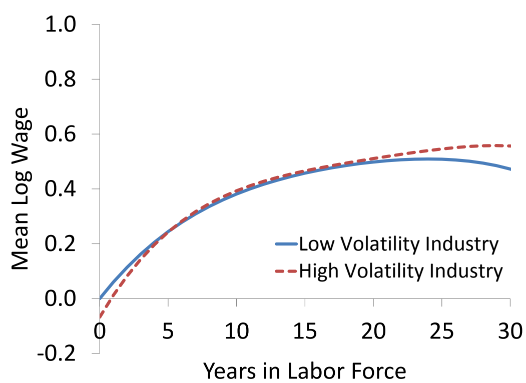

(b) Data: Unskilled Workers

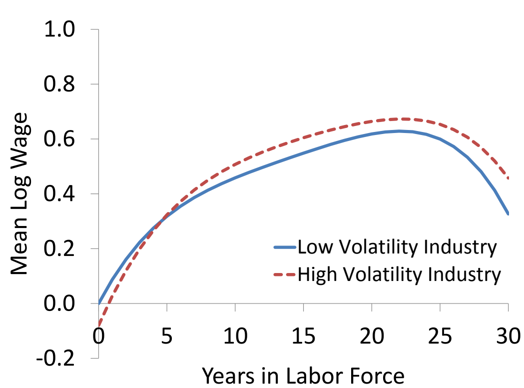

(c) Data: Skilled Workers

Figure 1.5: Expected Wage Profiles

2 Endogenous Credit Booms and Busts: Learning in a Model of Consumer Default

2.1 Introduction

Between 1984 and 2004 the U.S. economy experienced an explosive rise in consumer unsecured debt and bankruptcies. The unsecured debt-to-income ratio for U.S. households almost doubled over this period, increasing from 4.9% in 1984 to 9.1% in 2004. Perhaps more surprisingly, after remaining remarkably stable for nearly twenty-five years, the consumer bankruptcy filing rate more than quadrupled, rising from 1.6 per 1,000 adults in 1984 to 7.0 in 2004.[36] The secular rise in unsecured debt reversed during the recent financial crisis as the unsecured debt-to-income ratio fell by 21.7% between 2008 and 2011 and has yet to recover. This dramatic and persistent reduction in unsecured debt was accompanied by a substantial tightening in lending standards, a reduced willingness to lend by creditors, and lower demand for consumer credit by households.

In this paper we evaluate whether a heterogeneous agent model of optimal consumer default with learning and aggregate uncertainty can help account for these facts. In our framework, aggregate states differ by the mean and variance of the idiosyncratic endowment process. Households form beliefs about the probability of transitioning between aggregate states based on the realized sequence of aggregate shocks. Creditors form expectations about household default probabilities by observing the history of default rates conditional on the loan amount, aggregate state, and household's endowment at the date of origination. Based on these beliefs, households construct their optimal decision rules and creditors determine the appropriate default premium to charge on each loan contract.

We view learning as a natural modeling choice in response to a changing economic environment like that experienced in the U.S. since 1984. As Cogley and Sargent (2008) argue, if

[36] We stop our analysis in 2004 due to implementation of the Bankruptcy Abuse Prevention and Consumer Protection Act of 2005, which significantly increased the costs of declaring bankruptcy, thereby altering household incentives to both borrow and declare bankruptcy.

we assume that agents know the underlying parameters of our model with certainty, as is the standard assumption imposed by rational expectations, we are implicitly assuming that all learning is complete. Although this assumption may be innocuous and serve as a convenient simplification in many cases, there is substantial evidence that the post-1984 period was fundamentally different than that which preceded it.[37] Moreover, in the wake of the recent financial crisis, many observers are again wondering whether or not the underlying parameters of the economy have shifted. It is important to recognize that the first academic papers to document a reduction in aggregate volatility during the Great Moderation did not appear until the late 1990's. Given that it took quite a long time even for academic economists to begin to question whether or not the economic environment had fundamentally changed, we cannot reasonably expect the average agent in our model to know with certainty in 1984 one way or the other. We view learning as a natural response to this critique. It takes time for economic agents to change their beliefs, and we argue that we are better able to understand the behavior of households and creditors during this period by explicitly modeling their dynamic learning process.

While the theoretical model that we develop here is general enough to encompass a variety of learning algorithms, we focus on the case of constant gain learning for several reasons. First, this approach is consistent with actual creditor behavior during this period.[38] Second, it has been shown that constant gain learning is preferred to recursive least squares when agents suspect that the economy may be undergoing a period of structural change.[39] Third, constant gain learning is easier to implement and less computationally intensive than Bayesian learning.

[37] See, for example, Kim and Nelson (1999), McConnell and Perez-Quiros (2000), and Benati and Surico (2009).

[38] See Thomas (2000) for a discussion of credit industry practices during the 1980's and 1990's.

[39] As discussed in Adam, Marcet, and Nicolini (2008), constant gain learning is preferred in a changing economic environment because recent realizations are more informative about the data generating process. Therefore, households prefer to place more weight on recent realizations, as is true with constant gain learning.

As our main theoretical result, we demonstrate that when either households learn about aggregate state transition probabilities or creditors learn about household default probabilities, a rise in the unsecured debt-to-income ratio is the natural response of our economy to a sequence of favorable aggregate shocks like that experienced during the Great Moderation.[40] Using a calibrated version of our model that allows for both types of learning, we demonstrate that learning can explain most of the boom in consumer credit experienced during the Great Moderation, but is unable to produce a sizable increase in the bankruptcy rate. Allowing for a reduction in the costs of declaring bankruptcy over this period enables our model to match the rise in the unsecured debt-to-income ratio and bankruptcy filing rate observed during the Great Moderation.[41]

The intuition for our results is the following: Realizing a string of favorable aggregate shocks leads households to become more optimistic about the future as they discount the probability of transitioning to an unfavorable state. The perception of lower endowment uncertainty reduces households' precautionary savings motive, which leads to an increase in borrowing. Moreover, for any given endowment and loan size, the fraction of households that default is increasing in the variance and decreasing in the mean of the endowment process. Hence, a sequence of favorable aggregate shocks also results in a lower than expected default rate for any given debt contract. In response, creditors revise downward their expectations about default probabilities and reduce the default premium charged on debt contracts, reinforcing the rise in the household debt-to-income ratio.

The impact of learning on the bankruptcy filing rate during a period of reduced aggregate volatility is less clear. On one hand, since the likelihood that a household will default is increasing in their debt burden, an increase in borrowing leads to a higher incidence of

[40] By 'favorable aggregate shocks' we mean those states in which households' idiosyncratic endowment process has a relatively high mean and low variance. We will refer to states with a relatively low mean and high variance for the endowment process as unfavorable.

[41] Gross and Souleles (2002) and Livshits et al. (2010) emphasize the role of a reduction in the costs of declaring bankruptcy, broadly defined, in explaining the observed rise in bankruptcy filings during the Great Moderation.

default.[42] In addition, because households perceive lower future endowment volatility, the punishment from default, which we model as exclusion from credit markets, is believed to be less severe. This increases incentives to declare bankruptcy. On the other hand, a decrease in the default premium makes it easier for households to refinance their debt and avoid default. Lower interest rates also increase the benefit of having access to credit markets, which reduces the attractiveness of bankruptcy. Our quantitative exercises suggest that these changes in households' incentives to declare bankruptcy offset in the aggregate, leaving the bankruptcy filing rate largely unchanged. By allowing for a reduction in the costs of bankruptcy, our model generates a simultaneous rise in the bankruptcy filing rate and unsecured debt-to-income ratio that matches the data.

The fact that our model generates realistic movements in consumer debt and bankruptcies during the Great Moderation leads us to assess whether learning can aid in explaining the dramatic and sustained contraction in consumer debt observed during the financial crisis. To address this question, we subject our model to a sequence of aggregate shocks beginning in 2005 that mimic those realized during the crisis. We find that allowing households to learn about the probability of transitioning between aggregate states amplifies both the increase in the unsecured debt-to-income ratio prior to the crisis as well as the steepness of the decline in response to the crisis.[43] This credit crunch becomes increasingly severe and protracted when creditors also learn about household default probabilities.

Within our framework, we find that learning affects both the supply and demand for credit during the financial crisis. On the demand side, learning about the aggregate state

[42]Here we make an important distinction between a default rate and the incidence of default. A default rate is the fraction of households who find it optimal to default on a specific debt contract at a given date. The incidence of default, on the other hand, is the fraction of households who find it optimal to default across all available debt contracts at a given date. It is the incidence of default in our model that corresponds to the bankruptcy filing rate that we observe in the data.

[43]Although our model predicts an increase in the unsecured debt-to-income ratio between 2004 and the financial crisis, this measure remains relatively flat in the data during this period. We think several factors that we do not model in this paper can account for this discrepancy, among them: the implementation of bankruptcy reform in 2005 which dramatically altered household incentives to declare bankruptcy and undertake unsecured debt and a potential substitution away from unsecured debt and toward debt tied to housing during the housing boom.

during the financial crisis raises households' perceived probability of transitioning to an unfavorable aggregate state in the future. This increases households' precautionary savings motive and reduces their demand for credit. On the supply side, the financial crisis generates a substantial rise in default rates, causing creditors to raise the default premium on loan contracts, which we interpret as a tightening of lending standards. Our model is thus able to generate an endogenous credit cycle that strongly resembles that which the U.S. economy experienced during the recent financial crisis.

Our paper is closely related to several strands of literature. First, we contribute to the work on the rise in household debt and bankruptcies. The empirical literature has proposed many explanations for these findings. Boyes and Faith (1986) and Shepard (1984), for example, argue that changes in the U.S. consumer bankruptcy code made declaring bankruptcy more attractive to potential filers. Buckley and Brinig (1998), Gross and Souleles (2002), and Fay, Hurst, and White (2002), on the other hand, contend that the rise in defaults was primarily a result of a decline in the cost of filing for bankruptcy, either non-pecuniary or pecuniary in nature. Hacker (2006) and Barron, Elliehausen, and Staten (2000) argue that an increase in income volatility led more households into financial trouble; Warren and Tyagi (2003) highlight the role of greater idiosyncratic expense risk; and Barron and Staten (2003) cite credit market innovations that reduced the transaction costs associated with issuing debt.

In a recent paper, Livshits et al. (2010) evaluate the ability of several leading theories to quantitatively account for this experience in the context of an equilibrium model of consumer default. These authors point to a reduction in the transaction cost of issuing consumer credit and a simultaneous decline in the cost of filing for bankruptcy as the most likely explanation. While learning about default probabilities and changes in transaction costs can have similar implications for interest rates and household borrowing, we view learning as a simple method for endogenizing the changes in credit market conditions that exogenously varying transaction costs are intended to capture. Learning about default rates in our model has the additional feature that creditors adjust interest rates differentially across their portfolio of available loan contracts based on differences between the realized and expected default rate for each loan type. This implies that interest rates change more for debt contracts that experienced larger forecast errors. Moreover, we demonstrate that changes in beliefs about the probability of transitioning to different aggregate states is an important driver of aggregate debt statistics.

Most current models that are able to generate a credit crunch appeal to an exogenous change in borrowing constraints or the introduction of a financial wedge. For example, Guerrieri and Lorenzoni (2011) study the effects of an exogenous and unexpected permanent tightening in consumers' borrowing capacity on consumer spending in a heterogeneous agent, incomplete-markets model. Chatterjee and Eyigungor (2011b) introduce an exogenous financial wedge into a model of the housing market to generate a persistent decline in home prices following an unexpected and permanent increase in housing supply. Our work contributes to this rapidly expanding literature by demonstrating how learning can generate similar aggregate behavior in an endogenous fashion.

We also contribute to recent literature that employs statistical learning algorithms in quantitative settings.[44] For example, Eusepi and Preston (2011) introduce constant gain learning into a standard real business cycle model and find that it generates increased volatility in hours worked, thereby bringing the model's predictions closer to the data. Carceles-Poveda and Giannitsarou (2008) and Adam et al. (2008) use learning of this kind to shed new light on a variety of asset pricing puzzles. Our paper contributes to these findings by illustrating that learning is a quantitatively important factor in explaining unsecured debt dynamics since 1984, and it is also the first to consider statistical learning in a heterogeneous agent environment, or in a model of optimal default.

The remainder of this paper proceeds as follows. Section 1.2 introduces our full model which is later used to generate our quantitative results. Section 1.3 presents our analytical

[44]Evans and Honkapohja (2001) provide a comprehensive overview of the literature on statistical learning, its theoretical properties, and potential applications.

results along with a simplified version of the model that provides intuition for our quantitative results. In Section 1.4 we formally calibrate our model and quantify how much of the increase in consumer debt and bankruptcies over the Great Moderation can be explained by learning. In Section 1.5 we explore our model's implications for the recent financial crisis. Finally, Section 1.6 concludes.

2.2 Model

We consider a consumer defaultable debt model in the spirit of Livshits et al. (2010) and Chatterjee, Corbae, Nakajima, and Rios-Rull (2007). Time is discrete and infinite. The economy is populated by a measure one of infinitely lived households that receive a stochastic endowment y_t each period. The process from which this endowment is drawn depends on the realization of an aggregate state variable $s_t \in S = \{s^1, \ldots, s^N\}$, where S is a time-invariant set and s_t evolves according to a Markov process with transition matrix Π. The only asset is a one-period, unsecured and unconditional discount bond that trades at a price set by a pool of risk-neutral, perfectly competitive creditors.

Households choose whether or not to repay their debt each period. A defaulting household enters bankruptcy, which we model after Chapter 7 of the U.S. bankruptcy code.[45] If they choose not to repay their debt obligations, the household is said to default and is relieved of their outstanding debt. Default is punished with an endowment cost, a one-time utility cost, and restricted access to credit markets. In the period of default, the household can neither save nor borrow. In the period following default, the household's credit report is marked

[45] The U.S. bankruptcy code offers consumers two choices when filing for bankruptcy protection: Chapter 7 and Chapter 13. A household that chooses to file under Chapter 7 is relieved of all outstanding debt obligations in exchange for their assets net of any personal exemptions. A household that chooses to file under Chapter 13, on the other hand, agrees to pay back a portion of their outstanding debt obligations over a 3-5 year period in exchange for the ability to keep their assets. In either case, the household is not allowed to refile under the same chapter for a period of 6 years, and a record of their bankruptcy is maintained on their credit report for a period of 10 years. The conditions of default in our model are chosen to match Chapter 7 of the U.S. bankruptcy code, which accounts for approximately 70% of bankruptcy filings over the period under consideration. Moreover, given the choice between Chapters 7 and 13, a household would only choose Chapter 13 if they have assets that they would like to keep but would otherwise lose by filing under Chapter 7. Since there is only one asset in our model, a defaulting household will inevitably have a negative asset position, and therefore will always prefer to file under Chapter 7.

with a bankruptcy flag. Households with a bankruptcy flag are considered to be in a state of bad credit that persists for a random number of periods. While in bad credit standing, a household does not incur any additional costs and may save but cannot borrow.

2.2.1 Timing of Events

In any period t, the timing of events is as follows:

1. Households enter with prior beliefs about the aggregate state transition matrix, and creditors enter with prior beliefs about default probabilities.

2. The aggregate state s_t and idiosyncratic endowments y_t are realized.

3. Given their prior and the realized aggregate state, households form posterior beliefs about the aggregate state transition matrix.

4. Creditors announce a bond price schedule consistent with their beliefs.

5. With probability θ, households who are in bad credit standing have their bankruptcy flag removed and regain full access to credit markets.

6. Given their posterior beliefs and bond prices, households in good credit standing make default, consumption, and borrowing decisions, while households in bad credit standing make consumption and saving decisions.

7. Given their prior and the observed default rates, creditors form posterior beliefs about household default probabilities.

2.2.2 Household's Problem

As has become standard in the literature on statistical learning, we adopt the model of anticipated utility originally developed by Kreps (1998) and first used in applied work by Sargent (1999). In this framework, agents reoptimize at each point in time given their current beliefs. Cogley and Sargent (2008) demonstrate that when agents are not too risk averse, anticipated utility models closely approximate the results generated by models in which agents also learn, but are considered to be fully rational in the Bayesian sense.[46] Moreover, anticipated utility models have the advantage of being more tractable than models that use Bayesian learning rules. To our knowledge, this paper is the first to use anticipated utility in a model of optimal default.[47]

Each period households receive a stochastic endowment y, the log of which evolves according to the following first-order autoregressive process:

$$\log(y) = (1-\rho)\mu_s + \rho \log(y_{-1}) + \varepsilon \tag{2.1}$$

where $\varepsilon \sim N(0, \eta_s^2)$. The unconditional mean of the endowment process and the variance of the idiosyncratic endowment shock depend on the realization of the aggregate state s.

The household's state is composed of the debt with which it enters the period b (where we adopt the convention that $b > 0$ represents a household with positive assets, while $b < 0$ represents a household with negative assets, or positive debt), its idiosyncratic endowment realization in the current period y, and the realization of the aggregate state s, and is denoted by the triplet $(b, y; s)$. A household in good credit standing (G) observes the bond price schedule set by creditors and chooses whether to default (D) or repay their debt obligations (R):

$$V_t^G(b, y; s) \equiv \max_{R,D}\{V_t^R(b, y; s), V_t^D(y; s)\}, \tag{2.2}$$

where $V_t^D(y; s)$ represents the value of defaulting and $V_t^R(b, y; s)$ is the value associated with

[46]More precisely, when agents have constant relative risk aversion preferences with a coefficient of two or less, the predictions of a model in which agents use recursive least squares are nearly identical to those generated by a model in which agents use Bayesian learning.

[47]This paper is also the first, to our knowledge, to allow for aggregate uncertainty in a model of optimal default. We demonstrate that learning in an environment with aggregate uncertainty is the key feature that allows our model to closely match the data.

repaying their debt at date t.[48] Note that since a defaulting household is relieved of their debt obligations, the value of defaulting is independent of b.

If the household repays its debt, it then optimally chooses consumption and its asset position with which it leaves the period. The value of this option is given by:

$$V_t^R(b, y; s) = \max_{b'} \; u(c) + \beta \mathbb{E}_t \left[V_t^G(b', y'; s') | y; s \right] \tag{2.3}$$

subject to

$$c + q_t(b', y; s)b' = y + b$$

where $\mathbb{E}_t$ are household expectations conditional on their beliefs about the aggregate state transition matrix at date t and $q_t(b', y; s)$ is the bond price for a household leaving the period with assets b' with current endowment y in state s. A household that does not default remains in good credit standing and faces the same problem in the following period of whether or not to default and thus receives an expected continuation value of $\mathbb{E}_t \left[V_t^G(b', y'; s') | y; s \right]$.

If the household chooses to default, they are relieved of their outstanding debt obligations in exchange for an endowment cost and a one-time utility cost χ_t. The endowment cost represents a payment to satisfy the "good faith" requirement of the U.S. bankruptcy code for high-income households, and therefore, we assume that it is weakly increasing in the household's endowment.[49] The utility cost is intended to capture potential changes over time in the stigma attached to bankrupts. Households are also prohibited from saving in the period of default. Hence, a defaulting household simply consumes their endowment net of any bankruptcy costs. The value of defaulting in the current period is thus given by:

[48] The anticipated utility structure of the model implies that households reoptimize at each point in time given their current beliefs. For this reason, household decision rules and value functions, which depend on the household's beliefs about the aggregate state transition matrix at date t, are time-dependent and thus are appropriately labeled with time subscripts.

[49] We think this is an important feature of our model since otherwise a household with high income and a large amount of outstanding debt would have an incentive to game the system by filing for bankruptcy. This form of endowment cost is similar to that used by Arellano (2008) in a model of sovereign default.

$$V_t^D(y;s) = u(c) - \chi_t + \beta \mathbb{E}_t \left[V_t^B(0, y'; s') | y; s \right] \tag{2.4}$$

subject to

$$c = \min \{y, \psi \mathbb{E}[y|s]\}$$

and $V_t^B(b, y; s)$ is the value of a household that has a bankruptcy flag on their credit report and so is considered to be in bad credit standing.

Households in bad credit standing are restricted from borrowing.[50] But since the U.S. bankruptcy code does not prohibit asset accumulation after the discharge of debt, we allow households to save. Each period following default, the household has their bankruptcy flag removed and regains full access to credit markets with probability θ, while with probability $1 - \theta$ the bankruptcy flag remains on their credit report. The value of a household in this post-default state is given by:

$$V_t^B(b, y; s) = \max_{b' \geq 0} \ u(c) + \beta \mathbb{E}_t \left[\theta V_t^G(b', y'; s') + (1 - \theta) V_t^B(b', y'; s') | y, s \right] \tag{2.5}$$

subject to

$$c + q_t(b', y; s) b' = y + b.$$

2.2.3 Bond Prices

The bond price schedule is determined in equilibrium by the profit maximizing behavior of a pool of perfectly competitive, risk-neutral creditors that face an exogenously given, risk-free rate r. Creditors in our model face a proportional transaction cost $\tau > 0$ of making loans to households. One should think of τ as representing the cost to a lender of verifying a household's income prior to issuing a loan. The assumptions of risk neutrality and perfect

[50] Musto (2004) argues that creditors view default as an adverse signal about a household's future ability to repay their debt. Consequently, access to credit for households that have a bankruptcy flag on their credit report may be available on prohibitively tough terms or may not be available at all. Musto finds that this effect tends to last until the household's credit report is cleared of their bankruptcy flag, which occurs by law 10 years after the date at which their debt was discharged.

competition imply that creditors must earn zero expected profits on each credit contract they enter into with a household. Furthermore, the ability of creditors to price each loan based on its size, the household's income, and the aggregate state rules out cross-subsidization. As a result, bond prices fully reflect the expected default probability for a loan with these given characteristics. The bond price for a contract where $b' < 0$ is given by:

$$q_t(b', y; s) = \frac{1 - \tilde{D}_t(b', y; s)}{(1 + r)(1 + \tau)} \tag{2.6}$$

where $\tilde{D}_t(b', y; s)$ represents creditor beliefs about the default probability of a household borrowing b' with endowment y in state s at time t.[51] Creditors' expectations at time t incorporate information about realized default rates up to and including time $t - 1$. Note that this object is distinct from the household's decision of whether or not to default, which is described by an indicator function taking the value of 1 if the household defaults and 0 otherwise:

$$D_t(b, y; s) = \begin{cases} 1 & \text{if} \quad V_t^D(y; s) > V_t^R(b, y; s) \\ 0 & \text{otherwise} \end{cases} . \tag{2.7}$$

Since households that save will never find it optimal to default, they carry no default risk. Moreover, we assume there are no transaction costs ($\tau = 0$) associated with accepting deposits since income verification is unnecessary in this case. Thus, the bond price for a contract where $b' > 0$ is equal to $1/(1 + r)$.

2.2.4 Information and Learning

Households and creditors have incomplete information about the underlying model parameters. Each must use the information that they possess to form beliefs about these parameters. Households learn about the aggregate state transition matrix so that they can formulate optimal decision rules. Creditors learn about household default probabilities to ap-

[51] Throughout we will denote objects that are creditor or household beliefs with a tilde, while the true object will be presented without a tilde.

propriately price household debt contracts. Both agents use linear statistical learning rules when forming posterior beliefs given their prior and the realization of relevant economic variables. The details of these dynamic learning algorithms are outlined next.

Aggregate State Transition Probabilities

Households are uncertain about the transition probabilities governing the aggregate state. Let $\tilde{\Pi}_t$ represent household beliefs about the aggregate state transition matrix after the aggregate state is realized at time t. These beliefs may differ from the true aggregate state transition matrix Π. Given an initial prior $\tilde{\Pi}_0$, households learn over time about the aggregate state transition matrix by observing the realized sequence of aggregate states and using a linear updating rule to form their posterior beliefs.

Suppose that the observed transition at date t is from aggregate state s^i to s^j, and let $\tilde{\Pi}_t^k$ denote the k^{th} row of $\tilde{\Pi}_t$. If $\tilde{\Pi}_{t-1}$ is their prior belief about the aggregate state transition matrix at date t, then their posterior beliefs given the realized transition are:

$$\tilde{\Pi}_t^k = \begin{cases} \gamma_a 1^j + (1-\gamma_a)\tilde{\Pi}_{t-1}^k & \text{if } k = i \\ \tilde{\Pi}_{t-1}^k & \text{otherwise} \end{cases} \tag{2.8}$$

where 1^j is a row vector with a 1 as the j^{th} element and 0's elsewhere and γ_a is the gain parameter that governs the weight that households place on new information about the aggregate state transition matrix relative to their prior when forming their posterior beliefs. This updating rule implies that, in response to a transition from s^i to s^j, agents increase their beliefs about the probability of transitioning from s^i to s^j relative to transitioning from s^i to any other state. Since households receive no new information about transitions from states s^k for $k \neq i$, the corresponding rows of the transition matrix are not updated.

Default Probabilities

Creditors observe the aggregate state and the endowment of any household with whom they make a debt contract but do not observe the household's endowment in the following period nor do they know the parameters of their endowment process. While creditors can condition their loan contracts on all of the relevant state variables for the household, these assumptions imply that they are unable to compute the actual default probability for any given loan contract.

Creditors update their beliefs about default probabilities using the new information they obtain by observing actual default rates in the economy each period. When observed default rates differ from their expectations, creditors use their forecast errors to update their beliefs. Let $DR_t(b', y)$ represent the observed default rate at date t for households that borrowed an amount b' with endowment y at date $t-1$, and recall that $\tilde{D}_t(b', y; s)$ represents creditor beliefs about the default probability of a household that borrows b' with endowment y in state s at time t. Suppose that the aggregate state at date $t-1$ was s^i. For all b' and y, creditors' beliefs at date $t+1$ are then:

$$\tilde{D}_{t+1}(b', y; s^k) = \begin{cases} \gamma_d DR_t(b', y) + (1-\gamma_d)\tilde{D}_t(b', y; s^k) & \text{if } k = i \\ \tilde{D}_t(b', y; s^k) & \text{otherwise} \end{cases} . \tag{2.9}$$

The gain parameter γ_d governs the weight that creditors place on new information relative to their prior when forming their posterior beliefs. Given that creditors must announce a bond price schedule prior to household default decisions, the realized default rate at date t is the most recent default information available to creditors when they set bond prices at date $t+1$.

The fact that creditors learn about an endogenous object gives our model a self-referential property that operates in the following way: When a favorable (unfavorable) aggregate shock occurs, default rates are below (above) creditor expectations. Given their updating rule, expectations about default probabilities are revised downward (upward). In the following period, lower (higher) interest rates make it easier (harder) for a household to roll over its debt, thus leading to even lower (higher) default rates than expected. This mechanism imparts momentum into creditors' beliefs, amplifying the model's response to a sequence of favorable aggregate shocks.

2.2.5 Equilibrium

Definition *An equilibrium for this economy is sequences of household decision rules and beliefs* $b'_t(b, y; s)$, $D_t(b, y; s)$, *and* $\tilde{\Pi}_t$, *and creditor beliefs* $\tilde{D}_t(b', y; s)$, *such that, given initial beliefs for households and creditors* $\tilde{\Pi}_0$ *and* $\tilde{D}_0(b', y; s)$, *an initial distribution of households over bonds, endowments, and credit statuses* Φ_0, *learning rules, and sequences of bond prices* $q_t(b', y; s)$, *aggregate states* s_t, *and endowment shocks* y_t, *the decision rules solve each household's problem and bond prices maximize creditors' profits at every date* t.

2.3 Analytical Results

In this section we investigate the theoretical implications of learning by households about aggregate state transition probabilities and learning by creditors about household default probabilities. In order to both maintain tractability and to isolate the effects of each type of learning, we first consider the case in which only households learn and then consider the case in which only creditors learn.[52]

2.3.1 Household Learning

To understand how learning by households about aggregate state transition probabilities affects their consumption, savings, and default behavior, consider the case in which bond prices are held constant and $S \equiv \{c, e\}$, where 'c' represents a contraction and 'e' represents an expansion.[53] In addition, let the aggregate state transition matrix be given by:

[52]We analyze the effects of simultaneous learning by households and creditors later using a calibrated version of our model.

[53]Bond prices will remain constant in our economy if we set $\gamma_d = 0$, implying that creditors do not learn about household default probabilities.

$$\Pi = \begin{bmatrix} \pi_{ee} & 1-\pi_{ee} \\ 1-\pi_{cc} & \pi_{cc} \end{bmatrix},$$

where $\pi_{ij} = \Pr\{s_{t+1} = j | s_t = i\}$, and assume each household's idiosyncratic endowment takes a value $y \in \{y_L, y_H\}$, $y_L < y_H$, where

$$P^i = \begin{bmatrix} p^i_{HH} & p^i_{HL} \\ p^i_{LH} & p^i_{LL} \end{bmatrix},$$

and $p^i_{mj} = \Pr\{y' = j | y = m; s_{t+1} = i\}$. We consider the case in which the persistence of y_H and y_L are lower and higher, respectively, in contraction, $p^e_{HH} \geq p^c_{HH}$ and $p^e_{LL} \leq p^c_{LL}$. Finally, let $\tilde{\Pi}_0 = \Pi$ and suppose that $s_t = e$ for all $t \geq 0$.

At date $t = 1$, households update their beliefs about Π according to equation (1.8):

$$\tilde{\Pi}_1 = \begin{bmatrix} \gamma_a + (1-\gamma_a)\pi_{ee} & (1-\gamma_a)(1-\pi_{ee}) \\ 1-\pi_{cc} & \pi_{cc} \end{bmatrix}.$$

Given that $\gamma_a \in (0,1)$, $\gamma_a + (1-\gamma_a)\pi_{ee} = \pi_{ee} + \gamma_a(1-\pi_{ee}) > \pi_{ee}$ and $(1-\gamma_a)(1-\pi_{ee}) < 1-\pi_{ee}$. Therefore, households increase their perceived probability of remaining in state e and decrease their perceived probability of transitioning to state c as a result of realizing state e at date $t = 1$.

At date $t = 2$, households again update their beliefs about Π according to equation (1.8):

$$\tilde{\Pi}_2 = \begin{bmatrix} \gamma_a + (1-\gamma_a)[\gamma_a + (1-\gamma_a)\pi_{ee}] & (1-\gamma_a)^2(1-\pi_{ee}) \\ 1-\pi_{cc} & \pi_{cc} \end{bmatrix}$$

$$= \begin{bmatrix} (1-\gamma_a)^2\pi_{ee} + \gamma_a[1 + (1-\gamma_a)] & (1-\gamma_a)^2(1-\pi_{ee}) \\ 1-\pi_{cc} & \pi_{cc} \end{bmatrix}.$$

Repeating this procedure, it follows that at any date t households' beliefs about Π are given by:

$$\tilde{\Pi}_t = \begin{bmatrix} (1-\gamma_a)^t \pi_{ee} + \gamma_a \sum_{j=0}^{t-1} (1-\gamma_a)^j & (1-\gamma_a)^t (1-\pi_{ee}) \\ 1-\pi_{cc} & \pi_{cc} \end{bmatrix}$$

$$= \begin{bmatrix} (1-\gamma_a)^t \pi_{ee} + 1 - (1-\gamma_a)^t & (1-\gamma_a)^t (1-\pi_{ee}) \\ 1-\pi_{cc} & \pi_{cc} \end{bmatrix}.$$

Note that the rate at which households adjust their beliefs is increasing in the learning gain parameter γ_a.

Thus, given $s_t = e$ for all $t \geq 0$, we have the following result in the limit:

$$\lim_{t\to\infty} \tilde{\Pi}_t = \begin{bmatrix} 1 & 0 \\ 1-\pi_{cc} & \pi_{cc} \end{bmatrix}.$$

In response to a persistent sequence of expansionary aggregate shocks, households completely discount the possibility of transitioning to a contraction.

This result has important implications for consumption, savings, and default behavior. In particular, since expansions are associated with higher average income and lower income volatility, a string of favorable aggregate shocks leads households to become more optimistic about the future. The perception of lower endowment uncertainty reduces households' precautionary savings motive, and leads to an increase in borrowing.

2.3.2 Creditor Learning

In this section we explore the implications of creditors learning about default probabilities during an extended period of reduced aggregate volatility. To preserve analytical tractability, we consider a simplified version of our full model in which only creditors learn and assume that households do not learn about aggregate state transition probabilities. In addition, we assume that a defaulting household is forever restricted from borrowing in credit markets (i.e. $\theta = 0$).

Preliminary Results

Prior to introducing the example, we establish several preliminary results that will be useful for this section and the next. To start, let $\tilde{b}_t^{y,s}$ be the value of debt that makes a household with endowment y in state s indifferent between repaying its debt obligations and defaulting, so that $V_t^D(y;s) = V_t^R\left(\tilde{b}_t^{y,s}, y; s\right)$. Our first result establishes that households with debt greater than $\tilde{b}_t^{y,s}$ find it optimal to default.

Theorem 2.1 *A household with endowment y in state s finds it optimal to default if they have debt obligations $b < \tilde{b}_t^{y,s}$.*

Proof Notice that $V_t^D(y;s)$ is independent of b while $V_t^R(b,y;s)$ is increasing in b. Consider some $b < \tilde{b}_t^{y,s}$. Then $V_t^R(b,y;s) < V_t^R(\tilde{b}_t^{y,s}, y; s) = V_t^D(y;s)$. Hence, it is optimal for the household to default on their debt.

Our next result demonstrates that the default thresholds $\tilde{b}_t^{y,s}$ are decreasing in the bond price schedule. To start, we first establish that the household's problem is a contraction mapping and hence has a unique fixed point.

Lemma 2.2 *Define the operator T_q as follows:*

$$\left(T_q V_t^G\right)(b,y;s) = \max\left\{V_t^D(y;s), \max_{b'}\left\{u(y + b - q_t(b',y;s)b') + \beta\mathbb{E}\left[V_t^G(b',y';s')\right]\right\}\right\}.$$

T_q is a contraction mapping and hence there exists a unique fixed point $V_{t,q}^*(b,y;s)$.

Proof See Appendix A for proof.

Theorem 2.3 *The default thresholds $\tilde{b}^{y,s}$ decrease in response to an increase in the bond price schedule between periods t and $t+1$ (i.e. $q_{t+1}(b,y;s) \geq q_t(b,y,s)$ for all $\{b,y,s\}$).*

Proof See Appendix A for proof.

The intuition for this result is that as bond prices rise, interest rates fall by definition, and hence any household must be at least as well off.[54] Since defaulting households are restricted from borrowing forever in this example, the value of defaulting is independent of changes in the bond price schedule. Because the value of repaying debt is increasing in the bond price schedule and the value of default is independent of this object, for any given y and s a household must take on more debt to be indifferent between repaying their debt and defaulting. Therefore, the default thresholds must fall as a result of an increase in bond prices.

Example

In this section we provide a simple example to demonstrate how creditor learning in the presence of a sequence of favorable aggregate shocks causes interest rates to fall. As in the example with only household learning, consider the case where $S \equiv \{c, e\}$ and let Π and P^i be defined as in Section 1.3.1.

We will consider a sequence of shocks in which $s_t = e$ for $t \geq 0$. Suppose that creditors begin with beliefs that are consistent with the transition matrices of the true data generating process, Π and P^i. Let $\bar{b}_t \equiv \max_{y,s}\left\{\tilde{b}_t^{y,s}\right\}$ and $\underline{b}_t \equiv \min_{y,s}\left\{\tilde{b}_t^{y,s}\right\}$. Then from the perspective of a creditor, lending a household $b' \in [\bar{b}_t, 0)$ is risk free since the household will repay their debt in the following period with probability one. Hence, the corresponding bond prices are given by

$$q_0\left(b', y; e | b' \in [\bar{b}_0, 0)\right) = \frac{1}{(1+r)(1+\tau)}.$$

On the other hand, no creditor would ever lend a household $b' < \underline{b}_t$ since the household will default with probability one in the following period. Hence, the corresponding bond prices are given by

$$q_0\left(b', y; e | b' < \underline{b}_0\right) = 0.$$

[54] Recall that the interest rate paid on household savings is always equal to the risk-free rate since a household with $b \geq 0$ will never find it optimal to default.

Now consider a debt contract with a non-trivial default probability. In particular, suppose that $\tilde{b}_0^{L,c} = \max_{y,s}\left\{\tilde{b}_0^{y,s}\right\}$[55] and let $\hat{b}_t \equiv \max\{\tilde{b}_t^{H,e}, \tilde{b}_t^{L,e}, \tilde{b}_t^{H,c}\}$. Then lending a household $b' \in \left[\hat{b}_t, \tilde{b}_t^{L,c}\right)$ is risky since if the economy transitions to state c and the household receives endowment y_L in the following period, they will default. At date 0 creditors' expected probability of default is

$$\tilde{D}_0\left(b', y; e|b' \in [\hat{b}_0, \tilde{b}_0^{L,c})\right) = (1 - \pi_{ee})p_{yL}^c,$$

and the corresponding bond price is given by

$$q_0\left(b', y; e|b' \in [\hat{b}_0, \tilde{b}_0^{L,c})\right) = \frac{1 - (1 - \pi_{ee})p_{yL}^c}{(1 + r)(1 + \tau)}.$$

We now consider how these beliefs evolve when creditors learn. Conditional on $s_0 = e$, the household honored its debt obligation regardless of whether they receive y_L or y_H at $t = 0$ since $b' > \max\{\tilde{b}_0^{H,e}, \tilde{b}_0^{L,e}\}$. Hence, the realized default rate $DR_0(b', y|b' \in [\hat{b}_0, \tilde{b}_0^{L,c})) = 0$. Creditors observe this and update their beliefs as follows:

$$\tilde{D}_1\left(b', y; e|b' \in [\hat{b}_0, \tilde{b}_0^{L,c})\right) = (1 - \gamma_d)(1 - \pi_{ee})p_{yL}^c.$$

The bond price at date 1 is then given by

$$q_1\left(b', y; e|b' \in [\hat{b}_0, \tilde{b}_0^{L,c})\right) = \frac{1 - (1 - \gamma_d)(1 - \pi_{ee})p_{yL}^c}{(1 + r)(1 + \tau)}.$$

Given that $\gamma_d \in (0, 1)$,

$$\tilde{D}_1\left(b', y; e|b' \in [\hat{b}_0, \tilde{b}_0^{L,c})\right) < \tilde{D}_0\left(b', y; e|b' \in [\hat{b}_0, \tilde{b}_0^{L,c})\right)$$

[55] While this condition is not necessarily true in general, it is satisfied in our calibrated model.

and

$$q_1\left(b', y; e | b' \in [\hat{b}_0, \tilde{b}_0^{L,c})\right) > q_0\left(b', y; e | b' \in [\hat{b}_0, \tilde{b}_0^{L,c})\right).$$

Recall Theorem 1.3, which states that $\tilde{b}_t^{y,s}$ is decreasing in q. Hence, at date $t \geq 1$, either $b' \in [\hat{b}_t, \tilde{b}_t^{L,c})$, in which case the loan is risky, or $b' \geq \tilde{b}_t^{L,c}$, in which case the loan is risk free. Either way, as long as the economy remains in the expansion state, we have $DR_t(b', y | b' \in [\hat{b}_0, \tilde{b}_0^{L,c})) = 0$. Iterating on equation (1.9), it follows that

$$\tilde{D}_t\left(b', y; e | b' \in [\hat{b}_0, \tilde{b}_0^{L,c})\right) = (1-\gamma_d)^t (1-\pi_{ee}) p_{yL}^c$$

and

$$q_t\left(b', y; e | b' \in [\hat{b}_0, \tilde{b}_0^{L,c})\right) = \frac{1-(1-\gamma_d)^t (1-\pi_{ee}) p_{yL}^c}{(1+r)(1+\tau)}.$$

Thus, given $s_t = e$ for all t, we have the following results in the limit:

$$\lim_{t\to\infty} \tilde{D}_t\left(b', y; e | b' \in [\hat{b}_0, \tilde{b}_0^{L,c})\right) = 0,$$

and

$$\lim_{t\to\infty} q_t\left(b', y; e | b' \in [\hat{b}_0, \tilde{b}_0^{L,c})\right) = \frac{1}{(1+r)(1+\tau)}.$$

Creditors completely discount the probability of transitioning from expansion to contraction in the limit. This leads creditors to reduce the default premium charged on the bond to zero even though the household may default if the economy transitions to contraction and the household receives the low endowment shock.

2.3.3 Theoretical Results

In this section we prove that if creditors learn about default probabilities, then in response to a sequence of favorable aggregate shocks (1) bond prices and household debt will increase and (2) creditors' expectations and bond prices each converge in the limit.

Define the realized default rate at date t for households that borrowed an amount b' with endowment y in the previous period as

$$DR_t(b', y) \equiv \sum_{y' \in Y} \mathbb{I}(b' < \tilde{b}_t^{y',s'}) \Pr[y'|y; s'].$$

where $\mathbb{I}(\cdot)$ is the indicator function taking a value of 1 if the interior argument is true and 0 otherwise. The following theorem states that, given the learning algorithm used by creditors in our model, the realization of aggregate states for which the actual default rate is less (greater) than expected results in higher (lower) bond prices.

Theorem 2.4 *Let* $X_t(b', y) \equiv \{s' \in S : DR_t(b', y) \leq \tilde{D}_t(b', y; s)\}$. *If* $s_t \in X_t(b', y)$, *then* $\tilde{D}_{t+1}(b', y; s) \leq \tilde{D}_t(b', y; s)$ *and* $q_{t+1}(b', y; s) \geq q_t(b', y; s)$. *Otherwise,* $\tilde{D}_{t+1}(b', y; s) \geq \tilde{D}_t(b', y; s)$ *and* $q_{t+1}(b', y; s) \leq q_t(b', y; s)$.

Proof See Appendix A for the proof.

The following corollary presents the first of our two main theoretical results. It states that if the realized default rate is less than expected, household borrowing will increase on both the extensive and intensive margins.

Corollary 2.5 *Suppose* $s_{t-1} = s_t = s_{t+1} \equiv \hat{s}$ *where* $\hat{s} \in X_t(b', y)$ *for all* (b', y). *Let* i *index households. Then*

$$\int q_{t+1} b_{t+1}(i) di \leq \int q_t b_t(i) di,$$

and

$$\int \mathbb{I}(b_{t+1}(i) < 0) di \geq \int \mathbb{I}(b_t(i) < 0) di.$$

Proof Following Theorem 1.4, the bond price schedule increases in response to this sequence of shocks. Thus, the relative cost of borrowing (consumption today) declines. As a result, the income and substitution effects cause an increase in borrowing such that $q_{t+1}(b_{t+1}, y; \hat{s}) b_{t+1}(b, y; \hat{s}) \leq q_t(b_t, y; \hat{s}) b_t(b, y; \hat{s})$ for all b and y and each household. This establishes the former result, from which the latter directly follows.

We now establish what happens in the limit as the economy realizes an infinite sequence of aggregate shocks for which the actual default rate is less than the expected default rate at every date t.

Theorem 2.6 *Suppose* $s_t \in X_t(b', y)$ *for all* $t \geq 0$. *Then:*

1. $\lim_{t\to\infty} \tilde{D}_t(b', y; s) = \tilde{D}_\infty(b', y; s) \in \left[0, \tilde{D}_0(b', y; s)\right]$

2. $\lim_{t\to\infty} q_t(b', y; s) = q_\infty(b', y; s) \in [q_0(b', y; s), 1/((1 + r)(1 + \tau))]$.

Proof See Appendix A for the proof.

This result tells us that if the economy repeatedly experiences favorable aggregate shocks that produce default rates below expectations, then creditors' expectations and bond prices each converge in the limit. Moreover, we know that creditors' expectations are bounded above by their initial prior, while bond prices are bounded below by their initial value.

2.4 Quantitative Results

In this section we discuss our calibration exercise and the experiment that we conduct to determine how much of the simultaneous rise in consumer debt and bankruptcies between 1984 and 2004 can be accounted for by our model. To do so, we first construct a sequence of aggregate shocks through 2004 based on the unemployment rate and NBER recession dates. We then calibrate our model to match the observed unsecured debt-to-income ratio and bankruptcy filing rate in 1983 and simulate our model's response to the observed sequence of aggregate shocks between 1984 and 2004.

2.4.1 Calibration

We assume that each period in our model corresponds to one year. The aggregate state is discretized into four values by classifying the years 1890 through 2004 based on the unemployment rate (low or high) and NBER recession dates (expansion or contraction).[56] Using data from the Panel Study of Income Dynamics (PSID), Storesletten, Telmer, and Yaron (2004b) establish that income dispersion increases substantially during recessions relative to expansions. Conducting a similar analysis to those authors, we find that mean income tends to be high (low) when the unemployment rate is low (high). The set of aggregate states S thus contains four elements: expansion (e) or contraction (c), combined with either high (h) or low (l) mean income. Hence, $S = \{(e,h),(c,h),(c,l),(e,l)\}$. We construct the following transition matrix for the aggregate state by counting the transitions observed between 1890 and 1983 implied by our classification of years:

$$\Pi = \begin{bmatrix} 0.55 & 0.35 & 0.10 & 0.00 \\ 0.31 & 0.35 & 0.27 & 0.07 \\ 0.20 & 0.00 & 0.27 & 0.53 \\ 0.67 & 0.16 & 0.00 & 0.17 \end{bmatrix},$$

where $[\Pi]_{ij}$ represents the probability of transitioning from state i to state j.

We take the persistence and state-dependent standard deviation of the household's income process directly from Storesletten et al. (2004b), implying $\rho = 0.963$, $\eta_e = 0.088$ and $\eta_c = 0.162$. Hence, idiosyncratic income shocks are roughly twice as volatile during contractions than expansions. Our estimates for the state-dependent mean of the income process are $\mu_h = 7.55$ and $\mu_l = 7.51$.[57] Given that these values are in logs, we can conclude that the mean of the income process when the unemployment rate is low is about 4% higher than when the unemployment rate is high. We discretize the endowment process for each of the four aggregate states using the method employed by Tauchen and Hussey (1991).

The remaining parameters are chosen as follows. We set $\beta = 0.95$ and assume households have CRRA preferences with a coefficient of relative risk aversion $\sigma = 2$, as is standard in the

[56] Our unemployment rate series is constructed using data from Romer (1986) for the years 1890 to 1930, Lebergott (1964) for 1931 to 1940, and the Bureau of Labor Statistics from 1941 onward.

[57] See Appendix B for a detailed description of our estimation procedure.

literature. The risk-free interest rate is set to 1.7%, which is equal to the average real return on 1-year U.S. Treasury bills between 1984 and 2004. We set $\theta = 0.2$ to match an average exclusion from credit markets of six years. This implies that households in our model are, on average, able to refile for bankruptcy after six years, which is consistent with the U.S. bankruptcy code.

Finally, we choose values for ψ and τ that allow our model to most closely match the unsecured debt-to-income ratio and consumer bankruptcy filing rate in 1983. To do so we set $\tilde{D}_0(b', y; s)$ and $\tilde{\Pi}_0$ to be consistent with the true data generating process up to 1983 and then simulate the economy without household or creditor learning ($\gamma_a = \gamma_d = 0$) and assuming the utility cost of default is constant (normalized to 0) given the observed sequence of aggregate shocks from 1890 to 1983.[58] We repeatedly perform this exercise for a finite grid of ψ and τ and choose the pair of values that minimize a weighted sum of the square deviations from our targets. Our baseline parameterization is summarized in Table 2.1.

Table 2.1: Calibrated Parameters

Parameter	Value	Source / Target
σ	2.0	standard
β	0.95	standard
r	0.017	real return on 1 yr US T-bills (1983-2004)
θ	0.2	avg exclusion from credit markets of 6 yrs
ρ	0.963	Storesletten et al. (2004b)
η_e	0.088	Storesletten et al. (2004b)
η_c	0.162	Storesletten et al. (2004b)
μ_h	7.55	own estimate using PSID data
μ_l	7.51	own estimate using PSID data
χ_t	0.0	normalization
ψ	0.37	bankruptcy filing rate in 1983 of 0.16%
τ	0.007	debt-to-income ratio in 1983 of 4.9%

[58] This implies that household and creditor beliefs are consistent with the true data generating process in 1983.

2.4.2 Learning and the Rise in Unsecured Debt and Default

To analyze our model's ability to account for the rise in consumer debt and bankruptcy filing rates over this period, we must parameterize the dynamic learning process. Households and creditors are assumed to use constant gain learning and therefore assign a lower weight to past observations to protect themselves against the possibility of structural change. We think this form of learning is appropriate since it closely resembles how creditors actually behaved during this period. It has also been shown that constant gain learning results in more accurate forecasts than recursive least squares when agents are concerned about the potential for structural change and is more tractable than Bayesian learning while producing strikingly similar results.[59]

Throughout the 1980's and 1990's it was common for lenders to use linear models to evaluate the credit-worthiness of potential borrowers and to update the parameters of these models frequently – not less than once every two years – by re-running their regressions using the most up-to-date data on consumer default decisions. Since creditors updated their models of consumer default at least once every two years and discarded past observations to protect themselves against population drift, we set $\gamma_d = 0.5$, which represents a significant departure from rational expectations.[60] Since at this point we have little guidance regarding the choice of γ_a, we consider a range of values that encompass both small and large deviations from rational expectations.

Our experiment consists of the following procedure using an economy of 10 million households:

1. Set $\tilde{D}_0(b', y; s)$ and $\tilde{\Pi}_0$ to be consistent with the empirically observed data generating process up until 1983.

[59] Adam et al. (2008) make a convincing case for the use of constant gain learning rather than recursive least squares in the context of an asset pricing model, and Cogley and Sargent (2008) demonstrate that the results generated under statistical and Bayesian learning are nearly indistinguishable when agents are not too risk averse.

[60] To be precise, $\gamma_d = 0.5$ implies that creditors place almost zero weight on their 1984 prior beliefs at the end of the simulation period.

2. Simulate the economy without learning ($\gamma_a = \gamma_d = 0$) given the observed sequence of aggregate shocks from 1890 to 1983.

3. Then simulate the economy with learning ($\gamma_a \geq 0$, $\gamma_d = 0.5$) given the observed sequence of aggregate shocks from 1984 to 2004.[61]

4. Document the unsecured debt-to-income ratio and the bankruptcy filing rate in the economy in 2004.

The results of this experiment are summarized in Table 2.2. We find that the debt-to-income ratio is monotonically increasing in γ_a and that values between 0.15 and 0.20 match the actual data quite well. Conversely, the elasticity of the bankruptcy filing rate to changes in γ_a is negligible, as even large changes in this parameter have little, if any, impact on the bankruptcy filing rate in 2004. The case without learning refers to our simulation with $\gamma_a = 0$ and $\gamma_d = 0$. This scenario corresponds to the rational expectations equilibrium of our model under the belief that the Great Moderation was a sequence of lucky draws from an unchanged data generating process and provides a useful benchmark for understanding the effects of learning.[62] Note that in this case both the bankruptcy filing rate and debt-to-income ratio in 2004 are essentially unchanged from their respective values in 1983. These results suggest that while learning by households and creditors is able to account for the rise in the unsecured debt-to-income ratio, it is unable to explain the increase in bankruptcies.

In Table 2.3, we explore the sensitivity of our model's predictions to changes in γ_d holding γ_a fixed at a value of 0.20. These results suggest that the bankruptcy filing rate is relatively insensitive to changes in the rate of learning about default rates. Moreover, while γ_d has a measurable impact on the unsecured debt-to-income ratio, the elasticity of the debt-to-income ratio with respect to changes in γ_d is far smaller than that with respect to changes

[61] In this part of the experiment, creditors update their beliefs based on realized default rates, while households update their beliefs about the aggregate state transition matrix based on the realized transitions of the aggregate state.

[62] See Stock and Watson (2003), Sims and Zha (2006) and Gambetti, Pappa, and Canova (2008) for evidence of this "lucky draws" interpretation of the Great Moderation.

Table 2.2: The Impact of Learning

	2004 Bankruptcy Filing Rate	2004 Debt-to-Income Ratio
$\gamma_a = 0.00$	0.0017	0.051
$\gamma_a = 0.05$	0.0016	0.064
$\gamma_a = 0.10$	0.0017	0.078
$\gamma_a = 0.15$	0.0020	0.089
$\gamma_a = 0.20$	0.0018	0.092
$\gamma_a = 0.25$	0.0018	0.093
$\gamma_a = 0, \gamma_d = 0$	0.0016	0.047
Actual Data	0.0070	0.091

in γ_a, as can be seen by comparing the results between Tables 2.2 and 2.3.

Table 2.3: Sensitivity Analysis for $\gamma_a = 0.20$

	2004 Bankruptcy Filing Rate	2004 Debt-to-Income Ratio
$\gamma_d = 0.00$	0.0016	0.086
$\gamma_d = 0.25$	0.0017	0.089
$\gamma_d = 0.50$	0.0018	0.092

In conclusion, our analysis suggests that learning by households about the transition matrix for the aggregate state is an important factor in generating a rise in the unsecured debt-to-income ratio like that observed in the data. Moreover, while learning about default rates by creditors has an effect on the unsecured debt-to-income ratio, this impact is relatively small. Finally, neither type of learning is able to generate a meaningful increase in the bankruptcy filing rate over this period. This finding suggests that the factors that influence households' incentives to declare bankruptcy in the presence of learning offset in the aggregate, generating virtually no change in the bankruptcy filing rate in response to this sequence of favorable shocks. In the next section we reconcile this deficiency by allowing for a reduction in the utility cost of bankruptcy.

2.4.3 Utility Cost of Default

Empirical work by Gross and Souleles (2002) and quantitative analysis by Livshits et al. (2010) has suggested that a reduction in the stigma associated with declaring bankruptcy can account for the rise in the bankruptcy filing rate that occurred after 1984. In our model, a reduction in the stigma of filing for bankruptcy corresponds to a reduction in the utility cost of default, χ_t. In this section, we consider whether allowing for a reduction in the utility cost of default can help our model account for both the rise in the unsecured debt-to-income ratio and bankruptcy filing rate from 1984 to 2004.

We calibrate a reduction in the utility cost of default by matching the observed rise in the bankruptcy filing rate between 1984 and 2004. In particular, we parameterize the time trend of the utility cost in years $t \geq 1984$ as follows:

$$\chi_{t+1} = \chi_t + m(t - 1984).$$

Given our normalization $\chi_t = 0$ for all $t \leq 1983$, we simulate the model to find values for the slope parameter m and gain parameter γ_a that most closely match the unsecured debt-to-income ratio and bankruptcy filing rate observed in 2004.[63] This procedure implies that $m = -0.017$ and $\gamma_a = 0.12$ allow the model to exactly match our targets for the unsecured debt-to-income ratio and the bankruptcy filing rate in 2004.

2.5 Implications for the Financial Crisis

The recent financial crisis was characterized by a significant tightening of credit standards and a drastic and protracted reduction in unsecured household debt, which has yet to recover.[64] As of 2011, the unsecured debt-to-income ratio is more than 21% below its

[63] To reduce our computational burden, we assume that households believe the utility cost in year t will remain unchanged in all future periods.

[64] Although this paper focuses on unsecured debt, we believe learning can also help explain the dramatic expansion and collapse in mortgage debt over the past decade, which played a more central role in the recent

pre-crisis level.

The dramatic reduction in household unsecured debt since 2008 was influenced by both supply and demand factors. According to the Federal Reserve Board's Senior Loan Officer Opinion Survey, the recent crisis was characterized by a historically large and persistent tightening in lending standards in 2008 and 2009, a reduction in creditors' willingness to lend, and a weakened demand for consumer loans. For instance, more than 80% of surveyed creditors tightened standards on consumer loans and more than 60% expressed a decreased willingness to lend at the peak of the crisis.[65] At the same time, a majority of creditors reported weaker demand for consumer loans from 2007 through the end of 2010. Given that our model with learning is able to capture the secular trends in the consumer unsecured debt-to-income ratio and bankruptcy filing rate during the Great Moderation, in this section we analyze whether learning can help us understand the debt dynamics observed during the recent financial crisis.

We study our model's response to a sequence of aggregate shocks like those realized during the recent financial crisis. Starting with the distribution of households over states implied by our model in 2004, we fix the utility cost of default $\chi_t = \chi_{2004}$ for all $t \geq 2005$ and simulate through 2014 using the same approach as is described in the previous section.

The model implied time series for the unsecured debt-to-income ratio are depicted in Figure 2.1 for four different scenarios. Case 1 corresponds to our model with $\gamma_a = 0.12$ and $\gamma_d = 0.5$. In this scenario, the unsecured debt-to-income ratio rises by 8.6% between 2004 and 2007 and then falls by 10.5% between 2007 and 2011. This parameterization includes both the supply and demand factors that reduce aggregate borrowing that we mentioned in the introduction. Households demand less debt because they believe it is more likely the economy stays in a state with relatively low mean income and high idiosyncratic volatility. In addition, bankruptcy rates during the financial crisis exceed creditor expectations, causing

crisis.

[65]Similar patterns are seen for lending standards for only credit card loans, which are closely related to our focus on unsecured debt.

Figure 2.1: Model Implied Unsecured Debt-to-Income Ratio

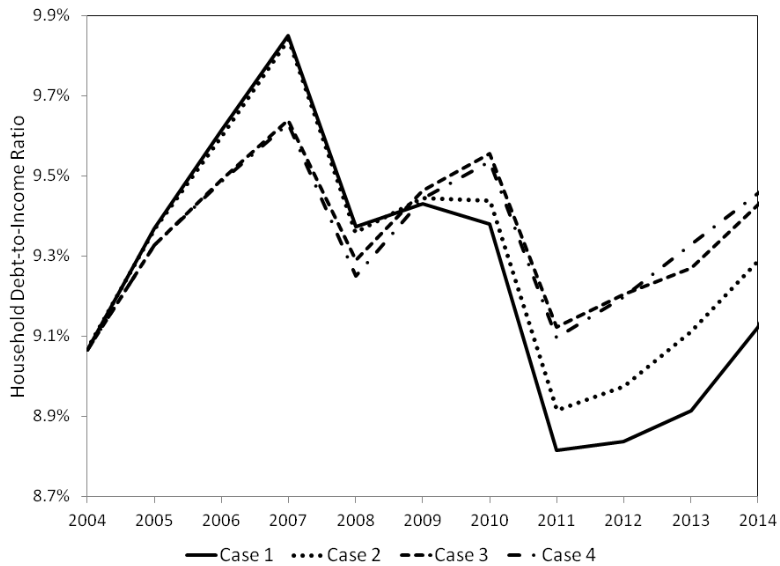

creditors to tighten lending standards by demanding a higher default premium. As a result of these forces, the model with both forms of learning is able to produce a steep and protracted decline in the household debt-to-income ratio like that seen in the data.

In Case 2 we leave $\gamma_a = 0.12$ and set $\gamma_d = 0.0$. This experiment isolates the impact of learning about the aggregate state transition matrix. Whereas the unsecured debt-to-income ratio rises by about the same amount as Case 1, the subsequent drop is less severe and the recovery is faster. This suggests that creditor learning about default probabilities has an asymmetric impact on household borrowing: its influence is minimal during credit expansions but substantial during collapses.[66] Therefore, it appears that learning by households about the aggregate state transition probabilities is most important for generating the initial boom phase and sharp reduction in lending following the crisis, while tighter lending standards amplify the depth and protracted nature of the financial crisis.

Case 3 corresponds to the scenario in which we leave $\gamma_d = 0.5$ and set $\gamma_a = 0.0$, while

[66] A recent paper by Bassett, Chosak, Driscoll, and Zakrajsek (2011) provides empirical evidence for this asymmetry, finding that a tightening shock to lending standards has a substantial impact on the macroeconomy while easing shocks have virtually no impact.

Case 4 considers the model's predictions when $\gamma_a = \gamma_d = 0$. The time series implied by these simulations are nearly indistinguishable, and the reduction in the unsecured debt-to-income ratio is far less severe than the scenarios that consider learning about the aggregate state. Without learning about the aggregate state, the unsecured debt-to-income ratio never falls below its level in 2004 during the financial crisis, which is at odds with the data. This implies that learning about default probabilities by creditors, on its own, is insufficient to generate a credit crunch.

While many existing studies explore the implications of an exogenous tightening in either transaction costs or credit constraints, our model generates a credit crunch through the endogenous response of both households and creditors to a sequence of unfavorable aggregate shocks.[67] In particular, a period of elevated idiosyncratic income volatility causes households to become more pessimistic about the future, increase their precautionary savings, and reduce borrowing. Increased income volatility also results in default rates on loan contracts in excess of lenders' expectations. In response, creditors increase the default premium charged on debt contracts. This credit tightening not only reduces households' incentives to borrow, but it also makes it more difficult for indebted households to roll over their debt, leading to a further increase in default rates. Hence, in our model, unfavorable aggregate shocks lead to an endogenous and persistent contraction in both the supply and demand for credit.

2.6 Conclusion

Is learning by households and creditors important for explaining the dynamics of household unsecured debt since the start of the Great Moderation? Our analysis suggests that it is. We develop a model of optimal default in which households and creditors learn about economic fundamentals in the presence of aggregate uncertainty. We show that in response to a sequence of favorable aggregate shocks, similar to those realized during the Great Moderation, households begin to discount the probability of transitioning to a recession and creditors

[67] See, for example, Guerrieri and Lorenzoni (2011) and Buera and Moll (2012).

reduce their default expectations. As a result, households' precautionary savings motive is reduced and interest rates fall, both of which lead to increased household borrowing.

We also demonstrate that learning can help explain the severe and protracted reduction in unsecured debt that occurred during the recent financial crisis. In response to lower mean income and elevated idiosyncratic income volatility during the crisis, households increase their beliefs that the economy will remain in an unfavorable state and reduce their demand for debt accordingly. The fact that default rates exceed creditor expectations during the crisis leads creditors to tighten lending standards. The result of these forces is that the overall decline in borrowing during the crisis is more severe and the economy takes significantly longer to return to pre-crisis debt levels in a model with learning. These facts are consistent with the debt dynamics observed during the recent financial crisis. Learning therefore appears to be a fruitful avenue for understanding credit booms that often accompany prolonged economic expansions, as well as the severity and persistence of financial crises.

2.7 Appendix A: Proofs

Lemma 1.2: We show that T_q is a contraction mapping by proving that it satisfies Blackwell's sufficient conditions for a contraction. In this lemma we drop the t subscripts on the value functions and household expectations for ease of notation.

- **Monotonicity:** Suppose $W^G(b,y;s) \leq V^G(b,y;s)$ for all $\{b,y,s\}$. Then:

$$\begin{aligned}(T_qV^G)(b,y;s) &= \max\left\{V^D(y;s), \max_{b'}\left\{u(y+b-q(b',y;s)b') + \beta\mathbb{E}\left[V^G(b',y';s')\right]\right\}\right\} \\ &\geq \max\left\{V^D(y;s), \max_{b'}\left\{u(y+b-q(b',y;s)b') + \beta\mathbb{E}\left[W^G(b',y';s')\right]\right\}\right\} \\ &= (T_qW^G)(b,y;s)\end{aligned}$$

- **Discounting:** Let $a \in \mathbb{R}_+$. Then

$$\begin{aligned}(T_q(V^G+a))(b,y;s) - (T_qV^G)(b,y;s) &= \max\{V^D(y;s), \max_{b'}\{u(y+b-q(b',y;s)b') \\ &\quad +\beta\mathbb{E}\left[V^G(b',y';s')+a\right]\}\} - (T_qV^G)(b,y;s) \\ &\leq \max\{V^D(y;s), \max_{b'}\{u(y+b-q(b',y;s)b') \\ &\quad +\beta\mathbb{E}\left[V^G(b',y';s')\right]\}\} + \beta a - (T_qV^G)(b,y;s) \\ &= \beta a\end{aligned}$$

Thus, the operator is a contraction mapping, and there exists a unique fixed point by the contraction mapping theorem. Denote the fixed point associated with the operator T_q as $V_q^*(b,y;s)$.

Theorem 1.3: We will show that if $q_{t+1}(b,y;s) \geq q_t(b,y;s)$ for all $(b,y;s)$ then $V_{q_{t+1}}^*(b,y;s) \geq V_{q_t}^*(b,y;s)$, i.e. that the fixed point under the $T_{q_{t+1}}$ operator is at least as large as the fixed point under the T_{q_t} operator for the entire state space. Since the value of default is invariant to the bond price schedule, this is equivalent to showing that the value of not defaulting is at least as large under q_{t+1} as under q_t for the entire state space.

Let $V^*_{q_t}(b, y; s)$ be the unique fixed point under q_t with associated policy functions $b^*_{q_t}(b, y; s)$ and $D^*_{q_t}(b, y; s)$. Applying the operator under q_{t+1} to this fixed point gives us:

$$(T_{q_{t+1}} V^*_{q_t})(b, y; s) = \max \left\{ V^D(y; s), \max_{b'} \left\{ u(y + b - q_{t+1}(b', y; s)b') + \beta \mathbb{E} \left[V^*_{q_t}(b', y'; s') \right] \right\} \right\}$$

$$\geq \max\{V^D(y; s), u(y + b - q_{t+1}(b^*_{q_t}(b, y; s), y; s) b^*_{q_t}(b, y; s))$$

$$+\beta \mathbb{E}[V^*_{q_t}(b^*_{q_t}(b, y; s), y'; s')]\}$$

$$\geq \max\{V^D(y; s), u(y + b - q_t(b^*_{q_t}(b, y; s), y; s) b^*_{q_t}(b, y; s))$$

$$+\beta \mathbb{E}[V^*_{q_t}(b^*_{q_t}(b, y; s), y'; s')]\}$$

$$= V^*_{q_t}(b, y; s)$$

Successively applying the operator $T_{q_{t+1}}$ gives a non-decreasing sequence of functions, all at least as large as $V^*_{q_t}(b, y; s)$, that converges to some limit – the fixed point under $T_{q_{t+1}}$: $V^*_{q_{t+1}}(b, y; s)$. Thus, $V^*_{q_{t+1}}(b, y; s) \geq V^*_{q_t}(b, y; s)$. Moreover, since $V^D_{q_{t+1}}(y; s) = V^D_{q_t}(y; s)$ for all t, we conclude $V^R_{q_{t+1}}(b, y; s) \geq V^R_{q_t}(b, y; s)$. As a result, $V^R_{q_{t+1}} \left(\tilde{b}^{y,s}_t, y; s \right) \geq V^R_{q_t} \left(\tilde{b}^{y,s}_t, y; s \right) = V^D(y; s)$, and the debt thresholds under q_{t+1} are no greater than the thresholds under q_t.

Theorem 1.4: Suppose at date t, $s_t = s^j \in X_t(b', y)$ is realized. Then learning implies

$$\tilde{D}_{t+1}(b', y; s) = \gamma_d DR_t(b', y) + (1 - \gamma_d) \tilde{D}_t(b', y; s).$$

Since $s^j \in X_t(b', y)$, $DR_t(b', y) \leq \tilde{D}_t(b', y; s)$. Thus,

$$\tilde{D}_{t+1}(b', y; s) = \gamma_d DR_t(b', y) + (1 - \gamma_d) \tilde{D}_t(b', y; s) \leq \gamma_d \tilde{D}_t(b', y; s) + (1 - \gamma_d) \tilde{D}_t(b', y; s).$$

Therefore, $\tilde{D}_{t+1}(b', y; s) \leq \tilde{D}_t(b', y; s)$ and $q_{t+1}(b', y; s) \geq q_t(b', y; s)$.

Now suppose at date t, $s_t = s^j \notin X_t(b', y)$ is realized. Then learning implies

$$\tilde{D}_{t+1}(b', y; s) = \gamma_d DR_t(b', y) + (1 - \gamma_d) \tilde{D}_t(b', y; s).$$

Since $s^j \notin X_t(b', y)$, $DR_t(b', y) \geq \tilde{D}_t(b', y; s)$. Thus,

$$\tilde{D}_{t+1}(b', y; s) = \gamma_d DR_t(b', y) + (1 - \gamma_d)\tilde{D}_t(b', y; s) \geq \gamma_d \tilde{D}_t(b', y; s) + (1 - \gamma_d)\tilde{D}_t(b', y; s).$$

Therefore, $\tilde{D}_{t+1}(b', y; s) \geq \tilde{D}_t(b', y; s)$ and $q_{t+1}(b', y; s) \leq q_t(b', y; s)$.

Theorem 1.6: Suppose initial beliefs $\tilde{D}_0(b', y; s_{-1})$ are such that there exists state k such that $DR_0(b', y) \leq \tilde{D}_0(b', y; s_{-1})$. Therefore, $X_0(b', y) \neq \emptyset$. By Theorem 1.4, if $s_0 \in X_0(b', y)$, $\tilde{D}_1(b', y; s_{-1}) \leq \tilde{D}_0(b', y; s_{-1})$ and $q_1(b', y; s_{-1}) \geq q_0(b', y; s_{-1})$.

Next, we show that if the initial set is non-empty, then it is non-empty for all t. Note that $\tilde{D}_1(b', y; s_{-1}) = \gamma_d DR_0(b', y) + (1 - \gamma_d)\tilde{D}_0(b', y; s_{-1}) \geq DR_0(b', y)$ since $s^k \in X_0(b', y)$. By Theorem 1.3, the debt thresholds are decreasing over time in response to higher bond prices, so that $DR_1(b', y) \leq DR_0(b', y)$. Therefore, $DR_1(b', y) \leq DR_0(b', y) \leq \tilde{D}_1(b', y; s_{-1})$ so that $X_1(b', y) \neq \emptyset$. An analogous argument shows that if $X_t(b', y) \neq \emptyset$ for any arbitrary t, then $X_{t+1}(b', y) \neq \emptyset$, and if we continue to draw from $X_t(b', y)$, $\tilde{D}_{t+1}(b', y; s_{-1}) \leq \tilde{D}_t(b', y; s_{-1})$ and $q_{t+1}(b', y; s_{-1}) \geq q_t(b', y; s_{-1})$ for all t.

Thus, $\tilde{D}_t(b', y; s_{-1})$ is a non-increasing sequence bounded below by 0. By the monotone convergence theorem, $\tilde{D}_t(b', y; s_{-1})$ must converge to some value. Call this limiting value $\tilde{D}_\infty(b', y; s_{-1}) \geq 0$. Similarly, since $q_t(b', y; s_{-1})$ is a non-decreasing sequence bounded above by $1/((1 + r)(1 + \tau))$, it too must converge to some limit $q_\infty(b', y; s_{-1})$.

2.8 Appendix B: Calibrating the Mean of the Endowment Process

To estimate μ_s, we first classify the years 1969 through 1991 as either high or low average income using annual changes in the national unemployment rate. If the unemployment rate increased by more than 1.3 percentage points, then the current year is classified as low average income. If the previous year is classified as low average income, then the current year is also classified as low average income if the decrease in the unemployment rate is less than 2/3 of the increase in the previous year. All other years are classified as high average income.

We follow Storesletten et al. (2004b) and construct a repeated panel from the PSID survey years 1968-1993. We extract income data, the age of the head, and the education level of the head from the PSID main family data files, along with the 1968 interview number and relationship to the head from the PSID individual data files, for all individuals across the PSID survey years 1968-1993. We then restrict our panel to include only those individuals who are members of, or are related to, a family that was included in the 1968 SRC cross-section sample. We define income to be the log of real income in 1968 dollars (deflated by the CPI) at the family level which is the sum of the head and wife's labor income, unemployment compensation, workers compensation, and help from relatives. Income attributed to the head of the household is then defined as total income divided by the number of persons in the family unit.

We select observations on individuals in each survey year into our panel if: (1) they are in the original sample in the previous year and the following year, (2) income is positive in the previous, current, and following year, (3) income growth rate is not less than 1/20 and not larger than 20 between the previous year and the current year or between the current year and the following year, and (4) the individual's age is between 22 and 60 years in the current year.

We then perform the following regression to isolate fixed effects associated with aggregate income, education, and age:

$$y_{it}^h = \theta_0 + \theta_1 \mathbb{D}_h + \theta_2 t + \theta_3^T \mathbf{x}_{it}^h + u_{it}^h,$$

where y_{it}^h is log (per capita) income (at the household level), $\mathbb{D}_h$ is a dummy variable that we set equal to 1 if the year is classified as low unemployment (high mean income), t is a time trend, and $\mathbf{x}_{it}^h$ is a vector composed of age, age squared divided by 100, age cubed divided by 10,000, and years of education completed for individual i of age h at date t.[68]

[68] This regression also identifies the idiosyncratic, uninsurable component of the income process, u_{it}^h, which

Our estimates for the state-dependent mean of the income process are thus given by $\mu_h = \theta_0 + \theta_1$ for the high mean income state and $\mu_l = \theta_0$ for the low mean income state.

Storesletten et al. (2004b) use to estimate ρ, η_e and η_c.

3 Bankruptcy Reform and the Housing Crisis

3.1 Introduction

Prior to 2005, the availability of debt relief through bankruptcy was widely known, the cost of filing was low, and little stigma was attached to those who filed. Bankruptcy was thus an attractive option for homeowners that wished to remain in their homes and could afford their mortgage payments if relieved of other debt obligations, such as credit card bills. This changed in 2005 as the Bankruptcy Abuse Prevention and Consumer Protection Act (BAPCPA) significantly raised the cost of filing and reduced the amount of debt that could be discharged. These changes to the bankruptcy code made it more difficult for struggling homeowners to loosen their budget constraints via bankruptcy, increasing the relative attractiveness of mortgage default. As a result, bankruptcy reform may have contributed to the severity of the housing crisis by inducing some homeowners to default that would have otherwise chosen to declare bankruptcy and keep their homes.

To understand exactly how the BAPCPA affected homeowners' incentives, consider a homeowner with negative home equity who, prior to the BAPCPA, could have had their unsecured debt discharged under Chapter 7 and remained in their home. With the introduction of the BAPCPA, this homeowner's ability to discharge their unsecured debt through Chapter 7 now depends on their income. In particular, if the homeowner has income above their state's median, they cannot file under Chapter 7 and are instead forced to file under Chapter 13 and enter into a repayment plan to which they must commit all of their non-exempt income for five years. Thus, bankruptcy became more costly for such a homeowner. If bankruptcy and mortgage default are substitutes, this higher cost will induce some households to default on their mortgage that would not have done so in the absence of the reform.[69] And because negative home equity is a necessary condition for mortgage default,

[69] We think of bankruptcy and mortgage default as being complements or substitutes just as we would any other goods. That is, they are substitutes (complements) if raising the cost of one increases (decreases) the incidence of the other.

the large decline in house prices that forced many homeowners underwater on their mortgage during the recent housing crisis may have amplified this rise in the mortgage default rate.[70]

Empirical work on the BAPCPA has reinforced this intuition. Li, White, and Zhu (2011), for example, argue that homeowners treated bankruptcy and mortgage default as substitutes in response to the BAPCPA, shifting from bankruptcy to default when the cost of the former rose. Using data on individual mortgages from LPS Analytics, these authors estimate that the BAPCPA increased the probability of default by 24% for prime borrowers and 14% for sub-prime borrowers with mortgages originated in 2004 and 2005. In a complementary study, Morgan, Iverson, and Botsch (2011) document a significant rise in the default rate of subprime mortgages in response to the BAPCPA. Although neither of these studies explicitly consider data from the housing market crash, their conclusions support the view that the BAPCPA may have increased the number of mortgage defaults during the housing crisis, thereby contributing to the severe and protracted decline in home prices.

Although this empirical work suggests that making bankruptcy more costly may have worsened the housing crash, theoretically this conclusion is ambiguous. While increasing the cost of filing for bankruptcy raises the relative attractiveness of mortgage default, rational mortgage lenders will respond by tightening lending standards on those households who are more likely to default to offset the potential for greater losses. Tighter lending standards, in turn, will tend to discourage these households from taking out a mortgage to purchase a home. Importantly, this effect is concentrated on households who bought homes in 2005 and 2006 – exactly those homeowners who are most likely to find themselves underwater as a result of a collapse in house prices – and works to reduce the mortgage default rate during the crisis. Given the presence of these opposing forces, the net impact on mortgage defaults could be either positive or negative depending on the relative magnitude of each effect.

In this paper, we quantify the effects of the BAPCPA on the housing market crash of 2007 using a quantitative-theoretic, equilibrium model of the U.S. housing market. In our

[70] If a household has positive home equity net of transaction costs, then selling their home and repaying their mortgage will always dominate the option to default.

framework, households optimally choose between renting and owning their housing space and can finance the purchase of a home by taking out a mortgage. Households interact in credit markets with rational lenders who provide unsecured credit and mortgage loans at terms that fully reflect the general equilibrium incentives each household has to renege on their obligations. Each period, homeowners optimally choose between remaining in or selling their home, filing for bankruptcy, defaulting on their mortgage, or simultaneously declaring bankruptcy and defaulting on their mortgage. Thus, our model is rich enough to determine whether tighter mortgage lending standards in the years prior to the crisis dominated the increased attractiveness of mortgage default during the crisis.

We calibrate our model to match salient characteristics of the U.S. economy prior to 2005 and then conduct several tests to ensure that our model adequately captures key empirical facts regarding the BAPCPA and the housing market crash. First, we discipline the model to match the empirical findings of Li et al. (2011) by calibrating the bankruptcy cost under reform to produce a rise in the mortgage default rate of 21.6% for new homeowners in response to the BAPCPA.[71] Next, we test whether the model produces a decline in house prices and a rise in mortgage default rates, on the order of that found in the data, in response to a housing crash. Following Chatterjee and Eyigungor (2011b), we model this crash as an unexpected increase in the economy's owner-occupied housing supply and find that our model is able to capture a decline in house prices and rise in mortgage default rates similar to the data. The model also replicates key dynamics in the bankruptcy filing rate, unsecured debt-to-income ratio, and price-rent ratio during the crash. The fact that our model is able to replicate these empirical facts gives us confidence about its implications for the counterfactual exercise that is central to our analysis. As our main quantitative experiment, we construct a counterfactual transition in the U.S. economy in which there is no bankruptcy reform in 2005 but the economy still undergoes a housing crisis in 2007. We then compare the data from this housing crisis to an economy that implemented bankruptcy

[71]Throughout we will refer to homeowners that purchased their home in the previous period as new homeowners.

reform in 2005.

Contrary to existing arguments in the empirical literature, our results suggest that the BAPCPA did not contribute significantly to the severity of the housing crisis. In particular, the mortgage default rate is only 2.7% higher in 2007 while the path of house prices during the crisis is virtually unaltered as a result of bankruptcy reform.

In our model, bankruptcy and mortgage default appear to be treated as substitutes by households in response to the BAPCPA, as its implementation leads to lower bankruptcy and higher mortgage default rates in the aggregate. Indeed, some households find it optimal to default on their mortgage in states where they would have optimally decided to declare bankruptcy in the absence of BAPCPA. Bankruptcy and mortgage default are substitutes for these households, and by making bankruptcy more costly to file, mortgage default becomes more likely. However, there are additional forces at work in our model that generate our results.

Prior to the BAPCPA, a household's home equity in excess of their state's homestead exemption would be paid to creditors in the event of bankruptcy. Non-exempt home equity thus served as collateral for unsecured debt contracts, leading to lower interest rates for homeowners with high home equity. Under the BAPCPA, homeowners with high income relative to their non-exempt home equity are forced into a repayment plan, meaning that their non-exempt home equity no longer serves as collateral for their unsecured debt obligations. For these households, interest rates on unsecured debt are now independent of their homeownership status, which reduces the benefits of homeownership. While homeowners with negative home equity are not directly impacted by this change, the continuation value of remaining in their home falls, inducing the marginal homeowner to default on their mortgage.[72]

Second, since the BAPCPA increased the cost of bankruptcy for high-income households,

[72]In order for default to be optimal in our model, a homeowner must not only have negative home equity (a necessary condition for default), but must also want to move. Homeowners want to move in our model because shocks to their income, assets or house size have made their current mortgage-house combination suboptimal.

making it less likely that these households will declare bankruptcy, unsecured creditors are able to offer lower interest rates and engage in more risky lending.[73] If we assume that a household must repay their unsecured debt if they default on their mortgage (and do not simultaneously declare bankruptcy), bankruptcy and mortgage default are complementary as the former reduces the costs associated with the latter. A rise in risky lending that leads to an increase in bankruptcy filings will thus also tend to cause an increase in the mortgage default rate.

In our model, rational mortgage lenders internalize these changes in homeowners' incentives and respond by tightening lending standards in the years prior to the housing crisis. Higher mortgage interest rates lead new homebuyers to choose smaller homes with lower initial loan-to-income, loan-to-value, and mortgage payment-to-income ratios, on average. These mortgage contracts are inherently less risky, making these new homebuyers far less likely to default on their mortgage during the crisis. This force offsets the increased attractiveness of default and ultimately drives our conclusion that the BAPCPA caused only a slightly higher default rate during the housing crisis and had minimal effect on the severity of the drop in house prices. Accounting for the general equilibrium response of unsecured debt and mortgage interest rates to changes in households' incentives is therefore crucial to adequately assess the impact of the BAPCPA on the housing crisis.

Finally, we use our framework to consider the impact of mortgage cram down, which has been extensively discussed in policy circles and in the academic literature, on the severity of the housing crisis. Under this policy, homeowners with negative home equity are able to treat the portion of their mortgage that exceeds the value of their home as unsecured debt, which can then be discharged through the bankruptcy process. The objective of this policy would be to reduce mortgage default rates and positively impact owner-occupied house prices through a feedback effect from defaults on prices. Our analysis suggests that this policy would have slightly reduced mortgage default rates during the recent housing

[73] Risky lending refers to unsecured debt contracts for which the household's bankruptcy decision in the following period is not trivial.

crisis, while producing a substantially higher bankruptcy rate and having minimal impact on aggregate house prices.

The remainder of the paper is structured as follows. In the next section we briefly discuss papers that are relevant to our current analysis. Section 2.3 then provides a detailed description of our full quantitative framework prior to bankruptcy reform. Next, Section 2.4 describes the BAPCPA and specifies how we model this reform in our quantitative analysis. The following section presents our parameterization and the model fit to the pre-crash period. Section 2.6 details our quantitative results, describing the effect of the BAPCPA on impact and during the housing crash and discussing the intuition for our findings. The impact of the hypothetical cram down program is then assessed in Section 2.7. Finally, Section 2.8 concludes.

3.2 Literature Review

Several recent papers aim to isolate and quantify the effects of the BAPCPA. Much of this work is empirical in nature.[74] As described in the introduction, the most relevant for our work are Li et al. (2011) and Morgan et al. (2011) who document that mortgage default rates increased in response to the BAPCPA. A primary benefit of our quantitative approach relative to their empirical analysis is that we are able to construct the counterfactual experiment that these authors envision.

In response to the recent housing crisis, there is a rapidly growing literature that aims to explain the rise in mortgage defaults and decline in house prices using quantitative models of the U.S. housing market. Corbae and Quintin (2011), for example, assess the importance of mortgage innovations, through the introduction of non-traditional mortgages, and conclude that this channel can explain approximately 40% of the rise in foreclosures during the crisis. Recent work by Jeske, Krueger, and Mitman (2011) evaluates the impact of in-

[74]Two notable exceptions are Chatterjee et al. (2007) and Li and Sarte (2006), who quantitatively analyze the impact of introducing means testing, in the spirit of the BAPCPA, on the consumer bankruptcy decision but do not consider the reform's impact on mortgage default decisions.

terest rate subsidies for the government sponsored enterprises on housing market outcomes. They determine that these subsidies substantially increase mortgage origination and lower aggregate welfare, but have little impact on default rates. Closely related to our study is that of Chatterjee and Eyigungor (2011b) who demonstrate that an unexpected increase in the supply of owner-occupied housing – along with frictions in the mortgage market and foreclosure delays – can go a long way toward explaining the sharp increase in foreclosures and precipitous drop in home prices during the housing crash.

Each of these quantitative studies, though, abstracts from unsecured credit and thus from the bankruptcy versus mortgage default decision. Mitman (2011) takes up this task and exploits variations in homestead exemptions and recourse laws across states to demonstrate that while bankruptcy rates are lower in states with higher homestead exemptions, foreclosure rates are higher. Mitman also examines his model's predictions for the long-run effects of the BAPCPA, but does not explore the implications of bankruptcy reform for the severity of the housing crisis, which is the primary focus of our analysis.

Moreover, Mitman (2011) models mortgages as one-period contracts and abstracts from the transaction costs associated with buying and selling a home. Although these assumptions improve analytical tractability and perhaps are appropriate for a steady state analysis, the inherent risks to both households and lenders in a long-term mortgage contract, such as changes in income and house prices, are of first-order importance for our dynamic analysis. It is therefore crucial that we model mortgages as long-term contracts and explicitly account for transaction costs to adequately assess how the BAPCPA impacted the subsequent housing crisis. To the best of our knowledge, ours is the first paper to allow for both short-term, unsecured debt and long-term, collateralized mortgage loans in a model of optimal consumer default.[75]

[75]In fact, the only other model, again to the best of our knowledge, to simultaneously consider short and long-term debt is Arellano and Ramanarayanan (2010), who consider unsecured debt instruments of different maturities in a sovereign default model.

In considering long-term debt, we build on the work introducing longer maturity bonds into models of sovereign default by Chatterjee and Eyigungor (2011a) and Hatchondo and Martinez (2009) and consumer default by Chatterjee and Eyigungor (2011b).

3.3 Model Economy

We consider an environment in which time is discrete and infinite. The economy is populated by a continuum of infinitely lived households, a pool of perfectly competitive, risk-neutral financial intermediaries, and a government. There is an exogenous and perfectly elastic supply of a homogeneous consumption good which is taken as the numeraire. The economy also has exogenous and perfectly divisible supplies of owner-occupied (K_t) and rental (H_t) housing space with prices P_t and R_t at date t, respectively. Households derive utility from consumption and the size of their housing space. Financial intermediaries accept deposits and offer competitively priced one-period unsecured debt contracts and multi-period mortgages, the latter of which households can use to help finance the purchase of housing space. The government levies income taxes on households but does not provide transfers or goods and services that affect the household's problem.

3.3.1 Households

Households are heterogeneous with respect to their homeownership status, house size k_t, mortgage payment x_t, assets a_t, endowment y_t, and credit status. We use $k_t = 0$ and $x_t = 0$ to denote a household that does not own a home and therefore does not have a mortgage.

Households in our model face three sources of uncertainty. First, each household receives an idiosyncratic and stochastic endowment y_t each period, the log of which evolves according to a first-order autoregressive process:

$$\log(y_t) = \rho \log(y_{t-1}) + \varepsilon_t$$

where $\varepsilon_t \sim N(0, \sigma)$ is i.i.d over time and across households. Second, owner-occupied housing is subject to idiosyncratic proportional depreciation shocks, δ_t, that are i.i.d. across households and time.[76] The value of this shock is given by:

[76]We introduce this feature to capture two important characteristics of the U.S. housing market: (1) homeowners occasionally choose to default on their mortgage obligations, and (2) homeowners move fre-

$$\delta_t = \begin{cases} \overline{\delta} & \text{with probability } \phi \\ 0 & \text{otherwise} \end{cases}$$

where $\overline{\delta} > 0$. A household that exits period t with a house of size k_t and experiences depreciation shock δ_{t+1} enters period $t+1$ with a house of size $k_{t+1} = (1 - \delta_{t+1})k_t$. Finally, households are subject to an idiosyncratic expense shock, e_t, which directly reduces the assets with which they enter the period.[77] This expense shock is also assumed to be i.i.d. across households and time, and its value is given by:

$$e_t = \begin{cases} \overline{e} & \text{with probability } \xi \\ 0 & \text{otherwise} \end{cases}$$

where $\overline{e} > 0$. A household that exits period t with assets a^*_{t+1}, which they have optimally chosen, and experiences expense shock e_{t+1} enters period $t+1$ with assets $a_{t+1} = a^*_{t+1} - e_{t+1}$.

Figure 3.1 depicts how households move between different homeownership and credit statuses in our model. For example, a household that enters the period as a homeowner with good credit can become (i) a *homeowner with bad credit* by declaring bankruptcy and having home equity less than the homestead exemption, (ii) a *renter with bad credit* by defaulting on their mortgage, declaring bankruptcy and defaulting on their mortgage, or declaring bankruptcy with home equity in excess of the homestead exemption, or (iii) a *renter with good credit* by selling their home.[78]

---

quently. Depreciation shocks create the potential for negative home equity, a prerequisite for mortgage default, in a steady state in which owner-occupied house prices are constant. These shocks also tend to result in a suboptimal combination of mortgage loan and house size given a household's assets and income, which is the main reason why homeowners choose to move in our model.

[77] Expense shocks are meant to capture unanticipated household expenses relating to medical expenses, divorce costs, unexpected births of children, among others, which are commonly cited by bankrupts as contributing to their decision to file. See Livshits, MacGee, and Tertilt (2007) and Livshits et al. (2010) for further discussion on the importance of expense shocks for the consumer bankruptcy decision.

[78] We will use the term "good credit" to mean a household that has access to credit markets and "bad credit" to mean a household that is excluded from credit markets due to a past bankruptcy filing and/or mortgage default.

Figure 3.1: Summary of Household's Problem

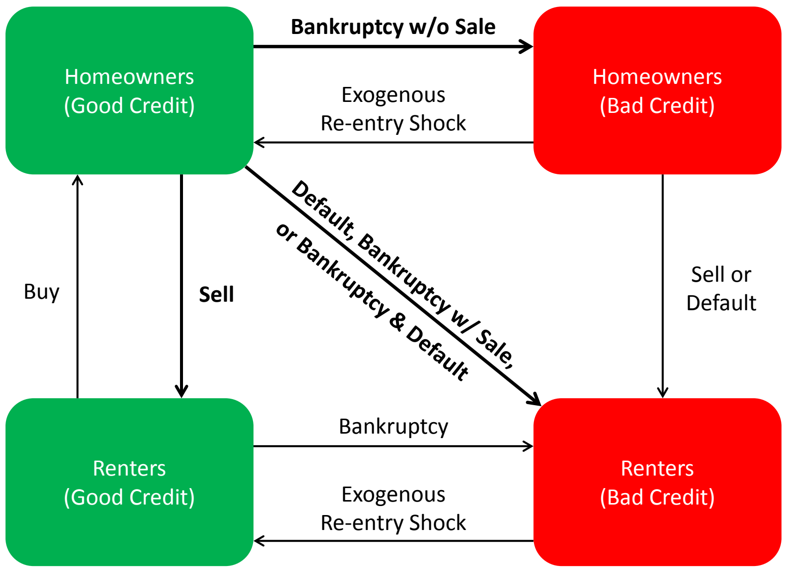

We begin by describing the decision problems for a homeowner with good credit because their decision to file for bankruptcy versus defaulting on their mortgage is affected by the BAPCPA and is thus the focus of our analysis.

Problem of a Homeowner with Good Credit

A homeowner with good credit must decide between making their mortgage payment and continuing as a homeowner (O_t), selling their home (S_t), defaulting on their mortgage (D_t), filing for bankruptcy (B_t), or both filing for bankruptcy and defaulting on their mortgage (BD_t). The value of having this decision is given by:

$$V_t(k_t, x_t, a_t, y_t) = \max_{O_t, S_t, D_t, B_t, BD_t} \{O_t(k_t, x_t, a_t, y_t), S_t(k_t, x_t, a_t, y_t),$$

$$D_t(a_t, y_t), B_t(k_t, x_t, y_t), BD_t(y_t)\}$$

The value associated with making their mortgage payment and continuing as a homeowner is

$$O_t(k_t, x_t, a_t, y_t) = \max_{c_t, a^*_{t+1}} u(c_t, k_t) + \beta \mathbb{E}_t[V_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1})|y_t]$$

subject to

$$c_t + q_t(k_t, x_t, a^*_{t+1}, y_t)a^*_{t+1} + x_t = y_t - g(x_t, a_t, y_t) + a_t$$

$$y_t - g(x_t, a_t, y_t) + a_t \geq x_t$$

$$k_{t+1} = (1 - \delta_{t+1})k_t, \;\; x_{t+1} = \mu x_t, \;\; a_{t+1} = a^*_{t+1} - e_{t+1}$$

The first constraint is the household's budget constraint, where $q_t(k_t, x_t, a^*_{t+1}, y_t)$ is the price of a one-period unsecured debt contract for a household with house size k_t, mortgage payment x_t, and endowment y_t that wishes to carry assets a^*_{t+1} into the following period. Here $g(x_t, a_t, y_t)$ represents the income tax levied by the government on a household with the given characteristics. The second constraint restricts the household from paying their mortgage with unsecured debt by ensuring that their mortgage payment does not exceed their after-tax income plus their resources from their bond holdings with which they entered the period, net of the expense shock. The final three constraints represent the laws of motion for the household's home size, mortgage payment, and assets. While we discuss in detail our assumptions about mortgage contracts in the following section, for now it suffices to convey that mortgage payments decay over time at the constant rate $\mu \in (0, 1)$. Moreover, while the household chooses assets a^*_{t+1} with which to exit the period, the assets with which they enter the following period depend on the realized expense shock e_{t+1}. Note that the expectation on the right hand side of the value function is taken with respect to all three sources of uncertainty: the household's next period endowment, depreciation shock, and expense shock.

If instead they choose to sell their home, they receive the proceeds from the sale $P_t k_t$

less a proportional transaction cost χ_S. The household must also repurchase their mortgage contract from the lender for an amount equal to the present value of the promised stream of decaying mortgage payments, discounted at the risk-free interest rate. We assume that the sale and purchase of housing space occurs at the beginning of each period, and therefore, the household must rent housing space in the current period. The value of selling is thus:

$$S_t(k_t, x_t, a_t, y_t) = \max_{c_t, a^*_{t+1}, h_t} u(c_t, h_t) + \beta \mathbb{E}_t[V_{t+1}(0, 0, a_{t+1}, y_{t+1})|y_t]$$

subject to

$$c_t + q_t(0, 0, a^*_{t+1}, y_t)a^*_{t+1} + R_t h_t = y_t - g(0, a_t, y_t) + a_t + P_t k_t(1 - \chi_S) - \left(1 + \frac{\mu}{r + 1 - \mu}\right) x_t$$

$$a_{t+1} = a^*_{t+1} - e_{t+1}$$

The household may also decide to default on their mortgage. In this case they are relieved of their mortgage payment but must relinquish their home to the lender. The household must also rent housing space in the current period and is temporarily excluded from credit markets but may save. We assume that households with bad credit re-enter credit markets with probability λ each period. Hence, the value of defaulting is

$$D_t(a_t, y_t) = \max_{c_t, a^*_{t+1} \geq 0, h_t} u(c_t, h_t) + \beta \mathbb{E}_t[\lambda V_{t+1}(0, 0, a_{t+1}, y_{t+1})$$

$$+(1 - \lambda) X_{t+1}(0, 0, a_{t+1}, y_{t+1})|y_t]$$

subject to

$$c_t + q_t(0, 0, a^*_{t+1}, y_t)a^*_{t+1} + R_t h_t = y_t - g(0, a_t, y_t) + a_t$$

$$a_{t+1} = a^*_{t+1} - e_{t+1}$$

where $X_{t+1}(0, 0, a_{t+1}, y_{t+1})$ is the value of being a renter with bad credit. Note that the value of defaulting is independent of k_t and x_t since the household loses their home and is relieved of their mortgage in the current period.

Alternatively, a household may choose to file for bankruptcy and have their unsecured debt obligations discharged in exchange for a one-time utility cost $\nu > 0$ and temporary exclusion from credit markets. In addition, a household that files for bankruptcy may face either a one-time endowment cost $\omega_t(y_t)$ or be forced to sell their home. Homeowners who declare bankruptcy and are forced to sell their home are allowed to retain any home equity up to the homestead exemption ζ and must rent housing space in the current period.[79] We therefore divide the value of bankruptcy into two distinct pieces:

(1) The household is forced to sell their home:

$$B_t(k_t, x_t, y_t) = \max_{c_t, a^*_{t+1} \geq 0, h_t} u(c_t, h_t) - \nu + \beta \mathbb{E}_t[\lambda V_{t+1}(0, 0, a_{t+1}, y_{t+1})$$

$$+(1-\lambda) X_{t+1}(0, 0, a_{t+1}, y_{t+1}) | y_t]$$

subject to

$$c_t + q_t(0, 0, a^*_{t+1}, y_t) a^*_{t+1} + R_t h_t = y_t - g(0, 0, y_t) + \zeta$$

$$a_{t+1} = a^*_{t+1} - e_{t+1}$$

Note that the value of filing for bankruptcy is independent of the household's debt since it is entirely discharged.

(2) The household is allowed to keep their home:

$$B_t(k_t, x_t, y_t) = \max_{c_t, a^*_{t+1} \geq 0} u(c_t, k_t) - \nu + \beta \mathbb{E}_t[\lambda V_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1})$$

[79] We state the conditions under which a household that declares bankruptcy is forced to sell their home in the following section.

$$+(1-\lambda)X_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1})|y_t]$$

subject to

$$c_t + q_t(k_t, x_t, a^*_{t+1}, y_t)a^*_{t+1} + x_t = y_t - g(x_t, 0, y_t) - \omega_t(y_t)$$

$$k_{t+1} = (1-\delta_{t+1})k_t,\ x_{t+1} = \mu x_t,\ a_{t+1} = a^*_{t+1} - e_{t+1}$$

Finally, we allow a household to simultaneously file for bankruptcy and default on their mortgage. The value of doing so is:

$$BD_t(y_t) = \max_{c_t, a^*_{t+1} \geq 0, h_t} u(c_t, h_t) - \nu + \beta \mathbb{E}_t[\lambda V_{t+1}(0, 0, a_{t+1}, y_{t+1})$$

$$+(1-\lambda)X_{t+1}(0, 0, a_{t+1}, y_{t+1})|y_t]$$

subject to

$$c_t + q_t(0, 0, a^*_{t+1}, y_t)a^*_{t+1} + R_t h_t = y_t - g(0, 0, y_t) - \omega_t(y_t)$$

$$a_{t+1} = a^*_{t+1} - e_{t+1}$$

Here the value of defaulting and filing for bankruptcy together only depends on the household's endowment since defaulting results in the loss of housing space and mortgage payment, while bankruptcy relieves the household of its unsecured debt obligations.

Problem of a Homeowner with Bad Credit

Now consider the problem of a household that owns their housing space but is excluded from credit markets. Such a household necessarily has filed for bankruptcy in the past and has not yet regained access to credit markets. The decision problem of this type of household is analogous to that presented above, except that they are restricted from borrowing in unsecured credit markets and hence will not declare bankruptcy. The household chooses whether to repay their mortgage and continue as a homeowner (O^X_t), sell their home (S^X_t), or default on their mortgage (D_t^X). Their optimal choice is the one with the highest value:

$$X_t(k_t, x_t, a_t, y_t) = \max_{O_t^X, S_t^X, D_t^X} \{O_t^X(k_t, x_t, a_t, y_t), S_t^X(k_t, x_t, a_t, y_t), D_t^X(a_t, y_t)\}.$$

The value of making their mortgage payment and continuing as a homeowner is given by

$$O_t^X(k_t, x_t, a_t, y_t) = \max_{c_t, a_{t+1}^* \geq 0} u(c_t, k_t) + \beta \mathbb{E}_t[\lambda V_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1})$$

$$+(1-\lambda) X(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1})|y_t]$$

subject to

$$c_t + q_t(k_t, x_t, a_{t+1}^*, y_t) a_{t+1}^* + x_t = y_t - g(x_t, a_t, y_t) + a_t$$

$$k_{t+1} = (1-\delta_{t+1}) k_t, \;\; x_{t+1} = \mu x_t, \;\; a_{t+1} = \max\{a_{t+1}^* - e_{t+1}, 0\}$$ [80]

If instead they choose to sell their home, they receive the proceeds from the sale $P_t k_t$ less a proportional transaction cost χ_S. The household must also repurchase their mortgage contract from the lender for an amount equal to the present value of the promised stream of decaying mortgage payments, discounted at the risk-free interest rate. Recall that a household that sells their home must rent housing space in the current period. The value of selling is then

$$S_t^X(k_t, x_t, a_t, y_t) = \max_{c_t, a_{t+1}^* \geq 0, h_t} u(c_t, h_t) + \beta \mathbb{E}_t[\lambda V_{t+1}(0, 0, a_{t+1}, y_{t+1})$$

$$+(1-\lambda) X_{t+1}(0, 0, a_{t+1}, y_{t+1})|y_t]$$

[80] For simplicity, we assume that the size of the expense shock is capped by the household's positive assets when a household is excluded from credit markets. This will ensure that excluded households never want to declare bankruptcy.

subject to

$$c_t + q_t(0,0,a^*_{t+1},y_t)a^*_{t+1} + R_t h_t = y_t - g(0,a_t,y_t) + a_t + P_t k_t(1-\chi_S) - \left(1 + \frac{\mu}{r+1-\mu}\right) x_t$$

$$a_{t+1} = \max\{a^*_{t+1} - e_{t+1}, 0\}$$

The household may also decide to default on their mortgage. In this case they are relieved of their mortgage payment but must relinquish their home to the lender and remain temporarily excluded from credit markets. They must also rent housing space in the current period. The value of defaulting is

$$D^X_t(a_t,y_t) = \max_{c_t,a^*_{t+1}\geq 0,h_t} u(c_t,h_t) + \beta\mathbb{E}_t[\lambda V_{t+1}(0,0,a_{t+1},y_{t+1})$$

$$+(1-\lambda)X_{t+1}(0,0,a_{t+1},y_{t+1})|y_t]$$

subject to

$$c_t + q_t(0,0,a^*_{t+1},y_t)a^*_{t+1} + R_t h_t = y_t - g(0,a_t,y_t) + a_t$$

$$a_{t+1} = \max\{a^*_{t+1} - e_{t+1}, 0\}$$

Problem of a Renter with Good Credit

Next, consider the decision problem faced by a household that does not own a home and is in good credit standing. This type of household must choose between purchasing housing space (O^R_t), continuing to rent (L^R_t), and filing for bankruptcy (B^R_t). Their optimal choice is the one with the highest value:

$$V_t(0,0,a_t,y_t) = \max_{O^R_t,L^R_t,B^R_t} \left\{O^R_t(a_t,y_t), L^R_t(a_t,y_t), B^R_t(y_t)\right\}.$$

Households can finance the purchase of housing space using a combination of savings and a mortgage. If the household decides to purchase a house of size k_t, commits to first mortgage payment x_t, chooses to carry assets a^*_{t+1} into the following period, and has endowment y_t, then the lender issues a mortgage with value $m_t(k_t, x_t, a^*_{t+1}, y_t)x_t$ to the household. We impose that the household must be able to afford the sum of the purchase price $P_t k_t$, a proportional moving cost χ_B, and their first mortgage payment x_t without the need to borrow in unsecured credit markets. The value of purchasing a home is thus:

$$O^R_t(a_t, y_t) = \max_{c_t, k_t, x_t, a^*_{t+1}} u(c_t, k_t) + \beta \mathbb{E}_t[V_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1})|y_t]$$

subject to

$$c_t + q_t(k_t, x_t, a^*_{t+1}, y_t)a^*_{t+1} + P_t k_t(1 + \chi_B) + x_t = y_t - g(x_t, a_t, y_t) + a_t + m_t(k_t, x_t, a^*_{t+1}, y_t)x_t$$

$$y_t - g(x_t, a_t, y_t) + a_t + m_t(k_t, x_t, a^*_{t+1}, y_t)x_t \geq P_t k_t(1 + \chi_B) + x_t$$

$$P_t k_t \geq m_t(k_t, x_t, a^*_{t+1}, y_t)x_t$$

$$k_{t+1} = (1 - \delta_{t+1})k_t,\ x_{t+1} = \mu x_t,\ a_{t+1} = a^*_{t+1} - e_{t+1}$$

where the third constraint restricts the household from taking out a mortgage that exceeds the value of the home.

If the household decides to repay their unsecured debt and continue renting housing space, the value is given by:

$$L^R_t(a_t, y_t) = \max_{c_t, a^*_{t+1}, h_t} u(c_t, h_t) + \beta \mathbb{E}_t[V_{t+1}(0, 0, a_{t+1}, y_{t+1})|y_t]$$

subject to

$$c_t + q_t(0, 0, a^*_{t+1}, y_t)a^*_{t+1} + R_t h_t = y_t - g(0, a_t, y_t) + a_t$$

$$a_{t+1} = a^*_{t+1} - e_{t+1}$$

Finally, the household can choose to file for bankruptcy subject to the same costs and penalties described above. The value of pursuing this option is

$$B_t^R(y_t) = \max_{c_t, a_{t+1}^* \geq 0, h_t} u(c_t, h_t) - \nu + \beta \mathbb{E}_t[\lambda V_{t+1}(0, 0, a_{t+1}, y_{t+1})$$

$$+(1-\lambda) X_{t+1}(0, 0, a_{t+1}, y_{t+1}) | y_t]$$

subject to

$$c_t + q_t(0, 0, a_{t+1}^*, y_t) a_{t+1}^* + R_t h_t = y_t - g(0, 0, y_t) - \omega_t(y_t)$$

$$a_{t+1} = a_{t+1}^* - e_{t+1}$$

Problem of a Renter with Bad Credit

Lastly, consider the decision problem of a household that does not own housing space and is excluded from credit markets. To (slightly) simplify our analysis, we restrict this type of household from purchasing a home, and hence, they must rent housing space until they regain access to credit markets. The problem of this type of household is:

$$X_t(0, 0, a_t, y_t) = \max_{c_t, a_{t+1}^* \geq 0, h_t} u(c_t, h_t) + \beta \mathbb{E}_t[\lambda V_{t+1}(0, 0, a_{t+1}, y_{t+1})$$

$$+(1-\lambda) X_{t+1}(0, 0, a_{t+1}, y_{t+1}) | y_t]$$

subject to

$$c_t + q_t(0, 0, a_{t+1}^*, y_t) a_{t+1}^* + R_t h_t = y_t - g(0, a_t, y_t) + a_t$$

$$a_{t+1} = \max\{a_{t+1}^* - e_{t+1}, 0\}$$

3.3.2 Financial Intermediaries

We assume that financial intermediaries are risk neutral and competitive. For simplicity, we consider a representative financial intermediary that accepts deposits, lends to households in unsecured credit markets, and sells mortgages to help households finance the purchase of owner-occupied housing space. The financial intermediary can also borrow or lend risk-free at the exogenously given interest rate r.

For computational tractability, we model mortgage contracts as perpetuities with payments that decay over time. In particular, when taking out a mortgage, the mortgagee agrees to the sequence of payments $\{x, \mu x, \mu^2 x, ...\}$, where $\mu \in (0,1)$, until they either default or sell their home. The decaying nature of mortgage payments allows households to gradually build home equity over time, even with a constant house price.

Consider a mortgage sold to a household planning to purchase a home of size k_t, with initial payment x_t, end of period assets a^*_{t+1}, and endowment y_t. The intermediary then disperses the amount $m_t(k_t, x_t, a^*_{t+1}, y_t)x_t$ to the household in the current period and receives the first payment x_t. If the household defaults in the following period, the intermediary takes control of the house and sells it through a foreclosure process, recovering a fraction $1 - \chi_S$ of its post-depreciation shock market value $P_{t+1}k_{t+1}$, where χ_S is a proportional transaction cost.

If the household decides to sell, they must repurchase their mortgage contract from the lender for an amount equal to the present value of the promised stream of decaying mortgage payments, discounted at the risk-free interest rate, or $(1 + \mu/(r + 1 - \mu))x_{t+1}$.

If the household declares bankruptcy, their unsecured debt obligations are discharged in exchange for temporary exclusion from credit markets, a one-time utility cost, and either a one-time endowment cost or the forced sale of their home.[81] If the home is liquidated as part of the bankruptcy proceedings, the intermediary receives the present value of the mortgage discounted at the risk-free interest rate. From the intermediary's perspective, bankruptcy in this case is equivalent to the sale of the home. On the other hand, if the household is allowed to keep their home, the intermediary receives the continuation value of the mortgage conditional on the household's choice of assets, realized endowment, depreciation shock, expense shock, and inability to borrow in unsecured credit markets.

[81] We will discuss the details pertaining to the U.S. bankruptcy code and its treatment of homeownership in the following section.

If the household neither defaults, sells, nor declares bankruptcy, the intermediary receives the continuation value of the mortgage conditional on the household's choice of assets, realized endowment, depreciation shock, and expense shock in the following period.

Let $\mathbb{D}_t(k_t, x_t, a_t, y_t)$ be an indicator function equal to 1 if a household with these characteristics finds it optimal to default at time t and 0 otherwise. Likewise, let $\mathbb{S}_t(k_t, x_t, a_t, y_t)$ be an indicator function equal to 1 if a household with these characteristics sells their home (either because they find it optimal to sell or because their home is liquidated during bankruptcy) and 0 otherwise, and similarly, define $\mathbb{B}_t(k_t, x_t, a_t, y_t)$ for a household that declares bankruptcy but is not forced to sell their home. The zero profit condition for this mortgage contract is then:

$$m_t(k_t, x_t, a^*_{t+1}, y_t)x_t = x_t$$

$$+\frac{1}{1+r+\alpha_t}\{\mathbb{E}_t[\underbrace{\mathbb{D}_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1})P_{t+1}k_{t+1}(1-\chi_S)}_{\text{Value if household defaults}}$$

$$+\underbrace{\mathbb{S}_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1})\left(1+\frac{\mu}{r+1-\mu}\right)x_{t+1}}_{\text{Value if house is sold}}$$

$$+\underbrace{\mathbb{B}_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1})m^X_{t+1}(k_{t+1}, x_{t+1}, a^*_{t+2}, y_{t+1})x_{t+1}}_{\text{Continuation value of mortgage after bankruptcy}}$$

$$+(1-\mathbb{D}_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1}))(1-\mathbb{S}_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1}))$$

$$\underbrace{(1-\mathbb{B}_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1}))m_{t+1}(k_{t+1}, x_{t+1}, a^*_{t+2}, y_{t+1})x_{t+1}}_{\text{Continuation value of mortgage without bankruptcy}}|y_t]\}$$

where α_t is a time-varying credit wedge and the expectation is taken over the realization of the household's next period endowment, depreciation shock, and expense shock.

Since the value of a mortgage today depends on its continuation value tomorrow if the household files for bankruptcy and is allowed to keep their home, creditors must also price mortgages to households that are excluded from credit markets even though such a mortgage is never actually sold in equilibrium.[82] A household that is excluded from credit markets, owns a home of size k_t, and has a mortgage with payment x_t will never file for bankruptcy (since they will not have any unsecured debt), but they may choose to sell their home or default on their mortgage. In addition, the household is allowed to re-enter credit markets with probability λ each period. Let $\mathbb{D}_t^X(k_t, x_t, a_t, y_t)$ and $\mathbb{S}_t^X(k_t, x_t, a_t, y_t)$ be indicator functions analogous to those described above but specific to households that are excluded from credit markets. The value of this mortgage contract is then

$$m_t^X(k_t, x_t, a_{t+1}^*, y_t)x_t = x_t$$

$$+\frac{1}{1+r+\alpha_t}\{\mathbb{E}_t[(1-\lambda)\{\underbrace{\mathbb{D}_{t+1}^X(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1})P_{t+1}k_{t+1}(1-\chi_S)}_{\text{Value of remaining excluded and defaulting}}$$

$$+\underbrace{\mathbb{S}_{t+1}^X(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1})\left(1+\frac{\mu}{r+1-\mu}\right)x_{t+1}}_{\text{Value of remaining excluded and selling}}$$

$$+(1-\mathbb{D}_{t+1}^X(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1}))$$

$$\underbrace{(1-\mathbb{S}_{t+1}^X(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1}))m_{t+1}^X(k_{t+1}, x_{t+1}, a_{t+2}^*, y_{t+1})x_{t+1}}_{\text{Value of remaining excluded and continuing mortgage}}\}$$

$$+\lambda\{\underbrace{\mathbb{D}_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1})P_{t+1}k_{t+1}(1-\chi_S)}_{\text{Value of re-entering credit markets and defaulting}}$$

$$+\underbrace{\mathbb{S}_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1})\left(1+\frac{\mu}{r+1-\mu}\right)x_{t+1}}_{\text{Value of re-entering credit markets and selling}}$$

$$+(1-\mathbb{D}_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1}))$$

$$\underbrace{(1-\mathbb{S}_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1}))m_{t+1}(k_{t+1}, x_{t+1}, a_{t+2}^*, y_{t+1})x_{t+1}}_{\text{Value of re-entering credit markets and continuing mortgage}}\}|y_t]\}$$

[82] We assume that households that are excluded from unsecured credit markets are also excluded from mortgage markets, and therefore, this type of mortgage is never sold to households in our model. Given our assumption of competitive financial intermediaries, though, one can think of an active secondary mortgage market in which this type of mortgage, along with all other active mortgages, are traded. It is this market in which the continuation value, or price, of mortgages such as this one are determined.

Together, these functional equations determine the profit maximizing, equilibrium mortgage contract pricing schedules $m_t(k_t, x_t, a^*_{t+1}, y_t)$ and $m^X_t(k_t, x_t, a^*_{t+1}, y_t)$.

The financial intermediary also offers one-period, unsecured, pure discount bonds which households cannot commit to repay. Suppose, for example, a household with house size k_t, mortgage payment x_t, and endowment y_t promises to repay an amount a^*_{t+1} in the following period. The intermediary then disperses the amount $q_t(k_t, x_t, a^*_{t+1}, y_t)a^*_{t+1}$ to the household in the current period. If the household does not declare bankruptcy in the following period, then the intermediary is repaid the amount a^*_{t+1} in full. On the other hand, if the household declares bankruptcy, the intermediary recovers an amount $\psi_{t+1}(k_{t+1}, x_{t+1}, a^*_{t+1}, y_{t+1})$ which depends on the household's characteristics and the current bankruptcy laws in place. The zero profit condition for this type of loan is:

$$q_t(k_t, x_t, a^*_{t+1}, y_t)a^*_{t+1} = \frac{1}{1+r}\{\mathbb{E}_t[\underbrace{(1 - \mathbb{B}_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1}))a^*_{t+1}}_{\text{Value of repaying loan}}$$

$$+ \underbrace{\mathbb{B}_{t+1}(k_{t+1}, x_{t+1}, a_{t+1}, y_{t+1})\psi_{t+1}(k_{t+1}, x_{t+1}, a^*_{t+1}, y_{t+1})}_{\text{Value of declaring bankruptcy}}|y_t]\}.$$

3.3.3 Government

There is also a government that levies income taxes on households. We include a government in our model to capture two of the primary financial benefits of homeownership in the U.S.: (1) the implicit rental income from homeownership is not taxed and (2) mortgage interest payments are tax deductible. While the former induces high-income households to purchase rather than rent their housing space, the latter gives an incentive for homebuyers to finance their purchase with debt rather than equity. For simplicity, we assume that government consumption does not provide any benefit to households and that tax revenues are not rebated to households.

The tax levied on each household (g) is modeled after the U.S. tax code. A household's taxable income i is the sum of their current endowment and interest on deposits less the greater of (i) their mortgage interest payment $\mu r x_t/(r+1-\mu)$ and (ii) the standard deduction s:

$$i(x_t, a_t, y_t) = y_t + r\max\{a_t, 0\} - \max\left\{\frac{\mu r x_t}{r+1-\mu}, s\right\}.$$

We assume that the tax rate $\tau(i(x_t, a_t, y_t))$ is weakly increasing in the household's taxable income. The tax levied on a household is then:

$$g(x_t, a_t, y_t) = \int_0^{i(x_t,a_t,y_t)} \tau(w)dw$$

and their after-tax income is given by $y_t - g(x_t, a_t, y_t)$.

3.3.4 Market Clearing

Let $\Phi_t(k_t, x_t, a_t, y_t, cs_t)$ represent the distribution of households over owner-occupied housing space, mortgage payments, assets, endowments, and credit statuses (cs_t) entering period t. The prices P_t and R_t adjust each period so that the aggregate demands for owner-occupied and rental housing space equal their exogenous supplies:

$$K_t = \int k_t(k_t, x_t, a_t, y_t)d\Phi_t(k_t, x_t, a_t, y_t, cs_t)$$

$$H_t = \int h_t(k_t, x_t, a_t, y_t)d\Phi_t(k_t, x_t, a_t, y_t, cs_t)$$

3.3.5 Equilibrium

An equilibrium in this economy is a sequence of prices $\{P_t, R_t, q_t, m_t, m_t^X\}$, exogenous sequences of owner-occupied and rental housing stocks $\{K_t, H_t\}$, sequences of household decision rules, and a sequence of distributions of households over states $\{\Phi_t\}$, such that, taking prices, the bankruptcy code, housing supplies, and the initial distribution of households over states Φ_0 as given, at each date t:

1. Households optimally solve their decision problems.

2. Creditors maximize profits.

3. Markets for owner-occupied and rental housing clear.

3.4 The BAPCPA

Prior to the BAPCPA, most households with significant unsecured debt obligations could benefit by filing for bankruptcy. Households who did not own a home were able to have all of their unsecured debt obligations extinguished in exchange for having a bankruptcy flag on their credit report for a period of 10 years.[83] We model this penalty as a one-time utility cost and temporary exclusion from credit markets, during which time households can neither borrow in unsecured credit markets nor purchase a home.[84] There are no other costs associated with declaring bankruptcy in this case, and unsecured creditors do not recover anything (i.e. $\omega_t(y_t) = 0$ and $\psi_t(k_t, x_t, a_t^*, y_t) = 0$).

The U.S. bankruptcy code provides exemptions that households can use to protect certain assets from seizure by creditors. The largest and most commonly used is an exemption for the home, which allows homeowners to keep their home equity up to a prespecified limit known as the homestead exemption. Homeowners with home equity less than the homestead exemption were allowed to keep their home and file under Chapter 7. Homeowners with home equity greater than the homestead exemption, on the other hand, were forced to sell their home and transfer all home equity in excess of the homestead exemption (non-exempt home equity) to their unsecured creditors.

[83] The presence of a bankruptcy flag on a household's credit report has been shown to severely restrict their access to credit (see Musto (2004)).

[84] The one-time utility cost is meant to capture the social stigma attached to bankrupts discussed extensively in the literature (see Fay et al. (2002) and Gross and Souleles (2002)).

In terms of our model, define the home equity of a household with house size k_t and payment x_t at time t as

$$HE_t(k_t, x_t) \equiv P_t k_t - \left(1 + \frac{\mu}{r + 1 - \mu}\right) x_t$$

and let ζ be the homestead exemption. Prior to the BAPCPA, a homeowner with $HE_t(k_t, x_t) \leq \zeta$ would be allowed to keep their home, while a homeowner with $HE_t(k_t, x_t) > \zeta$ would be forced to sell, raising an amount $P_t k_t$. Out of these funds, the mortgage lender would receive the present value of the promised stream of decaying mortgage payments, discounted at the risk-free interest rate, or $[1 + \mu/(r + 1 - \mu)]x_t$, the household would keep an amount equal to the homestead exemption ζ, and unsecured creditors would be paid all non-exempt home equity up to the original loan amount:

$$\psi_t(k_t, x_t, a_t^*, y_t) = \min\left\{|a_t^*|, \max\left\{HE_t(k_t, x_t) - \zeta, 0\right\}\right\}.$$ [85]

The BAPCPA made it more costly for households to declare bankruptcy. It raised the average total bankruptcy filing costs under both Chapter 7 and Chapter 13, capped the homestead exemption for households who have owned their home for less than 3-1/2 years, increased the number of years before a household could refile from six to eight, and introduced means testing that severely restricted high-income households' ability to benefit from bankruptcy. While all of these reforms clearly affect a homeowner's decision to file, we focus our attention on the effects of means testing as it is likely to have the largest impact on household behavior.

To illustrate the impact of means testing introduced under the BAPCPA, consider a household that either does not own a home or has home equity below the homestead exemption. The first step is to convert the household's income over the previous six months

[85] The fact that non-exempt home equity is seized by unsecured creditors during bankruptcy should lead to lower interest rates on borrowing in unsecured credit markets for these borrowers. We find this to be a quantitatively important benefit of homeownership which was reduced by the BAPCPA.

to an annualized basis and then compare it to the median income in their home state. If their income is less than the median, they are permitted to file under Chapter 7 and are unaffected by the reform. Conversely, if their income is above the median, the household may be forced to file under Chapter 13 and commit to a repayment plan.[86] In this case, the household's unsecured debt is discharged, but they are required to pay all non-exempt income to their creditors for a period of five years. For simplicity, we model this penalty as a one-time endowment cost. Specifically, if $\overline{y}$ is median income, then a household that declares bankruptcy, does not own a home or owns a home but has home equity less than the homestead exemption, and has endowment $y_t > \overline{y}$, is required to repay an amount

$$\omega_t(y_t) = \kappa(y_t - \overline{y})$$

in the current period to their unsecured creditors, in addition to facing the same one-time utility cost and temporary exclusion from credit markets discussed above. It follows in this case that,

$$\psi_t(k_t, x_t, a_t^*, y_t) = \min\left\{|a_t^*|, \kappa(y_t - \overline{y})\right\},$$

where the creditor's recovery amount is bounded above by the initial loan amount.

Now consider a household that owns their home and has home equity in excess of the homestead exemption. If the household's non-exempt home equity is greater than five times their non-exempt income, then the household is forced to sell their home and pay all non-exempt home equity to their unsecured creditors. Otherwise, the household is allowed to keep their home, but must pay all non-exempt income to their creditors for a period of five years. In both cases the household is subject to the same one-time utility cost and temporary exclusion from credit markets discussed above. In terms of our model, if

[86] This occurs if their income in excess of their exempt income, where exempt income includes the funds required for housing and transportation costs and personal expenses as well as additional amounts for their mortgage and car payments, exceeds $2,000. See Li et al. (2011) for a detailed description of how a household's non-exempt income is computed.

$$HE_t(k_t, x_t) - \zeta > 5(y_t - \overline{y})$$

then the household is forced to sell their home, raising an amount $P_t k_t$. Out of these funds, the mortgage lender receives the present value of the promised stream of decaying mortgage payments, discounted at the risk-free interest rate, or $[1 + \mu/(r + 1 - \mu)]x_t$, the household receives an amount equal to the homestead exemption ζ, and unsecured creditors are paid all non-exempt home equity up to the original loan amount:

$$\psi_t(k_t, x_t, a_t^*, y_t) = \min\left\{|a_t^*|, \max\left\{HE_t(k_t, x_t) - \zeta, 0\right\}\right\}.$$

On the other hand, if

$$HE_t(k_t, x_t) - \zeta \leq 5(y_t - \overline{y})$$

the household is required to repay an amount

$$\omega_t(y_t) = \kappa(y_t - \overline{y}),$$

in the current period to their unsecured creditors, in addition to facing the same one-time utility cost and temporary exclusion from credit markets discussed above. It follows in this case that,

$$\psi_t(k_t, x_t, a_t^*, y_t) = \min\left\{|a_t^*|, \kappa(y_t - \overline{y})\right\},$$

where, again, the creditor's recovery amount is bounded above by the initial loan amount.

Importantly, in this case the creditor's recovery amount only depends on the household's income and is independent of the household's homeownership status. While this household would have benefited from owning a home through lower interest rates on unsecured debt prior to the BAPCPA, this benefit is no longer available under the BAPCPA. Consequently, homeowners that currently have negative home equity and expect to have high income in the future, perceive a reduced benefit to future homeownership. This lower future benefit may induce some homeowners to default on their mortgage instead of staying in their home. Understanding this mechanism is important when discussing the effects of the BAPCPA on households' incentives to default.

Figure 3.2: Implementation of BAPCPA

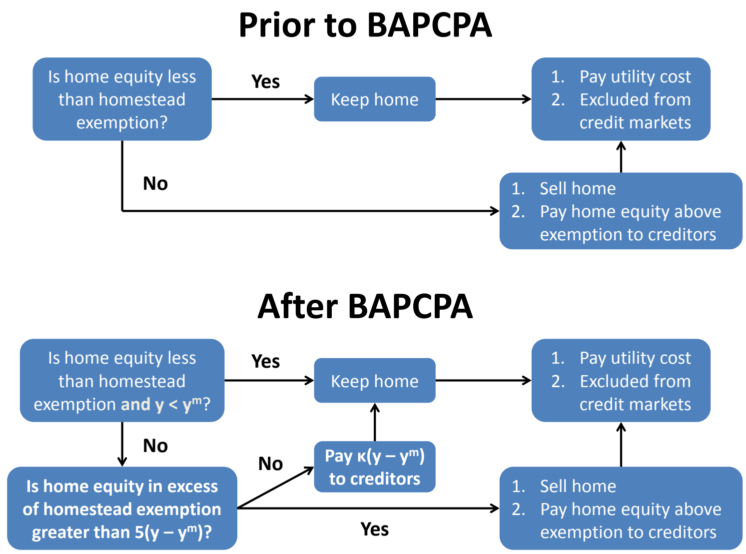

The effects of means testing implemented under the BAPCPA on the costs of bankruptcy faced by households in our model is depicted in Figure 3.2. Clearly, bankruptcy reform made filing for bankruptcy much more costly for high-income households. In the following sections we calibrate our model and quantify the effects of BAPCPA on the recent housing crisis.

3.5 Parameterization

We assume that each period in our model corresponds to one year. Many of our model parameters are common in the literature and can therefore be set outside of the model. We first discuss how these parameters are chosen and then describe how we calibrate the remaining model parameters.

The parameters governing households' stochastic first-order autoregressive endowment process are set to $\rho = 0.97$ and $\sigma = 0.129$, which are consistent with the findings of Storesletten et al. (2004b). We discretize this process with a 17-state Markov chain using Tauchen and Hussey (1991)'s method.

We assume that a household's flow utility at date t is given by:

$$u(c_t, h_t) = \frac{\left(c_t^{1-\theta} h_t^{\theta}\right)^{1-\gamma}}{1-\gamma}$$

where γ is a proxy for risk aversion and θ determines the share of income spent on housing space.[87] We set $\gamma = 2$, which is a standard value for this parameter used in the literature. Empirical work by Davis and Ortalo-Magne (2011) indicates that household's spend approximately 24% of their income on housing services, so we set $\theta = 0.24$.

The parameters related to the housing sector that are determined outside of our model are $(\chi_B, \chi_S, \overline{\delta}, \zeta)$. The proportional transaction costs for buying and selling are set to $\chi_B = 0.025$ and $\chi_S = 0.070$, respectively, which are in line with the values reported by Gruber and Martin (2003). Pennington-Cross (2006) finds that the value received from the sale of a foreclosed home is about 78% of the market value for a similar non-foreclosed home.[88] Since in our model a household that chooses to default on their mortgage also often has incurred the housing depreciation shock, we set $\overline{\delta}$ such that creditors receive 78% of the value of the pre-depreciation shock home after selling transaction costs. This implies that $\overline{\delta} = 0.15$.[89] Since our model is intended to represent the U.S. economy, we compute the average homestead exemption across states, where each state is weighted by its share

[87] Cobb-Douglas preferences imply that a renting household will choose to spend constant fractions of their wealth on non-durable consumption and housing services. Note that households do not derive any direct utility benefit from owning versus renting their housing space in this model.

[88] This finding is in line with estimates from other work, including Shilling, Benjamin, and Sirmans (1990), who find values that range in 22%-24%.

[89] More formally, the value to the creditor of a foreclosed home that received the depreciation shock is $(1-\overline{\delta})(1-\chi_S)P_t k_t = 0.85(0.925)P_t k_t = 0.786 P_t k_t$, matching the empirical literature.

of U.S. households.[90] Using data collected by Mitman (2011), we find the weighted average homestead exemption to be 1.10 times median household income. We normalize median income in our model to 1 and therefore set $\zeta = 1.10$.

The risk-free interest rate is set to 4% as is standard in the literature. The positive value for the expense shock $\overline{e}$ is set to 3.33 times median income, which is consistent with the findings of Livshits et al. (2007). The probability of re-entering credit markets after declaring bankruptcy or defaulting on a mortgage λ is set to 12%, implying that, on average, an excluded household re-enters credit markets after 8.5 years. Although households that declared bankruptcy during the pre-reform period were only restricted from refiling for 6 years, there is empirical evidence that filing for bankruptcy impacts a household's credit market status for as long as their credit score is adversely affected. Moreover, underwriting standards by the government-sponsored enterprises over this period suggest that access to mortgage markets is also similarly restricted after a bankruptcy or default.[91]

Finally, we calibrate the tax schedule. As in Chatterjee and Eyigungor (2011b) we assume that a household in our model files their taxes as married filing separately and calibrate the model's income tax schedule to match that of the U.S. economy in 1998. Table 3.1 presents the implied tax schedule, and we set the standard deduction $s = 0.1116$.

Table 3.1: Model Income Tax Brackets

Taxable Income (i)	Tax Rate (τ)
0 – 0.64	0.15
0.64 – 1.55	0.28
1.55 – 2.37	0.31
2.37 – 4.23	0.36
4.23 – ∞	0.396

[90]We exclude states with an infinite homestead exemption from this calculation.

[91]For example, Musto (2004) finds that households that declare bankruptcy face restricted access to credit markets at potentially prohibitively tough terms for 10 years after they file – at which point the bankruptcy flag is removed from their credit report. Defaulting on a mortgage also negatively impacts a household's credit score and thus their ability to borrow in unsecured credit markets (see Christie (2010) and Brevoort and Cooper (2010)).

The remaining parameters to be calibrated are the discount factor β, the utility cost of bankruptcy ν, the rate of decay for mortgage payments μ, the probability the household receives a depreciation shock ϕ, and the probability a household receives an expense shock ξ. These parameters are jointly calibrated to match the unsecured debt-to-income ratio, bankruptcy filing rate, percentage of homeowners with less than 30% home equity, mortgage default rate, and bankruptcy rate among new homeowners in the stationary distribution of the model prior to the BAPCPA.

Since these statistics are intended to capture a steady state in the U.S. housing market prior to the BAPCPA, we choose targets that predate the substantial rise in homeownership rates and house prices that corresponded with the housing boom in the mid-2000's. The target bankruptcy filing rate is set to 1.4%, which was the total bankruptcy filing rate in 2004 as reported by Li and White (2009). The percentage of homeowners with home equity less than 30% is taken from Chatterjee and Eyigungor (2011b), who in turn compute this number from the 1998 Survey of Consumer Finances. This value is set to 23.0%. The annual foreclosure rate according to the Mortgage Banker's Association was about 1.0%. However, using data from LPS Analytics between 2001 and 2003, Herkenhoff and Ohanian (2012) find that roughly 15% of homeowners entering the foreclosure process self-cure and remain in their home. Since defaulting on a mortgage is synonymous with losing the home through a foreclosure process in our model, we exclude such households from our target statistic, implying a mortgage default rate of 0.85%. The target bankruptcy rate for new homeowners is set to 0.57%, as reported in Li et al. (2011). The target unsecured debt-to-income ratio is set to 9.6%. This statistic is computed by constructing a revolving debt-to-income ratio measure from the Flow of Funds Accounts and adjusting this series with the historical spread between the unsecured and revolving debt-to-income ratios implied by Livshits et al. (2010). Finally, we choose to target a homeownership rate of 66.4%, which matches the ten-year average in the U.S. economy prior to 2003.

The joint calibration of these five parameters $(\beta, \nu, \mu, \phi, \xi)$ is achieved by conducting a grid search over the parameters, computing the stationary distribution of the economy for each set of parameters, and choosing the combination that minimizes a weighted sum of squared residuals between the empirical and model values for the target statistics.[92] Table 3.2 summarizes the parameter values.

Table 3.2: Parameterization

Parameter	Value	Source/Target
γ	2.0	Standard
θ	0.24	Davis and Ortalo-Magne (2011)
ρ	0.97	Storesletten et al. (2004b)
σ	0.129	Storesletten et al. (2004b)
r	0.04	Standard
χ_B	0.025	Gruber and Martin (2003)
χ_S	0.070	Gruber and Martin (2003)
$\bar{\delta}$	0.15	Pennington-Cross (2006)
ζ	1.10	Mitman (2011)
λ	0.12	Average Exclusion Period of 8.5 yrs
$\bar{e}$	3.33	Livshits et al. (2007)
β	0.936	Unsecured Debt-to-Income Ratio = 9.6%
R/P	0.052	Homeownership Rate = 66.4%
ν	1.6	Bankruptcy Filing Rate = 1.4%
μ	0.966	Fraction of HO with < 30% HE = 23.0%
ϕ	0.005	Mortgage Default Rate = 0.85%
ξ	0.004	New HO Bankruptcy Rate = 0.57%

Table 3.3 presents the calibration results and other relevant model statistics in the pre-reform stationary distribution of the model. The model is able to match the pre-reform empirical moments for the statistics targeted in our calibration exercise reasonably well. It also performs well in replicating several relevant statistics that are not targeted by our calibration exercise. Notably, the model replicates the home equity distribution rather well, only slightly underpredicting the fraction of homeowners with home equity less than 25 percent, in addition to matching the fraction of homeowners with less than 30 percent home equity.

[92]See Appendix A for this chapter for a detailed description of our algorithm to solve for the model's stationary distribution. When computing the pre-BAPCPA steady state, we also normalize the credit wedge to be equal to zero, i.e. $\alpha_t = 0$.

Table 3.3: Steady State Results

Statistic	Data	Model
Homeownership Rate*	66.4%	71.0%
Bankruptcy Filing Rate*	1.4%	1.7%
New Homeowner Bankruptcy Filing Rate*	0.57%	0.65%
Aggregate Mortgage Default Rate*	0.85%	0.90%
Homeowners with $<$ 25% Equity	19.0%	16.5%
Homeowners with $<$ 30% Equity*	23.0%	23.7%
Average Loan-to-Value at Origination	?	83.5%
Average Income of Homeowners to Renters	2.02	1.77
Average Annual Home Sales	4.3%	4.8%
Loan-to-Income Ratio	3.9	3.6
Unsecured Debt-to-Income Ratio*	9.6%	11.5%

* = Calibration Target

Matching this region of the home equity distribution is particularly important because it suggests that the fraction of homeowners that are pushed underwater on their mortgage by a drop in house prices similar to the recent housing crash is the same in the model and the data – a necessary feature of a model that quantitatively evaluates the effects of an unexpected housing crash.

The initial stationary distribution of the model is also consistent with the fact that homeowners have higher income, on average, than renters. The average income of homeowners relative to renters is 1.77 in our model, compared to 2.02 in the data.

Moreover, the pre-crisis stationary distribution is consistent with several statistics regarding the relative size of mortgages. In particular, our model matches the empirical loan-to-income value rather well and generates a loan-to-value at origination of 83.5%. Although we were unable to locate an analog to this statistic in the data, this value seems reasonable. Finally, the model implies that 4.8% of all owner-occupied houses are sold each year, which is in line with the ten-year average prior to 2003 as reported by the National Association of Realtors.

Now that we have calibrated the model and determined that it is able to match key empirical statistics in the pre-BAPCPA stationary distribution, we turn to the primary quantitative objective of this paper: assessing the impact of the BAPCPA on the U.S. economy during the recent housing crisis. The next section details how we use the model to make such an assessment and describes our quantitative results.

3.6 Quantitative Results

In this section we detail the quantitative experiment that we run to assess the impact of the BAPCPA on the housing market crash. The experiment is based on the economy experiencing two shocks: a bankruptcy reform shock in 2005 and a housing crisis in 2007. We then compute the perfect foresight transition path of the economy in response to the following sequences of events assuming each event is unanticipated by the agents in our model:

1. *Actual Timeline*: In 2005 the U.S. economy experiences an unexpected change to the bankruptcy code that mimics the BAPCPA and then suffers a housing crisis in 2007.[93]

2. *Counterfactual*: The U.S. bankruptcy code is not altered in 2005 but the economy does experience a housing crisis in 2007.

The ability to run counterfactual exercises that incorporate general equilibrium effects through housing, mortgage, and unsecured debt prices, to isolate the impact of the BAPCPA on the subsequent housing crisis is a key benefit of constructing a quantitative model like that presented in this paper.

We model the BAPCPA shock as an unexpected and permanent change to the U.S. bankruptcy code as outlined in Section 2.4: an introduction of the income and asset means testing consistent with this reform. Following Chatterjee and Eyigungor (2011b) we model

[93]The fact that we model bankruptcy reform as unexpected in 2005 seems reasonable given our annual calibration. Although this act was originally introduced in Congress in 1998, it gained little political support until Republican majorities increased in Congress in 2004. It was ultimately passed by the U.S. Congress on April 14, 2005 and signed into law by President Bush on April 20th of that same year. Its provisions affected bankruptcy filings on or after October 17, 2005.

the housing crisis as an unexpected increase in the owner-occupied housing supply in 2007. Unlike these authors, however, we assume that this shock is temporary and dissipates over time, which implies that the housing market eventually returns to a state consistent with the initial stationary distribution.[94] We find that a 4% shock to the supply of owner-occupied housing produces a decline in house prices similar in magnitude to the decline in the S&P Case-Shiller 20-City Home Price Index between January 2007 and January 2009. The size of our housing supply shock is also in line with the empirical estimates of excess housing supply reported by McNulty (2009).

To capture the fact that housing prices remain below their peak in 2006, we also include an exogenous credit wedge that increases the cost of issuing mortgages in the periods immediately following the housing market crash. A substantial rise in credit spreads during this period is documented in Hall (2011), who finds increases in various spreads on the order of 1.0 - 3.7%. In particular, we model this wedge as an additional spread – above the risk-free rate – that the creditor requires, represented by α_t. The effect of the credit wedge on the mortgage pricing equations is demonstrated in Section 2.3.2. To match the upper-end of Hall (2011)'s estimates, we set $\alpha_t = 0.035$ for $t = 2008, \ldots, 2012$, and then allow this wedge to slowly decline back to zero by 2020. In sum, these modeling assumptions imply that the housing supply and credit markets return to their initial standing by 2020.

While solving the counterfactual perfect foresight transition path is relatively straightforward – given that the economy only experiences one unexpected shock – solving the transition under the actual timeline is more complicated. To solve for this transition, we have to compute two different transitions, and then combine the results from each to form the actual sequence of events. First, we compute a transition path of our economy starting from the pre-BAPCPA stationary distribution that experiences a bankruptcy reform-only shock in 2005, and then transitions to the post-BAPCPA steady state from there. This transition gives us the model statistics for 2005 and 2006 as well as the distribution of households over

[94] In particular, we assume that the owner-occupied housing supply remains elevated between 2007 and 2012 and then declines to its original value by 2020.

states entering 2007. The second transition starts with this distribution and subjects the economy to an unexpected 4% increase in the owner-occupied housing supply. The transition from this shock to the steady state for the post-BAPCPA economy with an owner-occupied housing supply consistent with the initial stationary distribution is then computed. This transition provides us with the statistics for the economy from 2007 onward. Appendix A for this chapter presents a more detailed description of our model solution, including solving for the economy's stationary distribution given a fixed set of parameters, and also computing the perfect foresight transitions described in this section.

Prior to computing the transitions, we have to determine the value for κ, which controls the cost of filing for bankruptcy for high-income households that pass the income means test after bankruptcy reform. Recall that a household that files for bankruptcy and is forced into a repayment plan due to their high income must pay $\kappa(y_t - \bar{y})$ in the current period, where $\bar{y}$ is the economy's median income. We choose κ to match Li et al. (2011)'s findings that, on impact, bankruptcy reform increased the default probability of households that owned their home for less than three years by 21.6%.[95] To compute this statistic, we have to solve the entire perfect foresight transition of the economy in response to only the BAPCPA shock in 2005. We then find the value of κ that produces an increase in the mortgage default probability of new homeowners of 21.6%. A value of $\kappa = 1.0$ most closely matches this statistic.

3.6.1 BAPCPA and the Housing Crisis

We first consider the effect of the BAPCPA on the U.S. economy on impact when it was introduced in 2005. In Table 3.4 we compare the model implied statistics in 2005 with the BAPCPA to the statistics taken from the pre-reform stationary distribution, along with the percentage change in each of these statistics in response to the BAPCPA.

[95] We compute this number from their findings that the probability of defaulting on a prime mortgage – which represented 81% of outstanding mortgages – increased by 23.4%, and the probability of defaulting on a subprime mortgage increased by 13.9%.

Table 3.4: BAPCPA on Impact in 2005

Statistic	Initial Steady State	BAPCPA	BAPCPA / Initial Steady State
P	1.0	1.0	0.0%
R	0.0520	0.0517	-0.5%
Bankruptcy Rate	1.69%	1.68%	-1.0%
Mortgage Default Rate	0.90%	1.13%	25.5%
R/P	5.20%	5.17%	-0.5%
Unsecured Debt-to-Income	11.48%	11.53%	0.5%

Upon implementation, the BAPCPA reduces the bankruptcy rate and produces a substantially higher mortgage default rate, which is consistent with the empirical literature. In particular, the bankruptcy filing rate falls by 1.0% and the mortgage default rate increases by 25.5%. However, the BAPCPA had minimal impact on house and rental prices. The price of owner-occupied housing is unchanged and the rental price only declines by 0.5% in response to the reform.[96] In addition, by reducing households incentives to file for bankruptcy and increasing the expected recovery for creditors in the event that they do, the BAPCPA generates a rise in the amount of unsecured borrowing, evidenced by a 0.5% increase in the unsecured debt-to-income ratio.

We now turn to the primary quantitative question of this paper: To what extent did the BAPCPA impact the housing market crash? We begin by assessing the ability of our model to match the severity of the housing crisis. The model statistics in 2007 are presented in Table 3.5.[97]

Evident from this table is that our model produces a housing crash that looks very much like the data. Specifically, the unexpected supply shock generates a substantial decline in house prices, by 25.5%, and a quintupling in the mortgage default rate – from 0.9% in the pre-crisis steady state to 4.5% in 2007 – which remains elevated for several years following the

[96]The fact that the owner-occupied house price is unchanged despite a substantial rise in the mortgage default rate suggests that there may be minimal feedback from mortgage defaults on house prices. This finding is in line with Chatterjee and Eyigungor (2011b) and will be discussed in more detail in Section 2.7.

[97]See Appendix B for graphs depicting the transitions for relevant model statistics.

Table 3.5: The Housing Crisis

Statistic	Initial Steady State	Value in 2007
P	1.00	0.75
R	0.052	0.044
Bankruptcy Rate	1.7%	3.3%
Mortgage Default Rate	0.9%	4.5%
R/P	5.2%	5.9%
Unsecured Debt-to-Income	11.5%	10.4%

crash. By comparison, the S&P Case-Shiller 20–City Home Price Index fell by 27.6% between January 2007 and January 2009, while the adjusted annual foreclosure rate reported by the Mortgage Bankers Association reached 4.2% in 2008. Moreover, this crash is accompanied by a pronounced rise in bankruptcy filing rates and a severe and protracted decline in unsecured borrowing relative to income similar to those observed in the data during this period.

The model also captures the empirical fact that the rent-price ratio rose during the housing crash.[98] Our model predicts a 13.5% increase in this statistic from 2004 to 2007. Thus, to clear both the owner-occupied and rental housing markets in response to the housing supply shock, the owner-occupied house price must decline more relative to the rental price.

The fact that our model has quantitative predictions that are consistent with *both* Li et al. (2011)'s findings that bankruptcy reform caused mortgage default rates to rise by 21.6% for new homeowners *and* the response of the housing and unsecured debt markets to the recent housing crisis, gives us confidence in its implications for the impact of the BAPCPA on the severity of the housing crisis. To analyze this question, we compare the implications of our model under the actual and counterfactual timelines.

Table 3.6 compares the economy with bankruptcy reform to the counterfactual economy by contrasting the statistics in 2007 between the two simulations. These results suggest that bankruptcy reform had little impact on the severity of the housing crisis, producing only modestly higher bankruptcy and mortgage default rates in 2007. The aggregate mortgage

[98]See Davis, Lehnert, and Martin (2008) for quarterly data on this ratio.

default rate is 2.7% higher in the economy that underwent bankruptcy reform in 2005. However, the BAPCPA had little impact on the aggregate prices for owner-occupied and rental housing. In the next section we discuss some intuition for these findings.

Table 3.6: BAPCPA's Impact on the Housing Crisis in 2007

Statistic	BAPCPA	No Reform	BAPCPA / No Reform
P	0.745	0.745	0.0%
R	0.044	0.043	1.1%
Bankruptcy Rate	3.27%	3.22%	1.4%
Mortgage Default Rate	4.52%	4.40%	2.7%
R/P	5.88%	5.82%	1.1%
Unsecured Debt-to-Income	10.38%	10.51%	-1.2%

3.6.2 Discussion

Figure 3.3 depicts decision rules for a homeowner during the housing crisis in 2007 that is underwater on their mortgage both with (right figure) and without (left figure) the reform. This figure displays these decision rules fixing a household's house size and mortgage payment presented in the endowment (x-axis) - asset (y-axis) space. For ease of interpretation, assets and income are normalized by mean income, and negative values on the y-axis represent debt relative to mean income.

There are several facts about the impact of the BAPCPA on homeowner decisions that are evident from this figure. First, reform reduces the region in which continuing to own is optimal (yellow). With reform, the household now finds it optimal to default on their mortgage (orange) or to sell their home (green) in several regions where they find it optimal to continue to own their home in the absence of reform. Second, the introduction of the BAPCPA leads to a reduction in the region where it is optimal to declare bankruptcy for high-income households. By decreasing the probability the homeowner files for bankruptcy, the BAPCPA causes unsecured creditors to reduce interest rates on their lending, leading

Figure 3.3: BAPCPA and Homeowner Decision Rules in 2007

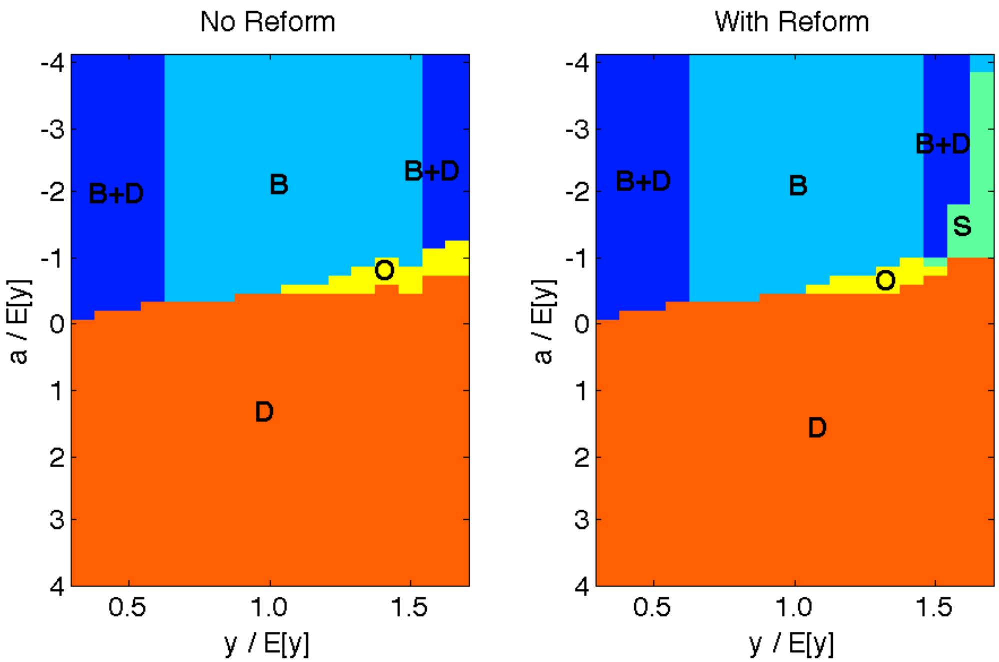

to an increase in risky lending for which bankruptcy-only and bankruptcy and default are non-trivial decisions.

Note that we can decompose the aggregate mortgage default rate into three components: households that simultaneously declare bankruptcy and default, non-excluded households that default on their mortgage, and excluded households that default on their mortgage. By considering these three components we can determine the quantitative importance of these effects in accounting for the higher mortgage default rate in the transition with the BAPCPA. We find that the reduced benefits to homeownership due to the BAPCPA, which lead to a rise in the number of non-excluded households that default on their mortgage, is the most quantitatively important factor. However, the rise in the number of households that declare bankruptcy and default on their mortgage due to the BAPCPA is also quantitatively important, while the impact of the BAPCPA on excluded homeowners is quantitatively insignificant. We now discuss these features in further detail.

Reduced Benefits to Homeownership

Prior to the BAPCPA, a homeowner with positive non-exempt home equity was able to borrow against that home equity in unsecured credit markets, as it served as collateral for unsecured debt in the event of bankruptcy. Consequently, homeowners with non-exempt home equity experienced a benefit of being able to borrow in unsecured credit markets at more favorable interest rates. This changed with the implementation of the BAPCPA and reduced the benefits of homeownership for some homeowners.

Consider a homeowner with income above the median ($y_t > \overline{y}$) and home equity above the exemption such that $5(y_t - \overline{y}) > HE_t(k_t, x_t) - \zeta > 0$. Table 3.7 depicts the amount recovered by creditors in the event of bankruptcy for a household with these characteristics if they are a homeowner or renter. Prior to the introduction of the BAPCPA, this household directly benefits by facing lower interest rates on unsecured borrowing by being a homeowner since, in the event of bankruptcy, the creditor recovers their non-exempt home equity (up to the face value of the bond). Following the introduction of the BAPCPA, creditor recovery no longer varies with the household's homeownership status, eliminating the benefit to this household of borrowing at lower interest rates because they own a home.

Table 3.7: Creditor Recovery and the BAPCPA

	Before Reform	After Reform
Homeowner	$HE_t(k_t, x_t) - \zeta$	$\kappa(y_t - \overline{y})$
Renter	0	$\kappa(y_t - \overline{y})$

This reduction in the benefit of homeownership is acute for homeowners with specific characteristics: relatively low current levels of home equity, expectations of high future income, and a desire to borrow in unsecured credit markets. We should expect to see households with these characteristics substituting away from homeownership in response to the BAPCPA. This intuition is reinforced by the decision rules depicted in Figure 3.4, as we see that homeowners with a higher current endowment now prefer to default on their mortgage (orange) or sell (green) as a result of the reform rather than continuing owning (yellow).

Figure 3.4: BAPCPA and Homeowner Decision Rules in 2007

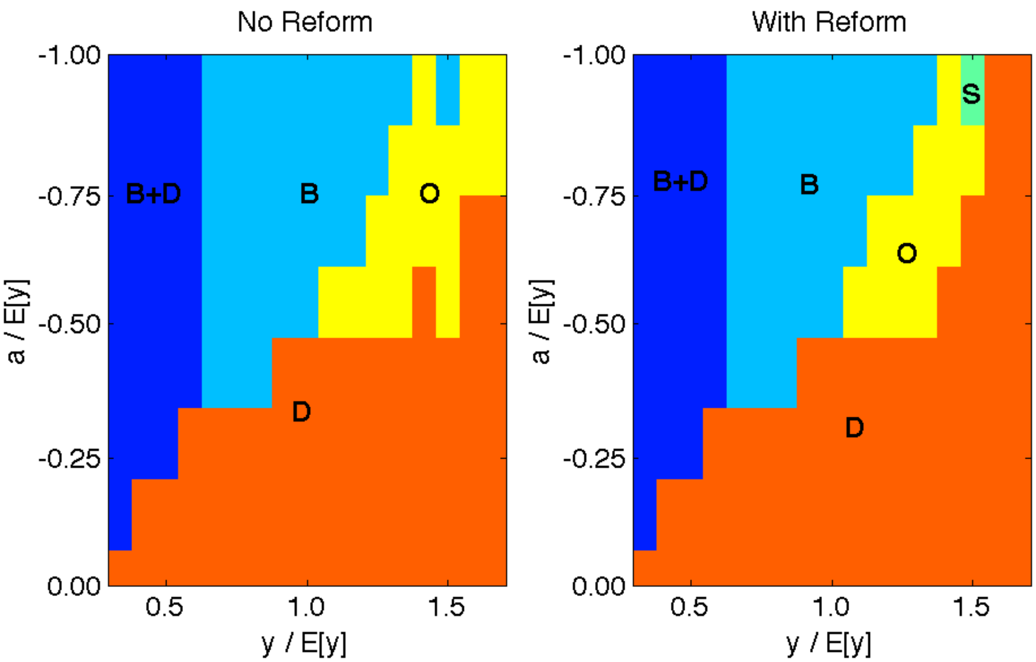

The expansion of the region in which households find it optimal to default on their mortgage is a quantitatively important determinant of the higher mortgage default rates observed with the BAPCPA compared to the economy that did not enact reform. At the peak of the crisis in 2007, an increase in the number of non-excluded households that default on their mortgage and do not declare bankruptcy accounts for 91.8% of the higher mortgage default rate with reform. Moreover, this channel accounts for more than 90% of the higher mortgage default rate with reform from 2007 through 2014, on average.

Thus, because the BAPCPA decreased the relative benefit of homeownership by reducing the dependence of unsecured interest rates on a household's homeownership status, underwater homeowners who would otherwise have decided to remain in their home now prefer to default on their mortgage. This effect tends to increase the mortgage default rate during the housing crisis.

Looser Unsecured Credit Lending Standards

The BAPCPA significantly increased the cost of filing for bankruptcy for high-income

Figure 3.5: BAPCPA and Homeowner Decision Rules in 2005

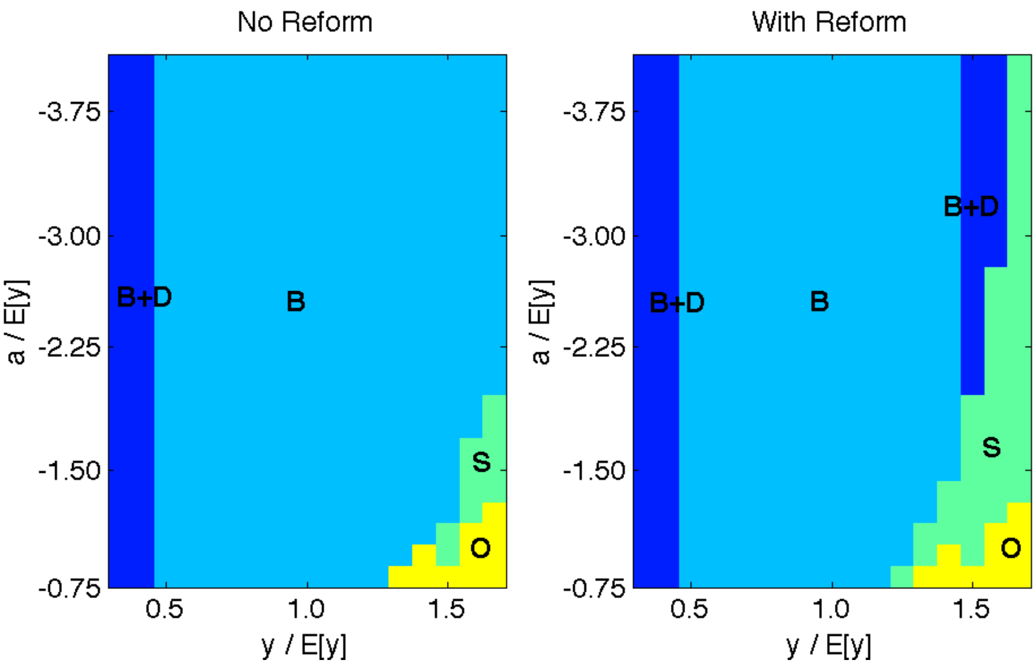

households. This led to a dramatic reduction in the region in which households find it optimal to declare bankruptcy, as seen in Figure 3.5. This figure depicts decision rules for a homeowner in 2005 in the case with and without bankruptcy reform. Prior to the BAPCPA, this household would declare bankruptcy with near certainty (i.e. across all endowments) for large amounts of debt. This high probability of bankruptcy leads to prohibitively high interest rates for households that desire to borrow that amount of debt. As a result, it is unlikely that households would choose to borrow an amount that causes them to declare bankruptcy and default.

With the implementation of the BAPCPA, the probability that a high-income household declares bankruptcy falls dramatically. Unsecured creditors thus expect to be repaid in full with a higher probability after the reform, leading to lower interest rates and increased lending to high-income households in regions in which they may declare bankruptcy and default on their mortgage in the following period (dark blue).[99]

[99]Recall that the creditor is repaid in full if the household sells their home. A reduction in the area in which households find it optimal to declare bankruptcy and an increase in the optimal sell region implies higher expected returns for the creditor.

An increase in the degree of lending for levels of debt where households are more likely to get mapped into a region where they find it optimal to simultaneously declare bankruptcy and default leads to higher mortgage default rates. Quantitatively, we find that this effect is important, although it is not as sizable as the higher default rates caused by the lower benefit of homeownership just discussed. Specifically, an increase in the number of households that simultaneously declare bankruptcy and default due to a rise in risky lending accounts for 8.2% of the difference between mortgage default rates with and without reform in 2007. This fraction is slightly higher than the average impact for this channel, which was 7.2% from 2007 through 2014. Therefore, the increased complementarity between bankruptcy and mortgage default that results from a rise in risky unsecured lending is quantitatively important in accounting for the higher mortgage default rate under BAPCPA.

Tighter Mortgage Lending Standards

An important feature of our model is how bankruptcy and default incentives are fully reflected in the terms at which households can borrow in credit markets. In response to the increased incentive for households to default on their mortgage or sell, mortgage lenders expect to receive a lower return on loans to new homebuyers.[100] To continue to break even in expectation despite these changing incentives, mortgage lenders must tighten their lending standards, which in our model is accomplished by raising interest rates on those households that are now more likely to either default or sell in the future as a result of the reform.

Table 3.8 depicts how several characteristics of new homebuyers change in response to the introduction of the BAPCPA. Most of these metrics move in very intuitive directions and imply a tightening of mortgage lending standards. For example, the average house size, initial loan-to-income, loan-to-value, and mortgage payment-to-income ratios all decline. Although the fact that the average income of new homeowners declines may at first appear to contradict a tightening of mortgage lending standards, this result is also intuitive. Since

[100]More specifically, lenders face higher credit and prepayment risk after the BAPCPA.

high-income households experience the higher incentive to default on their mortgage, in equilibrium, mortgage lenders tighten standards for high-income households relatively more in response to the BAPCPA, leading to lower average income for new homeowners.

Table 3.8: BAPCPA's Impact on Mortgage Lending Standards

Statistic	Change
House Size	-4.0%
Income	-1.4%
LTI	-12.8%
LTV	-0.7%
MTI	-4.0%

Tighter mortgage lending standards tend to reduce the probability households will find it optimal to default on their mortgage after the BAPCPA, offsetting some of the increased incentives homeowners have to default on their mortgage. On net, these effects nearly offset each other, implying a relatively small role for the BAPCPA in the severity of the housing crisis.

3.7 Mortgage Cram Down and the Housing Crisis

Given the depth and protracted nature of the recent housing crisis, there has been extensive discussions in policy and academic circles about policy initiatives aimed to reduce foreclosures and stabilize house prices. One such proposal, often referred to as mortgage cram down, has received attention in the economics and law literature.[101] Under this policy, homeowners with negative home equity are able to treat the portion of their mortgage that exceeds the value of their home as unsecured debt, which can then be discharged through the bankruptcy process. Thus, the value of the mortgage is reduced until the homeowner is no longer underwater, which should reduce the number of homeowners that find it optimal to default on their mortgage. To the extent that mortgage defaults have a feedback effect

[101]For examples, see White and Zhu (2008), Levitin (2009), and Scarberry (2010).

on prices, lower default rates may also positively impact house prices. In this section we use our framework to evaluate the impact of this policy on the severity of the housing crisis.

To implement this policy in our model, households that are underwater on their mortgage have their mortgage payments reduced until the present value of these payments is equal to the value of their home net of selling costs. Mortgage lenders internalize this policy change, recognizing that the stream of payments they expect to receive will be reduced in the following period if the household declares bankruptcy and qualifies for a mortgage cram down.

We then compute a perfect foresight transition in our economy under the following timeline: In 2005 the BAPCPA is implemented. The economy experiences an unexpected housing crisis in 2007, in which the time series for the supply of owner-occupied housing and credit wedges are the same as in the previous transitions. In addition, the mortgage cram down policy is unexpectedly and permanently implemented, allowing households with negative home equity to reduce their payments until they are no longer underwater from 2007 onward.

Table 3.9 compares statistics in 2007 for transitions for the model with and without the cram down policy.[102] These results suggest that the cram down policy generates a reduction in the mortgage default rate relative to an economy that did not enact this policy. Specifically, the mortgage default rate is 2.1% lower in the economy that enacted cram down. This reduction in mortgage defaults comes at the expense of a 6.2% higher bankruptcy rate. Notably, however, this policy has very little impact on aggregate prices.

To gain intuition for these results, we consider how homeowner decision rules are altered by this policy. Figure 3.6 depicts homeowner decision rules in 2007 with and without the cram down policy in effect. Evident from this figure is a dramatic reduction in the region where the household finds it optimal to default, either by simultaneously declaring bankruptcy and defaulting on their mortgage or by only defaulting on their mortgage. By

[102]Not presented in this table is the fact that the cram down policy is well-utilized by homeowners. Specifically, 18.3% of homeowners that declare bankruptcy in 2007 take advantage of the benefits of cram down, and this percentage rises over the next several years after the onset of the crisis.

Table 3.9: Cram Down's Impact on the Housing Crisis in 2007

Statistic	No Cram Down	Cram Down	Cram Down / No Cram Down
P	0.745	0.745	0.0%
R	0.044	0.043	-0.9%
Bankruptcy Rate	3.27%	3.47%	6.2%
Mortgage Default Rate	4.52%	4.43%	-2.1%
R / P	0.059	0.058	-0.9%

allowing homeowners to reduce their mortgage payments and no longer be underwater on their mortgage by declaring bankruptcy, the cram down policy incentivizes homeowners to declare bankruptcy in regions where they would find it optimal to default on their mortgage in the absence of the policy. Quantitatively, this produces a slightly lower mortgage default rate and a substantial increase in the bankruptcy filing rate during the peak of the crisis.

An inherent difficulty with the cram down policy is also evident from this figure. In particular, it is difficult to identify which households with negative home equity will actually default on their mortgage. Recall that while negative home equity is a necessary condition for mortgage default, it is not sufficient. Therefore, some households with negative home equity may take advantage of the lower mortgage payments under cram down even though they would not have defaulted on their mortgage in the absence of this program. In this case, the policy imposes losses on mortgage lenders and, possibly, unsecured creditors, without even reducing the number of mortgage defaults.

This can be seen in Figure 3.6, as some households that would have found it optimal to stay in their home and not declare bankruptcy in the absence of cram down, instead find it optimal to declare bankruptcy and take advantage of lower mortgage payments when cram down is available. This is true for the states that switch from own (yellow) to bankruptcy (light blue). The impact of cram down on the mortgage default rate is limited by the fact that homeowners have asymmetric information about their willingness to stay in their home when they have negative home equity. As a result, we find that cram down only produces a

Figure 3.6: Cram Down and Homeowner Decision Rules in 2007

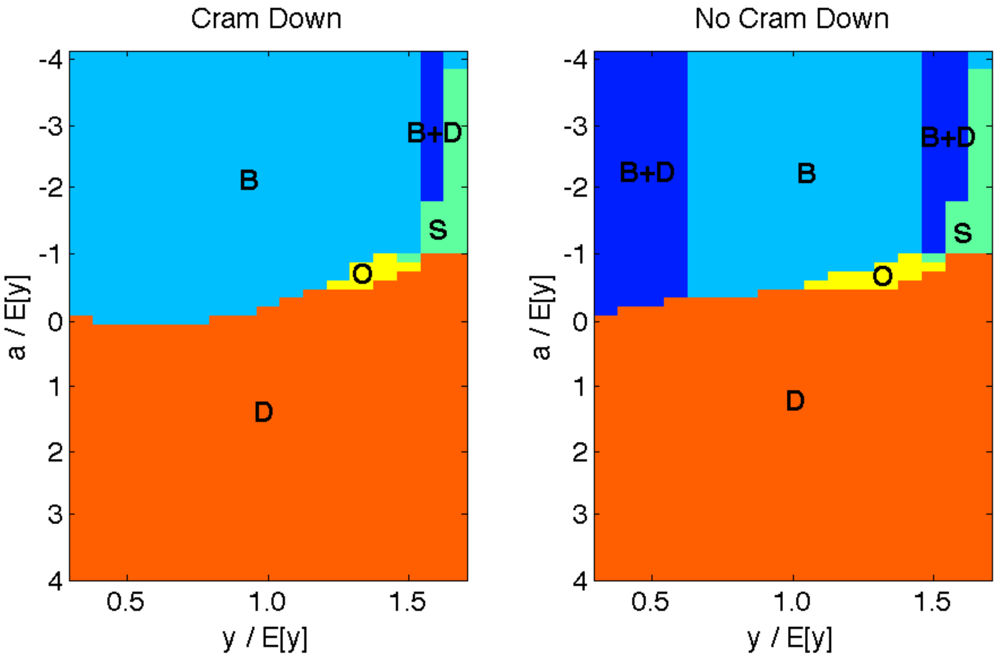

2.1% reduction in the mortgage default rate at the peak of the crisis.

To the extent that an objective of the cram down policy is to stabilize and raise owner-occupied home prices by decreasing the supply of housing produced by homeowners that default on their mortgage, our analysis suggests this policy would not have been successful. The owner-occupied price path for the model with cram down is virtually identical to an economy that does not enact this policy. And although rental prices fall more in the economy with cram down, due to the fact that demand for rental housing is lower because more homeowners choose to stay in their homes, this effect is also minimal. These results are in line with Chatterjee and Eyigungor (2011b), who find that the feedback effect of defaults on house prices is rather limited. In particular, these authors find that the drop in the owner-occupied house price at the peak of the crisis would still be 84% of the actual decline in house prices if mortgage defaults were completely disallowed. Therefore, our analysis suggests that although the mortgage cram down policy may have reduced mortgage default rates during the housing crisis, its effect on aggregate prices would have been minimal.

3.8 Conclusion

This paper investigates whether the BAPCPA of 2005 exacerbated the recent housing crisis in the context of a quantitative-theoretic, equilibrium model of unsecured debt and mortgage markets. We conclude that, although the BAPCPA did produce higher mortgage default rates, it had minimal effect on the severity of the housing crisis and virtually no impact on aggregate house prices.

Understanding how unsecured debt and mortgage prices respond to new incentives to declare bankruptcy and default in response to reform is key to our findings. In particular, the BAPCPA increased homeowner incentives to default by reducing the benefit homeowners derived from borrowing against their non-exempt home equity in unsecured credit markets, leading marginal homeowners to switch from owning to defaulting on their mortgage. Moreover, a rise in risky unsecured lending brought about by a reduction in the likelihood of bankruptcy for high-income households, increased the probability that a homeowner enters a region in which they find it optimal to declare bankruptcy and default. These incentives that tend to increase the mortgage default rate were offset by the fact that mortgage lenders tightened lending standards for new homeowners by requiring the household to purchase a smaller home and/or undertake a mortgage with a lower loan-to-income, loan-to-value, and mortgage payment-to-income ratio. Tighter mortgage standards, in turn, reduce the likelihood that households will find themselves in a position in which they prefer to default on their mortgage. On net, these incentives tend to offset, implying that the BAPCPA had a limited impact on the severity of the housing crisis.

3.9 Appendix A: Solution Algorithm

In this appendix we detail the solution algorithm for our model. We begin by describing how to solve for the steady state and then discuss how we solve for the perfect foresight transition paths of our economy under the actual sequence of events and the counterfactual sequence in which there is no bankruptcy reform.

3.9.1 Solving for the Stationary Distribution

Solving for the initial stationary distribution (i.e. prior to bankruptcy reform and the housing shock) of our economy entails fixing prices for owner-occupied and rental housing and solving the following fixed point problem in our economy without bankruptcy reform. To do so we first set $P = 1.0$ and $R = 0.052$, which is in line with the historical rental price to owner-occupied ratio in the U.S. economy from 1960 to 2000 documented in Davis et al. (2008). The solution algorithm is as follows:

1. Guess initial values for $V(k, x, a, y), X(k, x, a, y), q(k, x, a, y), m(k, x, a, y), m^X(k, x, a, y)$. Denote these initial guesses with a 0 subscript.

2. Taking these guesses as given, compute household optimal decision rules. From these optimal decisions, compute implied values for $V_1(k, x, a, y)$, $X_1(k, x, a, y)$, $q_1(k, x, a, y)$, $m_1(k, x, a, y)$, and $m_1^X(k, x, a, y)$ from the functional equations outlined in Section 2.3.

3. Compute the maximum of the absolute value of the differences between the initial guesses for these functions (denoted 0) and the implied values for these functions (denoted 1) given the initial guesses. If this maximum absolute difference is less than a pre-specified tolerance level, stop value function iteration, and we have found the fixed point of the operator. Conversely, if the maximum absolute difference exceeds the tolerance level, use the implied values computed in this step as the initial guess in step 2.

4. Iterate on 2 and 3 until the maximum difference is less than the tolerance level.

5. Once value function iteration is completed, we use the resulting household optimal decision rules to simulate an economy of 30 million households over 500 periods to compute the stationary distribution. The initial supplies for owner-occupied and rental housing are determined by setting these values equal to their respective demands implied by this initial stationary distribution. Label these initial housing supplies as K_0 and H_0 respectively.

Now we have the pre-reform, pre-housing shock stationary distribution of our economy. To compute the stationary distribution under changes in the bankruptcy code, we first set the supply of owner-occupied and rental housing equal to our desired values. Then, given initial guesses for P and R, we solve for the implied demand for both types of housing using the five steps just outlined. We then adjust P and R until both housing markets clear.

3.9.2 Solving for the Perfect Foresight Transition

This section details how we solve for the perfect foresight transition in our economy under both the actual sequence of events and the counterfactual sequence of events in which the economy experiences a housing crisis in 2007 – mimicked by an unanticipated increase in the supply of owner-occupied housing – but did not implement bankruptcy reform in 2005. We begin with the actual sequence of events, as this transition is more complicated than the counterfactual.

Actual Sequence of Events

Under this transition the economy experiences an unexpected change to the bankruptcy code in 2005 and an unexpected shock to the owner-occupied housing supply in 2007. In the first few years following the housing crash, mortgage lenders experience a credit wedge that raises the cost of issuing mortgages. To compute the perfect foresight transition under this sequence of events, we must actually compute two different transitions and then combine the results from each. Recall that we assume the following timeline:

1. 2004: Economy is in pre-bankruptcy reform steady state with housing supplies given by K_0 and H_0.

2. 2005: An unexpected and permanent change to the bankruptcy code occurs.

3. 2007: An unexpected increase in the supply of owner-occupied housing occurs, such that the supply of owner-occupied housing becomes $\tilde{K} = 1.04K_0$ in 2007. This elevated housing supply persists until 2012 and then slowly returns to its initial value by 2020.

4. 2008: A credit wedge equal to 0.035 raises the cost of issuing new mortgages. Like the housing supply shock, this wedge persists until 2012 and then returns to zero by 2020.

Thus, to correctly compute the economy's transition given this sequence of events, we must solve for two perfect foresight transitions. The first is for an economy that begins in the pre-bankruptcy reform steady state with housing supplies given by K_0 and H_0 and experiences an unexpected and permanent change to the bankruptcy code in 2005. From this transition we derive the relevant statistics for this economy in 2005 and 2006 in addition to the distribution of households entering 2007.

Next, we compute the transition for the post-bankruptcy reform economy that experiences a housing shock in 2007, beginning from the distribution implied by the bankruptcy reform-only transition in 2007, to the steady state for the post-bankruptcy reform economy with housing supplies equal to their initial values K_0 and H_0.

We now detail how we compute each of these transitions.

Reform-Only Transition

We assume that the economy takes $T = 40$ years to transition to its new steady state after experiencing an unexpected shock.[103] The steps for computing this transition are then:

[103]This assumption is confirmed if the economy has successfully transitioned to the terminal steady state

1. Using the algorithm outlined in Section 2.9.1, solve for the steady states of the economy both pre- and post-bankruptcy reform with the housing supplies equal to their initial values K_0 and H_0.

2. Set terminal values for $V_T(k, x, a, y)$, $q_T(k, x, a, y)$, $X_T(k, x, a, y)$, $m_T(k, x, a, y)$, and $m_T^X(k, x, a, y)$ equal to their values in the post-reform steady state.

3. Set $K_t = K_0$, $H_t = H_0$, and $\alpha_t = 0$ for all t.

4. Guess a sequence of owner-occupied house prices and rental prices $\{P_t, R_t\}_{t=1}^{39}$.

   (a) Use the decision rules and pricing functions from the post-reform steady state to back out the $t = 39$ pricing functions from the functional equations defining these pricing functions outlined in Section 2.3.2. Given these pricing functions and the guessed house prices, compute optimal household decisions for $t = 39$ under the assumption that bankruptcy reform is in place.

   (b) Repeat this step from $t = 38$ to $t = 1$, documenting household decision rules at each point in time along the transition, to compute the sequence of decision rules and pricing functions along the way.

5. Next, starting from the stationary distribution defined by the pre-bankruptcy reform economy, simulate the distribution of people each period given the sequences of decision rules determined in (b) from $t = 1$ to $t = 40$.

6. From the distribution of people, compute demand for owner-occupied and rental housing for each period, and compute excess demand for both types of housing at each point in time.

7. If excess housing demand and supply are below some pre-specified threshold for both owner-occupied and rental housing at each point in time along the transition, then we

in 40 years. If this is not the case, we increase T.

have successfully solved for the perfect foresight transition. If not, adjust P_t and R_t along the transition, increasing (decreasing) each slightly if the excess demand (supply) for that form of housing is too high at period t. Return to 4 with this new guess for the sequences of prices.

This algorithm gives us the perfect foresight transition of the economy under the assumption that the economy only experienced bankruptcy reform but no housing crisis. We use the statistics from the first two periods of this transition, corresponding to 2005 and 2006, in the final transition under the actual sequence of events.

Full Transition

To compute the full transition, however, the economy must experience a housing supply shock in 2007. To compute the statistics along the transition during and after the housing supply shock, we follow the algorithm just outlined, but use the value and pricing functions from the post-reform steady state as the terminal values (period T) and the distribution that is implied from the second period of the reform-only transition as the initial distribution entering 2007. We also use the time series for the housing supply and credit wedges outlined in Section 2.6.

Counterfactual

The counterfactual experiment assumes the following timeline:

1. 2004-2006: Economy is in pre-bankruptcy reform steady state with housing supplies given by K_0 and H_0.

2. 2007: An unexpected increase in the supply of owner-occupied housing occurs, such that the supply of owner-occupied housing becomes $\tilde{K} = 1.04K_0$ in 2007. This elevated housing supply persists until 2012 and then slowly returns to its initial value by 2020.

3. 2008: A credit wedge equal to 0.035 raises the cost of issuing new mortgages. Like the housing supply shock, this wedge persists until 2012 and then returns to zero by 2020.

To solve for the counterfactual transition we follow the detailed transition algorithm outlined in the preceding section, but use the steady state corresponding to an economy that does not undergo bankruptcy reform for the terminal values for value and pricing functions. The initial distribution of households is taken as the pre-reform stationary distribution with housing supplies given by K_0 and H_0.

3.10 Appendix B: Transitions

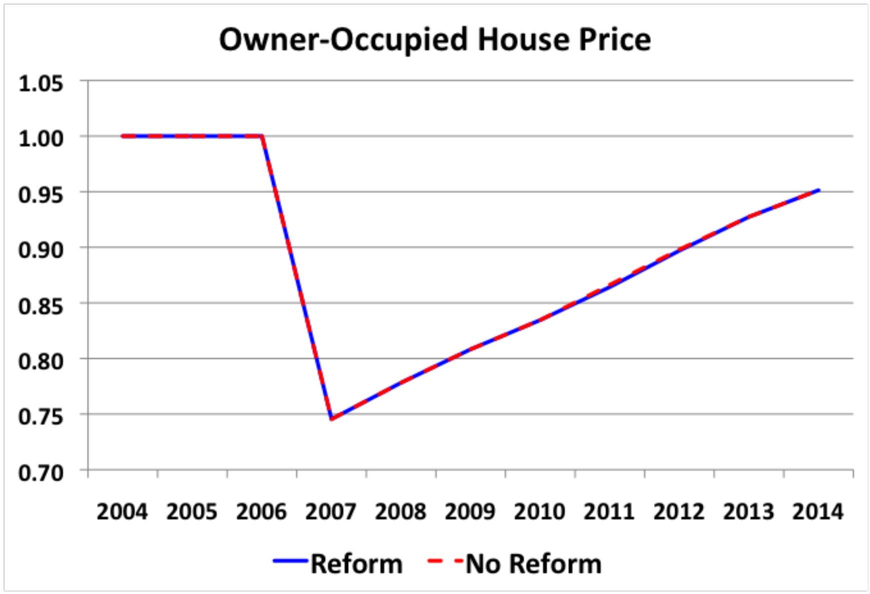

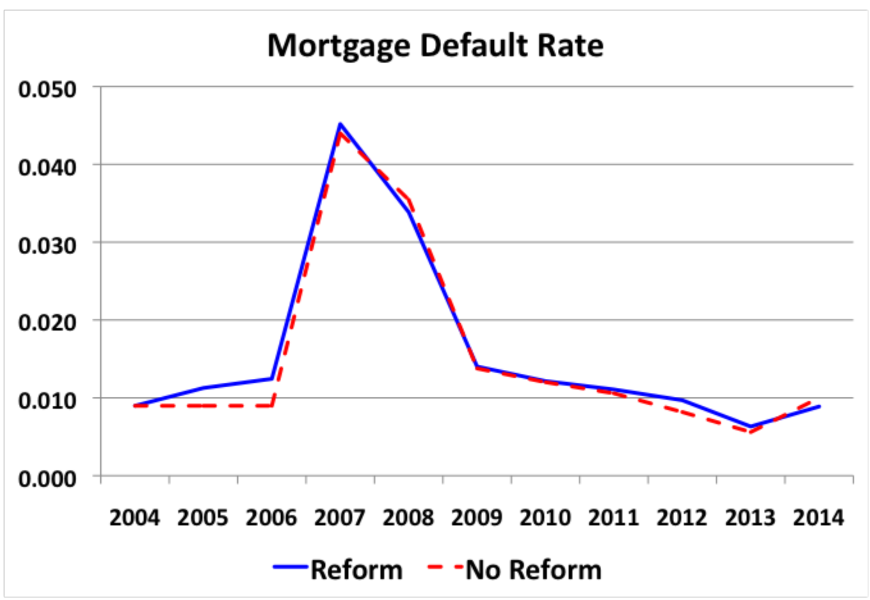

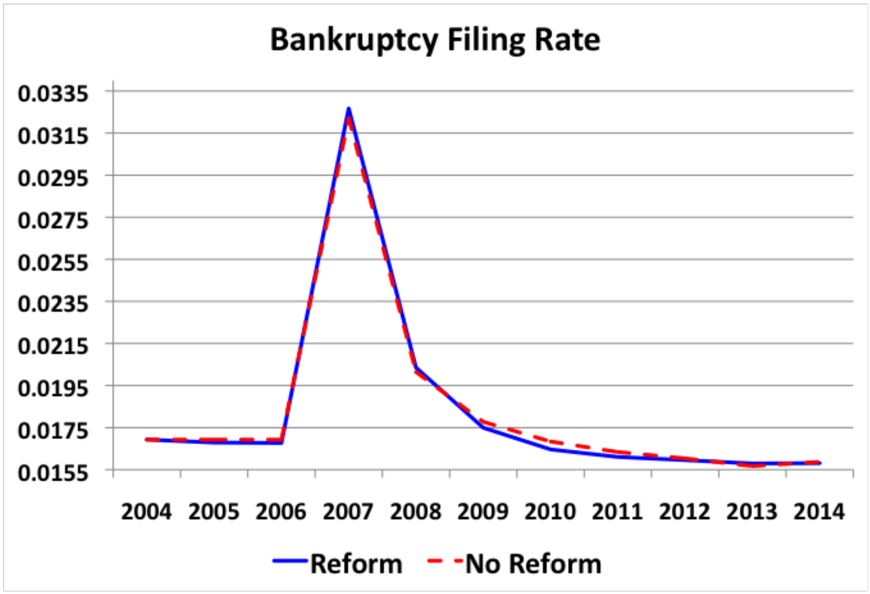
Bankruptcy Filing Rate
0.0335
0.0315
0.0295
0.0275
0.0255
0.0235
0.0215
0.0195
0.0175
0.0155
2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014
Reform
No Reform

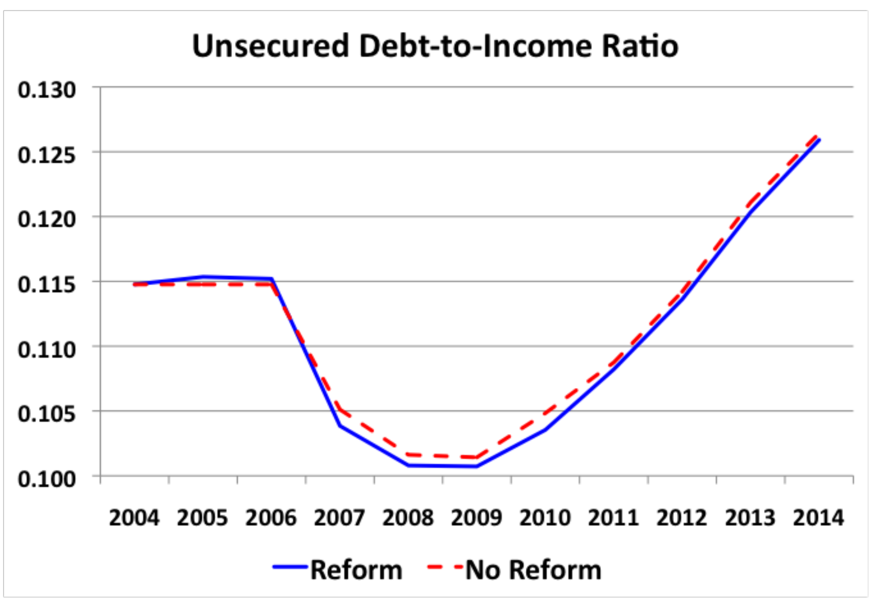
Unsecured Debt-to-Income Ratio
0.130
0.125
0.120
0.115
0.110
0.105
0.100
2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014
Reform
No Reform

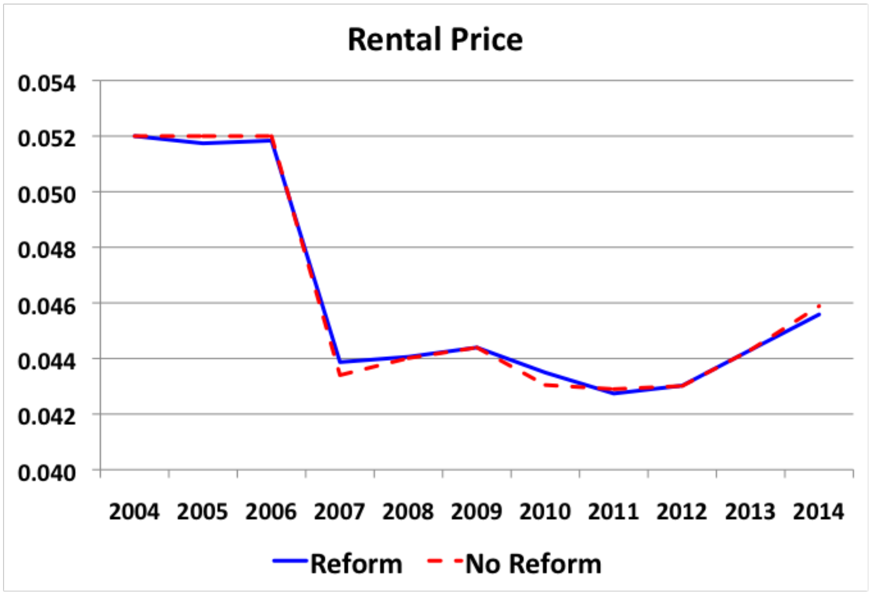
Rental Price
0.054
0.052
0.050
0.048
0.046
0.044
0.042
0.040
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
Reform
No Reform

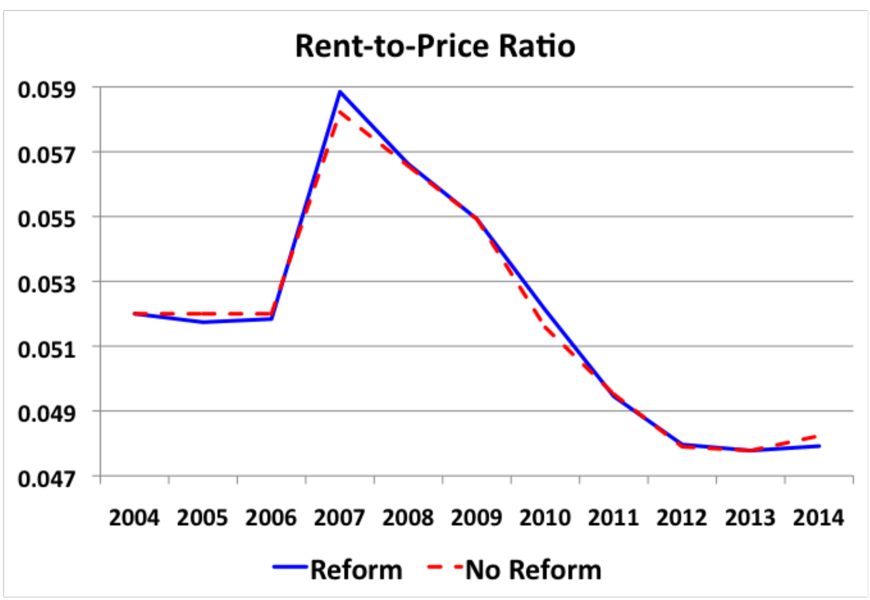
Rent-to-Price Ratio
0.059
0.057
0.055
0.053
0.051
0.049
0.047
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
Reform
No Reform

References

Abbring, J. H. and J. R. Campbell (2005, May). A firm's first year. Tinbergen Institute Discussion Papers 05-046/3, Tinbergen Institute.

Abowd, J. M. and O. Ashenfelter (1981). Anticipated unemployment, temporary layoffs, and compensating wage differentials. In *Studies in Labor Markets*, pp. 141–186. National Bureau of Economic Research, Inc.

Abowd, J. M. and D. Card (1989, March). On the covariance structure of earnings and hours changes. *Econometrica 57*(2), 411–45.

Adam, K., A. Marcet, and J. P. Nicolini (2008, February). Stock market volatility and learning. Working Paper Series 862, European Central Bank.

Arellano, C. (2008, June). Default risk and income fluctuations in emerging economies. *American Economic Review 98*(3), 690–712.

Arellano, C. and A. Ramanarayanan (2010). Default and the maturity structure in sovereign bonds. Technical report.

Artuç, E. and J. McLaren (2012, November). Trade policy and wage inequality: A structural analysis with occupational and sectoral mobility. Working Paper 18503, National Bureau of Economic Research.

Barron, J. M., G. Elliehausen, and M. E. Staten (2000). Monitoring the household sector with aggregate credit bureau data. *Business Economics*.

Barron, J. M. and M. Staten (2003). The value of comprehensive credit reports: Lessons from the U.S. experience summary.

Bassett, W. F., M. B. Chosak, J. C. Driscoll, and E. Zakrajsek (2011). Changes in bank lending standards and the macroeconomy. Working paper, Federal Reserve Board.

Benati, L. and P. Surico (2009, September). VAR analysis and the Great Moderation. *American Economic Review 99*(4), 1636–52.

Bound, J., C. Brown, G. J. Duncan, and W. L. Rodgers (1994, July). Evidence on the validity of cross-sectional and longitudinal labor market data. *Journal of Labor Economics 12*(3), 345–68.

Boyes, W. J. and R. L. Faith (1986, April). Some effects of the Bankruptcy Reform Act of 1978. *Journal of Law & Economics 29*(1), 139–49.

Brevoort, K. P. and C. R. Cooper (2010). Foreclosure's wake: The credit experiences of individuals following foreclosure. Technical report.

Buckley, F. H. and M. F. Brinig (1998, January). The bankruptcy puzzle. *Journal of Legal Studies 27*(1), 187–207.

Buera, F. J. and B. Moll (2012, January). Aggregate implications of a credit crunch. NBER Working Papers 17775, National Bureau of Economic Research, Inc.

Caju, P. D., G. Ktay, A. Lamo, D. Nicolitsas, and S. Poelhekke (2010, 04-05). Inter-industry wage differentials in E.U. countries: What do cross-country time varying data add to the picture? *Journal of the European Economic Association 8*(2-3), 478–486.

Carceles-Poveda, E. and C. Giannitsarou (2008, July). Asset pricing with adaptive learning. *Review of Economic Dynamics 11*(3), 629–651.

Carroll, C. D. and A. A. Samwick (1997, September). The nature of precautionary wealth. *Journal of Monetary Economics 40*(1), 41–71.

Chatterjee, S., D. Corbae, M. Nakajima, and J.-V. Rios-Rull (2007, November). A quantitative theory of unsecured consumer credit with risk of default. *Econometrica 75*(6), 1525–1589.

Chatterjee, S. and B. Eyigungor (2011a). Maturity, indebtedness, and default risk. Technical report.

Chatterjee, S. and B. Eyigungor (2011b). A quantitative analysis of the U.S. housing and mortgage markets and the foreclosure crisis. *Federal Reserve Bank of Philadelphia Working Paper*.

Christie, L. (2010, April 10). How foreclosure impacts your credit score. *CNNMoney.com*.

Cogley, T. (2002, March). Idiosyncratic risk and the equity premium: evidence from the consumer expenditure survey. *Journal of Monetary Economics 49*(2), 309–334.

Cogley, T. and T. J. Sargent (2008, 02). Anticipated utility and rational expectations as approximations of Bayesian decision making. *International Economic Review 49*(1), 185–221.

Corbae, D. and E. Quintin (2011, February). Mortgage innovation and the foreclosure boom.

Cubas, G. and P. Silos (2012, October). Comparative advantage and risk premia in labor markets. *Federal Reserve Bank of Atlanta Working Paper 2012-15*.

Dahl, M., T. DeLeire, and J. A. Schwabish (2011). Estimates of year-to-year volatility in earnings and in household incomes from administrative, survey, and matched data. *Journal of Human Resources 46*(4).

Davis, M. A., A. Lehnert, and R. F. Martin (2008, 06). The rent-price ratio for the aggregate stock of owner-occupied housing. *Review of Income and Wealth 54*(2), 279–284.

Davis, M. A. and F. Ortalo-Magne (2011, April). Household expenditures, wages, rents. *Review of Economic Dynamics 14*(2), 248–261.

Dickens, W. T. and L. F. Katz (1987). Inter-industry wage differences and theories of wage determination. NBER Working Papers 2271, National Bureau of Economic Research, Inc.

Dillon, E. (2012, July). Risk and return tradeoffs in lifetime earnings. mimeo, Arizona State University.

Dynan, K. E., D. W. Elmendorf, and D. E. Sichel (2012, July). The evolution of household income volatility. Technical report.

Eusepi, S. and B. Preston (2011, October). Expectations, learning, and business cycle fluctuations. *American Economic Review 101*(6), 2844–72.

Evans, G. and S. Honkapohja (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press.

Fay, S., E. Hurst, and M. J. White (2002, June). The household bankruptcy decision. *American Economic Review 92*(3), 706–718.

Feinberg, R. M. (1981). Earnings-risk as a compensating differential. *Southern Economic Journal*, 156–163.

Friedman, M. and S. Kuznets (1954). *Income from Independent Professional Practice*. NBER.

Gambetti, L., E. Pappa, and F. Canova (2008, 03). The structural dynamics of U.S. output and inflation: What explains the changes? *Journal of Money, Credit and Banking 40*(2-3), 369–388.

Gibbons, R. and L. F. Katz (1992, July). Does unmeasured ability explain inter-industry wage differentials? *Review of Economic Studies 59*(3), 515–35.

Gourieroux, C., A. Monfort, and E. Renault (1993, Suppl. De). Indirect inference. *Journal of Applied Econometrics 8*(S), S85–118.

Gross, D. B. and N. S. Souleles (2002, March). An empirical analysis of personal bankruptcy and delinquency. *Review of Financial Studies 15*(1), 319–347.

Gruber, J. and R. Martin (2003). Precautionary savings and the wealth distribution with illiquid durables. Technical report.

Guerrieri, V. and G. Lorenzoni (2011, November). Credit crises, precautionary savings, and the liquidity trap. NBER Working Papers 17583, National Bureau of Economic Research, Inc.

Guvenen, F. (2009, January). An empirical investigation of labor income processes. *Review of Economic Dynamics 12*(1), 58–79.

Hacker, J. S. (2006). *The Great Risk Shift*. Toronto: Oxford University Press.

Hall, R. E. (2011, 05). The high sensitivity of economic activity to financial frictions. *Economic Journal 121*(552), 351–378.

Hatchondo, J. C. and L. Martinez (2009, September). Long-duration bonds and sovereign defaults. *Journal of International Economics 79*(1), 117–125.

Heathcote, J., K. Storesletten, and G. L. Violante (2008). Insurance and opportunities: A welfare analysis of labor market risk. *Journal of Monetary Economics 55*(3), 501–525.

Heathcote, J., K. Storesletten, and G. L. Violante (2010, August). The macroeconomic implications of rising wage inequality in the united states. *Journal of Political Economy 118*(4), 681–722.

Herkenhoff, K. and L. Ohanian (2012). Foreclosure delays and unemployment. Technical report.

Hsieh, C.-T., E. Hurst, C. Jones, and P. Klenow (2012, May). The allocation of talent and u.s. economic growth. mimeo.

Jeske, K., D. Krueger, and K. Mitman (2011, October). Housing and the macroeconomy: The role of bailout guarantees for government sponsored enterprises. NBER Working Papers 17537, National Bureau of Economic Research, Inc.

Jovanovic, B. (1979, October). Job matching and the theory of turnover. *Journal of Political Economy 87*(5), 972–90.

Kambourov, G. and I. Manovskii (2008, February). Rising occupational and industry mobility in the united states: 1968-97. *International Economic Review 49*(1), 41–79.

Kambourov, G. and I. Manovskii (2012, September). A cautionary note on using (march) current population survey and panel study of income dynamics data to study worker mobility. *Macroeconomic Dynamics FirstView*, 1–23.

Kaplan, G. (2012, June). Moving back home: Insurance against labor market risk. *Journal of Political Economy 120*(3), 446–512.

Kim, C.-J. and C. R. Nelson (1999, November). Has the U.S. economy become more stable? A Bayesian approach based on a markov-switching model of the business cycle. *The Review of Economics and Statistics 81*(4), 608–616.

Kreps, D. (1998). Anticipated utility and dynamic choice. *Frontiers of Economic Theory*.

Krueger, A. B. and L. H. Summers (1988, March). Efficiency wages and the inter-industry wage structure. *Econometrica 56*(2), 259–93.

Lagakos, D. and G. L. Ordonez (2011). Which workers get insurance within the firm? *Journal of Monetary Economics 58*(6), 632–645.

Lebergott, S. (1964). *Manpower in Economic Growth.* New York: McGraw-Hill.

Leigh, J. P. (1983). Job choice across industries when earnings are uncertain. *Quarterly Review of Economics and Business 23*(3), 54–69.

Levitin, A. J. (2009, April). Resolving the foreclosure crisis: Modification of mortgages in bankruptcy. *Wisconsin Law Review*, 565–655.

Li, W. and P.-D. Sarte (2006, April). U.S. consumer bankruptcy choice: The importance of general equilibrium effects. *Journal of Monetary Economics 53*(3), 613–631.

Li, W., M. White, and N. Zhu (2011). Did bankruptcy reform cause mortgage default rates to rise? *American Economic Journal: Economic Policy*, 123–47.

Li, W. and M. J. White (2009, November). Mortgage default, foreclosure, and bankruptcy. NBER Working Papers 15472, National Bureau of Economic Research, Inc.

Livshits, I., J. MacGee, and M. Tertilt (2007, March). Consumer bankruptcy: A fresh start. *American Economic Review 97*(1), 402–418.

Livshits, I., J. MacGee, and M. Tertilt (2010, April). Accounting for the rise in consumer bankruptcies. *American Economic Journal: Macroeconomics 2*(2), 165–93.

Low, H., C. Meghir, and L. Pistaferri (2010, September). Wage risk and employment risk over the life cycle. *American Economic Review 100*(4), 1432–67.

McConnell, M. M. and G. Perez-Quiros (2000, December). Output fluctuations in the United States: What has changed since the early 1980's? *American Economic Review 90*(5), 1464–1476.

McNulty, J. E. (2009, October). The long-run demand for housing, the housing glut, and the implications for the financial crisis. *Business Economics 44*(4), 206–215.

Meghir, C. and L. Pistaferri (2004, 01). Income variance dynamics and heterogeneity. *Econometrica 72*(1), 1–32.

Miller, R. A. (1984, December). Job matching and occupational choice. *Journal of Politcal Economy 92*(6), 1086–1120.

Mitman, K. (2011, November). Macroeconomic effects of bankruptcy and foreclosure policies. Working papers.

Moffitt, R. A. and P. Gottschalk (2002, March). Trends in the transitory variance of earnings in the united states. *Economic Journal 112*(478), C68–C73.

Morgan, D. P., B. Iverson, and M. Botsch (2011). Subprime foreclosures and the 2005 bankruptcy reform. Technical report.

Murphy, K. and R. Topel (1987). Unemployment, risk, and earnings: testing for equalizing wage differences in the labor market. In K. Lang and J. Leonard (Eds.), *Unemployment and the structure of labor markets*. London: Basil Blackwell.

Musto, D. K. (2004, October). What happens when information leaves a market? Evidence from postbankruptcy consumers. *Journal of Business 77*(4), 725–748.

Neal, D. (1995, October). Industry-specific human capital: Evidence from displaced workers. *Journal of Labor Economics 13*(4), 653–77.

Parent, D. (2000, April). Industry-specific capital and the wage profile: Evidence from the national longitudinal survey of youth and the panel study of income dynamics. *Journal of Labor Economics 18*(2), 306–23.

Pennington-Cross, A. (2006). The value of foreclosed property. *Journal of Real Estate Research 28*(2), 193–214.

Romer, C. (1986). Spurious volatility in historical unemployment data. *The Journal of Political Economy 94*(1), pp. 1–37.

Sargent, T. J. (1999). *The Conquest of American Inflation*. Princeton University Press.

Scarberry, M. S. (2010). A critique of congressional proposals to permit modification of home mortgages in Chapter 13 bankruptcy. *Pepperdine Law Review 37*.

Shepard, L. (1984, October). Personal failures and the Bankruptcy Reform Act of 1978. *Journal of Law & Economics 27*(2), 419–37.

Shilling, J. D., J. D. Benjamin, and C. Sirmans (1990). Estimating net realizable value for distressed real estate. *Journal of Real Estate Research 5*(1), 129–140.

Sims, C. A. and T. Zha (2006, March). Were there regime switches in U.S. monetary policy? *American Economic Review 96*(1), 54–81.

Slichter, S. H. (1950). Notes on the structure of wages. *The Review of Economics and Statistics 32*(1), pp. 80–91.

Smith, A. (2009). *An Inquiry into the Nature and Causes of the Wealth of Nations. 1776* (Reprint ed.). Digireads.com Publishing.

Smith, Anthony A., J. (1993, Suppl. De). Estimating nonlinear time-series models using simulated vector autoregressions. *Journal of Applied Econometrics 8*(S), S63–84.

Smith, Anthony A., J. (2008). Indirect inference. In S. N. Durlauf and L. E. Blume (Eds.), *The New Palgrave Dictionary of Economics*. Basingstoke: Palgrave Macmillan.

Stock, J. H. and M. W. Watson (2003, August). Has the business cycle changed and why? In *NBER Macroeconomics Annual 2002, Volume 17*, NBER Chapters, pp. 159–230. National Bureau of Economic Research, Inc.

Storesletten, K., C. Telmer, and A. Yaron (2004a). Consumption and risk sharing over the life cycle. *Journal of Monetary Economics 51*(3), 609–633.

Storesletten, K., C. Telmer, and A. Yaron (2007, October). Asset pricing with idiosyncratic risk and overlapping generations. *Review of Economic Dynamics 10*(4), 519–548.

Storesletten, K., C. I. Telmer, and A. Yaron (2004b, June). Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy 112*(3), 695–717.

Tauchen, G. and R. Hussey (1991, March). Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. *Econometrica 59*(2), 371–96.

Thomas, L. C. (2000). A survey of credit and behavioural scoring: Forecasting financial risk of lending to consumers. *International Journal of Forecasting 16*(2), 149–172.

Vereshchagina, G. and H. A. Hopenhayn (2009, December). Risk taking by entrepreneurs. *American Economic Review 99*(5), 1808–30.

Warren, E. and A. W. Tyagi (2003). *The Two Income Trap: Why Middle-Class Mothers and Fathers Are Going Broke.* New York: Basic Books.

White, M. J. and N. Zhu (2008). Saving your home in Chapter 13 bankruptcy. NBER Working Papers 14179, National Bureau of Economic Research, Inc.