Financial Innovation, the Discovery of Risk, and the U.S. Credit Crisis[*]

Emine Boz
International Monetary Fund

Enrique G. Mendoza
University of Maryland and NBER

July 22, 2010

Abstract

Uncertainty about the riskiness of a new financial environment was an important factor behind the U.S. credit crisis. We show that a boom-bust cycle in debt, asset prices and consumption characterizes the equilibrium dynamics of a model with a collateral constraint in which agents learn "by observation" the true riskiness of the new environment. Early realizations of states with high ability to leverage assets into debt turn agents overly optimistic about the probability of persistence of a high-leverage regime. Conversely, the first realization of the low-leverage state turns agents unduly pessimistic about future credit prospects. These effects interact with the Fisherian deflation mechanism, resulting in changes in debt, leverage, and asset prices larger than predicted under either rational expectations without learning or with learning but without Fisherian deflation. The model can account for 69 percent of the rise in net household debt and 53 percent of the rise in residential land prices between 1997 and 2006, and it predicts a sharp collapse in 2007.

JEL Classification: *F41, E44, D82*
Keywords: credit crisis, financial innovation, imperfect information, learning, asset prices, Fisherian deflation

1 Introduction

A key factor behind the U.S. financial crisis was the large increase of household credit, residential land prices, and leverage ratios that preceded it. Between 1996 and 2006, the year in which the crisis started with the collapse of the sub-prime mortgage market, the net credit assets of U.S. households and non-profit organizations fell from -35 to -70 percent of GDP (see the top panel of Figure 1). By contrast, this ratio had remained very stable in the previous two decades. During 1996-2006, the market value of residential land as a share of GDP also surged, from about 45 percent to nearly 75 percent (see the bottom panel of Figure 1). Debt grew much faster than land values, however, because the ratio of net credit assets to the market value of residential land, a macroeconomic measure of the household leverage ratio, rose from 0.64 to 0.93 in absolute value.

Figure 1: Net Credit Market Assets and Value of Residential Land
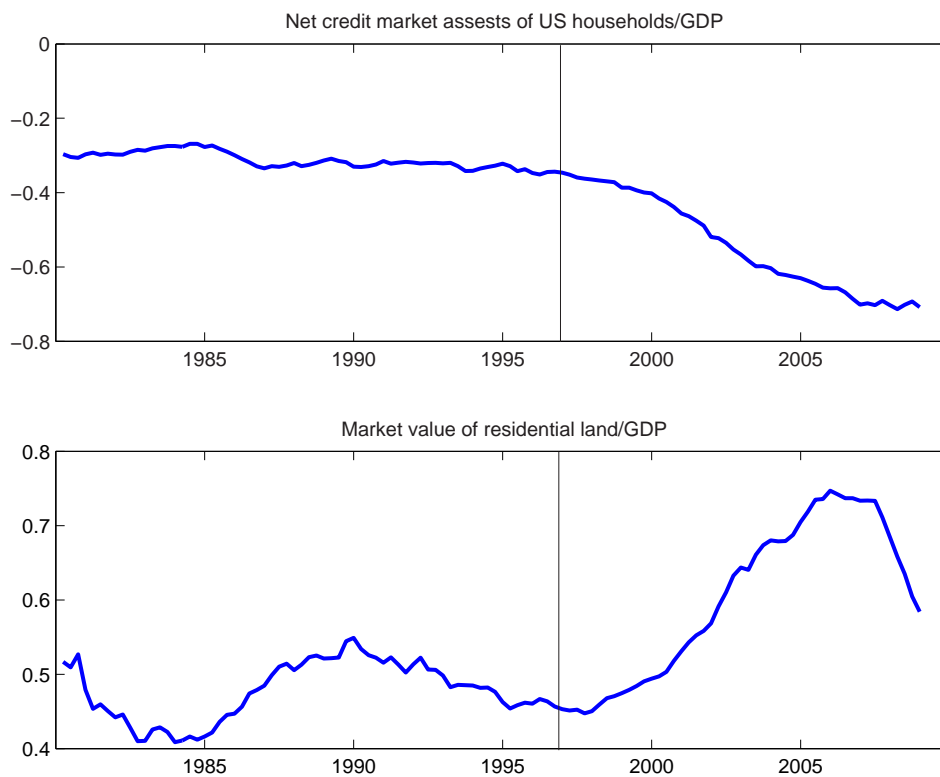
Notes: This figure plots the net credit market assets to GDP ratio for the U.S. households and non-profit organizations. Sources: Net Credit Market Assets: Flow of Funds Accounts of the U.S. provided by the Board of Governors of the Federal Reserve System. Value of Residential Land: Davis and Heathcote (2007).

As the timeline in Figure 2 shows, the rapid growth of household credit and leverage started with a period of significant financial innovation characterized by two central features: First, the introduction of new financial instruments that "securitized" the payment streams generated by a

1

wide variety of assets, particularly mortgages. Second, far-reaching reforms that radically changed the legal and regulatory framework of financial markets.

Figure 2: Timeline of Events During the Run-up to the U.S. Credit Crisis

2006 Peak of stock and housing markets	• **2008** Net credit assets-GDP bottoms
1999 Gramm-Leach-Bliley Act	• **2000** Commodity Futures Modernization Act
	• **1997** Issuance of the first CDS at JPMorgan
1995 Net credit assets-GDP starts falling	• **1995** New Community Reinvestment Act
1987 Issuance of the first CDO	
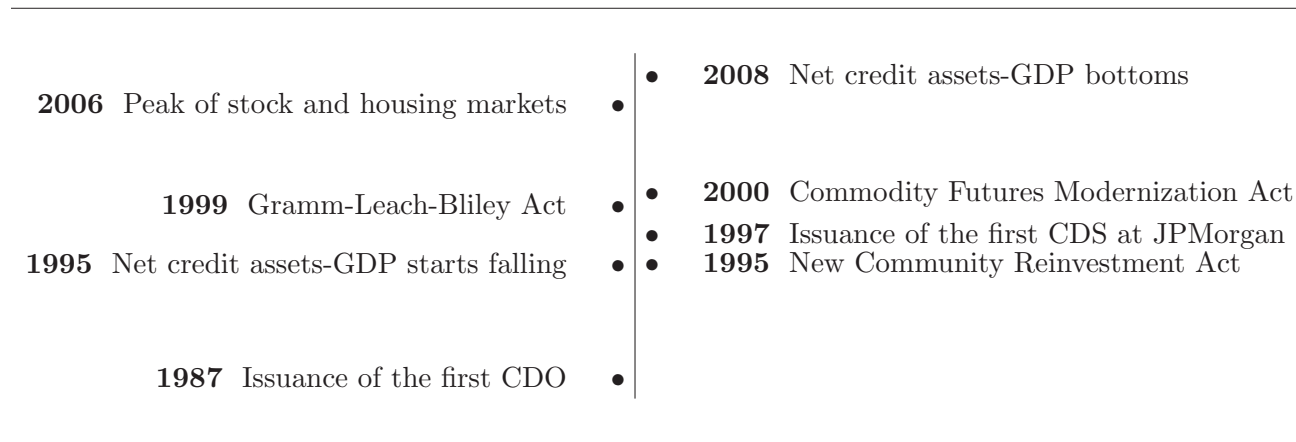
The gradual introduction of collateralized debt obligations (CDOs) dates back to the early 1980s, but the securitization boom that fueled the growth of household debt started in the mid 1990s with the introduction of collateralized mortgage obligations (CMOs) and insurance contracts on the payments of CDOs and CMOs known as credit default swaps (CDSs). In addition, synthetic securitization allowed third parties to trade these securities as bets on the corresponding income streams without being a party to the actual underlying loan contracts. By the end of 2007, the market of CDSs alone was worth about $45 trillion (or 3 times U.S. GDP).

The financial reforms introduced in the 1990s were the most significant since the Great Depression, and in fact aimed at removing the barriers separating bank and non-bank financial intermediaries set in 1933 with the Glass-Steagall Banking Act. Three Acts were particularly important for the housing and credit booms: The 1995 New Community Reinvestment Act, which strengthened the role of Fannie Mae and Freddie Mac in mortgage markets and facilitated mortgage securitization; the 1999 Gramm–Leach–Bliley Act, which removed the prohibition that prevented bank holding companies from owning other financial companies; and the 2000 Commodity Futures Modernization Act, which stipulated that financial derivatives such as CDSs would not be regulated as futures contracts, securities, or lotteries under any federal law.

We show in this paper that financial innovation of this magnitude can lead to a "natural" underpricing of the risk associated with the new financial environment, and that this can produce a surge in credit and asset prices, followed by a collapse. Undervaluing the risk was natural because of the lack of data on the default and performance records of the new financial instruments, and on the

stability of the financial system under the new regulatory framework. In line with this argument, the strategy of "layering of risk" justified the belief that the new instruments were so well diversified that they were virtually risk free. The latter was presumably being attained by using portfolio models that combined top-rated tranches of assets with tranches containing riskier assets–under the assumption that the risk of the assets was priced correctly. As Drew (2008) described it: *"The computer modelers gushed about the tranches. The layers spread out the risk. Only a catastrophic failure would bring the structure crashing down, and the models said that wouldn't happen."*

We recognize that several factors were at play in causing the credit boom that ended with the financial crash, including moral hazard in financial markets and rating agencies, reckless lending practices, growing global financial imbalances, and the lack of government supervision and regulation. In this paper, however, we focus exclusively on the role of financial innovation in an environment with imperfect information and imperfect credit markets, so we can show how these frictions alone can result in a pronounced credit boom-bust cycle. In particular, we propose a model in which the true riskiness of the new financial environment can only be discovered with time, and this learning process interacts with a collateral constraint that limits the debt of private agents not to exceed a fraction of the market value of their holdings of a fixed asset (i.e., land).[1]

Financial innovation is modelled as a structural change that increases the leverage limit, thus moving the economy to a "high-leverage" state. Agents know that in this new environment one of two financial regimes can materialize in any given period: one in which high ability to leverage continues, and one in which there is a switch back to the pre-financial-innovation leverage limit (the "low-leverage" state). They do not know the true riskiness of the new financial environment, because they lack data with which to estimate accurately the true regime-switching probabilities across high- and low-leverage states. They are Bayesian learners, however, and so they learn over time as they observe regime realizations, and in the long-run their beliefs converge to the true regime-switching probabilities. Hence, in the long-run the model converges to the rational expectations (RE) solution, with the risk of the financial environment priced correctly. In the short-run, however, optimal plans and asset prices deviate from the RE equilibrium, because beliefs differ from those of the RE solution, and this leads to a mispricing of risk.

---

[1]Following Davis and Heathcote (2007), we decided to focus on residential land and fluctuations in its price, instead of focusing on housing prices. Davis and Heathcote decomposed U.S. housing prices into the prices of land and structures, and found that between 1975 and 2006 residential land prices quadrupled while prices of physical structures increased only by 33 percent in real terms. Furthermore, land prices are about three times more volatile than prices of structures. Thus, land prices are more important than the prices of residential dwellings for understanding the evolution of housing prices.

Figure 3: Banks' Willingness To Lend
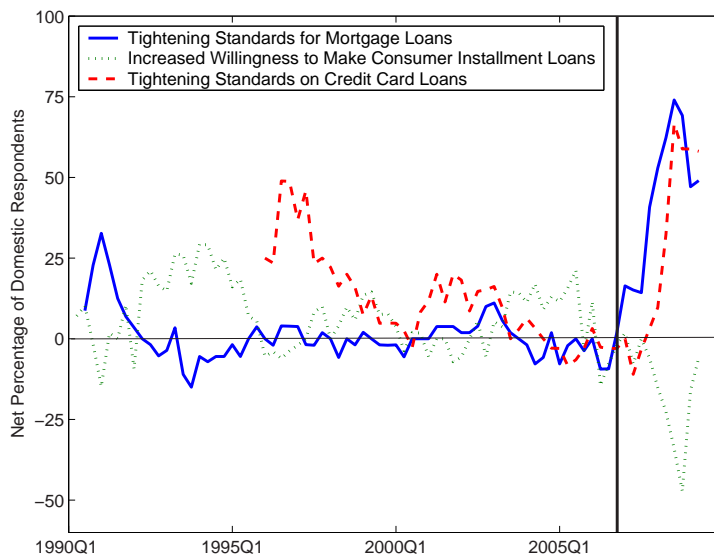
Notes: This figure plots the net percentage of domestic banks that reported tightening standards for mortgage loans and credit card loans; and increased willingness to make consumer installment loans. The banks can choose from five answers, 1) tightened significantly, 2) tightened somewhat, 3) remained unchanged, 4) eased somewhat, 5) eased significantly. Net percentages are calculated by subtracting the number of banks that chose 4 or 5 from those that chose 1 or 2, and then dividing by the total number of respondents. Source: Willingness to Lend Survey of the U.S., provided by the Board of Governors of the Federal Reserve System.

The collateral constraint introduces into the model the well-known Fisherian debt-deflation mechanism of financial amplification, but the analysis of the interaction of this mechanism with the learning dynamics is a novel feature of our work. In particular, the deviations of the agents' beliefs from the true RE regime-switching probabilities distort asset pricing conditions. The resulting over- or under-pricing of assets translates into over- or under-inflated collateral values that affect the debt-deflation dynamics.

Quantitative analysis shows that the process of discovery of risk in the presence of collateral constraints has important effects on macroeconomic aggregates, and leads to a period of booming credit and land prices, followed by a sharp, sudden collapse. We conduct an experiment calibrated to U.S. data in which we date the start of financial innovation in the first quarter of 1997 and the beginning of the financial crisis in the first quarter of 2007. Hence, from 1997 to the end of 2006 we assume that the economy experienced the high-leverage regime, followed by a switch to the low-leverage regime in the first quarter of 2007. The outstanding stock of net credit assets did not rise sharply then (see Figure 1), but the fraction of banks that tightened standards for mortgage and credit card loans jumped from nearly zero to over 50 percent (see Figure 3). The initial priors of the Bayesian learning process are calibrated to match observed excess returns on Fannie Mae

4

MBS at the beginning of 1997, and the high- and low-leverage limits are set equal to the observed leverage ratios before 1997 and at the end of 2006.

Under these assumptions, our model predicts that agents became optimistic about the probability of persistence of the high-leverage regime very soon after 1997, and remained so until they observed the switch to the low-leverage regime. During this "optimistic phase," debt, leverage and collateral values (i.e., land prices) rise significantly above what the RE equilibrium predicts.[2] In fact, the model accounts for 69 percent of the rise in net household debt and 53 percent of the rise in residential land prices during 1997-2006. Conversely, when agents observe the first realization of the low-leverage regime, they respond with a sharp correction in their beliefs and become unduly pessimistic, causing sharp downward adjustments in credit, land prices and consumption.[3]

The results also show that the interaction between the debt-deflation mechanism and the learning mechanism is quantitatively significant. The model predicts effects on debt and asset prices that are nearly twice as large when we allow for these two mechanisms to interact than when we remove either one.

Although we focus on the recent U.S. credit crisis and the financial innovation that preceded it, our framework applies to many episodes of credit booms and busts associated with large changes in the financial environment. It is well-known, for instance, that many of the countries to which the financial crisis spread after hitting the U.S. displayed similar pre-crisis features, in terms of a large expansion of the financial sector into new instruments under new regulations, and also experienced large housing booms (e.g., the United Kingdom, Spain, Iceland, Ireland). There is also evidence of a similar process at work before the Great Depression, specifically the securitization boom in the commercial mortgage market in the 1920s (Goetzmann and Newman (2010)). Moreover, Mendoza and Terrones (2008) found that 33 (22) percent of credit booms observed in the 1965-2006 period in developed (emerging) economies occurred after periods of large financial reforms. Looking at specific countries, the credit booms of Central and Eastern European transition economies in the aftermath of their financial liberalization, and those observed in the Baltic states in the mid 2000's, just before they entered the European Union, are very good examples.[4] In both cases there was

---

[2]The degree of optimism generated in the optimistic phase is at its highest just before agents observe the first realization of the low-leverage regime. This occurs because, when the new financial environment is first introduced, agents cannot rule out the possibility of the high-leverage regime being absorbent until they experience the first realization of the low-leverage state.

[3]The transition to the low-leverage regime is taken as given. One can think of it as being due to a disruption in financial intermediation as in Gertler and Kiyotaki (2010) that is not explicitly modelled in this paper.

[4]See Lipschitz, Lane, and Mourmouras (2002) for a discussion of capital flows to transition economies and the resulting policy challenges.

significant financial innovation, and since these countries had not liberalized financial markets or been in the EU before, there was little relevant history on which to base expectations.

We model learning following the approach proposed by Cogley and Sargent (2008). They offer an explanation of the equity premium puzzle by modelling a period of persistent pessimism caused by the Great Depression. They assume high and low states for consumption growth, with the true transition probabilities across these states unknown. Agents learn the true probabilities over time as they observe (without noise) the realizations of consumption growth. Similarly, in our setup, the true probabilities of switching across leverage regimes are unknown, and agents learn about them over time.

This paper is also related to the broader macro literature on the macroeconomic implications of learning. Most of this literature focuses on learning from noisy signals (see, for example, Blanchard, L'Huillier, and Lorenzoni (2008), Boz (2009), Boz, Daude, and Durdu (2008), Edge, Laubach, and Williams (2007), Lorenzoni (2009), Nieuwerburgh and Veldkamp (2006)). The informational friction in these models typically stems from signal extraction problems requiring the decomposition of signals into a persistent component and a noise component. The informational friction in models like ours and Cogley and Sargent (2008) is fundamentally different, because there is no signal extraction problem. Agents observe realizations of the relevant variables without noise. Instead, there is imperfect information about the true transition probability distribution of these variables. The financial innovations that led to the U.S. credit crisis provide a natural laboratory to study the effects of this class of learning models, because the new financial products clearly lacked the time-series data needed to infer the true probability of "catastrophic failure" of credit markets (i.e., the probability of switching to a low-leverage regime).

The credit constraint used in our model is similar to those widely examined in the macro literature on financial frictions and in the international macro literature on Sudden Stops. When these constraints are used in RE stochastic environments, precautionary savings reduce significantly the long-run probability of states in which the constraints are binding (see Mendoza (2010) and Durdu, Mendoza, and Terrones (2009)). In our learning model, however, agents have significantly weaker incentives for building precautionary savings than under rational expectations, until they attain the long-run in which they have learned the true riskiness of the financial environment. Since agents borrow too much during the optimistic phase, the economy is vulnerable to suffer a large credit crunch when the first switch to a regime with low leverage occurs.

Our credit constraint also features the "systemic credit externality" present in several models of financial crises. In particular, agents do not internalize the implications of their individual actions on credit conditions because of changes in equilibrium prices, and this leads to "overborrowing" relative to debt levels that would be acquired without the externality. The studies on overborrowing like those by Uribe (2006), Korinek (2008), and Bianchi (2009) explore whether credit externalities can generate excessive borrowing in decentralized equilibria relative to a constrained social optimum. Our paper makes two contributions to this line of research. First, we show that the discovery of risk generates sizable overborrowing (relative to the RE decentralized equilibrium), because of the unduly optimistic expectations of agents during the optimistic phase of the learning dynamics. This remains the case even in variants of our model with credit constraints that do not include the credit externality. Second, we provide the first analysis of the interaction between the credit externality and the underpricing of risk driven by a process of "risk discovery."

Our work is also related to the literature on credit booms. The stylized facts documented by Mendoza and Terrones (2008) show that credit booms have well-defined cyclical patterns, with the peak of credit booms preceded by periods of expansion in credit, asset prices, and economic activity followed by sharp contractions. Most of the models of financial crises, however, emphasize mechanisms that amplify downturns and explain crashes but leave booms unexplained. In this regard, our model aims to explain both the boom and the bust phase of credit cycles.

Finally, our paper is also related to some of the recent macro/finance literature on the U.S. credit crisis that emphasizes learning frictions and financial innovation, particularly the work of Howitt (2010) and Favilukis, Ludvigson, and Nieuwerburgh (2010). Howitt studies the interaction of expectations, leverage and a solvency constraint in a representative agent setup similar to ours, and he shows that adaptive learning about asset returns leads to periods of "cumulative optimism" followed by "cumulative pessimism," and this can lead to a crisis. Our analysis differs in that we study Bayesian learning, instead of adaptive expectations, and we model learning about the persistence of a financial regime, defined in terms of the maximum leverage ratio specified by a collateral constraint.[5] Favilukis, Ludvigson, and Nieuwerburgh (2010) analyze the macroeconomic effects of housing wealth and housing finance in a heterogenous agents, DSGE model with credit constraints. They study transition dynamics from an environment with high financial transaction

---

[5]There is also an interesting contrast across the two studies in terms of the motivation for focusing on learning to study the financial crisis. Howitt argues that the learning friction matters because agents learn in adaptive fashion about the behavior of asset returns, in a financial regime that is in fact unchanged, while we argue that it matters because agents learn gradually the true persistence of a new financial regime, while they have perfect information about the random process that drives dividends.

costs and tight credit limits to one with the opposite features. Our analysis has a similar flavor, but we focus on the effects of imperfect information and learning on macro dynamics, while they study a rational expectations environment.

The remainder of this paper proceeds as follows: Section 2 describes the model and the learning process. Section 3 examines the model's quantitative implications. Section 4 concludes.

2 A Model of Financial Innovation with Learning

We study a representative agent economy in which risk-averse individuals formulate optimal plans facing exogenous income fluctuations. The risk associated with these fluctuations cannot be fully diversified because asset markets are incomplete. Individuals have access to two assets: a non-state-contingent bond and an asset in fixed supply (land). The credit market is imperfect, because individuals' ability to borrow is limited not to exceed a fraction κ of the market value of their land holdings. That is, κ imposes an upper bound on the agents' leverage ratio.

The main feature that differentiates our model from typical macro models with credit frictions is the assumption that agents have imperfect information about the regime-switching probabilities that drive fluctuations in κ.[6] Specifically, we model a situation in which financial innovation starts with an initial shift from a low-leverage regime (κ^l) to a regime with higher ability to leverage (κ^h). Agents do not know the true regime-switching probabilities between κ^l and κ^h in this new financial environment. They are Bayesian learners, so in the long-run they learn these true probabilities and form rational expectations. In the short-run, however, they form their expectations with the posteriors they construct as they observe realizations of κ. Hence, they "discover" the true riskiness of the new financial environment only after they have observed a sample with enough regime realizations and regime switches to estimate the true regime-switching probabilities accurately.

We assume that the risk-free interest rate is exogenous in order to keep the interaction between financial innovation and learning tractable. At the aggregate level, this assumption corresponds to an economy that is small and open with respect to world capital markets. This is in line with recent evidence suggesting that in the era of financial globalization even the U.S. risk-free rate has been significantly influenced by outside factors, such as the surge in reserves in emerging economies and the persistent collapse of investment rates in South East Asia after 1998 (see Warnock and

---

[6]In previous work we studied a similar informational friction but in a setup in which the credit constraint does not depend on market prices. In that scenario, the distortions produced by the learning process in the aftermath of financial innovation do not interact with the credit externality present in the model we study here.

Warnock (2006), Bernanke (2005), Durdu, Mendoza, and Terrones (2009), Mendoza, Quadrini, and Rìos-Rull (2009)). Moreover, post-war data from the *Flow of Funds* published by the Federal Reserve show that, while pre-1980s the United States was in virtual financial autarky, because the fraction of net credit of U.S. nonfinancial sectors financed by the rest of the world was close to zero, about one half of the surge in net credit since the mid-1980s was financed by the rest of the world (see Mendoza and Quadrini (2010)). Alternatively, our setup can be viewed as a partial equilibrium model of the U.S. economy that studies the effects of financial innovation on household debt and residential land prices, taking the risk-free rate as given, as in Corbae and Quintin (2009) and Howitt (2010).

2.1 Agents' Optimization Problem

Agents act atomistically in competitive markets and choose consumption (c_t), land holdings (l_{t+1}) and holdings of one-period discount bonds (b_{t+1}), taking as given the price of land (q_t) and the gross real interest rate (R) so as to maximize a standard intertemporal utility function:

$$E_0^s \left[\sum_{t=0}^{\infty} \beta^t u(c_t) \right] \tag{1}$$

It is critical to note that E_t^s represents the expectations operator conditional on the representative agent's beliefs formulated with the information available up to and including date t. As we explain below, these beliefs will differ in general from the rational expectations formulated with perfect information about the persistence of the financial regime, which are denoted E_t^a.

The agents' budget constraint is:

$$c_t = z_t g(l_t) + q_t l_t - q_t l_{t+1} - \frac{b_{t+1}}{R} + b_t \tag{2}$$

Agents operate a concave neoclassical production function $g(l_t)$ subject to a stochastic TFP shock z_t. Since land is in fixed aggregate supply, a linear production technology could also be used. We will use the curvature of $g(l_t)$, however, to calibrate the model so that the condition that arbitrages returns across bonds and land is consistent with U.S. data on the risk-free interest rate and the value of residential land as a share of GDP (see Section 3 for details).

TFP shocks follow an exogenous discrete Markov process (which can be enriched to include also interest rate shocks). For these shocks, we assume that agents know their true Markov process

without informational frictions. That is they know the Markov transition matrix $\pi(z_{t+1} \mid z_t)$ and the corresponding set Z of M possible realizations of z at any point in time (i.e., $z_t \in Z = \{z_1 < z_2 < .... < z_M\}$). Alternatively, we could assume that TFP shocks are also affected by imperfect information.

Frictions in the credit market force agents to comply with a collateral constraint that limits the value of debt (given by b_{t+1}/R since $1/R$ is the price of discount bonds) to a time-varying fraction κ_t of the market value of their land holdings:

$$\frac{b_{t+1}}{R} \geq -\kappa_t q_t l_{t+1} \tag{3}$$

In this constraint, κ_t is a random variable that follows a "true" Markov process characterized by a standard two-point, regime-switching process with regimes κ^h and κ^l, with $\kappa^h > \kappa^l$, and transition probabilities given by $F^a = p^a(\kappa_{t+1} \mid \kappa_t)$.[7] The continuation transition probabilities are denoted $F_{hh}^a \equiv p^a(\kappa_{t+1} = \kappa^h \mid \kappa_t = \kappa^h)$ and $F_{ll}^a \equiv p^a(\kappa_{t+1} = \kappa^l \mid \kappa_t = \kappa^l)$, and the switching probabilities are $F_{hl}^a = 1 - F_{hh}^a$ and $F_{lh}^a = 1 - F_{ll}^a$. The long-run probabilities of the high- and low-leverage regimes are $\Pi^h = F_{lh}^a/(F_{lh}^a + F_{hl}^a)$ and $\Pi^l = F_{hl}^a/(F_{lh}^a + F_{hl}^a)$ respectively, and the corresponding mean durations are $1/F_{hl}^a$ and $1/F_{lh}^a$. The long-run unconditional mean, variance, and first-order autocorrelation of κ are:

$$E^a[\kappa] = (F_{lh}^a \kappa^h + F_{hl}^a \kappa^l)/(F_{lh}^a + F_{hl}^a) \tag{4}$$

$$\sigma^2(\kappa) = \Pi^h(\kappa^h)^2 + \Pi^l(\kappa^l)^2 - (E[\kappa])^2 \tag{5}$$

$$\rho(\kappa) = F_{ll}^a - F_{hl}^a = F_{hh}^a - F_{lh}^a \tag{6}$$

As explained earlier, agents know κ^h and κ^l but do not know F^a. Hence, they make decisions based on their individual beliefs characterized by E^s which evolve over time. Using μ to denote the Lagrange multiplier on the credit constraint, the agents' optimality conditions for bonds and land are:

$$u'(c_t) = \beta R E_t^s \left[u'(c_{t+1}) \right] + \mu_t \tag{7}$$

$$q_t(u'(c_t) - \mu_t \kappa_t) = \beta E_t^s \left[u'(c_{t+1}) \left(z_{t+1} g'(l_{t+1}) + q_{t+1} \right) \right]. \tag{8}$$

---

[7]One could also specify a continuous AR(1) process for κ such as $\kappa_t = m_t + \kappa_{t-1} + \epsilon_t$. The different regimes could be captured with a shift in the mean: $m \in \{m^h, m^l\}$ and the agents could learn about the process governing m. We conjecture that this specification would yield similar results as agents could turn optimistic about the persistence of the high mean regime for κ.

With the caveat that agents use subjective beliefs to form expectations, these conditions are standard from models with credit constraints. Following Mendoza (2010), we can show that the second condition implies a forward solution for land prices in which the future stream of land dividends is discounted at the stochastic discount factors adjusted for the shadow value of the credit constraint:

$$q_t = E_t^s \left[\sum_{j=0}^{\infty} \left(\prod_{i=0}^{j} M_{t+i}^{t+1+i} \right) z_{t+1+j} g'(l_{t+1+j}) \right], \qquad M_{t+i}^{t+1+i} \equiv \frac{\beta u'(c_{t+1+i})}{u'(c_{t+i}) - \mu_{t+i}\kappa_{t+i}} \qquad (9)$$

Defining the return on land as $R_{t+1}^q \equiv (z_{t+1} g'(l_{t+1}) + q_{t+1})/q_t$ and the period marginal utility of consumption as $\lambda_{t+1} \equiv \beta u'(c_{t+1})$, the excess return on land can be expressed as:

$$E_t^s \left[R_{t+1}^q - R \right] = \frac{(1 - \kappa_t)\mu_t - cov_t^s(\lambda_{t+1}, R_{t+1}^q)}{E_t^s [\lambda_{t+1}]} \qquad (10)$$

Thus, as in Mendoza (2010), the borrowing constraint enlarges the standard premium on land holdings, driven by the covariance between marginal utility and asset returns, by introducing direct and indirect effects. The direct effect is represented by the term $(1 - \kappa_t)\mu_t$. The indirect effect is represented by the fact that the credit constraint hampers the agents' ability to smooth consumption and hence reduces $cov_t^s(\lambda_{t+1}, R_{t+1}^q)$. Moreover, since the expected land returns satisfy $q_t E_t^s[R_{t+1}^q] \equiv E_t^s[z_{t+1} g'(l_{t+1}) + q_{t+1}]$, we can rewrite the forward solution for the agents' land valuation as:

$$q_t = E_t^s \left[\sum_{j=0}^{\infty} \left(\prod_{i=0}^{j} \left(\frac{1}{E_t^s[R_{t+1+i}^q]} \right) \right) z_{t+1+j} g'(l_{t+1+j}) \right]. \qquad (11)$$

The expressions in (10) and (11) imply that the collateral constraint lowers land prices because it increases the rate of return at which future land dividends are discounted. Note also that land valuations are reduced at t not just when the constraint binds at t, which increases $E_t^s[R_{t+1}^q]$, but also if agents expect (given their beliefs F^s) that the constraint can bind at any future date, which increases $E_t^s[R_{t+1+i}^q]$ for some $i > 0$. Thus this expression suggests that the learning process and the collateral constraint interact in an important way. For instance, suppose the credit constraint was binding at t, pessimistic beliefs such that agents expect higher future land returns because of tight credit conditions will depress more current land prices, which will tighten more the collateral constraint.

The economy has a fixed unit supply of land, hence market clearing in the land market implies that the land holdings of the representative agent must satisfy $l_t = 1$ for all t, and the rest of the

equilibrium conditions reduce to the following:

$$u'(c_t) = \beta R E_t^s \left[u'(c_{t+1}) \right] + \mu_t \tag{12}$$

$$q_t(u'(c_t) - \mu_t \kappa_t) = \beta E_t^s \left[u'(c_{t+1}) \left(z_{t+1} g'(1) + q_{t+1} \right) \right] \tag{13}$$

$$c_t = z_t g(1) - \frac{b_{t+1}}{R} + b_t \tag{14}$$

$$\frac{b_{t+1}}{R} \geq -\kappa_t q_t 1 \tag{15}$$

2.2 General Features of the Learning Setup

Following Cogley and Sargent (2008), our learning setup features two-point passive learning without noise, so that the belief transition probability matrix denoted by $p_t^s(\kappa_{t+1} \mid \kappa_t)$ converges to its true value $p_t^s(\kappa_{t+1} \mid \kappa_t) \to p^a(\kappa_{t+1} \mid \kappa_t)$ for sufficiently large t. With this setup, agents learn about the transition probability matrix only as they observe realizations of the shocks. In particular, they learn about the true regime-switching probabilities of κ only after observing a sufficiently long set of realizations of κ^h and κ^l.[8]

This learning setup fits nicely our goal of studying a situation in which financial innovation represents an initial condition with a state κ^h but with imperfect information about the true riskiness of this new environment. Agents are ignorant about the true transition distribution of κ, since there is no data history to infer it from. Over time, if they observe a sequence of realizations of κ^h for a few periods, they build optimism by assigning a probability to the possibility of continuing in κ^h that is higher than the true value. We refer to this situation as the "optimistic phase." Such optimism by itself is a source of vulnerability, because it is quickly reversed into a "pessimistic phase" with the opposite characteristics as the first few realizations of κ^l hit the economy. In addition, during the optimistic phase, the incentives to build precautionary savings against the risk of a shift in the ability to leverage are weaker than in the long-run RE equilibrium. This increases the agents' risk of being caught "off-guard" (i.e., with too much debt) when the first shift to the low-leverage regime occurs.

Modeling imperfect information in this fashion implies a deviation from rational expectations. This is a key feature of our model, because it highlights the role of the short history of a new financial regime in hampering the ability of agents to correctly assess risk. This approach seems

---

[8]Time alone does not determine how fast agents learn. The order in which κ realizations, and switches between realizations, occur also matters.

better suited for studying the role of financial innovation in causing the financial crisis, as opposed to a standard signal extraction problem with noisy signals and rational expectations.

Since κ is exogenous, we are modeling a passive learning structure from *and* about exogenous variables, which facilitates significantly the analysis and numerical solution of the model. In particular, it allows us to split the analysis into two parts. The first part uses Bayesian learning to generate the agents' sequence of posterior density functions $\{f(F^s \mid \kappa^t)\}_{t=1}^T$. Each of these density functions (one for each date t) is a probability distribution over possible Markov transition matrices F^s. Since agents do not know the true transition matrix F^a, the density function changes with the sequence of realizations observed up to date t (i.e., $\{\kappa^t, \kappa^{t-1}, ..., \kappa^1\}$ where $\kappa^t = (\kappa_t, \kappa_{t-1}, ..., \kappa_1)$) and with the initial date-0 priors, as we explain below. If date T is high enough to accomodate sufficient sampling of regime switches across κ^h and κ^l, the sequence $\{f(F^s \mid \kappa^t)\}_{t=1}^T$ converges to a distribution with all its mass in F^a. In other words, beliefs converge to the true values, so that in the long-run the model converges to the RE equilibrium.

The second part of the analysis characterizes the model's equilibrium by combining the sequences of posterior densities obtained in the first part, $\{f(F^s \mid \kappa^t)\}_{t=1}^T$, with chained solutions from a set of "conditional" optimization problems. These problems are conditional on the posterior density function of F^s that agents form with the history of realizations up to each date t. This approach takes advantage of the fact that, because of the passive learning, agents do not benefit from trying to improve their inference about the regime switching probabilities by "experimenting" using their optimization problems. As a result, the evolution of beliefs can be analyzed separately from the agents' optimal consumption and savings plans. The remainder of this Section examines in more detail the Bayesian learning setup and the construction of the model's equilibrium.

2.3 Learning and the Sequence of Beliefs

The learning framework takes as given an observed history of realizations of T periods of leverage regimes, denoted by κ^T, and a prior of F^s for date $t = 0$, $p(F^s)$, and it yields a sequence of posteriors over F^s for every date t, $\{f(F^s \mid \kappa^t)\}_{t=1}^T$.[9] At every date t, from 0 to T, the information set of the agent includes κ^t as well as the possible values that κ can take (κ^h and κ^l).

Agents form posteriors from priors using a beta-binomial probability model. Since agents know the realization vector of κ, information is imperfect only with regard to the Markov transition matrix across $\kappa's$, and, because κ can only take two values, this boils down to imperfect information about

---

[9]In describing the learning problem, we employ the notation used by Cogley and Sargent (2008).

the continuation probabilities F_{hh}^a and F_{ll}^a. The other two elements of the transition matrix of κ are recovered using $F_{ii}^a + F_{ij}^a = 1$, where F_{ij}^a denotes the true probability of switching from state i to state j.

The agents' posteriors about F_{hh}^s and F_{ll}^s have Beta distributions as well, and the parameters that define them are determined by the number of regime switches observed in a particular history κ^t. As in Cogley and Sargent (2008), we assume that the priors for F_{hh}^s and F_{ll}^s included in $p(F^s)$ are independent and determined by the number of transitions assumed to have been observed prior to date $t = 1$. More formally,

$$p(F_{ii}^s) \propto (F_{ii}^s)^{n_0^{ii}-1}(1 - F_{ii}^s)^{n_0^{ij}-1} \tag{16}$$

where n_0^{ij} denotes the number of transitions from state i to state j assumed to have been observed prior to date 1.

The likelihood function of κ^t conditional on F_{hh}^s and F_{ll}^s is obtained by multiplying the densities of F_{hh}^s and F_{ll}^s:

$$f(\kappa^t \mid F_{hh}^s, F_{ll}^s) \propto (F_{hh}^s)^{(n_t^{hh}-n_0^{hh})}(1 - F_{hh}^s)^{(n_t^{hl}-n_0^{hl})}(1 - F_{ll}^s)^{(n_t^{lh}-n_0^{lh})}(F_{ll}^s)^{(n_t^{ll}-n_0^{ll})}. \tag{17}$$

Multiplying the likelihood function by the priors delivers the posterior kernel:

$$k(F^s \mid \kappa^t) \propto (F_{hh}^s)^{(n_t^{hh}-1)}(1 - F_{hh}^s)^{(n_t^{hl}-1)}(1 - F_{ll}^s)^{(n_t^{lh}-1)}F_{ll}^{s(n_t^{ll}-1)}, \tag{18}$$

and dividing the kernel using the normalizing constant $M(\kappa^t)$ yields the posterior density:

$$f(F^s \mid \kappa^t) = k(F^s \mid \kappa^t)/M(\kappa^t) \tag{19}$$

where

$$M(\kappa^t) = \iint (F_{hh}^s)^{(n_t^{hh}-1)}(1 - F_{hh}^s)^{(n_t^{hl}-1)}(1 - F_{ll}^s)^{(n_t^{lh}-1)}(F_{ll}^s)^{(n_t^{ll}-1)} \, dF_{hh}^s dF_{ll}^s.$$

The number of transitions across regimes is updated as follows:

$$n_{t+1}^{ij} = \begin{cases} n_t^{ij} + 1 & \text{if } \kappa_{t+1} = \kappa^j \text{ and } \kappa_t = \kappa^i, \\ n_t^{ij} & \text{otherwise.} \end{cases}$$

14

Note that the posteriors are of the form $F_{hh}^s \propto Beta(n_t^{hh}, n_t^{hl})$ and $F_{ll}^s \propto Beta(n_t^{lh}, n_t^{ll})$. That is, the posteriors for κ^h only depend on n_t^{hh} and n_t^{hl} and not on the other two counters, n_t^{lh} and n_t^{ll}, and the posteriors for κ^l only depend on n_t^{lh} and n_t^{ll}. This is important because it implies that the posteriors of F_{hh}^s change only as n_t^{hh} and n_t^{hl} change, and this only happens when the date-t realization is κ^h. If, for example, the economy experiences realizations $\kappa = \kappa^h$ for several periods, agents learn only about the persistence of the high-leverage regime. They learn nothing about the persistence of the low-leverage regime. To learn about this, they need to observe realizations of κ^l. Since in a two-point, regime-switching setup persistence parameters also determine mean durations, it follows that both the persistence and the mean durations of leverage regimes can be learned only as the economy actually experiences κ^l and κ^h.[10]

Figure 4: Evolution of Beliefs
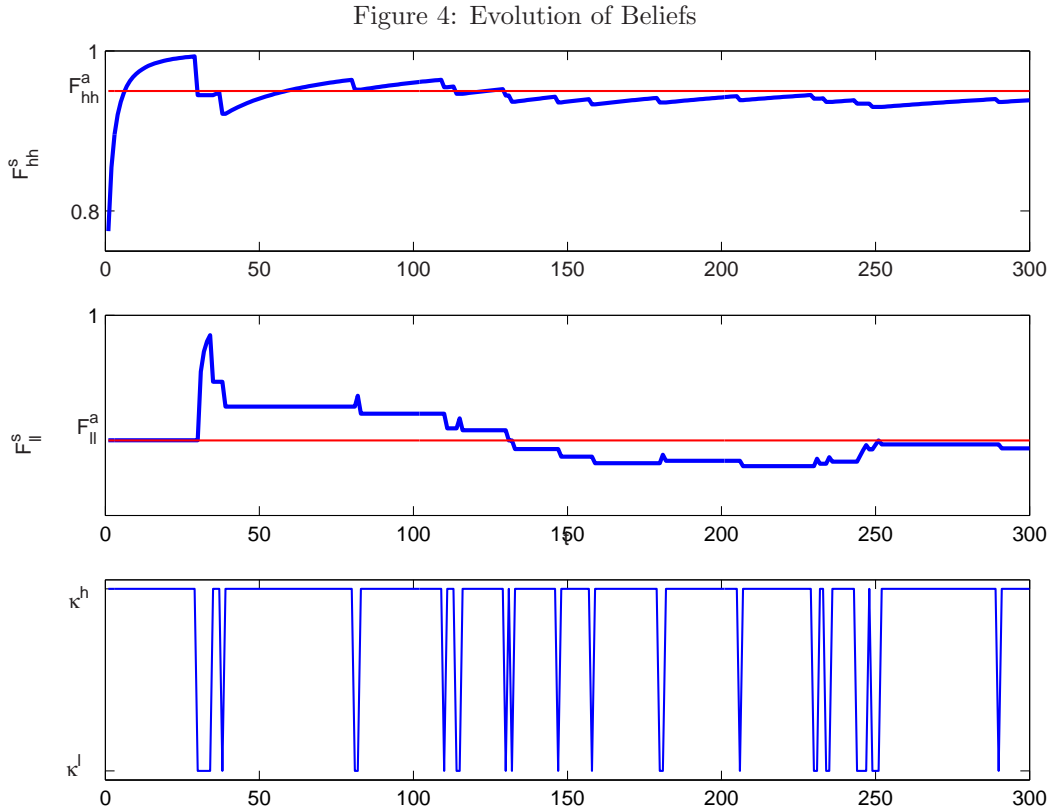
Notes: This figure plots the evolution of beliefs about F_{hh}^s (top panel), F_{ll}^s (middle panel), and the associated realizations of κ (lower panel). The horizontal red lines in the upper two panels mark the true values of the corresponding variables.

_____

[10]If priors, as well as F_{hh}^a and F_{ll}^a, are correlated, learning would likely be faster, because agents would update their beliefs about both F_{hh}^s and F_{ll}^s every period, instead of updating only one or the other depending on the regime they observe. But this is akin to removing some of the informational friction by assumption. In an extreme case, imagine that $F_{hh}^a = F_{ll}^a$ and that the agents know about this property of the model. In this case, agents know an important characteristic of the transition probability matrix (i.e. that is symmetric), which weakens the initial premise stating that they do not know any of its properties. Agents would learn about the persistence of both regimes no matter which one they observe.

15

We illustrate the learning dynamics of this setup by means of a simple numerical example. We choose a set of values for F_{hh}^a, F_{ll}^a, and initial priors, and then simulate the learning process for 300 quarters (75 years) using a hypothetical sequence of κ realizations produced by a stochastic simulation of the true Markov-switching process. The results are plotted in Figure 4, which shows the time paths of the conditional averages of F_{hh}^s and F_{ll}^s based on the beliefs formed at each date t in the horizontal axis, after observing the corresponding κ_t shown in the bottom panel. The true regime-switching probabilities are set to $F_{hh}^a = 0.95$ and $F_{ll}^a = 0.5$. These values are used only for illustration purposes (they are not calibrated to actual data as in the solution of the full model in Section 3). In addition, the date-0 priors are set to $F_{hh}^s \sim Beta(0.7, 0.7)$ and $F_{ll}^s \sim Beta(0.7, 0.7)$. With these priors, the agents update their beliefs about the persistence of the high-leverage regime to around 0.78 after observing $\kappa_1 = \kappa^h$.

The most striking result evident in Figure 4 is that financial innovation can lead to significant underestimation of risk. Specifically, the initial sequence of realizations of κ^h observed until just before the first realization of κ^l (the "optimistic phase") generates substantial optimism. In this example, the optimistic phase covers the first 30 periods. The degree of optimism produced during this phase can be measured by the difference between the conditional expectation based on the date-t beliefs, F_{hh}^s, and the corresponding rational expectations value, F_{hh}^a (shown in the horizontal line of the top panel). As the Figure shows, the difference grows much larger during the optimistic phase than in any of the subsequent periods. For example, even though the economy remains in κ^h from around date 40 to date 80, the magnitude of the optimism that this period generates is smaller than in the initial optimistic phase. This is because it is only after observing the first switch to κ^l that agents rule out the possibility of κ^h being an absorbent state. As a result, F_{hh}^s cannot surge as high as it did during the initial optimistic phase. Notice also that the first realizations of κ^l generate a "pessimistic phase," in which F_{ll}^s is significantly higher than F_{ll}^a, so the period of unduly optimistic expectations is followed by a period of unduly pessimistic expectations.

Figure 4 also illustrates the fact that the beliefs about the average duration of κ^h (κ^l) are updated only when the economy is in the high- (low-) leverage state. This is evident, for example, in the initial optimistic phase (the first 30 periods), when F_{ll}^s does not change at all. As explained above, for the agents to learn about the duration of the high- (low-) leverage regime, the economy needs to actually be in that regime. This feature of the learning dynamics also explains why in this example F_{hh}^s converges to F_{hh}^a faster than F_{ll}^s. Given that the low-leverage regime is visited much

less frequently, it takes longer for the agents to learn about its persistence, and therefore F_{ll}^s takes longer to converge to F_{ll}^a.

2.4   Recursive Competitive Equilibrium

We define the model's competitive equilibrium in recursive form. Since in the quantitative analysis we solve the model by policy function iteration on the equilibrium conditions (12)-(15), we formulate the recursive equilibrium using these conditions instead of a Bellman equation (Appendix A describes the solution method in detail). The state variables in the recursive equilibrium are defined by the triple (b, z, κ).

The solution strategy works by breaking down the problem into a set of conditional optimization problems (COP) that are conditional on the beliefs agents have each period. We add time indices to the policy and pricing functions in the recursive equilibrium so as to identify the date of the beliefs that match the corresponding COP. It is critical to note that this solution strategy works because the law of iterated expectations holds with passive Bayesian learning (see Appendix B in Cogley and Sargent (2008)). This is important because, in solving each COP, agents take into account that they are in a learning environment, so they form expectations about the future, including future $\kappa's$ *and* the associated future beliefs, conditional on the information and beliefs they have in the current planning period. Since the law of iterated expectations holds, however, $E_t^s[E_{t+1}^s(x_{t+2})] = E_t^s[x_{t+2}]$.

Consider first the date-1 COP. At this point agents have observed κ_1, and use it to update their beliefs. Thus, we pull $f(F^s \mid \kappa^1)$ from the sequence of posterior density functions solved for in the previous subsection. This is the first density function in the sequence $\{f(F^s \mid \kappa^t)\}_{t=1}^T$, and it represents the first posterior about the distribution of F^s that agents form, given that they have observed κ_1 and given their initial priors. The solution to the date-1 COP is given by policy functions $(b_1'(b, z, \kappa), c_1(b, z, \kappa), \mu_1(b, z, \kappa))$ and a pricing function $q_1(b, z, \kappa)$ that satisfy the

equilibrium conditions (12)-(15) rewritten in recursive form:

$$u'(c_1(b, z, \kappa)) = \beta R \left[\sum_{z' \in Z} \left(\int E_1^s[u'(c_1(b', z', \kappa')) \mid f(\kappa' \mid \kappa^1, F^s)] f(F^s \mid \kappa^1) dF^s \right) \pi(z' \mid z) \right] + \mu_1(b, z, \kappa)$$

(20)

$$q_1(b, z, \kappa)(u'(c_1(b, z, \kappa)) - \mu_1(b, z, \kappa)\kappa) = $$

(21)

$$\beta \left[\sum_{z' \in Z} \left(\int E_1^s[u'(c_1(b', z', \kappa')) \left(z'g'(1) + q_1(b', z', \kappa') \right) \mid f(\kappa' \mid \kappa^1, F^s)] f(F^s \mid \kappa^1) dF^s \right) \pi(z' \mid z) \right]$$

$$c_1(b, z, \kappa) = zg(1) - \frac{b_1'(b, z, \kappa)}{R} + b$$

(22)

$$\frac{b_1'(b, z, \kappa)}{R} \geq -\kappa q_1(b, z, \kappa)1$$

(23)

where the expectations inside the integrals in (20) and (21) are of the form $E_1^s[x \mid f(\kappa' \mid \kappa^1, F^s)] \equiv \sum_{i=l}^{h} \Pr(\kappa' = \kappa^i \mid \kappa^1, F^s)x$, for a random variable x that depends on κ'. These expectations taken by themselves are analogous to those we would calculate if we were solving a standard RE model using an Euler-equations method with an invariant Markov transition function F^s. Since agents here do not know F^a, however, the expectations in (20) and (21) also integrate over $f(F^s \mid \kappa^1)$. The time subscripts that index the policy and pricing functions indicate the date of the beliefs use to form the expectations (which is also the date of the most recent observation of κ, which is date 1 in this case).

Recall that Equations (20)-(23) incorporate the market-clearing condition in the land market, which requires $l = 1$. Moreover, given (20)-(21), the pricing function $q_1(b, z, \kappa)$ satisfies the asset pricing equation (11).

It is critical to note that solving for date-1 policy and pricing functions means solving for a full set of optimal plans over the entire (b, z, κ) domain of the state space and conditional on $f(F^s \mid \kappa^1)$. Thus, we are solving for the optimal plans agents "conjecture" they would make over the infinite future acting under the beliefs given by $f(F^s \mid \kappa^1)$. This COP remains recursive, particularly in terms of forming expectations about future variables and beliefs, because the law of iterated expectations still holds. For characterizing the "actual" equilibrium dynamics to match against the data, however, the solution of the date-1 COP determines optimal plans for date 1 only.

Generalizing the date-1 problem we can define COPs for all subsequent dates $t = 2, ..., T$ using the corresponding density function $f(F^s \mid \kappa^t)$ for each date t in the sequence of beliefs solved for earlier. This is crucial because $f(F^s \mid \kappa^t)$ changes as time passes and each subsequent κ_t is observed,

reflecting the passive Bayesian learning, which implies that the policy and pricing functions that solve each COP also change. Thus, history matters for the "full solution" of the model because assuming different histories κ^t yields different densities $f(F^s \mid \kappa^t)$, and hence different sets of policy functions. If at any two dates t and $t+j$ we give the agents the same values for (b, z, κ), they in general, will not choose the same bond holdings for the following period because $f(F^s \mid \kappa^t)$ and $f(F^s \mid \kappa^{t+j})$ will differ.

The solution to the date-t COP is given by policy functions $(b'_t(b, z, \kappa), c_t(b, z, \kappa), \mu_t(b, z, \kappa))$ and a pricing function $q_t(b, z, \kappa)$ that satisfy the model's equilibrium conditions:

$$u'(c_t(b, z, \kappa)) = \beta R \left[\sum_{z' \in Z} \left(\int E_t^s[u'(c_t(b', z', \kappa')) \mid f(\kappa' \mid \kappa^t, F^s)] f(F^s \mid \kappa^t) dF^s \right) \pi(z' \mid z) \right] + \mu_t(b, z, \kappa) \tag{24}$$

$$q_t(b, z, \kappa)(u'(c_t(b, z, \kappa)) - \mu_t(b, z, \kappa)\kappa) = \tag{25}$$

$$\beta \left[\sum_{z' \in Z} \left(\int E_t^s[u'(c_t(b', z', \kappa')) \left(z'g'(1) + q_t(b', z', \kappa') \right) \mid f(\kappa' \mid \kappa^t, F^s)] f(F^s \mid \kappa^t) dF^s \right) \pi(z' \mid z) \right]$$

$$c_t(b, z, \kappa) = zg(1) - \frac{b'_t(b, z, \kappa)}{R} + b \tag{26}$$

$$\frac{b'_t(b, z, \kappa)}{R} \geq -\kappa q_t(b, z, \kappa)1 \tag{27}$$

We can now define the model's recursive equilibrium as follows:

Definition Given a T-period history of realizations $\kappa^T = (\kappa_T, \kappa_{T-1}, ..., \kappa_1)$, a recursive competitive equilibrium for the economy is given by a sequence of policy functions $[b'_t(b, z, \kappa), c_t(b, z, \kappa), \mu_t(b, z, \kappa)]_{t=1}^T$ and pricing functions $[q_t(b, z, \kappa)]_{t=1}^T$ such that: (a) $b'_t(b, z, \kappa), c_t(b, z, \kappa), \mu_t(b, z, \kappa)$ and $q_t(b, z, \kappa)$ solve the date-t COP conditional on $f(F^s \mid \kappa^t)$; (b) $f(F^s \mid \kappa^t)$ is the date-t posterior density of F^s determined by the Bayesian passive learning process summarized in Equation (19).

Intuitively, the complete solution of the recursive equilibrium is formed by chaining together the solutions for each date-t COP. That is, the equilibrium dynamics at each date $t = 1, ...T$ for a particular history κ^T are given by $\left[b'_t(b, z, \kappa), c_t(b, z, \kappa), \mu_t(b, z, \kappa), q_t(b, z, \kappa), f(F^s \mid \kappa^t) \right]_{t=1}^T$. At each date in this sequence, $b'_t(b, z, \kappa), c_t(b, z, \kappa), \mu_t(b, z, \kappa), q_t(b, z, \kappa)$, are the recursive functions that solve the corresponding date's COP using $f(F^s \mid \kappa^t)$ to form expectations. Hence, the sequence of equilibrium decision rules for bond holdings that the model predicts for dates $t = 1, ..., T$ would be obtained by chaining the relevant decision rules as follows: $b_2 = b'_1(b, z, \kappa)$, $b_3 = b'_2(b, z, \kappa)$, ..., $b_{T+1} = b'_T(b, z, \kappa)$.

3 Quantitative Analysis

In this section we calibrate the model to U.S. data and study its quantitative predictions for the following financial innovation experiment: At $t = 1$, financial innovation begins with the first realization of κ^h, followed by an optimistic phase in which the same regime continues for J periods. At date $J+1$ the first realization of κ^l occurs, and the financial regime remains in state κ^l from dates $J + 1$ to T. In short, the experiment assumes the sequence of realizations $\kappa_t = \kappa^h$ for $t = 1, ..., J$ and $\kappa_t = \kappa^l$ for $t = J + 1, ..., T$.

3.1 Baseline Calibration

The functional forms for preferences and technology are standard: $u(c_t) = \frac{c_t^{1-\sigma}}{1-\sigma}$ and $g(l_t) = l_t^\alpha$. The calibration requires setting values for the parameters $(\alpha, \beta, \sigma, R)$, the Markov process for z, and the parameters of the learning setup, which include κ^h, κ^l, n_0^{hh}, n_0^{hl}, n_0^{ll}, n_0^{lh}, J and T. We propose a set of baseline calibration parameters based on U.S. data, and later we conduct sensitivity analysis to evaluate the robustness of the results to changes in the baseline calibration.

We calibrate the model to a quarterly frequency at annualized rates. The beginning of the financial innovation experiment is dated as of 1997Q1. This is in line with the observations that 1997 was the year in which the first CDS was issued at JPMorgan and the first publicly-available securitization of loans under the New Community Reinvestment Act took place. Moreover, 1997 is also the year in which the net credit assets-GDP ratio shown in Figure 1 started on its declining trend. We date the start of the financial crisis as of 2007Q1, to match the early stages of the subprime mortgage crisis after the Fall of 2006. This is in line with the observation that the net fraction of banks reporting tighter standards for mortgage loans jumped significantly to 16 percent in 2007Q1, as shown in Figure 3.[11]

The above timing assumptions imply that the first realization of κ^h occurred in 1997Q1 and κ^h continued to be observed trough 2006Q4. Hence, the optimistic phase lasts $J = 40$ quarters (10 years). The first realization of κ^l occurred in 2007Q1 and κ^l continued to be observed through 2008Q4. Thus, the learning period has a total length of $T = 48$ quarters, in which the first 40 realizations of κ are κ^h and the remaining 8 are κ^l.

---

[11]These dates are broadly in line with those assumed by Campbell and Hercowitz (2009), who study the welfare implications of a transition from a high- home-equity-requirement regime to a low-equity-requirement regime. They assume that the former corresponds to the 1982-1994 period, while the latter starts in 1995.

In the pre-financial-innovation period, before 1997, we assume that there was only one financial regime with $\kappa = \kappa^l$, and hence the only source of uncertainty were TFP shocks. The values of $(\alpha, \beta, \sigma, R)$ and κ^l are then set so that the model's stochastic stationary state under these assumptions is consistent with various averages from U.S. data from the pre-financial-innovation period.

We set the real interest rate to 2.66 percent annually. This is the average ex-post real interest rate on U.S. three-month T-bills during the period 1980Q1-1996Q4.

Our data proxy for b are the net credit market assets of U.S. households and non-profit organizations in the *Flow of Funds* dataset, and our proxy for ql is the series on the value of residential land estimated by Davis and Heathcote (2007) and kindly updated by them through 2009Q1. These are the data plotted in Figure 1 as shares of GDP. The 1980Q1-1996Q4 average ratios of the value of residential land and net credit market assets relative to GDP are 0.477 and -0.313 respectively. The two ratios are fairly stable around these averages throughout the 1980Q1-1996Q4 period, in contrast with the sharp trends they display after 1996.

As described in the Introduction, we construct a macro estimate of the household leverage ratio, or the loan-to-value ratio, by dividing net credit market assets by the value of residential land. Then, we set the value of κ^l by combining the 1980Q1-1996Q4 average of this ratio with the calibrated value of R (assuming also that the collateral constraint was binding in the pre-financial-innovation era). This yields $\kappa^l = 0.659/1.0266 = 0.642$. The fact that net credit assets and land values were fairly stable prior to 1997, as shown in Figure 1, supports the idea of using this constant value of κ^l to characterize the pre-financial-innovation regime.

The value of κ^h is set equal to the 2006Q4 leverage ratio, hence $\kappa^h = 0.926$. This represents the largest leverage ratio that the economy was able to support in the new financial regime just before the financial crisis hit. Note, however, that κ^h does not always bind in the new regime. First, as the economy moves from the pre-financial-innovation regime to the regime with stochastic κ, agents build up debt over time, and hence the equilibrium leverage ratio does not jump to its new ceiling immediately as the new regime begins. Second, the new regime features two possible realizations of κ that are occasionally binding, so κ^h only binds with some probability in the long-run.[12]

The value of σ is set to $\sigma = 2.0$, the standard value in quantitative DSGE models, and β is set so that the pre-financial-innovation model matches the observed standard deviation of consumption

---

[12]Our calibrated values of κ^h and κ^l are in line with the parameter values that Favilukis, Ludvigson, and Nieuwerburgh (2010) chose to calibrate their collateral constraint (0.75 in their tight credit regime and 1 in their loose credit environment).

relative to output over the 1980Q1-1996Q4 period, which is 0.8. This procedure yields $\beta = 0.91$. Notice that, given the calibrated value of R, the rate of time preference exceeds the real interest rate (i.e., $\beta R = 0.934 < 1$). This is important because it ensures the existence of an ergodic distribution of bond holdings given that asset markets are incomplete. Intuitively, this occurs because of the interaction between the precautionary savings motive, which pushes for increasing bond holdings, and the incentive to tilt consumption towards the present, and hence reduce bond holdings, because $\beta R < 1$. Consumption tilting and precautionary savings will also play a key role later in our analysis of the macro dynamics induced by financial innovation.

Using the 1980Q1-1996Q4 average of the value of residential land to GDP, the value of R, and the condition that arbitrages the returns on land and bonds, which follows from the optimality conditions (12)-(13), we obtain and implied value for α.[13] This yields $\alpha = 0.0251$.

We normalize mean output to 1 (since $L = 1$ and the unconditional mean of z also equals 1), and calibrate the model so that the observed average pre-financial-innovation ratios of consumption ($\bar{c}$) and bonds ($\bar{b}$) to GDP are consistent with the resource constraint in the average of the stochastic stationary state for that financial regime.[14] The observed average ratio of net credit assets to GDP in the 1980Q1-1996Q4 period yields $\bar{b} = -0.313$. In the case of the consumption-GDP ratio, the data show a slight trend, so we use the last observation of the pre-financial-innovation regime (1996Q4). This implies $\bar{c} = 0.670$. To make these values of $\bar{b}$ and $\bar{c}$ consistent with the resource constraint in the average of the stochastic steady state, we need to take into account the fact that investment and government absorption are included in the data but not in the model. To adjust for this discrepancy, we introduce a fixed, exogenous amount of autonomous spending A, so that the long-run average of the resource constraint is $1 = \bar{c} + A - \bar{b}(R-1)/R$. Given $\bar{b} = -0.313$, $\bar{c} = 0.6707$ and $R = 1.0266$ the value of A follows as a residual $A = 1 - \bar{c} + \bar{b}(R-1)/R = 0.321$.

The Markov process for z is set to approximate an AR(1) process $(\ln(z_t) = \rho \ln(z_{t-1}) + e_t)$ fitted to HP-filtered real U.S. GDP per capita using data for the period 1965Q1-1996Q4. The estimation yields $\rho = 0.869$ and $\sigma_e = 0.00833$, which imply a standard deviation of TFP of $\sigma_z = 1.68$ percent. We use Tauchen and Hussey (1991)'s quadrature method to construct a Markov approximation to this process assuming a vector of 9 realizations. The transition probability matrix and realization vector are available on request.

---

[13]Since the model with a single financial regime set at κ^l (i.e., the pre-financial-innovation regime) yields a collateral constraint that is almost always binding and a negligible excess return on land, we use the approximation $E[R^q] \approx R$, and then conditions (12) and (13) imply: $\alpha = (ql/zl^\alpha)[R - 1 + \beta^{-1}(1 - \beta R)(1 - \kappa^l)]$

[14]Consumption and GDP data are from the *International Financial Statistics* of the IMF.

Table 1: Parameters

β	Discount factor (annualized)	0.91
σ	Risk aversion coefficient	2.0
c	Consumption-GDP ratio	0.670
A	Lump-sum absorption	0.321
r	Interest rate (annualized)	2.660
ρ	Persistence of endowment shocks	0.869
σ_e	Standard deviation of TFP shocks	0.008
α	Factor share of land in production	0.025
L	Supply of land	1.0
κ^h	Value of κ in the high securitization regime	0.926
κ^l	Value of κ in the low securitization regime	0.642
F_{hh}^a	True persistence of κ^h	0.964
F_{ll}^a	True persistence of κ^l	0.964
n_0^{hh}, n_0^{hl}	Priors	0.3

The counters of the beta-binomial distribution that determine the initial priors of F_{hh}^s and F_{ll}^s are calibrated as follows: First, for simplicity, we impose the symmetry condition $n_0 = n_0^{hl} = n_0^{hh} = n_0^{ll} = n_0^{lh}$, so that there is only one counter to calibrate. Second, we calibrate n_0 so that the model matches an estimate of the observed excess return on land relative to the risk free rate for 1997Q2, which corresponds to the one-period-ahead expected excess return in the first date of the financial innovation experiment (date 1 in the experiment corresponds to 1997Q1). The data proxy for this excess return is the 1997Q2 spread on the Fannie Mae residential MBS with 30-year maturity over the T-bill rate. This excess return was equal to 47.6 basis points.[15] The model matches this excess return with $n_0 = 0.3$.

Since the model calibration is at a quarterly frequency, ideally we would like to use excess returns for securities with a quarterly maturity. However, residential MBSs do not have such short-term maturities, because the underlying assets tend to be long-term mortgages. Still, using the spread for the 30-year Fannie Mae MBS is useful because it actually makes it harder for the model to generate optimism. This is because securities with a quarterly maturity would likely have sharply

---

[15]The source of this excess return quote is Bloomberg. One complication that arises with using the 30 year MBS is the prepayment risk that tends to widen spreads. We use "option-adjusted spreads" from the same source that are adjusted for prepayment risk. The unadjusted spread is 117.6 basis points. We use the adjusted spread since we do not explicitly model prepayment risk, and hence we cannot expect the model to capture the portion of the spread due to this risk.

lower spreads than the 30-year securities, and thus the latter can be taken as an upper bound for the more accurate spreads. But higher spreads imply higher values of n_0, which weaken the mechanism generating optimism and pessimism in the learning process. Thus, by calibrating the priors to match the excess returns of the 30-year MBS we are looking at a "lower bound" of the optimism that the model can generate.

Figure 5 shows the density functions of the initial priors of F_{hh}^s and F_{ll}^s for Beta distributions with three different (n_0^{ii}, n_0^{ij}) pairs. The $Beta(0.3, 0.3)$ distribution corresponds to the baseline calibration. In this case, the priors have a U-shaped distribution with most of the mass concentrated around 0 and 1. Since this case assumes $n_0 = n_0^{ii} = n_0^{ij}$, the distribution is symmetric with a mean of 0.5 (and a variance of 0.156). By contrast, consider the $Beta(1, 1)$ distribution, which implicitly assumes that at least one observation of switch and continuation of each κ regime has been observed. This distribution also has a mean of 0.5, but the distribution is uniform over the (0,1) range, and it has a much lower variance than the $Beta(0.3, 0.3)$ distribution (0.083 v. 0.156).

Figure 5: Beta Distributions
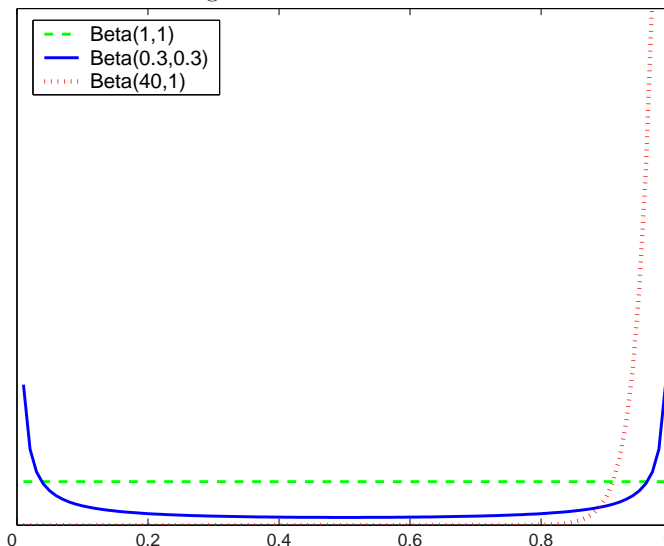
Notes: This figure plots the probability density function of the $Beta$ distribution with different n_0^{ii} and n_0^{ij} where $Beta(n_0^{ii}, n_0^{ij})$.

Figure 5 also plots the $Beta(40, 1)$ distribution, which matches the beliefs about F_{hh}^s that the learning process generates at period 40 of the financial innovation experiment. At this point, agents have observed 40 transitions from κ^h to κ^h and thus form beliefs characterized by a distribution that is highly skewed to the right, with most of the mass concentrated around 1. This reflects the high degree of optimism that the learning process can create.

We illustrate further how the initial priors yield optimistic and pessimistic beliefs by studying the evolution of the conditional means of the $Beta$ distributions of F_{hh}^s and F_{ll}^s over time as the sequence of realizations of κ is observed. Setting symmetric initial priors for F_{hh}^s and F_{ll}^s with a low value of n_0, as with the $Beta(0.3, 0.3)$ baseline, implies that agents conjecture that there are four "most likely" scenarios before the first realization of κ is observed: *a)* Both the high- and low-leverage regimes are extremely persistent, i.e., $F_{hh}^s \approx 1$ and $F_{ll}^s \approx 1$; *b)* The high-leverage regime is extremely persistent and the economy switches to the low-leverage regime rarely and for a short time, i.e., $F_{hh}^s \approx 1$ and $F_{ll}^s \approx 0$; *c)* The low-leverage regime is extremely persistent and the high-leverage regime occurs rarely and has a short duration, i.e., $F_{hh}^s \approx 0$ and $F_{ll}^s \approx 1$, *d)* Neither regime is persistent and the economy constantly moves between the two, i.e., $F_{hh}^s \approx 0$ and $F_{ll}^s \approx 0$. After observing the first few realizations of κ^h, however, the agents can rule out scenarios *c)* and *d)*.

Figure 6: Evolution of Beliefs under Alternative Priors
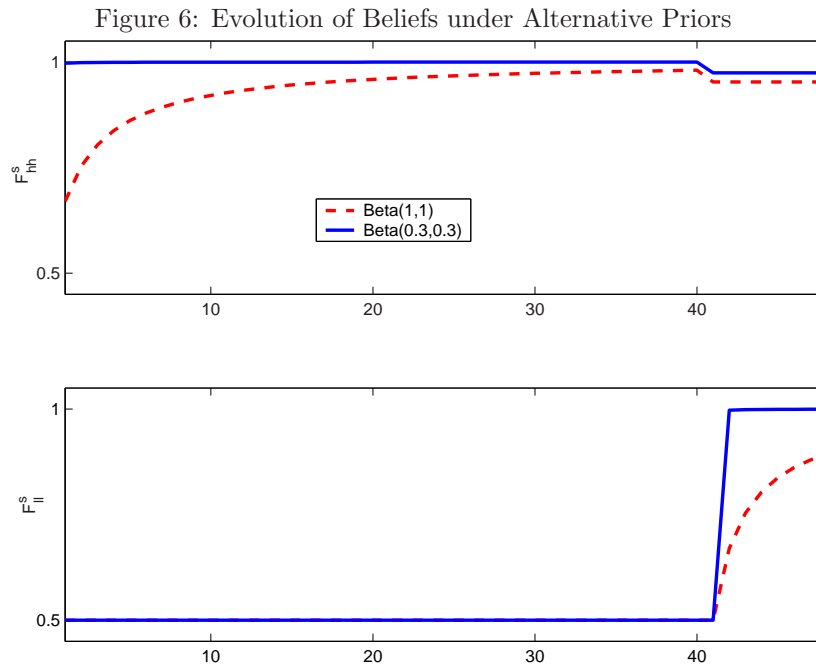
Notes: $Beta(0.3, 0.3)$ corresponds to our baseline calibration while $Beta(1, 1)$ shows beliefs if the priors were uniformly distributed.

Figure 6 plots the conditional averages of F_{hh}^s and F_{ll}^s corresponding to the beliefs in each of the 48 periods of the learning experiment, using both $Beta(0.3, 0.3)$ and $Beta(1, 1)$ as date-0 priors. The plots start at date 1 after the first realization κ^h has been observed. The bottom panel shows that, as discussed before, unless the economy switches to the low-leverage state, the agents cannot learn about the persistence of that state. Hence, their beliefs about this state remain unchanged at the initial prior of 0.5 for the first 40 periods. In contrast, the top panel shows that observing

the long spell of κ^h leads agents to update their beliefs about the persistence of this regime, and they become optimistic very quickly. In the baseline $Beta(0.3, 0.3)$ case, the average F_{hh}^s moves very close to 1 in just one quarter, while with the $Beta(1, 1)$ priors the buildup of optimism is more gradual, but still after 8 quarters the average F_{hh}^s approaches 90 percent. This rapid adjustment of beliefs also occurs with the surge of pessimism that follows the first observation of κ^l in period 41: With $Beta(0.3, 0.3)$ priors, agents update their average average F_{ll}^s from 0.5 to almost 1 in period 41, and with the $Beta(1, 1)$ priors the change is slower but again by period 48, the mean of F_{ll}^s is approaching 90 percent.

It is important to note that neither $Beta(0.3, 0.3)$ or $Beta(1, 1)$ introduce bias in the initial priors in favor of optimism or pessimism. This differs from the approach followed by Cogley and Sargent (2008), who studied the implications of inducing initial pessimism into the agents' priors. In our calibration, agents are not optimistic prior to period 1 because $Beta(0.3, 0.3)$ yields initial beliefs with average continuation probabilities of $F_{hh}^s = F_{ll}^s = 0.5$. This $Beta$ distribution does imply that agents' initial beliefs consider as "most likely" one of the four initial scenarios a)-d) mentioned above (i.e., they believe that either the switches in κ will be infrequent, as in scenarios a)-c), or that there will be a switch every period, as in scenario d)), but there is no initial bias in favor of either κ^h or κ^l.

At this point we have calibrated all of the parameters that are needed for solving the model. Notice in particular that the true transition probability matrix of κ (F^a) is not needed. Solving the COPs of the agents requires the sequence of beliefs about the transition probability matrix ($\{f(F^s \mid \kappa^t)\}_{t=0}^T$), which is determined with the parameters we already set. Still, calibrating the true transition probability matrix is necessary if we want to evaluate the macroeconomic effects of the imperfect information friction by comparing the solutions against the standard RE solution in which the "true" transition matrix of the financial regimes is known.

We calibrate F_{hh}^a so that the mean duration of high-leverage regimes is in line with the estimated duration of credit boom episodes in industrial economies, which Mendoza and Terrones (2008) estimated at about 7 years. This implies $F_{hh}^a = 0.964$. With this calibration of F_{hh}^a and conditional on observing κ^h at date 1, the probability of observing κ^h the following 39 periods is 0.241. Thus, the "true" probability of observing the long spell of κ^h that we assume in our financial innovation experiment, and that produces substantial optimism, is about 1/4. We assume a symmetric process by setting $F_{ll}^a = 0.964$.

An interesting implication of calibrating the "true" process of κ in this way is that the model features a long-run credit cycle consistent with average duration features of actual credit cycles, so that agents eventually learn that the economy will display a credit cycle with the duration and frequency that is typical of industrial countries. In the short-run, however, their expectations can deviate sharply from these regularities, along the lines of Reinhart and Rogoff (2009) "this-time-is-different" argument.

3.2 Quantitative Findings

The quantitative analysis is based on four sets of results derived from the numerical solutions: Long-run distributions of bond positions, forecast functions of macroeconomic aggregates, average changes in these aggregates at the "turning points" of the learning experiment, and expected excess returns. We compare the results of the learning model with the RE model (i.e., a model which retains the collateral constraint with its credit externality and debt deflation mechanism, but does not have a learning friction) and with a fixed price-learning (FPL) model in which land in the collateral constraint is valued at a constant price set to the long-run average (i.e., a model that retains the learning friction but removes the credit externality and the Fisherian deflation channel). In this case, the collateral constraint becomes $\frac{b_{t+1}}{R} \geq -\kappa_t \bar{q} l_{t+1}$.

3.2.1 Ergodic distributions

Figure 7 plots "conjectured" ergodic distributions of b for dates 1, 8, 40, 41 and 48 in the learning model and the true ergodic distribution of the RE model. We label the former as "conjectured" because the actual ergodic distribution of the learning model is the same as the one of the RE model, since in the long-run agents learn the true regime-switching process F^a, and thus the long-run equilibrium is the same as under rational expectations. The "conjectured" ergodic distributions for the other dates in the learning experiment are the agents' projections, or conjectures, of what the long-run equilibrium would look like if they forecast the dynamics of the financial regime using their current beliefs (i.e., these distributions assume that the corresponding period's beliefs about the transition probability matrix of κ are correct). Appendix B provides details on the computation of long-run distributions in the learning and RE models.

The conjecture that the beliefs are correct is clearly not valid in the learning model's equilibrium dynamics, because in the short-run beliefs do change from one period to the next and deviate sharply from the true transition probability matrix. Plotting the conjectured and RE long-run distributions

is useful, however, for illustrating the impact of the optimism and pessimism driving the model's dynamics on the agents "willingness" to borrow or save.

Figure 7: Ergodic Distributions of Bond Holdings
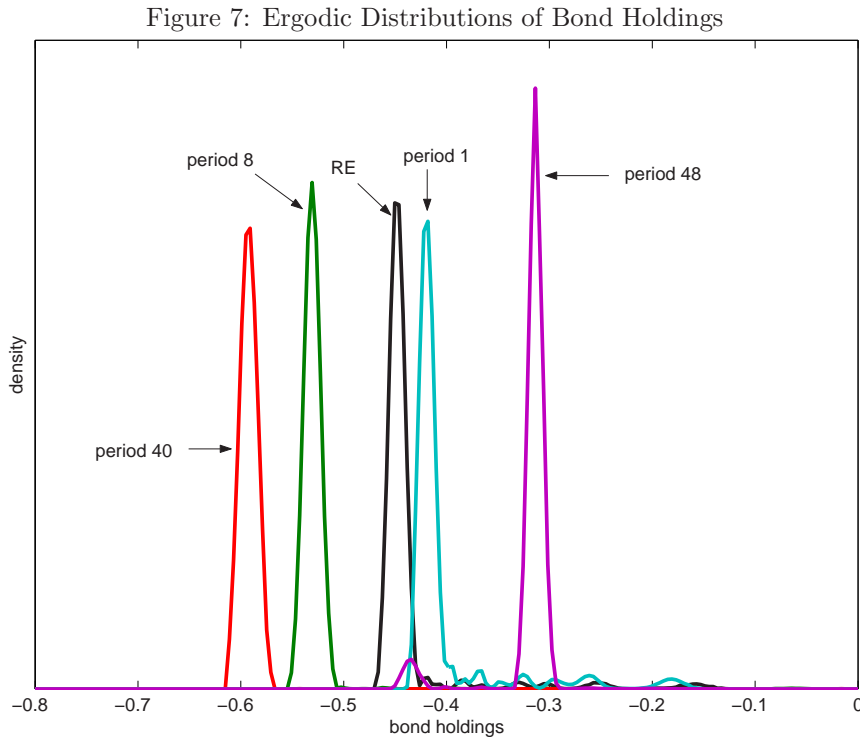
Notes: This figure plots the ergodic distribution of bond holdings implied by the learning model in periods 1 (initial period), 8, 40 (peak of optimism), 41, and 48 as well that of the rational expectations model marked by "RE."

Consider first the conjectured distribution for period 1. Recall that the mean of bond holdings pre-financial-innovation was -0.31, so we can tell that already by period 1 agents conjecture that the support of the long-run distribution of bonds will shift to the left (i.e. support higher debt levels), but not yet by as much as in the long-run RE equilibrium.

Compare now the RE ergodic distribution with the conjectured ergodic distribution for period 40. Large debt ratios (bond holdings in the interval [-0.62, -0.48]) are never a long-run equilibrium outcome in the RE model, but they are projected to be on the basis of the agents' period-40 beliefs. But this happens also much earlier than date 40 because, as shown before, it takes observing only the first few realizations of κ^h for agents to turn very optimistic relative to the RE transition probabilities. By period 8 agents already conjecture that debt positions in the [-0.58,-0.50] range are probable long-run equilibria, while in the RE ergodic distribution they have zero probability.

As optimism builds, the highest debt conjectured to have positive long-run probability rises, and the mass of probability assigned to debt levels larger than the largest debt under rational

expectations also rises. This process peaks at the peak of the optimistic phase in date 40. In short, during this phase, agents are willing to "overborrow" (take on more debt at the left tail of the conjectured ergodic distributions of b) than what is ever optimal in the RE model, and "undersave" (build less precautionary savings, or conjecture they can attain a lower average of b) than what is optimal in the RE model. When the first realization of κ^l hits and the pessimistic phase starts, the opposite effects take over and they peak at date 48. By then, agents are "underborrowing" and "oversaving" substantially.

3.2.2   Forecast functions

Forecast functions are useful for illustrating the model's equilibrium dynamics during the 48 periods of the learning experiment. Recall that in the calibration we defined a trajectory of realizations κ^T with $T = 48$, the first 40 are equal to κ^h and the last 8 are equal to κ^l, and we used the calibrated values of the learning process to compute the sequence of beliefs $\left[f(F^s \mid \kappa^t) \right]_{t=1}^{48}$. Given this information, we solved the COPs of each of these 48 periods to obtain bond decision rules associated with each set of beliefs. Given these bond decision rules, we constructed forecast functions that trace the dynamics of the expected values of the endogenous variables along the model's equilibrium path by chaining the decision rules that correspond to each period's beliefs and each realization of κ.

Intuitively, the algorithm that computes the forecast functions uses a law of motion for the evolution of the probability of the economy being in each triple (b, z, κ) as we move from date 0 to date 48. This law of motion can be computed for any triple of initial conditions (b, z, κ) in the state space, but we are interested in the triple that approximates the state of the U.S. economy in 1996Q4 (i.e. the initial conditions at the beginning of date 1 in the financial innovation experiment). Thus, we start at date 1 with all the probability concentrated in the coordinate of initial conditions (b_1, z_1, κ^h) where $b_1 = -0.345$ (the observed net credit assets as a share of GDP for 1996Q4) and $z_1 = 1$. Then, for each subsequent date, the value of κ is set to the corresponding realization in the κ^T sequence (κ^h for $t = 2, ..., 40$ and κ^l for $t = 41, ..., 48$), the transitions across values of z are computed with the Markov process of z, and the transitions across points in the state space of b are governed by the policy function $b'_t(b, z, \kappa)$ of the date-t COP. The procedure is similar to the standard forecast functions of a RE model, except that the policy function is time-varying because it varies with each set of beliefs in the sequence $\left[f(F^s \mid \kappa^t) \right]_{t=1}^{48}$ (see Appendix B for details).

29

Figure 8: Forecast Functions
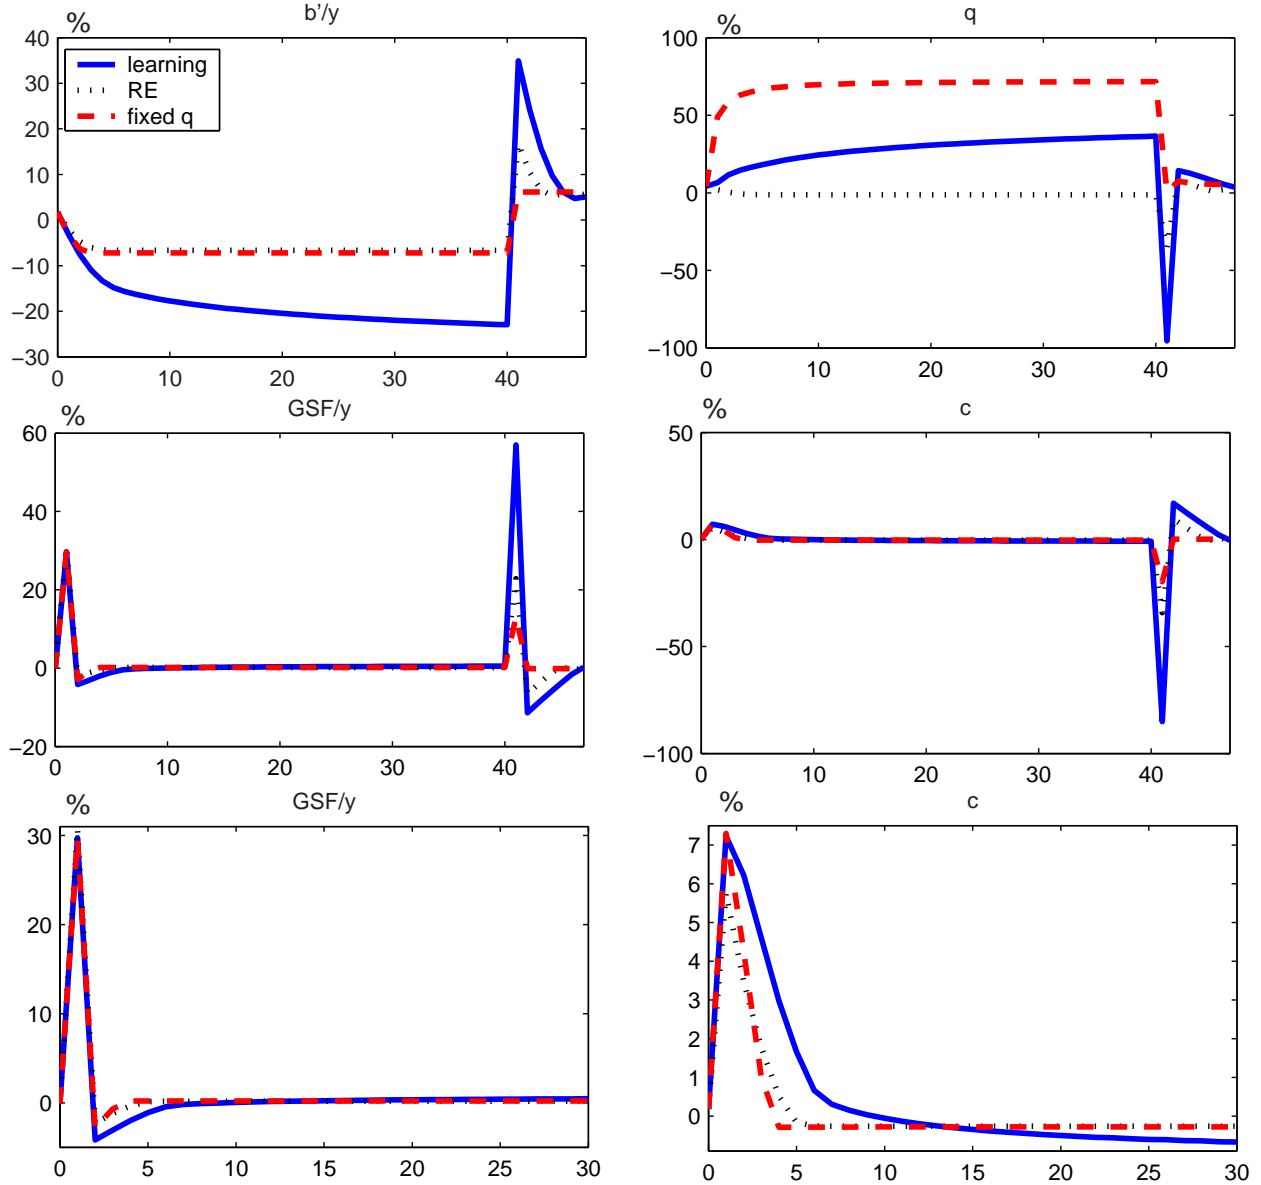
Notes: This figure plots the forecast functions of bond holdings-output ratio, price of land, consumption, and gross saving flow-output ratio (GSF/y) as percentage deviations from their long run means implied by the rational expectations scenario. GSF/y is calculated as $((b'/R) - b)/y$. Realizations of κ are as described in the text, κ^h in the first 40 periods and κ^l in the remaining 8. Date-0 b'/y is the 1996:Q4 observation from data (since debt data are end of period basis), so that the date-1 b'/y is the first endogenous choice of b' under κ^h, given an initial state determined by the data point from 1996:Q4. "Fixed q" refers to the scenario with the asset price on the right hand side of the credit constraint fixed at 0.456 which is the long run average of prices.

Figure 8 plots the forecast functions for the choice of bond holdings as a share of output (b'/y), consumption, the price of land, and the savings rate (GSF/y) as percent deviations from their long-run means in the learning, RE and FPL models. Recall that long-run means in the learning and RE models are identical because the ergodic distributions of the two are the same. The solid (blue) lines correspond to the learning model, the dashed (red) lines are for the FPL model, and the dotted (black) lines represent the RE model. Note that even the RE model generates some dynamics in this exercise, because the initial condition b_1 is not the long-run mean of the new financial regime with stochastic κ, and also because we are using a particular time series of realizations of κ (instead of averaging across possible κ realizations at each date t). Thus, these forecasts functions are conditional on the particular history κ^T that we assumed.

The forecast functions for bonds in the top-left panel show that during the optimistic phase agents consistently borrow more in the learning model than in both the RE and the FPL models. In the first two periods after financial innovation is introduced, the three models predict similar debt dynamics, but after that the optimism and the debt-deflation feedback loop at work in the learning model produce a much larger decline in bond holdings, while the bond dynamics in the RE and FPL models are similar.[16] Bond holdings as a share of output decline by as much as 23 percentage points below the long-run average at the peak of optimism of the learning model in period 40. These dynamics are in line with the downward trend in net credit market assets observed in the data. Interestingly, the combination of the learning friction and the debt-deflation channel delivers a much stronger decline in assets than the alternatives that retain only one of the two mechanisms. In the RE model there is no buildup of optimism to push for overborrowing, and in the FPL model there is no endogenous feedback from higher land prices into higher collateral and thus higher borrowing ability.

The switch to the pessimistic phase in period 41 brings about a large correction in bond holdings, which bounce back about 58 percentage points in the learning model. An adjustment of this magnitude is an equilibrium outcome, despite CRRA preferences and incomplete markets, because the arrival of the first realization of κ^l at date 41 is almost like a large, unexpected shock, in the sense that by date 40 agents believed that the state κ^h in which they had been living was almost absorbent (i.e., the average of F_{hh}^s at date 40 was very close to 1). Moreover, this large κ

---

[16]This occurs because (a) in the FPL model the price is fixed at the long-run average that is very similar to the RE model, and (b) the RE price displays very small deviations from its long-run average. As a result, and since the values of κ are the same in both models, the debt allowed by the collateral constraint in both models is about the same.

shock triggers a large Fisherian deflation effect (see below), which contributes to enlarge the debt adjustment. Bond holdings also jump up in the RE and FPL models, because of the switch from κ^h to κ^l in a state in which the collateral constraint was binding. But the adjustments are much smaller. The debt reversal in the RE case is about half the size of that in the learning model, and it reflects the effect of the debt-deflation mechanism in the absence of a switch to pessimistic beliefs. The FPL model yields the smallest correction, which isolates the effect of the switch to pessimistic beliefs without amplification due to the Fisherian deflation channel.

As agents overborrow during the optimistic phase in the learning model, they also bid more aggressively for the risky asset. This increases the price of land significantly, as shown in the top-right plot of Figure 8. This contrasts with the RE case, in which the price of land *declines* relative to the pre-financial-innovation price that prevailed at date 0. This occurs because the price of land in the RE model is falling to a lower long-run average in the financial innovation regime. In turn, the mean price of land in the RE model (with stochastic κ) is lower than in the pre-financial-innovation regime (with a constant κ^l) because, even though agents know the true distribution of κ, they now face uncertainty with regards to κ, since it is now a random variable. Hence, financial innovation implies a higher mean κ but also a higher variance of κ. The former enables the agents to borrow more, and therefore demand more of the risky asset and bid up its price, but the latter gives them an incentive to hold less of the risky asset, because the new financial environment is riskier and they are risk averse. We find that, if the gap between κ^h and κ^l is small, the "mean effect" dominates leading to higher land prices in the RE model, but as the gap widens, the "variance effect" becomes stronger and the mean land price in the RE model is lower than in the pre-financial-innovation equilibrium (as is the case in our baseline calibration).

The FPL model generates a larger asset price boom during the optimistic phase and a smaller price crash compared with the other two models. This is because the FPL model rules out the Fisherian deflation by construction, and hence at date 41 the downward spiral on land prices, collateral values, and debt that is at work in the learning model is not active. Moreover, the fixed land price for collateral valuation also serves as a limited asset price guarantee, which produces a larger price boom during the optimistic phase than in the learning and RE models. The guarantee is limited in the sense that it is not a guarantee on the price at which land is traded, but only on the price at which land is valued for collateral. Accordingly, the FPL model produces a smaller reversal in debt in period 41, as agents are able to borrow more than in the other two models

because of the constant land price for collateral. For the same reason, the larger land price increase in the optimistic phase does not feed back into a large debt expansion.

The center and bottom panels of Figure 8 show the forecast functions for the savings rate and consumption. Because of the large magnitude of the changes that occur at date 41, we split these plots into two pairs. The center pair shows dynamics for the full 48 periods of the experiment. The bottom pair shows only the first 30 periods. In studying these plots, it is critical to recall that the forecast functions show the effects of the learning experiment on macro variables for the given history of realizations κ^T, averaging across TFP shocks, and starting from average productivity and the 1996Q4 observation of b.

Consider first what consumption dynamics would look like in a perfect-foresight model where we switch from the constrained pre-financial innovation steady state with κ^l to a hypothetical financial innovation deterministic steady state for a regime with κ^h. The two steady states are well-defined because $\beta R < 1$, and hence the steady state of bonds is $b = -\kappa q(\kappa)$, where $q(\kappa)$ is the steady state land price, which is increasing in κ.[17] Thus, the increase in κ yields a lower steady state for b (higher steady state debt) because both κ and $q(\kappa)$ increase. But higher steady state debt means lower steady state consumption, since the non-financial wealth of the economy is invariant to changes in κ and the debt has to be serviced. Thus, financial innovation tilts the time profile of consumption. On impact, when κ is first increased, and for a few periods after that, consumption rises above the pre-financial-innovation level, but then it drops monotonically until it reaches its new steady state below the pre-financial-innovation level. This consumption tilting effect is also at work in the stochastic model, but is weaker because of the precautionary savings motive, which implies a smaller decrease in bond holdings and a smaller drop in consumption.

Now consider the forecast functions of the learning, RE and FPL models showing consumption dynamics in the center- and bottom-right plots of Figure 8. The fact that the dynamics for the first 40 periods are similar in all three models indicates that the consumption tilting effect dominates those dynamics. This is because consumption converges quickly to its new long-run average (which is identical in the learning and RE models, and very similar in the FPL model). There is over-consumption in the learning model relative to the RE and FPL models in the early stages after the switch to κ^h, because of the larger increase in debt (i.e., decline in bond holdings). In the first two periods, consumption is about the same in all three models, but the overconsumption in the

---

[17]The closed-form solution for the steady state land price with the collateral constraint binding is: $q(\kappa) = (\alpha\beta)/[\beta(R-1) + (1-\beta R)(1-\kappa)]$.

learning model is clear between the 3rd and 10th periods. After period 10, however, the dynamics driven by consumption tilting dominate in all three models. Consumption then remains smooth (as we are averaging across TFP and keeping κ constant at κ^h), until we arrive at date 41 and κ switches to κ^l.

At date 41, as explained earlier, the κ switch is almost like a large, unexpected shock in the learning and FPL models. In the learning model, which also has the Fisherian deflation, this produces a dramatic collapse in consumption. This is in line with the findings in Mendoza (2010) and Mendoza and Smith (2006), showing that in Fisherian deflation models there are equilibria outside the ergodic distribution of wealth, where the economy could land as a result of unexpected shocks, in which the impact response of consumption can be around -80 percent. In those models, however, precautionary savings and perfect information about the Markov processes of shocks rule out consumption drops of that magnitude from the equilibrium dynamics, while in our model the learning friction allows us to support them as short-run equilibria.

The RE and FPL models also produce large consumption declines when the economy switches to κ^l, but both are significantly smaller than in the learning model. In the RE model this is again because precautionary savings and the lack of overborrowing prevented a large accumulation of debt in the optimistic phase. In the FPL model the smaller consumption drop occurs because there is no Fisherian deflation of collateral values, which yields the smallest correction in debt, and hence implies the smaller consumption drop.

Figure 9 shows the evolution of the shadow value of the collateral constraint, expressed in terms of an implicit endogenous interest rate premium that measures the difference between the stochastic intertemporal marginal rate of substitution in consumption $(u'(c_t)/\beta E_t^s[u'(c_{t+1})])$ and the real interest rate R. Using condition (12), we can express this premium as $\mu_t/(u'(c_t) - \mu_t)$. Thus, if the constraint does not bind, there is no interest rate premium, and when it binds the premium rises as the constraint becomes more binding.

The dynamics of the interest premium confirm our previous argument stating that, when financial innovation starts, the constraint becomes nonbinding, and then it begins to bind after some time. In particular, in the baseline learning model the constraint begins to bind after period 5. Then it increases monotonically, at a decreasing rate, to converge to about 6 percent at the peak of the optimistic phase. In contrast, the FPL model generates a larger premium of up to 7 percent, while the RE model generates a premium of just above 2 percent in the optimistic phase. This is natural because in the FPL model rising land prices do not contribute to relax the borrowing limit,

and in the RE model the constraint is less binding because individuals desire to save more with rational expectations than with optimistic beliefs.

Figure 9: Interest Rate Premium
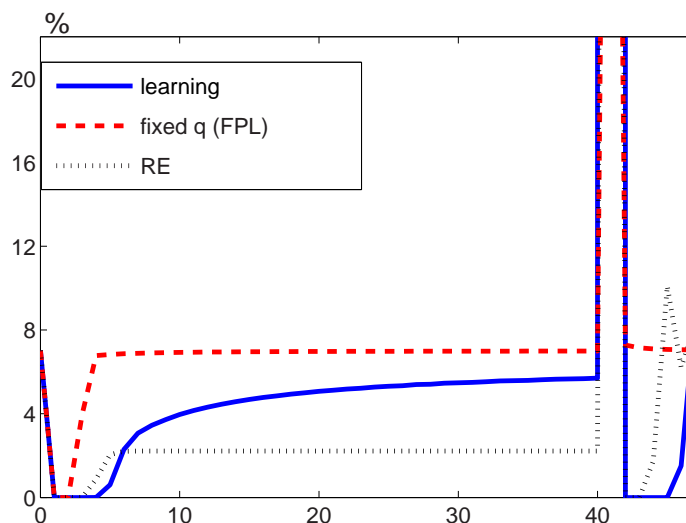
Notes: This figure plots the interest rate premium that represents the difference between the intertemporal MRS and the risk-free rate, which can be simplified to $\frac{\mu_t}{u'(c_t) - \mu_t}$. The premium in period 41 is 7064, 199, and 80 percent in baseline, RE, and FPL scenarios, respectively.

When the switch to the pessimistic phase takes place at date 41, there is a large jump in interest premia, in line with the large reversals in debt and consumption. The correction is the largest in the learning model, followed by the RE model, and the FPL model last. After date 41, however, the constraint becomes nonbinding for 4 periods in the learning model and for 1 period in RE. Afterwards the interest premia become positive, hovering around 7 percent in the FPL model throughout the rest of the experiment and gradually increasing to reach 8.8 and 8.2 percent in learning and RE models respectively.

3.2.3 Turning Points

Table 2 lists short-run changes in average bond-output ratios and land prices, calculated with the data used to construct the forecast functions, at the key turning points: the peak of optimism at date 40 relative to the pre-financial-innovation initial conditions, and at the end of the learning experiment relative to the peak of optimism (which we label as *financial crisis*). The figures shown in this table are the differences in the levels of b/y and q projected by the forecast functions, but not expressed in deviations from long-run means (as was the case in the plots of Figure 8).

Table 2: Average Changes at the Turning Points

	Data	RE	FPL	Baseline
Peak of Optimism:				
$E[(b/y)_{40} - (b/y)_0]$	-0.355	-0.083	-0.087	-0.246
$E[(ql/y)_{40} - (ql/y)_0]$	0.280	-0.025	0.307	0.147
Financial Crisis:				
$E[(b/y)_{48} - (b/y)_{40}]$	0.023	0.122	0.133	0.278
$E[(ql/y)_{48} - (ql/y)_{40}]$	-0.149	0.013	-0.303	-0.146

Note: Data column reports the difference between 2006Q4 and 1996Q4 observations in the top panel and the difference between 2008Q4 and 2006Q4 observations in the bottom panel. In columns 2-4 the realizations of κ are set to the path described in the text. Period 0 in all three scenarios corresponds to the 1996Q4 data observations, which are the initial conditions. qL/y is the aggregate market value of residential land divided by output.

This table illustrates two main results. First, the learning model can explain a large fraction of the observed increase in debt and land prices before the financial crisis. Second, the learning model generates significantly more debt in the optimistic phase than the RE or FPL models, and a much larger land price increase than the RE model.

The learning model can explain 69 percent of the increase in household debt observed in the data, since in the model the b/y ratio falls by almost 25 percentage points v. 36 percentage points in the data. The decline in bond holdings in the learning model is about 16 percentage points of GDP larger than in the RE or FPL models. The comparison with the RE model shows again that financial innovation, when agents are uncertain about the true nature of the new financial environment, produces significant overborrowing relative to what RE predicts. The comparison with the FPL model shows, also in line with our previous findings, that the interaction of the learning friction with the debt-deflation channel has significant quantitative implications for the size of the credit boom that the model can produce.

Comparing the changes in land prices, we find that the baseline learning model accounts for about 53 percent of the land price boom observed in the data (the increase in the price of land in the model reaches almost 15 percentage points at date 40, v. 28 percentage points in the data). In line with what we noted in the comparison of forecast functions, the RE model yields a slight fall, instead of an increase, in the price of land during the optimistic phase, and the FPL model generates a larger price increase than the learning model.

Consider now the changes in bond holdings and land prices during the financial crisis. The baseline learning model generates a large debt reversal of about 28 percentage points (and this after an even larger reversal between periods 40 and 41, as shown in Figure 8). By contrast, in the data the correction was only 2.3 percentage points. The model clearly overestimates the reversal in debt, but part of the discrepancy is due to the fact that the bond in the model is a one-period bond while the average maturity of household debt data is significantly higher than a quarter. This makes an important difference because in the model agents repay and re-finance their debt every period, but in the data this is not the case, particularly with long-term debt contracts such as 30-year mortgages. Hence, in the model, a switch to the low-leverage regime leads to an abrupt decline in debt, while in the data this is not going to affect immediately the outstanding stock of long-term debt.

The model does a good job at matching the observed decline in land prices during the financial crisis (14.6 and 14.9 percentage points in model and data respectively). This is after an initial price collapse between periods 40 and 41 that is significantly larger than what the model predicts between periods 40 and 48. In contrast, the FPL model now produces a larger price decline, about twice as large as in the data, and the price change in the RE model is again small and in the opposite direction from both the learning and FPL models.

3.2.4 Excess Returns on Land

Next we investigate the projections of excess land returns that underlie the discounting of future land dividends in land pricing, in order to illustrate further the agents' perception of the riskiness of land during the optimistic and pessimistic phases. Figure 10 plots the $t + j$-period-ahead expected excess returns for $j =1,..,50$ periods ahead of three initial dates $t =1$, 40, and 41. These are expectations that agents form looking into the future given beliefs and decision rules as of periods 1, 40 or 41. In each scenario, we set the initial state of b to the mean bond holdings predicted by the forecast function in Figure 8 for the corresponding date, κ to its corresponding value in the history κ^t, and TFP to its mean value.

Focusing on expected excess returns projected as of date 1, in the top-left panel of Figure 10, the excess returns in the RE model exceed those of the baseline learning setup for the first 13 periods. Staring from period 14, this ordering reverses. This pattern justifies the result showing that the land price at date 1 is lower in the RE model than in the baseline learning model (because agents in the latter expect relatively lower excess land returns in the first 13 periods, which carry

more weight in discounting land dividends–and recall that land dividends are simply driven by the exogenous TFP process, which is the same in all three models we are comparing). The FPL model yields expected excess returns that lie significantly below both the RE and baseline models, and this is consistent with the higher date-1 land price produced by the FPL model (see Figure 8). The FPL model has lower excess returns because the removal of the debt-deflation channel weakens the direct and indirect effects of the collateral constraint on excess returns shown in Equation (10).

Figure 10: Excess Returns
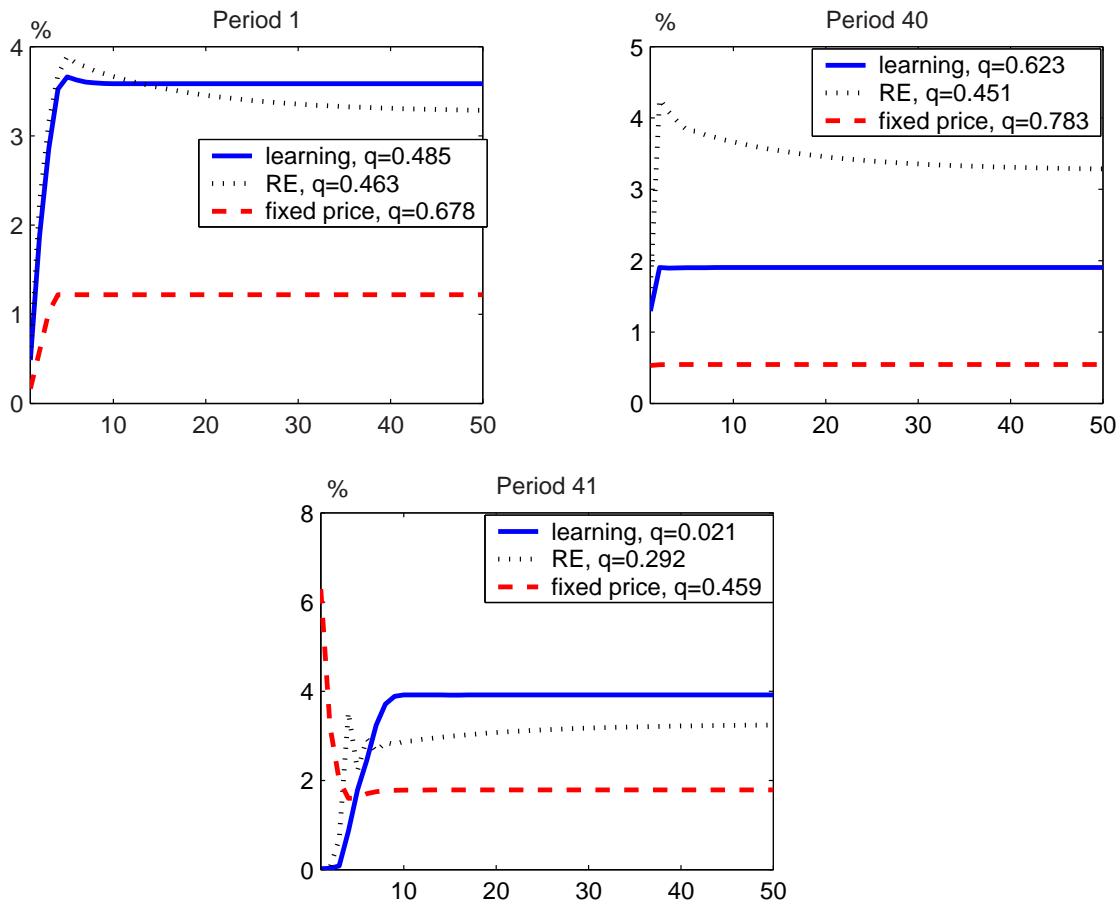
Notes: Expected excess returns for 50 periods ahead of initial dates $t =1$, 40, and 41, computed using the beliefs and associated equilibrium pricing function of the each date's optimization problem. The expected returns are conditional on the bond holdings predicted for each initial date by the forecast functions of Figure 8, the mean value of TFP ($z = 1$), and the value of kappa indicated in the history of realizations for each date t.

As agents turn very optimistic after observing the high-leverage regime persist for 40 periods, they become more willing to hold risky land at lower excess returns, and thus expected excess returns ahead of date 40 are lower in all three models relative to what they were ahead of date 1. Moreover, projected returns in the baseline learning model become significantly smaller than

in the RE model (see the top-right panel of Figure 10), and the FPL model predicts significantly smaller excess returns than both the RE and the learning model. This is because in the FPL model beliefs turn as optimistic as in the learning model, but the removal of the debt-deflation mechanism reduces risk premia on holding land.

In period 41, when the switch to the low-leverage regime takes place, the ordering of projected excess returns across the three models reverses (see the bottom panel of Figure 10). Excess returns in the RE model are projected initially to be higher than in the baseline learning model, but beyond seven periods ahead of date 41 the learning model predicts higher excess returns. Moreover, now the FPL model is the one that predicts the sharpest increase in excess returns in the first periods ahead of date 41, and then converges to significantly lower excess returns than both the RE and learning models. This pattern of projected excess returns is in line with the ranking of asset prices shown for date 41 in Figure 8, in which the fall in the price of land is the largest in baseline learning model, followed by the RE model, and the FPL model last.

3.2.5 Sensitivity Analysis

We now conduct a sensitivity analysis to study how changing the model's key parameters alters the magnitude of the turning point effects we just discussed. We focus on changes in the initial priors, the values of κ^h and κ^l, the discount factor and the interest rate. Columns 3 to 8 of Table 3 report the results, and Column 2 includes the results for the baseline learning model. Note that, in general, the parameterizations that generate larger booms during the optimistic phase also generate larger busts in the financial crisis.

Table 3: Sensitivity

	Baseline	$n_0 = 1$	$n_0 = 0.1$	$\kappa^l = 0.7$	$\kappa^h = 0.9$	$\beta = 0.92$	$R = 1.02$
$E[(b/y)_{40} - (b/y)_0]$	-0.246	-0.095	-0.276	-0.250	-0.236	-0.250	-0.259
$E[(ql/y)_{40} - (ql/y)_0]$	0.147	-0.011	0.180	0.151	0.154	0.150	0.165
$E[(b/y)_{48} - (b/y)_{40}]$	0.278	0.145	0.307	0.226	0.267	0.256	0.266
$E[(ql/y)_{48} - (ql/y)_{40}]$	-0.146	-0.015	-0.179	-0.111	-0.154	-0.111	-0.123

Note: The first column reproduces the baseline scenario results. The exercise conducted is the same as that explained in the note for Table 2 with different parameter values as indicated in column headings.

The third column of the table shows the results resetting the initial priors to $n_0 = 1$, which corresponds to the case with uniformly-distributed priors. In this scenario, debt dynamics are qualitatively the same as in the baseline scenario but the debt buildup is smaller (10 percentage points v. 25 in the baseline). The price of land, however, is about 1 percentage point lower in period 40 than in period 1, which is sharply at odds with the nearly 15 percentage points increase produced in the baseline case. Moreover, we found that with uniform priors, the land price follows a u-shaped trajectory in the optimistic phase, instead of the monotonically increasing path displayed in the baseline case. The reason for this different trajectory is that, throughout the optimistic phase, the beliefs about the persistence of the κ^h regime with the uniform priors are always lower than those in the baseline. As explained before, the means of the two distributions of priors are the same (0.5), but, with the baseline priors, agents turn optimistic more quickly after starting to observe κ^hs. The mean beliefs are the same in period 0, but the differences in the shape of the distributions of priors induce agents to become more optimistic, and faster, in the baseline scenario.

The more gradual buildup of optimism with the uniform priors affects the relative magnitude of the effects of higher mean and higher variance of κ on land prices post financial innovation. This explains the u-shaped trajectory of prices under uniform priors, because the effect of higher variance dominates that of higher mean, causing a decline in the price of land, until agents have turned optimistic enough. As agents observe more κ^hs, and sufficient optimism builds up, the higher mean dominates the higher variance, but under uniform priors this requires a longer sequence of κ^h than in the baseline case. Hence, if we look at an optimistic phase of more than 40 periods with the uniform priors, we again find that at the peak of the optimistic phase the price of land would be higher than in period 1.

The financial crisis effects on land prices and debt are also much weaker under the uniform priors than in the baseline scenario, again because of the more gradual adjustment of the priors (now in the switch to pessimistic beliefs and the buildup of pessimism). Debt rises by 15 percentage points instead of 28, and the price of land falls only by 1.5 percentage points, instead of nearly 15 percentage points.

Reducing the initial priors to $n_0 = 0.1$, which moves the distribution further away from uniform priors than in the baseline case, produces larger debt and land price booms. The size of the debt buildup is about 28 percentage points and the boom in the price of land reaches 18 percentage points. With these priors, the *Beta* distribution of initial priors has even more mass concentrated around 0 and 1 than in the baseline *Beta*(0.3, 0.3) case. Consequently, when agents observe the first

40

κ^h they turn more optimistic than in the baseline case, and hence they borrow more and demand more of the risky asset. Similar effects are at work, but in the opposite direction, in the pessimistic phase, and hence with $n_0 = 0.1$ we find a larger increase in bonds and a large drop in land prices in the financial crisis.

In contrast with what we found when changing initial priors, the magnitude of the movements in debt and land prices at the turning points in the baseline scenario are largely robust to changes in the values of κ^h, κ^l, β and R around their baseline values. Increasing κ^l (see the fifth column of Table 3) increases the size of the debt buildup and land price boom slightly. These slight changes occur, even though we still have the same sequence of 40 realizations of κ^h at the same value as in the baseline, because agents take into account the fact that with the higher value of κ^l the low-leverage regime is not as low as it was in the baseline. This results in both a higher mean and a lower variance of κ, which support both larger debt and higher land prices.

Reducing κ^h to 0.9 (sixth column of Table 3) reduces the size of the debt buildup slightly, because of the tighter credit constraint in the high-leverage regime that the lower κ^h represents. The land price boom is slightly larger, however, because the lower κ^h again reduces the variability of κ. Thus, both higher κ^l and lower κ^h increase land prices more than the baseline because both reduce the variance of κ, while they change the debt boom in opposite directions because higher κ^l increases the mean of κ but lower κ^h reduces it (relative to the baseline case).

The last two columns of Table 3 show the results of the sensitivity analysis for higher β and lower R. These two cases show similar implications, with both of them generating slightly larger debt and land price booms than in the baseline case during the optimistic phase. In contrast, the financial crisis effects with both higher β and lower R are slightly weaker than in the baseline.

4 Conclusion

The recent financial crisis in the United States was preceded by a decade of fast growth in household debt, residential land prices, and leverage ratios, accompanied by far-reaching financial innovation in terms of the introduction of new financial products and changes in the legal and regulatory framework of financial markets. In this paper, we argued that financial innovation in an environment with imperfect information and credit frictions was a central factor behind the credit and land price booms that led to the crisis. To make this point, we examined the interaction between financial

innovation, learning, and a collateral constraint in a stochastic equilibrium model of household debt and land prices.

We used the model to study the quantitative implications of an experiment calibrated to U.S. data in which financial innovation begins with a switch to a high-leverage regime, but agents do not know the true regime-switching probabilities across high- and low-leverage regimes. Agents are Bayesian learners, however, so in the long-run, after observing a long history of realizations of leverage regimes, they learn the true regime-switching transition probabilities. The collateral constraint introduces Irving Fisher's classic debt-deflation amplification mechanism, providing a vehicle for the waves of optimism and pessimism produced by Bayesian learning to have amplification effects on debt accumulation and land prices.

In our setup, a buildup of optimism is a natural consequence of financial innovation, because agents start without a sufficiently long time series of data to correctly evaluate the riskiness of the new financial environment. Calibrating the leverage regimes to data on the ratio of net household debt to residential land values, and the initial priors to the excess returns on the 30-year Fannie Mae MBS in early 1997, Bayesian learning predicts that agents would turn very optimistic, very quickly between the mid-1990s and the mid-2000s, after observing only a few quarters of the high-leverage regime.

The debt-deflation channel plays an important role because, as optimism built up and land prices rose, the agents' ability to borrow also grew. Similarly, when optimism turned to pessimism, after the first observation of the low-leverage regime, which we dated to the beginning of 2007, after the beginning of the sub-prime mortgage crisis in the Fall of 2006, the debt-deflation channel amplified the reversals in debt and in asset prices. This occurs because fire-sales of land drive down land prices and reduce the agents' ability to borrow.

The interaction of the learning friction and the debt-deflation mechanism generates a substantial amount of overborrowing, which accounts for over two thirds of the massive drop in net credit assets of U.S. households, and over one half of the boom in residential land prices, observed between 1997 and 2006. Moreover, the model also predicts a credit crunch, a crash in land values, a collapse in consumption and a surge in private savings when the economy experiences the first realization of the low-leverage regime. In contrast, the size of the debt and price booms, and the subsequent collapses, are significantly smaller in variants of the model that remove the learning friction (i.e., assuming rational expectations so that agents know the true regime-switching process of leverage

42

regimes) or the debt-deflation mechanism (i.e., leaving the learning friction in place but assuming that land prices used to value collateral are fixed at their long-run average).

Our work has important implications for the ongoing debate on financial reforms to prevent future financial crises. First, since by definition the true riskiness of a truly brand-new financial regime with new securities and new regulations cannot be correctly evaluated when the new regime starts, and little or no data is available on its performance, exposure to the credit boom-bust process we studied in this paper comes along with the potential benefits of financial innovation. Hence, close supervision of financial intermediaries in the early stages of financial innovation is critical. For example, capital requirements can be used to limit overborrowing, but they have to be used carefully, because tight regulatory constraints introduce additional distortions that undermine the benefits of financial innovation.

Second, our work suggests that there are limitations to the benefits of taxes or fees designed to manage "macroprudential risk." The argument for these taxes is that they work to correct the overborrowing that results from a systemic credit externality, by which agents fail to internalize the effect of their individual actions on the market prices that determine borrowing ability (by, for example, affecting the value of collateral assets). We showed here, however, that overborrowing can also result from optimistic beliefs, in our case due to imperfect information about the persistence probability of a high-leverage regime under a new financial environment. Assuming that policy-makers are as uninformed as households about how financial markets will perform after radical structural changes, taxes on debt can address overborrowing due to the credit externality, but cannot address overborrowing due to optimistic beliefs. Still, we do find that the overborrowing effect of learning without the credit externality is about one half of that produced when both learning and the credit externality are at work.

Third, the prospect of a new round of radical changes tightening the legal and regulatory framework of financial markets in industrial countries, which will affect the types of securities that will be available and the size of the markets in which they will trade, implies a new large financial innovation shock. Hence, agents once again will have to evaluate the riskiness of the new financial environment with subjective beliefs based on imperfect information. Thus, the risk exists that a few years of slow credit growth and poor performance in asset markets can lead to the buildup of pessimistic expectations that will hamper the recovery of financial markets.

References

Bernanke, Ben, 2005, "The Global Saving Glut and the U.S. Current Account Deficit," Sandridge lecture, Virginia Association of Economics.

Bianchi, Javier, 2009, "Overborrowing and Systemic Externatilities in the Business Cycle," mimeo, University of Maryland.

Blanchard, Olivier J., Jean-Paul L'Huillier, and Guido Lorenzoni, 2008, "News, Noise and Fluctuations: An Empirical Exploration," mimeo, MIT.

Boz, Emine, 2009, "Can Miracles Lead to Crises? The Role of Optimism in Emerging Market Crises," *Journal of Money, Credit, and Banking*, vol. 41, pp. 1189–1215.

Boz, Emine, Christian Daude, and Bora Durdu, 2008, "Emerging Market Business Cycles Revisited: Learning About the Trend," International Finance Discussion Paper 927, Federal Rreserve Board.

Campbell, Jeffrey R., and Zvi Hercowitz, 2009, "Welfare Implications of the Transition to High Household Debt?" *Journal of Monetary Economics*, vol. 56, pp. 1–6.

Cogley, Timothy, and Thomas J. Sargent, 2008, "The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression?" *Journal of Monetary Economics*, vol. 55, pp. 454–478.

Corbae, Dean, and Erwan Quintin, 2009, "Mortgage Innovation and the Foreclosure Boom," mimeo, University of Texas at Austin.

Davis, Morris A., and Jonathan Heathcote, 2007, "The Price and Quantity of Residential Land in the United States," *Journal of Monetary Economics*, vol. 54, pp. 2595–2620.

Drew, Jill, 2008, "The Frezy," *the Washington Post*, 12/16/2008.

Durdu, Ceyhun Bora, Enrique G. Mendoza, and Marco Terrones, 2009, "Precautionary Demand for Foreign Assets in Sudden Stop Economies: An Assessment of the New Merchantalism," *Journal of Development Economics*, vol. 89, pp. 194–209.

Edge, Rochelle M., Thomas Laubach, and John C. Williams, 2007, "Learning and Shifts in Long-Run Productivity Growth," *Journal of Monetary Economics*, vol. 54, pp. 2421–2438.

Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh, 2010, "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium," mimeo.

Gertler, Mark, and Nobuhiro Kiyotaki, 2010, "Financial Intermediation and Credit Policy in Business Cycle Analysis," mimeo, NYU and Princeton University.

Goetzmann, William N., and Frank Newman, 2010, "Securitization in the 1920's," Working Paper 15650, NBER.

Korinek, Anton, 2008, "Excessive Dollar Borrowing in Emerging Markets: Balance Sheet Effects and Macroeconomic Externalities," mimeo, University of Maryland.

Lipschitz, Leslie, Timothy Lane, and Alex Mourmouras, 2002, "Capital Flows to Transition Economies: Master or Servant?" Working Paper 02/11, International Monetary Fund.

Lorenzoni, Guido, 2009, "A Theory of Demand Shocks," *American Economic Review*, vol. 99, pp. 2050–2084.

Mendoza, Enrique G., 2010, "Sudden Stops, Financial Crises and Leverage: A Fisherian Deflation of Tobin's Q," *American Economic Review*, forthcoming.

Mendoza, Enrique G., and Vincenzo Quadrini, 2010, "Financial Globalization, Financial Crises and Contagion," *Journal of Monetary Economics*, vol. 57, pp. 24–39.

Mendoza, Enrique G., Vincenzo Quadrini, and Josè-Vìctor Rìos-Rull, 2009, "Financial Integration, Financial Development, and Global Imbalances," *Journal of Political Economy*, vol. 117, pp. 371–416.

Mendoza, Enrique G., and Katherine A. Smith, 2006, "Quantititive Implications of a Debt-Deflation Theory of Sudden Stops and Asset Prices," *Journal of International Economics*, vol. 70, pp. 82–114.

Mendoza, Enrique G., and Marco Terrones, 2008, "An Anatomy of Credit Booms: Evidence From Macro Aggregates and Micro Data," Working Paper 14049, NBER.

Nieuwerburgh, Stijn Van, and Laura Veldkamp, 2006, "Learning Asymmetries in Real Business Cycles," *Journal of Monetary Economics*, vol. 53, pp. 753–772.

Reinhart, Carmen, and Kenneth Rogoff, 2009, *This Time is Different: Eight Centuries of Financial Folly* (Princeton University Press).

Tauchen, George, and Robert Hussey, 1991, "Quadrature-based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," *Econometrica*, vol. 59, pp. 371–396.

Uribe, Martin, 2006, "On Overborrowing," *American Economic Review Papers and Proceedings*, vol. 96, pp. 417–421.

Warnock, Francis, and Veronica Cacdac Warnock, 2006, "International Capital Flows and U.S. Interest Rates," International Finance Discussion Paper 840, Federal Rreserve Board.

Appendix A   Solution Method (NOT FOR PUBLICATION)

We solve the model by an Euler equation method that combines price and policy function iterations using the land pricing equation and the general equilibrium conditions (12)-(15). By proceeding in this way, instead of solving the agents' Bellman equation, we avoid using aggregate states and iterations to converge on the representative agent condition matching individual and aggregate laws of motion for bond holdings.

The algorithm proceeds in these steps:

1. Define a history of realizations κ^T and calculate the sequence of posteriors $\{f(F^s \mid \kappa^t)\}_{t=1}^T$.

2. Take κ_1 and the date-1 posterior $f(F^s \mid \kappa^1)$ from the sequence in Step 1.

3. Guess a land pricing function $q_1(b_t, z_t, \kappa_t)$ and solve for the date-1 equilibrium conditions (12)-(15) using the posterior density function, $f(F^s \mid \kappa^1)$, and a policy function iteration algorithm.

4. Use the resulting policy functions $[b_1'(b, z, \kappa), c_1(b, z, \kappa), \mu_1(b, z, \kappa)]$ from Step 3 and the asset pricing equation (11) to compute a new pricing function $\hat{q}_1(b_t, z_t, \kappa_t)$. Note that we can use the current beliefs in computing this forward solution because the Law of Iterated Expectations still holds.

5. Compare $\hat{q}_1(b_t, z_t, \kappa_t)$ and $q_1(b_t, z_t, \kappa_t)$, if they satisfy a convergence criterion then the decision rules $[b_1'(b, z, \kappa), c_1(b, z, \kappa), \mu_1(b, z, \kappa)]$ and the pricing function $q_1(b_t, z_t, \kappa_t)$ are the solutions of the date-1 COP. If not, construct a new guess of the pricing function using a Gauss-Siedel rule and return to Step 3.

6. Move to the date-2 with history κ^2 and posterior $f(F^s \mid \kappa^2)$, taken from Step 1, and return to Step 3 in order to solve for the date-2 COP. Repeat for each date-t history κ^t and posterior $f(F^s \mid \kappa^t)$ for $t = 1, ...T$ solving each time for the corresponding date-t COP.

The passive Bayesian learning has important implications that can be useful in implementing the above algorithm:

1. The solutions to each date-t COP are not functionally related (i.e., solving a particular date-t problem does not require knowing anything about the solution for any other date). Thus,

the model can be solved by solving each date-t COP independently.[18] Still, we can speed convergence if, whenever $||f(F^s \mid \kappa^{t+j}) - f(F^s \mid \kappa^t)||$ is small enough under some metric, we use for the date $t+j$ COP initial guesses given by the date-t COP.

2. If $j \leq T$ is large enough for $f(F^s \mid \kappa^{t+j})$ to converge to F^a (for some convergance criterion), the solutions for all dates $t \geq j$ collapse to a standard recursive RE equilibrium using the true Markov process F^a.

3. Since the full equilibrium solution of the intertemporal sequence of allocations and prices from dates 1 to T is obtained by chaining the solutions of each date-t COP (for $t = 1, ..T$), one can think of solving the recursive equilibrium for a set of different histories $[\kappa_i^T]_{i=1}^I$, each supporting a different sequence of posterior densities $f(F^s \mid \kappa_i^t)_{t=1}^T$. We consider only one history κ^T because we take the stance that the financial innovation experiment we look at in the data can be represented by a particular history κ^T, intended to match the observed financial regimes between 1997 and 2007. The alternative would be to generate a set of I "potential" histories $[\kappa_i^T]_{i=1}^I$, which could be done using the true Markov process F^a, solve the model for each, and then take averages across these different solutions at each date t.

Appendix B   Computation of Ergodic Distributions, Forecast Functions, Excess Returns (NOT FOR PUBLICATION)

B.1 Ergodic Distribution and Forecast Functions under Rational Expectations

Define the *date t probability distribution* over bonds, productivity and collateral coefficients in the RE model as $\lambda_t(b, z, \kappa)$. The law of motion that governs the evolution of this distribution over time is:

$$\lambda_{t+1}(b', z', \kappa') = \sum_z \sum_\kappa \sum_{\{b: b' = g(b, z, \kappa)\}} \lambda_t(b, z, \kappa) \pi(z' \mid z) p(\kappa' \mid \kappa)$$

where $g(b, z, \kappa)$ is the policy function that sets the optimal decision rule for bonds, $\pi(z' \mid z)$ is the Markov transition probability for productivity shocks, and $p(\kappa' \mid \kappa)$ is the true Markov transition probability of κ (with the two Markov processes assumed to be independent). The unconditional limiting distribution of bonds, productivity and collateral coefficients is given by $\lambda(b, z, \kappa)$, and it represents the fixed point of the above law of motion. The algorithm computes the ergodic

---

[18]This fact can be used to develop a strategy to speed up the full solution of the model, because in a computer with n number of cores, we can solve n COPs for n different dates simultaneously.

distribution exactly in this way, by performing iterations of the law of motion until $\lambda_t(b, z, \kappa)$ and $\lambda_{t+1}(b', z', \kappa')$ satisfy a convergence criterion.

Forecast functions are averages of the model's endogenous variables computed at each date t using the corresponding distribution $\lambda_t(b, z, \kappa)$, starting from any initial condition (b_0, z_0, κ_0) with distribution $\lambda_0(b_0, z_0, \kappa_0) = 1$. By construction, just like iterations on the above law of motion of probabilities converge to the long-run distribution, forecast functions converge to unconditional long-run averages computed with the ergodic distribution, regardless of the initial conditions (as long as the ergodic distribution itself is unique and invariant).

Given $\lambda_t(b, z, \kappa)$, the *date-t conditional probability distribution* over κ^i for $i = h, l$ is defined as follows:

$$\tilde{\lambda}_t(b, z \mid \kappa^i) = \frac{\lambda_t(b, z, \kappa^i)}{\sum_b \sum_z \lambda_t(b, z, \kappa^i)}$$

Conditional forecast functions are averages for the models endogenous variables computed at each date t using the corresponding distribution $\tilde{\lambda}_t(b, z \mid \kappa^i)$. By construction, as $\lambda_t(b, z, \kappa^i) \rightarrow \lambda(b, z, \kappa)$, the date-$t$ conditional distribution $\tilde{\lambda}_t(b, z \mid \kappa^i)$ converges to the corresponding long-run conditional distribution $\tilde{\lambda}(b, z \mid \kappa^i)$. Moreover, conditional forecast functions of any endogenous variable converge to the corresponding conditional long-run average.

B.2 Forecast Functions in the Learning Model

The learning model has dynamics in the beliefs about the transition probability matrix of κ, and hence the RE definitions of conditional and unconditional forecast functions do not apply. Intuitively, one can construct a set of forecast functions and ergodic distributions by using the corresponding date$-t$ beliefs to form all the expectations about future states. In light of this, we define forecast functions in the learning model by averaging only over productivity shocks and tracking the decision rules produced at each date by the corresponding set of beliefs and the corresponding conditional optimization problem. Specifically, we compute forecast functions in the learning model as follows: Take as given (b_0, z_0, κ_0), then the relevant values of the forecast function of bonds in a

learning period of length T with a sequence of realizations $[\kappa_t]_{t=0}^T$ are:

$$\hat{b}_1 = E\left[b_1 \mid (b_0, z_0, \kappa_0), f(F^s \mid \kappa^0)\right] = h_0(b_0, z_0, \kappa_0; f(F^s \mid \kappa^0))$$

$$\hat{b}_2 = E\left[b_2 \mid b_0, f(F^s \mid \kappa^1)\right] = \sum_{z^1} \sum_{\{b_1:b_2=h_1(b_1,z_1,\kappa_1)\}} \pi(z^1 \mid z_0) h_1\left(b_1, z_1, \kappa_1; f(F^s \mid \kappa^1)\right)$$

$$\hat{b}_3 = E\left[b_3 \mid b_0, f(F^s \mid \kappa^2)\right] = \sum_{z^2} \sum_{\{b_2:b_3=h_2(b_2,z_2,\kappa_2)\}} \pi(z^2 \mid z_0) h_2\left(b_2, z_2, \kappa_2; f(F^s \mid \kappa^2)\right)$$

$$......$$

$$\hat{b}_{T+1} = E\left[b_{T+1} \mid b_0, f(F^s \mid \kappa^T)\right] = \sum_{z^T} \sum_{\{b_T:b_{T+1}=h_T(b_T,z_T,\kappa_T)\}} \pi(z^T \mid z_0) h_T\left(b_T, z_T, \kappa_T; f(F^s \mid \kappa^T)\right)$$

where $\pi(z^t \mid z_0) = \pi(z_t \mid z_{t-1})\pi(z_{t-1} \mid z_{t-2})...\pi(z_1 \mid z_0)$ is the probability of a particular history of realizations of productivity up to date t (for $t \geq 0$), $\left[f(F^s \mid \kappa^t)\right]_{t=0}^T$ is the sequence of beliefs, and $h_t\left(b_t, z_t, \kappa_t; f(F^s \mid \kappa^t)\right)$ is the optimal decision rule for bonds determined by the date-t COP using the date-t beliefs and evaluated for the states (b_t, z_t, κ_t). Note that because of the recursive structure of the $\hat{b}_t's$, the expectations that form these forecast functions are conditional not just on date-0 states (i.e., (b_0, z_0, κ_0),), but on the history of realizations $[\kappa_t]_{t=0}^T$ and the history of beliefs $\left[f(F^s \mid \kappa^t)\right]_{t=0}^T$.

The equivalent objects to compare with in the rational expectations model are:

$$\tilde{b}_1 = E\left[b_1 \mid (b_0, z_0, \kappa_0)\right] = g(b_0, z_0, \kappa_0)$$

$$\tilde{b}_2 = E\left[b_2 \mid b_0\right] = \sum_{z^1} \sum_{\{b_1:b_2=g(b_1,z_1,\kappa_1)\}} \pi(z^1 \mid z_0) g\left(b_1, z_1, \kappa_1\right)$$

$$\tilde{b}_3 = E\left[b_3 \mid b_0\right] = \sum_{z^2} \sum_{\{b_2:b_3=g(b_2,z_2,\kappa_2)\}} \pi(z^2 \mid z_0) g\left(b_2, z_2, \kappa_2\right)$$

$$......$$

$$\tilde{b}_{T+1} = E\left[b_{T+1} \mid b_0\right] = \sum_{z^T} \sum_{\{b_T:b_{T+1}=g(b_T,z_T,\kappa_T)\}} \pi(z^T \mid z_0) g\left(b_T, z_T, \kappa_T\right)$$

We can also express the forecast functions of the learning model with a slight modification of the treatment used under rational expectations. Define the probability distribution of TFP shocks and bond holdings at date t in the learning model as $\chi_t(b, z)$. The law of motion for the evolution of this probability over time, given the history κ^T of realizations of the leverage regimes, is defined as follows:

$$\chi_{t+1}(b', z') = \sum_b \sum_z \chi_t(b, z)\pi(z' \mid z) I_t(b', b, z, \kappa_t)$$

where $I_t(b', b, z, \kappa_t)$ is a binary indicator such that $I_t(b', b, z, \kappa_t) = 1 \leftrightarrow b' = h_t\left(b, z, \kappa_t; f(F^s \mid \kappa^t)\right)$ and zero otherwise.

At date-0, for example, we have $\chi_0(b_0, z_0) = 1$ for the particular initial conditions (b_0, z_0), and $\chi_0(b, z) = 0$ for all other pairs (b, z). We also have that:

$$I_t(b', b, z, \kappa_t) = \begin{cases} 1 & \text{if } b' = h_t\left(b, z, \kappa_t; f(F^s \mid \kappa^t)\right), \\ 0 & \text{otherwise.} \end{cases}$$

We could add the indicators for all other possible initial conditions, but since they satisfy $\chi_0(b, z) = 0$ they wash out from the law of motion from date 0 to 1. Hence we get $\chi_1(h_0\left(b_0, z_0, \kappa_0; f(F^s \mid \kappa^0)\right), z')$ $= \pi(z' \mid z_0)$ for each z' and zero otherwise (for all pairs (b', z') such that $b' \neq h_0\left(b_0, z_0, \kappa_0; f(F^s \mid \kappa^0)\right)$).

Now we can compute the expected bonds chosen at date 1 for beginning of period 2 as: $\hat{b}_2 = E\left[b_2 \mid b_0, f(F^s \mid \kappa^1)\right] = \sum_b \sum_z \chi_1(b, z) h_1\left(b, z, \kappa_1; f(F^s \mid \kappa^1)\right)$. At this point we can add over all values of bonds in the state space because the probabilities already have incorporated the information relevant for the "correct" bond positions that the system can land on in period 2 in the learning model.

Alternatively, we can define the probability law of motion as:

$$\chi_{t+1}(b', z') = \sum_z \sum_{\{b: b' = h_t(b, z, \kappa_t; f(F^s \mid \kappa^t))\}} \chi_t(b, z) \pi(z' \mid z)$$

In writing it this way, we take out the indicator function but keep track of only the relevant initial states that can land in each b' by constraining the set of $b's$ over which the summation is taken.

B.3 Expected Returns j Periods Ahead of Date t

Choose an initial triple (b_t, z_t, κ_t) with initial bond holdings set to $b_t = \hat{b}_t$. t is the period for which we are going to calculate the sequence of expected returns j periods ahead. $\hat{b}_t$ stands for the mean bond holdings at period t obtained from the forecast functions. z_t is set equal to 1. κ_t is set to its value used in the forecast function calculations for the corresponding period t. We calculate expected returns for any date $t+1+j$ as of date t. This calculation involves a numerator with the sum of dividends and price of date $t+1+j$, $[q_t(b_{t+j+1}, z_{t+j+1}, \kappa_{t+j+1}) + d(z_{t+j+1})]$, and a numerator with the price as of date $t+j$, $q_t(b_{t+j}, z_{t+j}, \kappa_{t+j})$, all of which are projected as of the initial date t.

51

We proceed in two steps. First, we calculate the probability tree of possible states in which the economy can land conditional on the initial triple (b_t, z_t, κ_t) up to J periods ahead. The events that we are capturing in this probability tree are the combinations of TFP and κ shocks. Second, we construct the $E_t^s[R_{t+1+j}^q]$ sequence for $j = 0, 1, ..., J$. Finally, as a cross check we recover the asset price in state (b_t, z_t, κ_t) of date t, i.e., $q_t(b_t, z_t, \kappa_t)$, using the sequence $E_t^s\left(\frac{1}{E_{t+j}^s[R_{t+1+j}^q]}\right)$ to recalculate the date-t price as the present discounted value of dividends discounted by expected returns.

In the first step to calculate the probability tree we put all the mass on the initial state that we are conditioning our calculations on for $j = 0$. In other words,

$$\lambda_t^t(b_t, z_t, \kappa_t) = 1.$$

Going forward these distributions evolve according to:

$$\lambda_{t+j+1}^t(b_{t+j+1}, z_{t+j+1}, \kappa_{t+j+1}) = \sum_{z_{t+j+1}} \sum_{\kappa_{t+j+1}} \sum_{b_{t+j+1} \in H_{t+j+1}} \lambda_{t+j}^t(b_{t+j}, z_{t+j}, \kappa_{t+j}) \pi(z_{t+j+1} \mid z_{t+j}) p_t^s(\kappa_{t+j+1} \mid \kappa_{t+j})$$

for $j = 0, 1, ..., J$ where p_t^s stands for the subjective transition probability matrix of κ corresponding to period t. H_{t+j+1} is the set of bolding holdings chosen conditional on a triple $(b_{t+j}, z_{t+j}, \kappa_{t+j})$, which is defined as $H_{t+j+1} = \{b_{t+j+1} : b_{t+j+1} = h_t\left(b_{t+j}, z_{t+j}, \kappa_{t+j} \mid f(F^s \mid \kappa^t)\right)\}$. The superscript t of λ_{t+j+1}^t highlights the fact that this is the date-$t + j$ element for the law of motion that started with initial conditions $\lambda_t^t(b_t, z_t, \kappa_t)$ as of date t, so that the probabilities are conditional on date t.

In the second step, to compute the expected returns, we first take the date $t + j$ element of the sequence of $\lambda's$, $\lambda_{t+j}^t(b_{t+j}, z_{t+j}, \kappa_{t+j})$. Intuitively, this is the equilibrium probability of landing in a particular state $(b_{t+j}, z_{t+j}, \kappa_{t+j})$ in period $t + j$, j periods ahead of the initial period. We then compute expected returns for any $t + 1 + j$ conditional on date t as:

$$E_t^s[R_{t+1+j}^q] = \sum_{z_{t+j+1}} \sum_{\kappa_{t+j+1}} \sum_{b_{t+j+1} \in H_{t+j+1}} \sum_{z_{t+j}} \sum_{\kappa_{t+j}} \lambda_{t+j}^j(b_{t+j}, z_{t+j}, \kappa_{t+j}) \pi(z_{t+j+1} \mid z_{t+j})$$
$$\times p_t^s(\kappa_{t+j+1} \mid \kappa_{t+j}) \frac{q_t(b_{t+j+1}, z_{t+j+1}, \kappa_{t+j+1}) + d(z_{t+j+1})}{q_t(b_{t+j}, z_{t+j}, \kappa_{t+j})}$$

where $d(z) = zg'(l)$. Note that $E_t^s[R_{t+1+j}^q]$ is in fact $E_t^s[R_{t+1+j}^q](b_{t+j}, z_{t+j}, \kappa_{t+j})$. In other words, the one period ahead expected returns depend on the date-j triple $(b_{t+j}, z_{t+j}, \kappa_{t+j})$.

To confirm that the calculations in the first two steps are correct, in the third step we recalculate $q_t(b_t, z_t, \kappa_t)$ as the sum of expected present discounted value of future dividends where discounting

is done using the equity returns (see Equation 11 in the text):

$$q_t(b_t, z_t, \kappa_t) = \sum_{j=0}^{J} \sum_{z_{t+j+1}} \sum_{\kappa_{t+j+1}} \sum_{b_{t+j} \in H_{j+1}} \sum_{z_{t+j}} \sum_{\kappa_{t+j}} \lambda_{t+j}^j(b_{t+j}, z_{t+j}, \kappa_{t+j}) \pi(z_{t+j+1} \mid z_{t+j})$$

$$\times p_t^s(\kappa_{t+j+1} \mid \kappa_{t+j}) \left(\prod_{i=0}^{j} \left(\frac{1}{E_t^s[R_{t+1+i}^q]} \right) \right) d(z_{t+j+1}).$$

To discount date-$t+j+1$ dividend, we divide it by the sum of all one-period-ahead expected returns up to that date. The calculation of expectations in this step utilizes the probability tree computed in the first step.