UC San Diego
UC San Diego Electronic Theses and Dissertations

Title
Essays on Macroeconomic Crises

Permalink
https://escholarship.org/uc/item/5g08345q

Author
Lind, Nelson R

Publication Date
2017

Peer reviewed|Thesis/dissertation

eScholarship.org
Powered by the California Digital Library
University of California

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Essays on Macroeconomic Crises

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

in

Economics

by

Nelson R. Lind

Committee in charge:

Professor Valerie Ramey, Chair
Professor Gordon Hanson
Professor Ivana Komunjer
Professor Natalia Ramondo
Professor Rossen Valkanov
Professor Johannes Wieland

2017

Copyright

Nelson R. Lind, 2017

All rights reserved.

The dissertation of Nelson R. Lind is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2017

DEDICATION

To Rebecca, John, John-Alex, and Alton.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I thank Valerie Ramey for her support and guidance throughout graduate school, and thank Gordon Hanson for inspiration to broaden my research. I also thank Darrel Cohen, Ivana Komunjer, David Lagakos, Marc Muendler, and Johannes Wieland for invaluable advice and countless discussions, and thank Davide Debortoli, Giacomo Rondina, and Irina Telyukova for mentorship during my first few years at UC San Diego. Finally, I thank Lars, Julie, Monica, Jason, Kate, John, and Roy for many years of friendship that made this possible.

Chapter 2 has benefited from discussions with Brent Bundick, Jean Barthélemy, Davide Debortoli, Andrew Foerster, James Hamilton, Matteo Iacoviello, Alejandro Justiniano, Marc Muendler, Valerie Ramey, Irina Telyukova, and Johannes Wieland. I also thank seminar participants at the Banque de France, CEF 2014, Midwest Macro Spring 2014, NBER DSGE Methods 2014, and UC San Diego for comments and suggestions.

Chapters 3 and 4 have benefited from discussions with Stefania Albanesi, Roy Allen, Mitch Downey, Daniel Greenwald, Gordon Hanson, Elizabeth Hastings, Alejandro Justiniano, Aubhik Khan, David Lagakos, Mark Machina, Marc Muendler, Tommaso Porzio, Giorgio Primiceri, Garey Ramey, Valerie Ramey, John Rehbeck, Achim Schmillen, John Shea, Joel Sobel, Julia Thomas, Alexis Toda, Joel Watson, and Johannes Wieland. I also thank seminar participants at FRB San Francisco, Claremont McKenna College, University of Pennsylvania, Ohio State University, University of Maryland, Columbia University, UT Austin, UC Santa Cruz, UC San Diego, and FRB Kansas City for comments and suggestions.

VITA

2008-2009	General Course, London School of Economics, London
2010	B. A. in Economics and Mathematics, St. Olaf College, Northfield
2010-2011	Visiting Scholar, University of California, Berkeley
2011-2017	Graduate Teaching Assistant, University of California, San Diego
2017	Ph. D. in Economics, University of California, San Diego

ABSTRACT OF THE DISSERTATION

Essays on Macroeconomic Crises

by

Nelson R. Lind

Doctor of Philosophy in Economics

University of California, San Diego, 2017

Professor Valerie Ramey, Chair

This dissertation examines economies that may occasionally enter periods of crisis. I first develop a model of asset pricing in the presence of frictions to financial intermediation. This model generates recurrent financial crises due to its strong non-linear dynamics. Next, I develop methodological tools for analyzing these types of highly non-linear dynamic equilibrium models. I then apply these tools to a theory of housing boom-bust cycles driven by endogenous shifts in lending standards.

Chapter 1 introduces a model of asset pricing in the presence of agency frictions between savers and financial intermediaries. This model can generate asymmetric price movements where asset values suddenly collapse during a financial crisis. During normal

times, intermediaries arbitrage away excess returns on assets and traditional asset pricing conditions hold. During a financial crisis, the net worth of intermediaries limits their ability to borrow from savers and they are unable to arbitrage away excess returns. Since their net worth depends on realized asset values, collapsing prices further tighten borrowing constraints leading to a large and sudden collapse in asset values.

In chapter 2, I introduce a local approach to solving highly non-linear models, generalizing perturbation to handle the class of piecewise smooth rational expectations models. First, I formalize the notion of an endogenous regime by introducing a regime-switching equilibrium (RSE) concept. This framework uses non-linear model features to explain macroeconomic regime changes, and makes the distribution of the regime an equilibrium object instead of imposing an external regime-switching structure. Then, I demonstrate how to apply perturbation within a slackened model, approximate the policy functions associated with a given belief about the regime, and solve for the equilibrium regime distribution using backwards induction. This approach (1) accounts for expectational effects due to the probability of regime change; (2) provides a framework for modeling regime-switching from first principles; and (3) connects macroeconomic theory to reduced-form regime-switching econometric models.

Chapter 3 develops a theory of housing boom-bust cycles driven by endogenous shifts in lending standards. The key friction is asymmetric information about default risk, implying that the economy occasionally and endogenously switches between two credit regimes. These regimes differ by whether or not lenders use income verification to screen the marginal homebuyer. A switch from the “screening” regime to the “pooling” regime leads to rapid home price appreciation, a collapse in down-payment requirements, and a reduction in income documentation — consistent with the shift in lending standards during the US housing boom. The episode ends in a foreclosure crisis once fundamentals revert and the screening regime returns. The theory predicts patterns for debt accumulation and mortgage spreads consistent with existing micro evidence.

Chapter 1

Asset Price Booms and Financial Crises

Financial crises often occur after boom-bust cycles in asset markets. This fact begs the question, do good times set the stage for a crisis?

In this chapter, I present a model of asset pricing in which a productivity boom can push the economy into a state of fragility where a deterioration in fundamentals triggers a financial crisis. A rise in asset prices leads to increased leverage in the financial sector. A crisis acts to resolve banker balance sheets, and — in the absence of government intervention — necessarily follows after a sufficiently long boom.

In particular, I consider asset pricing in the environment of Bocola (2016) extended to incorporate Arrow securities. The key friction in this economy comes from an agency problem between asset owners (intermediaries) and their creditors. Following Gertler and Karadi (2011), there is a moral hazard problem where intermediaries may divert assets and default on their debt. Usually, this agency problem has no bite, the financial sector efficiently intermediates funds, and a traditional no-arbitrage condition prices assets. In contrast, intermediaries will be unable to arbitrage away excess returns whenever the agency problem limits the flow of credit. In such a crisis, asset values will be low because the net worth of intermediaries constrains their ability to borrow from households, reducing the demand for assets.

The underlying fundamentals of the economy generate fluctuations in asset prices. Due to frictions limiting banker re-capitalization, financial sector leverage increases while asset prices are high. Over a short boom, leverage does not increase too much and the economy does not risk entering a crisis once fundamentals revert. However, a sufficiently long boom can push the economy into a fragile state where incentive constraints for bankers will bind after a decline in productivity. The crisis persists until re-capitalization of bankers reduces leverage and incentive constraints no longer bind. In this way, financial crises are the private sector's method of recapitalizing banks, and necessarily occur over random business cycles.

This framework provides a baseline asset pricing model in the presence of occasional endogenous financial crises[1]. It abstracts from feedback between the financial sector to the real economy. Due to its simplicity, it is straightforward to solve for equilibrium using global methods. The perturbation-based solution method developed in Chapter 2 allows for more complicated models built on this framework which incorporate macro-financial linkages.

1.1 Model

The economy consists of production firms and a representative household whose members are either workers or bankers. Production firms use labor and land to produce a homogenous final good. They need funds to finance their purchase of land and issue contingent claims to bankers. Bankers have the ability to intermediate funds between

[1] Also see Prestipino (2014).

workers and firms. Firms would not be able to raise funds without these intermediation services because workers do not have direct access to financial markets.

The representative household is endowed with a unit of labor and supplies a fraction H_t of their time in a competitive labor market. They save via risk-free banker deposits D_{t+1} with return of R_t, and consume the final good. Given flow utility from consumption of $u(C_t)$ and discount factor of β, they solve the following problem.

$$\begin{aligned} \max_{\{C_t, H_t, D_{t+1}\}_{t=0}^{\infty}} \quad & \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(C_t) \\ \text{s.t.} \quad & C_t + \frac{1}{R_t} D_{t+1} \leq W_t H_t + D_t + T_t, \quad H_t \leq 1 \end{aligned}$$

Their budget constraint states that their expenditure on consumption and savings must be covered by labor income, their beginning of period deposits, and net transfers of T_t to and from household members who work as bankers. The key assumption is that the household cannot freely choose these transfers and the only way to channel additional funds to bankers is through deposits.

Given no disutility from working, the household supplies their full unit of labor inelastically. Their inter-temporal decisions can be captured through the following stochastic discount factor.

$$\Lambda_{t,t+1} \equiv \beta \frac{u'(C_{t+1})}{u'(C_t)}$$

The household supplies funds to bankers up until the following pricing condition holds.

$$\frac{1}{R_t} = \mathbb{E}_t \Lambda_{t,t+1} \tag{1.1}$$

The return on risk-free debt must equal the present value of an additional unit of income next period.

Taking as given this stochastic discount factor, bankers accumulate wealth on behalf of the household. These bankers will transfer funds back to the household when they lose access to the financial market. Exit and entry is stochastic with household members randomly gaining and losing access to the intermediation technology. Between periods, bankers exit with chance $1-\psi$ and become workers. An equal measure of workers randomly gain access to the financial market so that the fraction of bankers in the overall household stays constant. The household transfers a fraction ω of the assets of exiting bankers to seed the activity of entering bankers.

Consider household member i with access to financial markets at time t and accumulated net worth of n_{it}. Access to the financial market allows this banker to invest in Arrow securities. Let the pricing kernel associated with these securities be $Q_{t,t+1}$ so that the cost of purchasing a random return of $a_{i,t+1}$ is $\mathbb{E}_t Q_{t,t+1} a_{i,t+1}$.

Bankers finance the purchase of securities out of their net worth and may additionally raise funds by taking in risk-free deposits. However, their ability to issue deposists is limited by the following agency problem.

During each period, a banker can choose to liquidate a fraction η of their portfolio and transfer the value of these assets to the representative household. In doing so, they default on their deposits and exit the financial market. Total deposits of an individual banker are publicly verifiable and so workers will be willing to make deposits only if a banker has no incentive to divert funds. This agency problem implies the following incentive constraint.

$$\eta \mathbb{E}_t Q_{t,t+1} a_{i,t+1} \leq V_t(n_{it}) \tag{IC}$$

Here, the function $V_t(n)$ is the state-t contingent value of being a banker with net worth of n and not diverting assets. This incentive constraint states that the value of diverting assets and exiting the market must be less than the value of continuing as a banker.

Bankers accumulate wealth in the financial market until they exogenously exit and transfer their wealth to the household. The wealth accumulation problem of the banker is:

$$
\begin{aligned}
V_t(n_{it}) = \max_{a_{i,t+1},d_{i,t+1}} \quad & \mathbb{E}_t \Lambda_{t,t+1}[(1-\psi)n_{i,t+1} + \psi V_{t+1}(n_{i,t+1})] \\
\text{s.t.} \quad & \eta \mathbb{E}_t Q_{t,t+1} a_{i,t+1} \leq V_t(n_{it}) && \text{(IC)} \\
& \mathbb{E}_t Q_{t,t+1} a_{i,t+1} \leq n_{it} + \frac{1}{R_t} d_{i,t+1} && \text{(BC)} \\
& n_{i,t+1} = a_{i,t+1} - d_{i,t+1}
\end{aligned}
$$

Bankers choose a state-contingent portfolio with return $a_{i,t+1}$ and a level of deposits $d_{i,t+1}$ to maximize the present value of their future net worth. They value future net worth using the stochastic discount factor of the household. The objective of this program follows from the dynamics of entry and exit. With chance $1-\psi$ the banker will exit and transfer funds of $n_{i,t+1}$ to the household and with chance ψ they will continue having access to the financial market. Continued financial market access has value of $V_{t+1}(n_{i,t+1})$. In addition to the IC constraint ensuring that they will not divert funds, the banker's wealth accumulation is subject to a budget constraint, BC. This budget constraint states that expenditure on securities can be no greater than the sum of net worth and funds raised by issuing deposits. The final constraint states that the banker's future net worth is equal to the payoff from their Arrow securities net of their outstanding deposits[2].

As in Bocola (2016), the value function is linear in net worth due to the linearity of this decision problem. This result enables a simple characterization of banker choices and aggregation.

Proposition 1.1.1 *The banker's problem has the following solution. The value function*

[2] Note that if bankers have limited liability then we must also impose that $a_{i,t+1} \geq d_{i,t+1}$ in all possible future states to ensure that the banker has no incentive to default ex-post.

is linear, $V_t(n) = v_t n$. The random coefficient v_t satisfies

$$v_t = \frac{\mathbb{E}_t \Lambda_{t,t+1}[(1-\psi) + \psi v_{t+1}]R_t}{1-\mu_t}$$

where μ_t is the multiplier on the IC constraint. The multiplier satisfies

$$\mu_t = \max\left\{1 - \frac{\mathbb{E}_t \Lambda_{t,t+1}[(1-\psi) + \psi v_{t+1}]R_t N_t}{\eta X_t}, 0\right\}$$

where

$$N_t \equiv \int_0^1 n_{it} \mathbf{1}\{i \text{ is a banker at time } t\} di$$

is aggregate banker net worth and

$$X_t \equiv \mathbb{E}_t Q_{t,t+1} A_{t+1}, \quad A_{t+1} \equiv \int_0^1 a_{i,t+1} \mathbf{1}\{i \text{ is a banker at time } t\} di$$

is the market value of all Arrow securities held by bankers.

Proof. The proof follows identically to in Appendix A of Bocola (2016) with modified notation to account for the full set of Arrow securities. ■

It is worth briefly discussing the coefficient v_t. This variable represents the marginal value of net worth to bankers — the value of additional internal funds. Since the value of funds transferred to the households is 1, bankers will want to retain funds so long as $v_t \geq 1$. Note that if $v_{t+1} = 1$ and the IC constraint is slack (so that $\mu_t = 0$) we get

$$v_t = \mathbb{E}_t \Lambda_{t,t+1}[(1-\psi) + \psi v_{t+1}]R_t = \frac{\mathbb{E}_t \Lambda_{t,t+1}[(1-\psi) + \psi]}{\mathbb{E}_t \Lambda_{t,t+1}} = 1$$

where I used the deposit supply condition (1.1) to eliminate R_t. Under these conditions, the value of internal funds is the same as transferring funds to the household. In contrast, if either $\mu_t > 0$ or $\mathbf{P}_t[v_{t+1} > 1] > 0$, then internal funds are more valuable than funds transferred to the household. In the presence of aggregate risk, the value of internal funds will be larger than funds transferred to the household, and so bankers will accumulate wealth until they are forced to exogenously exit.

Given this characterization of banker wealth accumulation, consider a banker's portfolio choice problem — their choice of state-contingent payoffs from Arrow securities. In each possible future state, they must be indifferent between increasing or decreasing their payoff. As a result the following indifference condition must hold in all future states.

$$\eta Q_{t,t+1}\mu_t = \Lambda_{t,t+1}[(1-\psi)+\psi v_{t+1}](1-R_t Q_{t,t+1}) \tag{1.2}$$

The left hand side is the cost of tightening the IC constraint by increasing the payoff $a_{i,t+1}$ in a particular future state. The right hand side is the state-by-state excess return on securities accounting for the marginal value of net worth across states.

This arbitrage condition pins down the pricing kernel, $Q_{t,t+1}$.

Proposition 1.1.2 *The equilibrium asset pricing kernel is*

$$Q_{t,t+1} = \frac{\Lambda_{t,t+1}[(1-\psi)+\psi v_{t+1}]}{\eta\mu_t + \Lambda_{t,t+1}[(1-\psi)+\psi v_{t+1}]R_t}.$$

Proof. Re-arrange condition (1.2). ■

The pricing kernel is influenced by banker entry and exit dynamics and the agency problem between workers and bankers. The term $(1-\psi)+\psi v_{t+1}$ captures the influence of entry and exit, and the multiplier on the IC constraint captures the agency problem. When financial intermediaries are not net-worth constrained so that IC constraint is slack and $\mu_t = 0$, we have $Q_{t,t+1} = 1/R_t$. All states of the world are discounted by the risk-free rate. In contrast, when net worth constrains the borrowing of financial intermediaries ($\mu_t > 0$), the pricing kernel decreases across all states and $Q_{t,t+1} < 1/R_t$.

We can now use this pricing kernel to price risky assets backed by issuance of

Arrow securities. In particular, firms issue Arrow securities at time t and use the proceeds to finance the purchase of land. The purchase price of a unit of land is P_t. They use land and labor to produce final goods in period $t+1$ before selling their land.

The profit maximization problem of firms is:

$$\begin{aligned} \max_{L_t,H_{t+1},A_{t+1}} \quad & \mathbb{E}_t Q_{t,t+1}A_{t+1} - P_t L_t \\ \text{s.t.} \quad & Y_{t+1} \leq Z_{t+1}L_t^{\alpha}H_{t+1}^{1-\alpha} \\ & A_{t+1} \leq P_{t+1}L_t + Y_{t+1} - W_{t+1}H_{t+1}. \end{aligned}$$

The firm's time-t profits consists of revenue from arrow securities net of expenditure on land. The first constraint is the firm's production technology. The second constraint is a collateral constraint. The firm can only issue securities whose payoffs are covered by its land holdings and by production revenue net of labor costs.

The budget constraint and collateral constraint must both bind and so we can reduce to the following problem in land and hours.

$$\max_{L_t,H_{t+1}} \mathbb{E}_t Q_{t,t+1}[P_{t+1}L_t + Z_{t+1}L_t^{\alpha}H_{t+1}^{1-\alpha} - W_{t+1}H_{t+1}] - P_t L_t$$

The first order conditions for this problem state that labor earns its marginal product

$$W_{t+1} = (1-\alpha)\frac{Y_{t+1}}{H_{t+1}}$$

and the price of land equals its expected discounted resale value plus marginal product

$$P_t = \mathbb{E}_t Q_{t,t+1}\left[P_{t+1} + \alpha\frac{Y_{t+1}}{L_t}\right].$$

This last condition allows us to use the pricing kernel to price land.

1.1.1 Equilibrium

First consider clearing in factor markets. Since labor supply is inelastic and equal to 1, wages are a fraction of output

$$W_t = (1-\alpha)Y_t.$$

For simplicity, I normalize the supply of land in the economy to 1. The equilibrium price of land is then

$$P_t = \mathbb{E}_t Q_{t,t+1}\left[P_{t+1} + \alpha Y_{t+1}\right].$$

Note that, all else equal, the price of land falls during a financial crisis because an increase in μ_t depresses the pricing kernel. When agency problems limit the flow of credit, intermediaries value the future less, which reduces the value of land.

Then, since labor and land are both in unit supply, productivity determines production of final goods.

$$Y_t = Z_t L_t^{\alpha} H_t^{1-\alpha} = Z_t$$

We effectively have an exogenous level of output given an exogenous process for productivity.

Next consider clearing in the final goods market. Since household consumption is the only use for final goods, the market clearing condition is

$$Y_t = C_t.$$

Total production of final goods must equal household consumption. As a result, the household stochastic discount factor is

$$\Lambda_{t+1} = \beta \frac{u'(Y_{t+1})}{u'(Y_t)}$$

The household's marginal rate of substitution between t and $t+1$ is effectively exogenous due to the exogeneity of productivity.

Finally, clearing in the financial market requires that bankers hold all securities issued by firms. As a result, the aggregate payoff from banker portfolios must equal

$$\int_0^1 a_{i,t+1} \mathbf{1}\{i \text{ is a banker at time } t\} \mathrm{d}i = A_{t+1} = P_{t+1}L_t + Y_{t+1} - W_{t+1}H_{t+1} = P_{t+1} + \alpha Y_{t+1}.$$

The banker's payoff equals the liquidation value of land plus land's share of output. This result implies that the aggregate initial value of banker portfolios is then

$$V_t = E_t Q_{t,t+1} A_{t+1} = \mathbb{E}_t Q_{t,t+1} \left[P_{t+1} + \alpha Y_{t+1}\right] = P_t.$$

The financial sector's assets have value equal to the land held by firms to be used in production at time $t+1$. The amount of deposits they must raise to finance these assets is

$$D_{t+1} = R_t(P_t - N_t).$$

Deposits equal the excess of the value of land over the banker's net worth.

Net worth evolves as

$$N_t = \psi(P_t + \alpha Y_t - D_t) + \omega(P_t + \alpha Y_t)$$

The total payoff from banker securities in t is $P_t + \alpha Y_t$. Each banker pays back their depositors giving net payoff of $P_t + \alpha Y_t - D_t$. Then a fraction $1 - \psi$ of these bankers exit delivering wealth to the household. In turn, the household delivers funds to new bankers in amount $\omega(P_t + \alpha Y_t)$. The total net worth of bankers operating in period t is then the net worth of the surviving bankers and the funds used to seed entering bankers.

Note that the complementarity condition for the IC constraint is

$$\mu_t \geq 0 \quad \perp \quad \eta P_t \leq v_t N_t.$$

When the IC constraint binds, the banker's net worth will pin down the value of land. Otherwise we have $\mu_t = 0$, the pricing kernel reduces to $1/R_t$, and the price of land satisfies

$$P_t = \frac{1}{R_t}\mathbb{E}_t[P_{t+1} + \alpha Y_{t+1}].$$

This condition corresponds to a traditional asset pricing condition for an asset with dividend of αY_{t+1} given a cost of funds of R_t. When agency problems do not limit the flow of credit, intermediaries arbitrage away excess returns on risky assets. Otherwise, intermediary net worth pins down the price of land.

1.2 Equilibrium Dynamics

I now examine the dynamics of land prices in this economy. The economy occasionally experiences financial crises where the flow of credit from workers to bankers stalls (the IC constraint binds) and the price of land collapses.

These asymmetric dynamics come from the non-linearity of the IC constraint. Since the model has two state variables — productivity and debt — it can be easily solved using global methods. I approximate policy functions via linear interpolation and then solve the model using time iteration.

1.2.1 Parameterization

For simplicity, I consider a model parameterization where productivity follows a two state Markov chain. Times are either bad or good with productivity (equal to output) either low or high:

$$Z_B = 1 - \varepsilon \quad \text{and} \quad Z_G = 1 + \varepsilon$$

The parameter $\varepsilon \geq 0$ determines the volatility of the economy. For stochastic simulations I set $\varepsilon = 0.01$.

The evoluation of productivity is is governed by the following transition matrix:

$$\Gamma = \begin{bmatrix} \rho & 1-\rho \\ 1-\rho & \rho \end{bmatrix}$$

The parameter ρ is the persistence of productivity, and I set $\rho = 0.95$.

For comparability to the literature, I choose the fraction of assets used to seed bankers and the survival rate of bankers following Gertler and Karadi (2011). These values are $\omega = 0.002$ and $\psi = 0.972$.

I choose the fraction of assets that can be diverted to ensure that the economy occasionally enters the crisis regime. Specifically, I choose η such that the economy is in the normal regime in its deterministic steady state (with $\varepsilon = 0$), but close to the boundary of regime change. This ensures that in the stochastic model (when $\varepsilon = 0.01$), the economy will sometimes enter the crisis regime during bad times. Specifically, I set $\eta = 0.412$ which is about 10% larger than the value of 0.381 using in Gertler and Karadi (2011).

I set the discount factor to 0.99 to generate a quarterly risk-free rate of about 4% in steady state. I set the risk aversion parameter to $\sigma = 2$. Without loss of generality, I set the land share to 1.

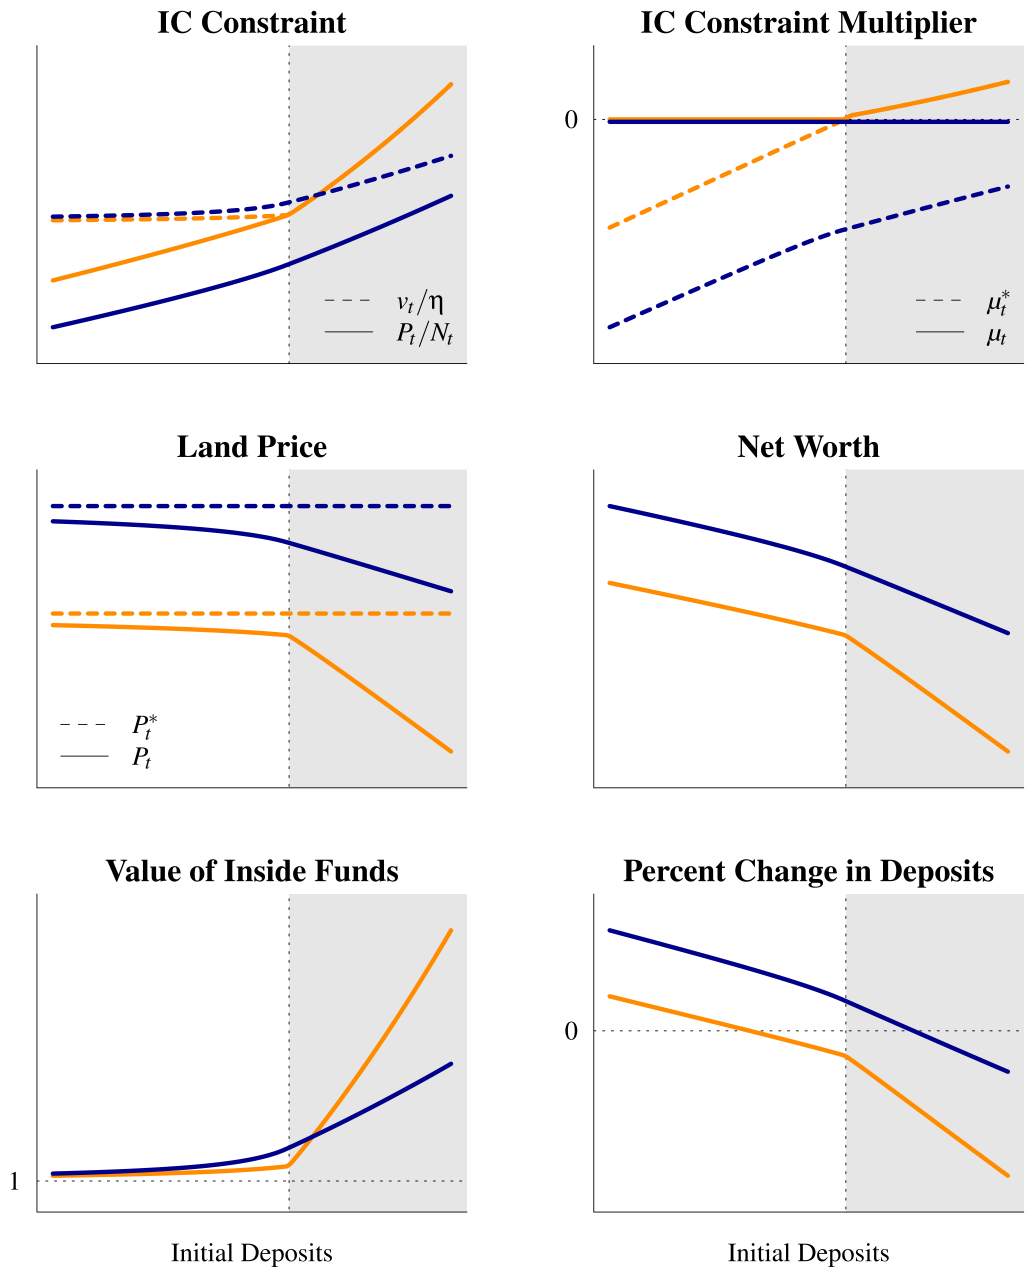

Figure 1.1: Policy functions across initial deposit levels in bad (orange) and good (blue) times. Shaded region indicates when a crisis occurs in bad times.

1.2.2 Numerical Results

Figure 1.1 shows equilibrium outcomes as a function of the beginning-of-period banker deposits (debt). Orange lines correspond to when productivity is low (bad times), and blue lines correspond to when productivity is high (good times). The top left diagram depicts the IC constraint and the top right diagram depicts the IC constraint multiplier (both μ_t and μ_t^*). Note that the constraint is

$$\eta P_t \leq v_t N_t \quad \Leftrightarrow \quad \frac{P_t}{N_t} \leq \frac{v_t}{\eta}.$$

The constraint puts an upper bound of v_t/eta on the banker's asset-to-net-worth ratio of P_t/N_t. The dashed lines on the top left show the upper bound while the solid lines show P_t/N_t. We can see that the IC constraint is slack and the multiplier is zero for all levels of deposits when times are good. For high levels of deposits, the constraint binds in the low productivity state of the world (depicted in orange). On the top right, we can see that when IC constraint is slack, we have $\mu_t^* < 0$ and $\mu_t = 0$. When it binds we have $\mu_t = \mu_t^* \geq 0$. The shaded region of each diagram corresponds to when the IC constraint binds in the low productivity state.

The next row shows the equilibrium land price and banker net worth. Both outcomes are decreasing in banker initial deposits and decrease rapidly in deposits during crises if times are bad and also in the absence of a crisis if times are good.

To interpret this result, it is useful to define a counterfactual frictionless land price as the price implied by no arbitrage between land and deposits.

$$P_t^* = \frac{1}{R_t}\mathbb{E}_t[P_{t+1}^* + \alpha Y_{t+1}] \quad \Longrightarrow \quad P_t^* = \sum_{j=1}^{\infty} \frac{1}{R_{t,t+j}} \alpha Y_{t+j}$$

with $R_{t,t+j} \equiv \Pi_{i=0}^{j-1} R_{t+i}$. This price is the fundamental price of the flow of payoffs from land given opportunity cost of finds captured by R_{t+j} for $j = 0, 1, \ldots$. Note that this price is complete a function of exogenous productivity, and does not depend on banker deposits.

If financial crises never occurred, so that $\mu_t = 0$, then the land price would equal this fundamental price. Departures from this level quantify the effect of potential crises on the price of land. The dashed lines in the middle left diagram show P_t^*.

Define the deviation of the frictionless price from the true price $\Delta_t \equiv P_t^* - P_t$. Combining the definition of P_t^* with the land pricing condition implies that the deviation satisfies

$$\Delta_t = \frac{1}{R_t}\mathbb{E}_t\Delta_{t+1} + \mathbb{E}_t(R_t^{-1} - Q_{t,t+1})[P_{t+1} + \alpha Y_{t+1}]$$

and so, letting $R_{t,t+j} \equiv \Pi_{i=0}^{j-1} R_{t+i}$,

$$\Delta_t = \sum_{j=0}^{\infty} \frac{1}{R_{t,t+j}}\mathbb{E}_t(R_{t+j}^{-1} - Q_{t+j,t+1+j})[P_{t+1+j} + \alpha Y_{t+1+j}].$$

Note that since $\mu_{t+j} = 0$ implies that $R_{t+j}^{-1} = Q_{t+j,t+1+j}$ and $\mu_{t+j} > 0$ implies that $R_{t+j}^{-1} > Q_{t+j,t+1+j}$, the terms in this series are zero in all future periods where there is not financial crises, and positive whenever there is a crisis. As a result, the deviation of frictionless land prices from true land prices depends on anticipation of future financial crises. As the probability of future crises gets small, we must have $\Delta_t \to 0$, and as future crises become more likely Δ_t will increase.

We can see this result in the middle left diagram. Δ_t corresponds to the gap between the dashed line (equal to P_t^*) and their corresponding solid line (equal to P_t). As initial deposits gets larger the economy gets closer to the crisis region. In this region, the economy will enter a crisis if productivity ends up being low. Therefore, the likelihood of a crisis occurring in the future is increasing in deposits. As a result, the deviation of the land price from the fundamental price is increasing in banker initial deposits because of anticipation effects.

Note that this effect is active in all states of the world, not just if the economy is currently in a crisis. As a result, the mere possibility of financial crises in the future tends to depress land prices and make them dependent on the balance sheet of financial intermediaries. That said, if the economy is currently in a crisis, then the binding IC constraint directly impacts land prices. P_t rapidly falls and Δ_t rapidly increases if banker debt is large — leading to a tight IC constraint.

The balance sheet of financial intermediaries depends on realized home prices and their initial debt burden. At low land prices, interedmiaries have low net worth (middle right diagram), which leads to binding incentive constraints. In those states of the world where bankers have a high amount of debt, their internal funds are scarce because they have low net worth. The bottom left diagram shows the value of inside funds. Funds delivered to the household have value of 1, and inside funds are more valuable than outside funds in all states of the world. Net worth relaxes incentive constraint by alleviating the agency problem between workers and bankers. When a banker has a low debt burden, they are more likely to have a high net worth in the future. The anticipation of future states of the world where the banker has little access to borrowing makes inside funds valuable today, and makes their value increase as the banker's balance sheet condition worsens (due to high liabilities in the form of outstanding deposits).

Finally, the bottom right diagram shows the growth rate of deposits, allowing us to visualize the evolution of banker deposits. When deposits are low, bankers accumulate additional deposits over time and do so at a faster rate when times are good. The two deposit levels at which deposit growth equals zero are the long-run deposits levels for each productivity state. Deposits converge to these levels over time while productivity stays constant. While times remain good, banker deposits growth is high, and, absent a deterioration in fundamentals, they end up with high leverage — a large stock of outstanding deposits. Note that this long-run level is in the crisis region. After a long period of high productivity, the economy converges to a fragile state where a decline in productivity will trigger a financial crisis. Following a crisis, bankers begin a process of deleveraging and will converge towards a lower level of deposits that is outside of the

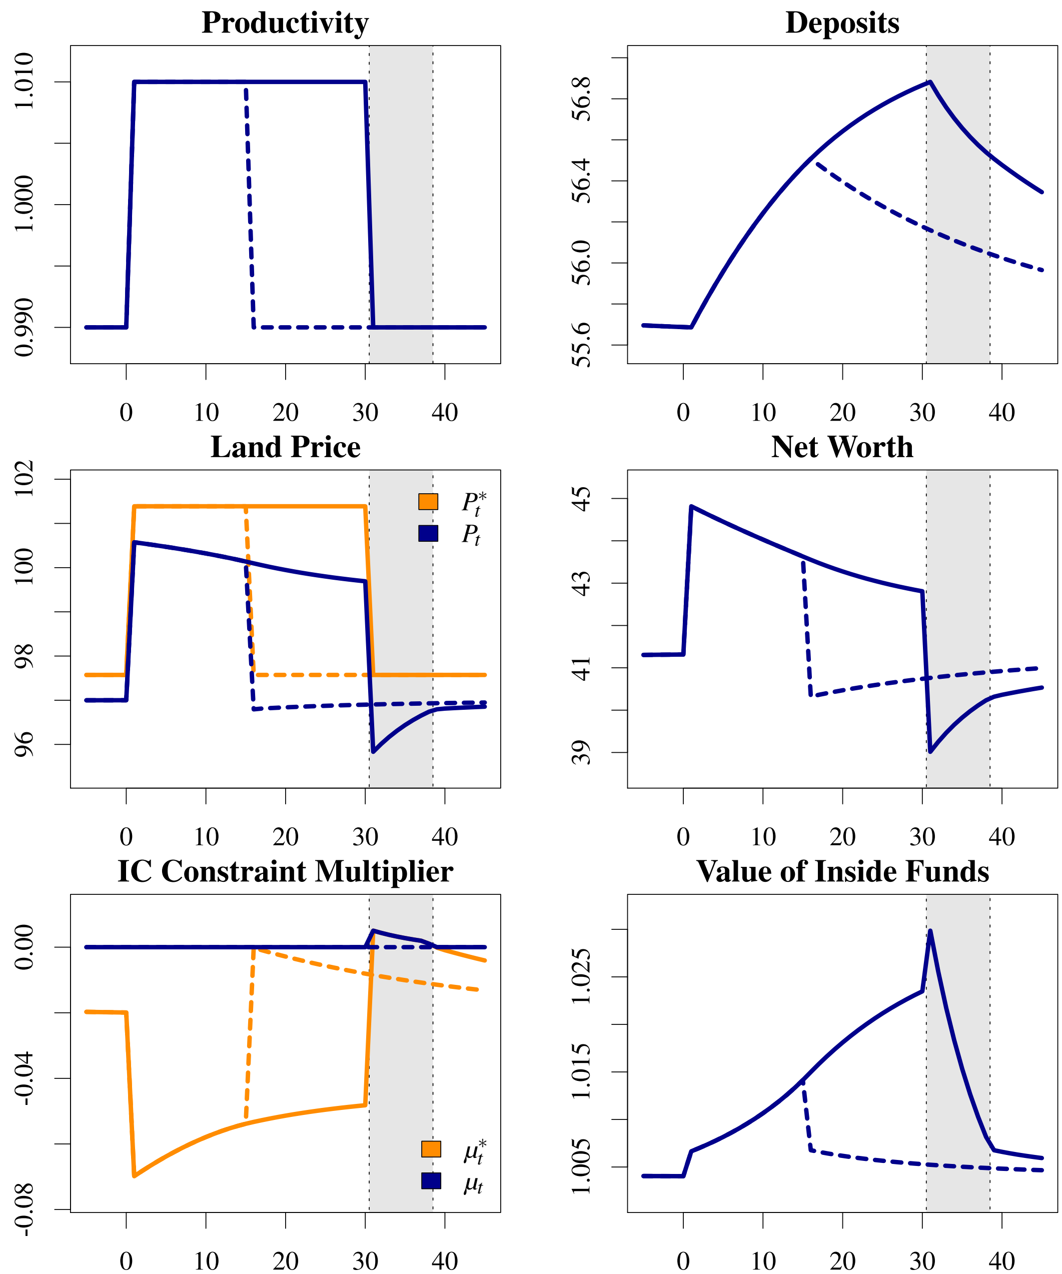

Figure 1.2: Short boom (dashed) versus long boom (solid). Shaded region indicates the crisis following the long boom.

crisis region.

Figure 1.2 shows these dynamics. I start the economy at a low productivity level at the associated long-run level of deposits. Then, in period 1, productivity increases. This generates boom in the asset market which increases intermediary net worth. Since some bankers exit and deliver their net worth to the representative household each period, the total net worth of the financial sector is steadily declining in the absence of new funds from the household. New bankers are only seeded with a portion of the wealth of exit bankers and so they must raise deposits to be able to invest in assets. Since land prices are high, they must use a high level of deposits. The result is that the banking system steadily accumulates debt while high productivity keeps land prices elevated.

As bankers accumulate additional liabilities, the economy moves towards the crisis region. Anticipating of a future crisis then depresses current land prices since land prices are expected to fall in a crisis.

The diagram shows two scenarios. The dashed line shows the dynamics of the economy if the productivity boom is short (15 periods) while the solid line shows the case of a long productivity boom (30 periods). In the first case, bankers do not accumulate enough leverage to push the economy into the crisis regime. The collapse in productivity pushes down the land price, which reduces banker net worth, but not by enough to make the IC constraint bind and trigger a crisis. In contrast, a crisis does occur after a long productivity boom. In a long boom, bankers accumulate enough debt to make the economy fragile, pushing it into the crisis region.

By comparing the two scenarios we can visualize the amplification generated by the binding IC constraint. The land price falls in both cases, but in the later case it over shoots. The binding IC constraint makes bankers unable to raise additional funds to purchase assets, reducing asset demand. This depresses land prices, which further reduces banker net worth. This feedback generates amplification.

The crisis persists while banker debt remains high. They begin a deleveraging process which moves the economy back out of the crises region over time. Low net worth bankers exit and are replaced with new bankers with seed funds from the household. This slow re-capitalization of the banking sector reduces leverage over time.

In period 38, the crisis ends. At this point the land price has recovered significantly and is almost back to its original level. However, note that throughout land prices are uniformly below their fundamental value. The deviation of the land price (the gap between the orange and blue lines in the middle left diagram) is a function of how far the economy is from the crisis regime. The deviation is increases during the productivity boom as bankers accumulate debt and then declines again once the crisis occurs and bankers reduce their leverage through recapitalization.

Because a crisis only occurs after a sufficiently long boom in productivity, they are relatively rare events. Figure 1.3 shows a simulation of the economy over 200 quarters (50 years). In this simulation, four crises occur. The first is brief and comes after two separate productivity booms. The first boom is long enough to significantly elevate banker deposits but not long enough to push the economy into the crisis regime. The second boom is brief, but further pushes up debt and brings the economy just into the crisis regime. However, the crisis in brief because the combined effect of the two productivity booms doesn't lead to a very high level of deposits. The first financial crisis is minor and the banking sector is quickly re-capitalized.

In contrast, the second crisis and third crises occur when banker debt is high after a very long productivity boom. The second crisis is brief because productivity improves. But this improvement in fundamentals is short, and the third crisis has a long period of deleveraging. During this period, productivity does not improve and banker debt steadily is reduced through re-capitalization. Once this long crisis ends, the land price recovers to almost its long-run (low productivity) level. The fourth crisis similarly occurs following a long-productivity boom, and involves a similar protracted period of banking sector recapitalization.

What we see from this simulation is that the duration of a crisis depends on whether or not productivity improves and how much debt bankers have accumulated.

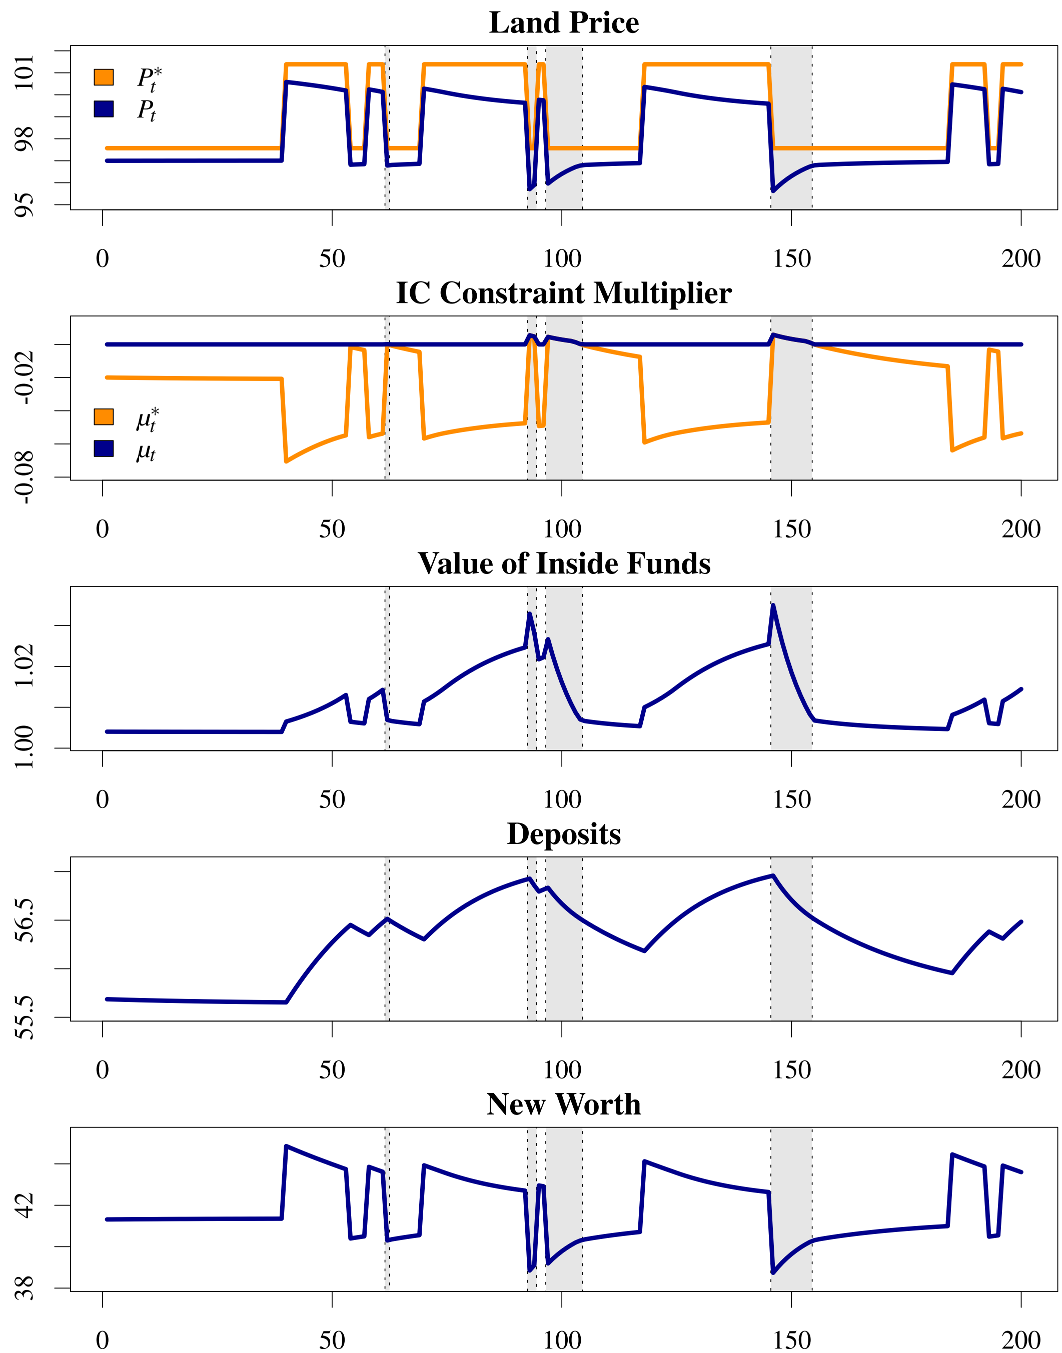

Figure 1.3: Simulated economy over 50 years (200 quarters). Shaded regions indicate the crisis regime. Initialized at bad-time stochastic steady state in $t = 0$

Note that an intermittent boom in productivity does end the crisis in the short term, but does not re-capitalize the banks. Ultimately, the economy will go through a long deleveraging period if bankers have accumulated a large amount of debt. In this model, a crisis is a necessary adjustment period after a long productivity boom during which the financial sector become highly levered.

1.3 Conclusion

This chapter has outlined a baseline model of endogenous financial crises. They key feature is an occasionally binding incentive constraint that represents an agency problem between workers and bankers. In this model, anticipation of future crises depresses land prices below their fundamental value, and the size of the discount depends on the underlying balance sheets of bankers. When bankers have a high amount of debt, they are likely to be net worth constrained in the near future. Since land prices fall during a crisis, this elevated crisis risk puts downward pressure on the current land price. However, since there is always some chance of a crisis in the long-run, land values are always depressed relative to a frictionless world.

The dynamics of crises follow from the dynamics of banker balance sheets. Bankers accumulate debt which land prices are large and so protracted productivity booms lead to increased debt and an economy that is susceptible to crises. The high leverage of the banking sector gets resolved through a process of recapitalization, and the length of a crisis depends on how much debt bankers accumulated during the preceding productivity boom. In this way, a crisis is an adjustment period that necessarily follows after a long period of relative prosperity in the economy.

The lessons of this model came from a full global solution for equilibrium. The dynamics surrounding endogenous financial crisis come from fully capturing the strong non-linearity of the IC constraint. Global methods typically depend on having a few states variables, and in this simple model we only needed to keep track of banking sector debt and an exogenous productivity shock. But to look at the macroeconomic implications of credit crunches and asset price busts — say in a model with endogenous investment that changes future labor productivity — we need to be able to solve larger models. The next chapter introduces a perturbation method that can handle the non-linearity in the IC constraint, and softens the curse of dimensionality so that we can solve extensions of the model incorporating feedback between the financial and real sectors of the economy.

Appendix

1.A Steady State

To solve for steady state we need to solve the following system of equations.

$$
\begin{aligned}
R &= 1/\beta \\
(1-\psi R)N &= [\omega-(R-1)\psi]P+(\psi+\omega)\alpha Y \\
D &= R(P-N) \\
\mu &= \max\left\{1-\frac{[(1-\psi)+\psi v]N}{\eta P},0\right\} \\
P &= \frac{\beta[(1-\psi)+\psi v]}{\eta\mu+[(1-\psi)+\psi v]}\left[P+\alpha Y\right] \\
(1-\mu)v &= (1-\psi)+\psi v.
\end{aligned}
$$

If $\mu=0$ then $v=1$ and $P=\alpha Y/(1-\beta)$, $N=[(\omega-(R-1)\psi)P+(\psi+\omega)\alpha Y]/(1-\psi R)$. Then we must have

$$1\leq\frac{[(1-\psi)+\psi v]N}{\eta P}$$

Therefore we have an upper bound on values of η consistent with the IC constraint being slack in steady state:

$$1\leq\frac{[(1-\psi)+\psi v]N}{P}$$

If this condition is violated then we must have $\mu > 0$.

1.B Model Equations and Solution

The model's control variables are: $R_t, P_t, \mu_t, N_t, v_t$. There is one endogenous state of D_t and one exogenous state of Y_t. The equations pinning down the controls and endogenous state are

$$
\begin{aligned}
\frac{1}{R_t} &= \mathbb{E}_t \beta \frac{u'(Y_{t+1})}{u'(Y_t)} \\
N_t &= \psi(P_t + \alpha Y_t - D_t) + \omega(P_t + \alpha Y_t) \\
D_{t+1} &= R_t(P_t - N_t) \\
\mu_t &= \max\left\{1 - \frac{\mathbb{E}_t \beta \frac{u'(Y_{t+1})}{u'(Y_t)}[(1-\psi) + \psi v_{t+1}]R_t N_t}{\eta P_t}, 0\right\} \\
P_t &= \mathbb{E}_t \frac{\beta \frac{u'(Y_{t+1})}{u'(Y_t)}[(1-\psi) + \psi v_{t+1}]}{\eta \mu_t + \beta \frac{u'(Y_{t+1})}{u'(Y_t)}[(1-\psi) + \psi v_{t+1}]R_t} [P_{t+1} + \alpha Y_{t+1}] \\
v_t &= \frac{\mathbb{E}_t \beta \frac{u'(Y_{t+1})}{u'(Y_t)}[(1-\psi) + \psi v_{t+1}]R_t}{1 - \mu_t}.
\end{aligned}
$$

Note that the first equation pins down R_t without needing to know any other equilibrium outcomes. To solve for a temporary equilibrium at time t given future policy functions, I guess a value for P_t. Then calculate N_t from the second equation and D_{t+1} from the third. Given P_t, N_t, and D_{t+1} we can then calculate μ_t from the fourth equation. Finally, calculate the residual in the fifth equation. This sequence of calculations implies a map from P_t to a residual in the fifth equation. Use a non-linear solver to find P_t that set the residual in the fifth equation to zero. Given this value of P_t, calculate v_t. This procedure solves for a temporary equilibrium for the control vector $(R_t, P_t, \mu_t, N_t, v_t)$ given state (Y_t, D_t) and policy functions governing next period outcomes.

I solve the model using time iteration (recursively solving for temporary equilibrium and updating policy functions), and approximate policy functions via linear interpolation on a discrete grid. I assume that output follows a three state Markov process and keep track of outcomes on an evenly space grid for debt.

Chapter 2

Regime Switching Perturbation for Non-linear Equilibrium Models

2.1 Introduction

Many salient macroeconomic problems are inherently non-linear[1]. Quantitatively realistic dynamic stochastic general equilibrium (DSGE) models often require many state variables, and estimation based on global solution techniques is typically infeasible due to the curse of dimensionality[2]. Instead, macroeconomists tend to rely on local methods to study medium and large scale models since they are fast and standardizable[3]. But since these methods linearize equilibrium conditions, they are invalid for studying highly

[1]For instance, highly non-linear features arise in models of sovereign default (Arellano (2008), Mendoza and Yue (2012), Adam and Grill (2013)), the zero lower bound on interest rates (Fernandez-Villaverde et al. (2012), Aruoba and Schorfheide (2012), and Nakata (2012)), sudden stops (Mendoza (2010),Bianchi and Mendoza (2013),Korinek and Mendoza (2013)), and financial crises (Sannikov and Brunnermeier (2012), He and Krishnamurthy (2013), Gertler and Kiyotaki (2013), and Boissay et al. (2013)).

[2]See Gust et al. (2012) for a successful application of a global method in a model with three endogenous and three exogenous state variables.

[3]See Jin and Judd (2002) and Schmitt-Grohe and Uribe (2004) for standard perturbation methods. Dynare (Adjemian et al. (2014)) provides a standardized solver. For recent local/perturbation methods for solving models with occasionally binding constraints see Brzoza-Brzezina et al. (2013) for a penalty function approach and Guerrieri and Iacoviello (2014) for a perfect foresight approach.

non-linear models[4].

Focusing on the class of piecewise smooth dynamic stochastic general equilibrium models, I introduce a generalized perturbation technique to approximate regime-switching representations of rational expectations equilibria. By overcoming the high cost of global methods and the limited applicability of perturbation methods, this regime-switching perturbation method enables the study of general equilibrium models that were previously considered intractable.

This method solves dynamic stochastic general equilibrium models whose structural equations are piece-wise smooth[5]. This class includes business cycle models incorporating policy constraints like the zero-lower-bound (Fernandez-Villaverde et al. (2012); Aruoba and Schorfheide (2012); Nakata (2012); Gavin et al. (2013)), models of crises based on collateral constraints (Mendoza and Smith (2006); Mendoza (2010); Korinek and Mendoza (2013)), and models of crises due to market breakdown (Boissay et al. (2013)). The solution method only relies on the piecewise smooth structure of the model – which is known directly from its equations – and does not rely on any information about how this non-linearity is reflected in equilibria.

A regime-switching characterization of rational expectations equilibrium underlies the solution method. This alternative regime-switching equilibrium (RSE) concept uses regime-conditional solution functions and specifies the distribution of the regime conditional on the state of the economy as an equilibrium object. I show that it is equivalent to rational expectations equilibrium (REE) because both definitions of equilibrium make identical in-equilibrium predictions. This result justifies using regime-switching representations of equilibrium and provides a theoretical underpinning for regime-switching perturbation.

[4]Braun et al. (2012) show that economic conclusions can be very incorrect if log-linearization techniques are used inappropriately in models incorporating the zero lower bound.

[5]Throughout I describe a real analytic function as a smooth function.

Introducing an endogenous regime is useful in models with strong non-linearities because the regime variable can absorb non-smooth features. Once the model's kinks and discontinuities are captured by the regime, the model is everywhere smooth which enables a perturbation based solution method. Global approaches for solving models incorporating the zero-lower-bound (ZLB) on interest rates (Gust et al. (2012); Aruoba and Schorfheide (2012); Fernandez-Villaverde et al. (2012)) also benefit from conditioning the solution function on whether or not the ZLB binds, and I formalize this technique via an endogenous regime variable.

This idea is also central to the perfect foresight perturbation approach introduced by Guerrieri and Iacoviello (2014). This important contribution demonstrates how local methods when coupled with a regime structure can successfully approximate equilibria in ZLB models. However, the perfect foresight approach means that the probability of transitioning between regimes does not influence the ultimate solution. The RSE concept introduced in this paper incorporates the distribution of the regime explicitly as an equilibrium object, which allows a researcher to account for expectational effects associated with regime changes when solving the model with perturbation.

Prior solution approaches for DSGE models incorporating regime change specify the regime distribution externally either as an exogenous Markov-switching process (as in Farmer et al. (2009); Bianchi (2012, 2013); Foerster (2013); Foerster et al. (2013)) or as a switching process that depends on the economy's state in a known way (as in Davig and Leeper (2006) or Barthélemy and Marx (2013)). This paper builds on this literature by showing how to allow the evolution of the regime to come endogenously from the model structure.

Instead of specifying how regimes evolve ex-ante, economists can explain the underlying reason for regime change using first principles. First, the RSE concept formalizes how to incorporate a regime as an equilibrium outcome. Then, by tying the realization of the regime in an RSE to non-linearity in theoretical models, regime changes can be explained as coming from shifts in the fundamental structure of the economy. The regime depends on economic fundamentals, and the distribution governing regime changes is an equilibrium object. Previous regime-switching approaches emerge as the special case when evolution of the regime can be determined ex-ante to solving the model[6].

Incorporating an endogenous regime requires modifying typical perturbation procedures. Calculating derivatives at a single steady state is no longer sufficient. Instead I show how to approximate RSEs relative to regime dependent reference points. The procedure uses an augmented model which nests the original model and a slack model via a perturbation parameter. The reference points solve the slack model, enabling an application of the implicit function theorem to approximate solutions to the original model. This procedure is closely related to solution methods for Markov-switching models. I show that solving for a first order RSE can be accomplished by adapting techniques for Markov-switching rational expectations models such as those in Farmer et al. (2011), Bianchi and Melosi (2012), Barthélemy and Marx (2013), Cho (2013), and Foerster et al. (2013).

To solve for how the regime relates to the state of the economy, I use backwards induction. The perturbation step of the solution procedure relies on a guess of the regime distribution, but the distribution is itself an equilibrium outcome. Since the regime is identified with the structure of the model – in particular in what region current outcomes are realized – the distribution is restricted by the rational expectations requirement. I show that a simple closed-form expression updates the distribution to be consistent with solution function approximations.

The equilibrium representation and solution method provides a connection to

[6]See appendix 2.E.

linear regime-switching models used routinely in time-series econometrics. In particular, a first order approximation of a regime-switching equilibrium implies a linear state-space system for observable data. This result suggests that macroeconomic processes which are well characterized by Markov-switching econometric models can be modeled as resulting from regime-switching equilibria of non-linear DSGE models. Due to this connection, the method provides a step towards empirical studies based on highly non-linear models – which are typically infeasible since estimation exacerbates the computational burden of solving for equilibria.

Section 2 provides intuition for the technique by presenting a univariate model. I solve a special case of the model directly, and show how the solution can be re-cast within a regime-switching structure. Finally, regime-switching perturbation is applied to solve the model exactly. Section 3 presents the general case. I introduce the class of piecewise smooth rational expectations models, the notion of regime-switching equilibrium, and the general solution procedure. This section also contains equivalence results justifying the method. Section 4 returns to the example model and compares global solutions with the solutions calculated using regime-switching perturbation. Section 5 concludes.

2.2 Example: A ZLB Model

This section works through a motivating example to fix ideas. The model incorporates the zero-lower-bound (ZLB) on interest rates into a simple Fisherian model of inflation determination. A special case of the model has closed-form solutions which take a piecewise linear form. I show how to re-express these equilibria in a regime-switching linear form. In this representation, the regime variable indicates whether or not the ZLB binds and the conditional distribution of the regime is an equilibrium object. Finally, I show how to solve for the solution functions in this regime switching representation via perturbation.

The model consists of the Fisher equation relating nominal and real interest rates, and a ZLB-constrained Taylor rule with interest rate persistence. This basic structure can be derived from a standard neoclassical growth model with money and nominal bonds[7].

$$\begin{aligned} r_t &= i_t - \mathbb{E}_t \pi_{t+1} \\ r_t &= (1-\rho_r)\bar{r} + \rho_r r_{t-1} + \varepsilon_t,\ \varepsilon_t \overset{iid}{\sim} \mathcal{N}(0, \sigma^2) \\ i_t^* &= (1-\rho_i)\bar{r} + \rho_i i_{t-1}^* + \theta \pi_t \\ i_t &= \max\{0, i_t^*\} \end{aligned} \tag{2.1}$$

The fisher equation defines the real interest rate r_t as the nominal interest rate i_t net of expected inflation $\mathbb{E}_t \pi_{t+1}$. The real rate evolves as an exogenous AR(1) process with persistence of ρ_r. Monetary policy has a desired nominal interest rate i_t^* and sets the nominal interest rate i_t equal to this target whenever feasible. Due to the zero lower bound, the central bank is constrained to setting the nominal rate at zero whenever the desired rate is negative.

The desired rate is determined by a weighted average of all past deviations of inflation from the central bank's zero inflation target. As a result, i_t^* can be interpreted as the central bank's accumulated commitment to make up for past deviations of inflation from this target. For example, if inflation has been below target historically and the ZLB is binding, i_t^* represents commitments to keep the interest rate low for an extended period.

I assume that a Taylor-type principle holds and the reaction coefficient θ is positive and greater than $1-\rho_i$. Under this assumption, the model has two steady states. The first corresponds to when policy achieves its zero inflation target. The second is the Friedman steady state where the rate of deflation is equal to the average real interest rate. Linear approximations based around either steady state are only valid if the process for r_t

[7]See appendix 2.F.

is sufficiently bounded to never push equilibrium outcomes across the kink generated by the ZLB. Since the interesting case is precisely when the ZLB may occasionally bind, standard perturbation is an inappropriate solution method.

2.2.1 Closed-Form Solution With $\rho_r = 0$

For the remainder of this section, I assume that the real interest rate has no persistence: $\rho_r = 0$. Under this assumption, the model has a closed-form solution.

To solve the model, first reduce to a single equation in the desired rate by eliminating the real rate, the realized nominal rate, and inflation. This reduction gives a single non-linear expectational difference equation in the desired rate:

$$\mathbb{E}_t i^*_{t+1} = (1-\rho_i)\bar{r} + \max\{\rho_i i^*_t, (\rho_i+\theta)i^*_t\} - \theta r_t \tag{2.2}$$

The max operator arises from the ZLB constraint. When the desired rate is below zero, the ZLB binds, and the right hand side has a slope of $\rho_i < 1$. When policy is not constrained by the ZLB, the slope is $\rho_i + \theta > 1$. Figure 2.1 depicts this equation and shows the two steady states of the model for comparison.

The only fundamental state variable in equation (2.2) is the exogenous shock, r_t. A minimum-state-variable equilibrium is a map $r_t \mapsto g(r_t) = i^*_t$ to determine the desired rate. Since r_t is *iid* and the future desired rate is determined by the same function in-equilibrium, the expected desired rate is constant over time: $\mathbb{E}_t[g(r_{t+1})] \equiv i^{*e}$. Inverting the right hand side of (2.2) gives a piecewise linear expression for the desired rate:

$$i^*_t = g(r_t) = \min\left\{\frac{i^{*e} - (1-\rho_i)\bar{r} + \theta r_t}{\rho_i}, \frac{i^{*e} - (1-\rho_i)\bar{r} + \theta r_t}{\rho_i + \theta}\right\} \tag{2.3}$$

For any given level of i^{*e}, today's desired rate is determined by one of two linear

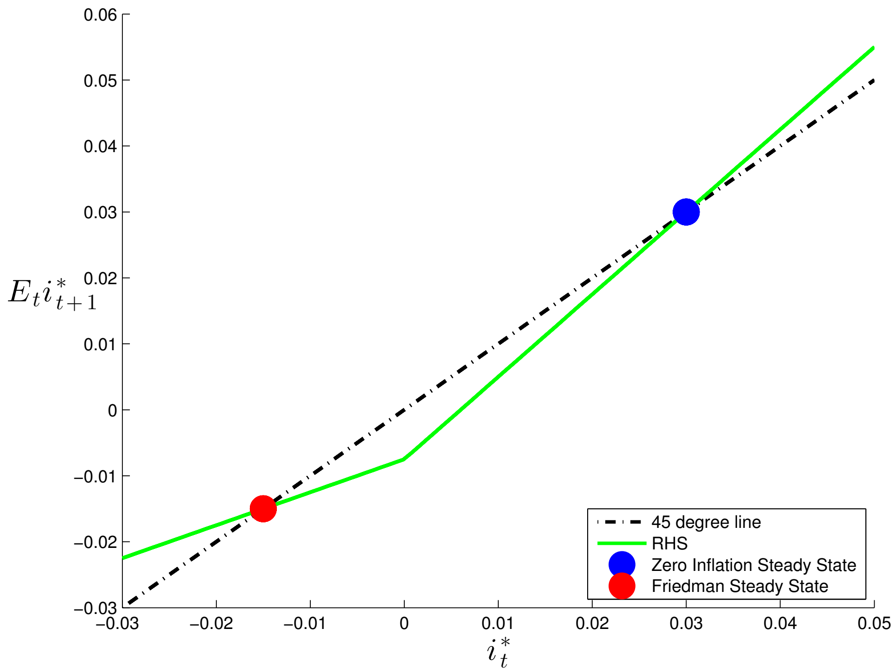

Figure 2.1: Depiction of The Fisherian Model: Equation (2.2)

functions. The first corresponds to when policy is constrained by the ZLB, and the second corresponds to when policy is unconstrained. Regime-switching, whether or not the ZLB binds, is then just between one or the other entry.

In turn, since $i^{*e} = \mathbb{E}_t\left[g(r_{t+1})\right]$, rational expectations requires that the expected desired rate must satisfy the fixed point condition:

$$i^{*e} = \int \min\left\{\frac{i^{*e} - (1-\rho_i)\bar{r} + \theta r_{t+1}}{\rho_i}, \frac{i^{*e} - (1-\rho_i)\bar{r} + \theta r_{t+1}}{\rho_i + \theta}\right\} \frac{1}{\sigma}\phi\left(\frac{r_{t+1} - \bar{r}}{\sigma}\right) dr_{t+1} \tag{2.4}$$

where ϕ denotes the standard normal pdf. Figure 2.2 illustrates this fixed point condition.

This equation has two solutions given the strict concavity of the min function and the fact that the right hand side limits to a linear functional form as $i^{*e} \to \pm\infty$. These two solutions are represented by the intersection of the right hand side of equation (2.4) with

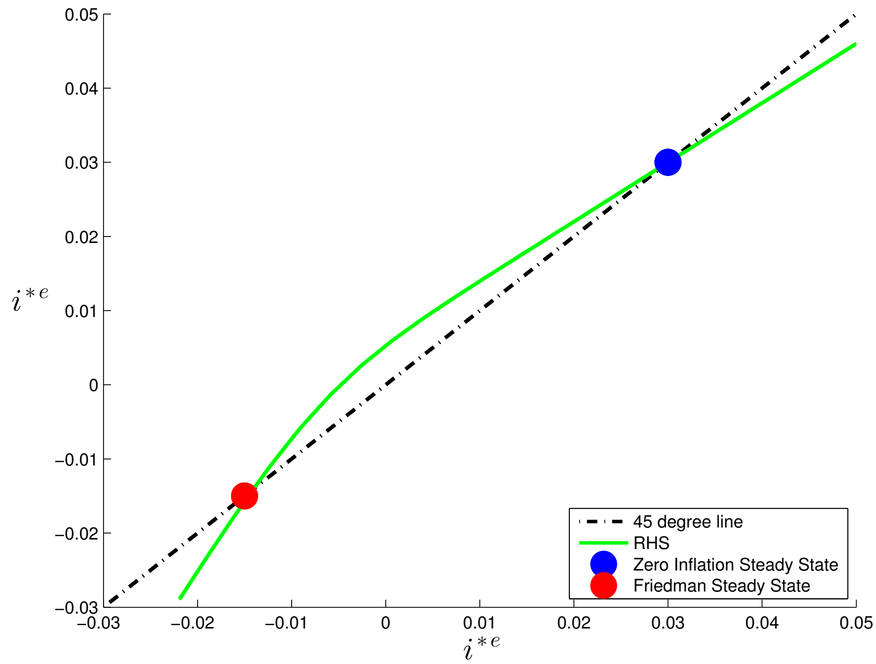

Figure 2.2: Fixed Point Condition For Rational Expectations: Equation (2.4)

the 45-degree line in figure 2.2. There are then precisely two equilibria corresponding to each of the two values for i^{*e} which solve equation (2.4). Since the expected desired rate converges to a steady state of the model as the volatility of the real rate shock is shutdown, these two equilibria can be associated with the two steady states of the model.

2.2.2 A Regime-Switching Characterization of Equilibrium

Structural equation (2.2) and the solution function (2.3) both suggest organizing our solution approach based on whether or not the desired rate is above or below zero. To formalize this idea, I define a regime-switching equilibrium concept where the solution is conditioned on whether or not the ZLB binds. Instead of representing equilibrium with a single solution function, I use two. One function determines outcomes when policy is constrained by the ZLB and the other determines outcomes when policy is not constrained by the ZLB. Note that distinguishing between these cases does not require any knowledge of the equilibria of the model and only uses the information contained in the model's structural equations.

Define a regime variable s_t which either takes a value of c when policy is *constrained* or u when policy is *unconstrained*. Condition outcomes on this regime variable by defining two function g_c and g_u which determine the desired rate depending on the prevailing regime. Since the regime is an endogenous outcome, its probability law is an equilibrium object. Denote the distribution of s_t conditional on r_t by $\Pi_s(r) \equiv \mathbb{P}[s_t = s \mid r_t = r]$.

To ensure that the regime variable correctly indicates when the ZLB binds, this conditional distribution must be consistent with the equilibrium process for the desired rate. In particular, $s_t = c$ can only occur when the desired rate is below zero and $s_t = u$ can only occur when the desired rate is above zero. This means that the events $\{s_t = c \text{ and } i_t^* > 0\}$ and $\{s_t = u \text{ and } i_t^* \leq 0\}$ must occur with probability zero and are off-equilibrium path events.

In turn, if regime $s_t = s$ does occur with positive probability given that $r_t = r$, the model equation must be satisfied:

$$\int \sum_{s'=c,u} g_{s'}(r')\Pi_{s'}(r')\sigma^{-1}\phi(r'/\sigma)\mathrm{d}r' \\ = (1-\rho_i)\bar{r} + \mathbf{1}\{s_t = c\}\rho_i g_s(r) + \mathbf{1}\{s_t = u\}(\rho_i + \theta)g_s(r) - \theta r \quad (2.5)$$

where I have used that $i_t^* \leq 0 \implies s_t = c$ and $i_t^* > 0 \implies s_t = u$. Note that this equation only must hold for (s_t, r_t) pairs which occur in-equilibrium.

The regions of the solution functions which never occur in-equilibrium are unrestricted in principle. However, re-writing the model as the system (2.5) has implicitly

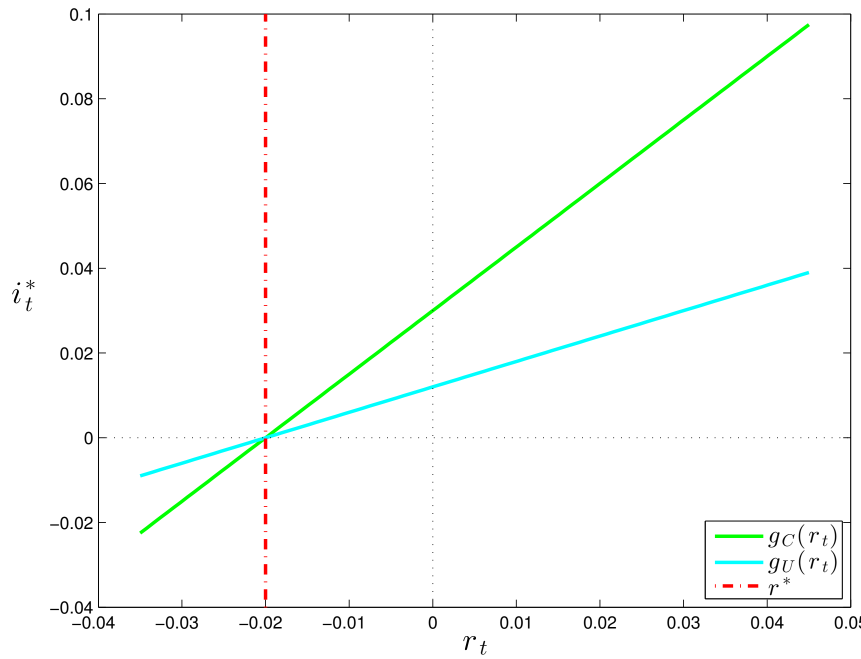

Figure 2.3: Regime-Switching Equilibrium: Functions in (2.6) and Threshold in (2.7)

pinned down the solution functions:

$$\begin{aligned} g_c(r) &= \tfrac{\theta}{\rho_i}(r - r^*) \\ g_u(r) &= \tfrac{\theta}{\rho_i+\theta}(r - r^*) \end{aligned} \tag{2.6}$$

where

$$r^* \equiv \frac{(1-\rho_i)\bar{r}}{\theta} - \int \sum_{s'=c,u} \frac{g_{s'}(r')}{\theta} \Pi_{s'}(r') \frac{1}{\sigma} \phi\left(\frac{r'-\bar{r}}{\sigma}\right) \mathrm{d}r' \tag{2.7}$$

These solution functions and the threshold r^* are depicted in figure 2.3. In effect, off-equilibrium path outcomes are pinned down via the linear extension of the entries of the max operator in the solution given in equation (2.3).

The distribution of the regime must be consistent with these functions. Given the form of g_c and g_u, the constrained regime can only occur when the real rate is below the threshold r^* and the unconstrained regime can only occur when the real rate is above r^*. Therefore, this switching threshold is a summary statistic for the equilibrium regime distribution:

$$\Pi_s(r) = \begin{cases} 1\{r \leq r^*\} & \text{if } s = c \\ 1\{r > r^*\} & \text{if } s = u \end{cases} \tag{2.8}$$

Using (2.8) in (2.7), implies that the switching threshold must be a fixed point of:

$$r^* = \frac{(1-\rho_i)r}{\theta} - \int \max\left\{\frac{1}{\rho_i}(r - r^*), \frac{1}{\rho_i + \theta}(r - r^*)\right\} \frac{1}{\sigma}\phi\left(\frac{r' - \bar{r}}{\sigma}\right) \mathrm{d}r' \tag{2.9}$$

This condition is equivalent to the fixed point condition (2.4) with $r^* = \theta^{-1}[(1-\rho_i)r - i^{*e}]$.

The regime-switching equilibrium concept represents the model solution using a regime-switching linear structure. Instead of solving for a kinked function (equation (2.3)), it is sufficient to solve for two linear functions (equation (2.6)). Given these functions, we then find a fixed point in the distribution Π via the summary statistic r^*.

Although this regime-switching equilibrium concept is a change of perspective, it is equivalent to the usual notion of a rational expectations equilibrium. By construction, it generates identical outcomes. The regime variable effectively absorbs all of the non-differentiability in the model, enabling the use of linear solution functions. As the next section shows, a perturbation method can solve for these well behaved functions.

2.2.3 Solution Using Regime-Switching Perturbation

The regime-switching perturbation method finds solution functions associated with any given conditional distribution for the regime, and uses backwards induction to solves for the equilibrium regime distribution.

The algorithm begins with a guess of the regime distribution: $\Pi = \hat{\Pi}^{(0)}$. At iteration n, perturbation delivers the solution functions $\hat{g}_c^{(n)}$ and $\hat{g}_u^{(n)}$ associated with the distribution $\hat{\Pi}^{(n-1)}$. The collection $\{\hat{g}_c^{(n)}, \hat{g}_u^{(n)}, \hat{\Pi}^{(n-1)}\}$ implies conditional probabilities for the events $\{i_t^* \leq 0\}$ and $\{i_t^* > 0\}$. These conditional probabilities imply a new estimate $\hat{\Pi}^{(n)}$. Iterating on these two steps can be interpreted as performing backward induction. Proposition 2.2.2 shows that the resulting sequence of regime distributions converges to an equilibrium distribution.

The central step to this procedure is the calculation of the solution function estimates, $\hat{g}_c$ and $\hat{h}_u$. In standard perturbation, solution functions are approximated by applying the implicit function theorem at a steady state of the model. The steady state provides a reference point central to solving the model.

In the current example, regime changes force the desired rate to switch between lying above and below zero. The use of a single steady state as a reference point is impossible. To resolve this issue, I use the support points of a stochastic process driven by the regime variable as the necessary regime-specific reference points.

Choose points $\tilde{i}_c^* < 0$ and $\tilde{i}_u^* > 0$ to serve as regime conditional reference points for the desired rate. Importantly these points must lie on either side of zero so that local information is obtained from both sides of the kink arising from the ZLB. Also, choose two reference points for the real rate: $\bar{r}_c$ and $\bar{r}_u$[8]. Given these points and a fixed regime distribution, define a slackness term as:

$$\Delta_{s,s'} = \tilde{i}_{s'}^* - (1-\rho_i)\bar{r} - \max\left\{\rho_i \tilde{i}_s^*, (\rho_i + \theta)\tilde{i}_s^*\right\} + \theta \bar{r}_s$$

The expected value of this slackness term is the residual in the model equation under the assumption that outcomes are determined as $(i_t^*, r_t) = (\tilde{i}_{s_t}^*, \bar{r}_{s_t})$ and $s_t \sim \int \Pi_s(\bar{r} + \varepsilon')\mathrm{d}F(\varepsilon')$.

[8] In this example, these points can be arbitrary. Generally, it is helpful to choose them based on estimates of the regime-conditional means of the state variables of the model

Next, augment the model with this slackness term using a nesting parameter $\eta \in [0,1]$:

$$\begin{aligned} 0 &= \mathbb{E}_t\left[i^*_{t+1} - (1-\rho_i)\bar{r} - \max\left\{\rho_i i^*_t, (\rho_i+\theta)i^*_t\right\} + \theta r_t - (1-\eta)\Delta_{s_t,s_{t+1}}\right] \\ r_{t+1} &= (1-\eta)\bar{r}_{s_{t+1}} + \eta(\bar{r}+\varepsilon_{t+1}),\ \varepsilon_{t+1} \sim \mathcal{N}(0,\sigma^2) \\ s_{t+1} \mid \varepsilon_{t+1} &\sim \Pi_s(\bar{r}+\varepsilon_{t+1}) \end{aligned} \tag{2.10}$$

The first equation arises from appending the slack term $(1-\eta)\Delta_{s,s'}$ to equation (2.2). The second equation specifies a distorted version of the real rate process. The evolution of the real rate is now a convex combination of the reference point $\bar{r}_s$ and the true process of the shock. Now both the regime variable and the Gaussian innovation directly drive the real rate. Finally, the last line specifies that the regime is driven by the true real rate process and not the modified real rate process.

This augmented model reduces to our original model when $\eta = 1$ because the residual term drops out and the future real rate shock is no longer distorted towards the reference point $\bar{r}_{s'}$. In the opposite case with $\eta = 0$, the continuous shock does not impact the future value of the real rate which takes values of $\bar{r}_c$ and $\bar{r}_u$. In this slack model, the residual term enters fully and ensures that the reference points $\tilde{i}^*_c$ and $\tilde{i}^*_c$ solve the model at the state-space points $\bar{r}_c$ and $\bar{r}_u$. The slackened model distorts agent beliefs in such a way to ensure that the chosen reference points solve the model.

Given fixed values for $\{\Pi, \tilde{i}^*_c, \tilde{i}^*_u, \bar{r}_c, \bar{r}_u\}$, a solution to this augmented model is a pair of functions $g_c(r,\eta)$ and $g_u(r,\eta)$ such that model equation (2.10) is satisfied when $i^*_t = g_{s_t}(r_t,\eta)$ and $i^*_{t+1} = g_{s_{t+1}}(r_{t+1},\eta)$. Note that these solutions now depend on the nesting parameter η.

Since the augmented model nests the original model, evaluating the solution functions $g_c(r,\eta)$ and $g_u(r,\eta)$ at $\eta = 1$ gives the solution functions of the original model for a given conditional distribution Π. On the other extreme, when $\eta = 0$ the augmented model reduces to the slack model. By construction, this slack model has the chosen reference point as solution. The desired rate $\tilde{i}_c^*$ solves the augmented model at the point $(\bar{r}_s, 0)$:

$$\begin{aligned} \tilde{i}_c^* &= g_c(\bar{r}_c, 0) \\ \tilde{i}_u^* &= g_u(\bar{r}_u, 0) \end{aligned}$$

Just as a deterministic steady state is the reference point for standard perturbation, the stochastic process $(\tilde{i}_{s_t}^*, \bar{r}_{s_t})$ serves as a reference process. Now, an application of the implicit function theorem based on the state-space points $\bar{r}_c$ and $\bar{r}_u$ (see appendix 2.A) delivers the partial derivatives of the augmented model's solution functions.

Calculating these derivatives leads to a first order approximation:

$$g_s(r, \eta) \approx \tilde{i}_s^* + \frac{\partial}{\partial r} g(\bar{r}_s, 0)(r - \bar{r}_s) + \frac{\partial}{\partial \eta} g(\bar{r}_s, 0)\eta$$

Setting $\eta = 1$ gives an approximate solution to the original model:

$$g_s(r, 1) \approx \left[\bar{i}_s^* - \frac{\partial}{\partial r} g(\bar{r}_s, 0)\bar{r}_s + \frac{\partial}{\partial \eta} g_s(\bar{r}_s, 0)\right] + \frac{\partial}{\partial r} g_s(\bar{r}_s, 0) r \equiv \hat{g}_s(r) \tag{2.11}$$

Note that, the derivative with respect to η adjusts the constant term to account for the level shift introduced by slackening the model.

In this example, the solution is linear so this approximation is in fact exact when Π is an equilibrium conditional distribution. To arrive at this result, I use the following lemma which specifies the constants in this approximation associated with any arbitrary choice of Π.

Lemma 2.2.1 *For a given conditional distribution* Π*, not necessarily an equilibrium distribution, the constants in the linear approximation* (2.11) *associated with* Π *are*

($s = u, c$)

$$\left[\tilde{i}_s^* - \frac{\partial}{\partial r} g(\bar{r}_s, 0)\bar{r}_s + \frac{\partial}{\partial \eta} g_s(\bar{r}_s, 0)\right] = a_s(\Pi) = -b_s r^*(\Pi)$$

$$\frac{\partial}{\partial r} g_s(\bar{r}_s, 0) = b_s = \begin{cases} \frac{\theta}{\rho_i} & s = c \\ \frac{\theta}{\rho_i + \theta} & s = u \end{cases}$$

where

$$r^*(\Pi) = \frac{\frac{1-\rho_i}{\theta}\bar{r} - \mathbb{E}_{r' \sim \mathcal{N}(\bar{r}, \sigma^2)}\left[\Pi_c(r')\frac{1}{\rho_i}r' + \Pi_u(r')\frac{1}{\rho_i+\theta}r'\right]}{1 - \mathbb{E}_{r' \sim \mathcal{N}(\bar{r}, \sigma^2)}\left[\Pi_c(r')\frac{1}{\rho_i} + \Pi_u(r')\frac{1}{\rho_i+\theta}\right]}$$

Proof. See appendix 2.A. ■

The slope constant b_s is the same slope constant in the true solution function given by equation (2.6). The constant term depends on the regime distribution through the threshold $r^*(\Pi)$. This threshold is the point in the state space where both $\hat{g}_c$ and $\hat{g}_u$ equal zero[9].

This lemma characterizes the solution delivered by using perturbation to approximate g_c and g_u holding Π fixed. However, if Π is an equilibrium regime distribution, then the threshold $r^*(\Pi)$ must be a threshold which satisfies the fixed point condition (2.9). It immediately follows that this perturbation method delivers the exact solution, as summarized in the following proposition:

Proposition 2.2.1 *If Π is an equilibrium distribution in an RSE of model* (2.1) *(so that it takes the form given in equation* (2.8)*), then for each regime the first order approximation* (2.11) *is equal to the unique linear regime-conditional solution function associated with Π.*

Proof. See appendix 2.B. ■

[9]It is also a sufficient statistic for the conditional regime distribution implied by these functions

This result implies that if we can solve for Π, we can solve the model. It begs the question, using perturbation to calculate solution functions, can backwards induction solve for an equilibrium value for Π? The answer to this question is yes.

To see this result, consider an update of the regime distribution to Π' given a previous guess of Π. If the regime is drawn according to Π and the desired rate is determined by the associated approximation $\hat{g}_{s_t}(r_t)$, then the desired rate is below zero if and only if $r_t \leq -a_{s_t}(\Pi)/b_{s_t} = r^*(\Pi)$. This implies that the conditional probability of the desired rate ending up below zero is $\sum_{s=u,c} \Pi_s(r_t)\mathbf{1}\{r_t \leq r^*(\Pi)\} = \mathbf{1}\{r_t \leq r^*(\Pi)\}$. This probability is the true chance that policy is constrained. Similarly, the desired rate is above zero if and only if $r_t > r^*(\Pi)$ so the probability of being unconstrained is $\sum_{s=u,c} \Pi_s(r_t)\mathbf{1}\{r_t > r^*(\Pi)\} = \mathbf{1}\{r_t > r^*(\Pi)\}$. Therefore, the only consistent belief is

$$\Pi'_s(r) = \begin{cases} \mathbf{1}\{r \leq r^*(\Pi)\} & \text{if } s = c \\ \mathbf{1}\{r > r^*(\Pi)\} & \text{if } s = u \end{cases} \tag{2.12}$$

This equation specifies an updating equation for beliefs which takes as input a conditional distribution Π and returns an updated conditional distribution Π'.

Backward induction based on this updating rule is guaranteed to generate a sequence of conditional distributions which converges to an equilibrium conditional distribution:

Proposition 2.2.2 *For sufficiently small* $\sigma > 0$ *and for almost every initialization of the conditional distribution* $\Pi^{(1)}$*, the sequence* $\{\Pi^{(n)}\}_{n=1}^{\infty}$ *generated by recursion on the backward induction rule* (2.12) *converges to a conditional distribution* Π^* *associated with a regime-switching equilibrium of model* (2.1).

Proof. See appendix 2.C. ■

This backwards induction procedure solves the model to arbitrary precision. The regime-switching perturbation method combines backwards induction with perturbation to solve jointly for regime-conditional solution functions and the distribution governing the evolution of the regime.

This example demonstrates how regime-switching perturbation works within a simple model. By introducing a slack term into the model structure, we can choose regime-specific reference points at which to apply the implicit function theorem and approximate the model solution. In fact, regime-switching perturbation solves the model exactly because the true solution has a linear regime-switching representation.

2.3 The General Case

This section examines the concept of regime-switching equilibrium, and regime-switching perturbation. The first sub-section specifies the general framework and the central assumption defining the class of piecewise smooth rational expectations models. The second sub-section provides equilibrium definitions and an equivalence result that justifies focusing on regime-switching representations of rational expectations equilibria. The third sub-section examines the details of the method, and the fourth sub-section summarizes the algorithm.

2.3.1 General Framework

Suppose a macroeconomic theory implies a model for a control variable vector $Y_t \in \mathcal{Y} \subset \mathbb{R}^{n_Y}$ and state variable vector $X_t \in \mathcal{X} \subset \mathbb{R}^{n_X}$ which takes the form

$$\begin{aligned} 0 &= \mathbb{E}_t f(Y_{t+1}, Y_t, \tilde{X}_t, X_t) \\ X_{t+1} &= \tilde{X}_t + \Sigma \varepsilon_{t+1},\ \varepsilon_{t+1} \overset{iid}{\sim} F \end{aligned} \tag{2.13}$$

The vector $\tilde{X}_t$ represents the pre-determined part of the state variables. The *iid* innovation, ε_{t+1}, introduces a stochastic component to the state vector via the impact matrix Σ. Note that Σ may be singular so that some state-variables are entirely pre-determined. I will refer to $\mathcal{X}$ as the state space, $\mathcal{Y}$ as the control space, and $\mathcal{Y} \times \mathcal{X}$ as the outcome space of the model.

For instance, consider a real business cycle model with irreversible investment. Letting C_t, I_t, K_t, N_t, z_t denote consumption, investment, capital, labor, and log total factor productivity respectively, the model's equilibrium conditions are:

$$
\begin{aligned}
C_t^{-\tau} &\leq \Lambda_t \perp I_t \geq 0 \\
\Lambda_t &= \beta \mathbb{E}_t C_{t+1}^{-\tau} \left[1 - \delta + \alpha e^{z_{t+1}} \left(\frac{K_{t+1}}{N_{t+1}} \right)^{\alpha - 1} \right] \\
\frac{C_t^{\tau}}{1 - N_t} &= (1 - \alpha) e^{z_t} \left(\frac{K_t}{N_t} \right)^{\alpha} \\
e^{z_t} K_t^{\alpha} N_t^{1-\alpha} &= C_t + I_t \\
K_{t+1} &= (1 - \delta) K_t + I_t \\
z_t &= \rho z_{t-1} + \varepsilon_t
\end{aligned}
$$

Here, Λ_t denotes the marginal utility from an increase in savings while ε_t is an exogenous iid innovation to the log of total factor productivity.

The first equation defines a complementarity condition for whether or not investment is strictly positive or constrained to equal zero. Investment is constrained when the marginal utility from savings is less than the marginal utility from consumption.

The model fits into the form in (2.13) as

$$0 = \mathbb{E}_t \underbrace{\begin{bmatrix} C_t^{-\tau} - \min\left\{\Lambda_t, \left(e^{z_t} K_t^{\alpha} N_t^{1-\alpha}\right)^{-\tau}\right\} \\ \Lambda_t - \beta C_{t+1}^{-\tau}\left[1 - \delta + \alpha e^{z_{t+1}} \left(\frac{K_{t+1}}{N_{t+1}}\right)^{\alpha-1}\right] \\ \tilde{X}_{1,t} - (1-\delta)K_t - e^{z_t} K_t^{\alpha} N_t^{1-\alpha} + C_t \\ \tilde{X}_{2,t} - \rho z_t \\ X_{1,t} - K_t \\ X_{2,t} - z_t \end{bmatrix}}_{f(Y_{t+1}, Y_t, \tilde{X}_t, X_t)}$$

$$\begin{bmatrix} X_{1,t+1} \\ X_{2,t+1} \end{bmatrix} = \begin{bmatrix} \tilde{X}_{1,t} \\ \tilde{X}_{2,t} \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \varepsilon_{t+1}$$

Here, the control variable vector is taken to be the full vector of outcomes:

$$Y_t = [C_t, \Lambda_t, K_t, N_t, z_t]'$$

and definitions for X_t and $\tilde{X}_t$ are appended to the model. To get this formulation, re-write the complementarity condition using a min operator, combine the resource constraint and capital law of motion to eliminate I_t, append the equation $\tilde{X}_t = [(1-\delta)K_t + e^{z_t} K_t^{\alpha} N_t^{1-\alpha} - C_t, \rho z_t]'$ to the model, and append the identity $X_t = [K_t, z_t]'$. The law of motion for the state variables comes from replacing $(1-\delta)K_t + e^{z_t} K_t^{\alpha} N_t^{1-\alpha} - C_t$ in the capital law of motion with $\tilde{X}_{1,t}$, and replacing ρz_t in the law of motion for productivity with $\tilde{X}_{2,t}$.

Since investment is constrained to be non-negative (irreversible investment), the first model equation has a kink. Like in the previous ZLB example, this kink introduces a strong non-linearity into the model. To allow for models with these types of non-smooth features, I make the following assumption on f:

Assumption 2.3.1 *The function $y \mapsto f(y', y, x', x)$ is piece-wise real analytic for each* *value of* $(y', x', x) \in \mathcal{Y} \times \mathcal{X} \times \mathcal{X}$.

Specifically, there is a partition of the control space[10] *into S regions,* $\{\mathcal{Y}_s\}_{s=1}^{S}$, *and a collection of real analytic functions* $\{f_s\}_{s=1}^{S}$ *so that*

$$f(y', y, x', x) = \sum_{s=1}^{S} f_s(y', y, x', x) \mathbf{1}\{y \in \mathcal{Y}_s\} \tag{2.14}$$

This assumption defines a class of piecewise smooth dynamic stochastic general equilibrium models. The model equation f is smooth almost everywhere, and once the control space is broken into the partition $\{\mathcal{Y}_s\}_{s=1}^{S}$ it is smooth on each piece.

Note that in any given model, the function f is known. For instance, in the ZLB model the desired rate must lie in $\mathbb{R}$ and the model is smooth after splitting into pieces $(-\infty, 0]$ and $(0, \infty)$, corresponding to when the ZLB does and does not bind, respectively. In the irreversible investment RBC model, the control space is $\mathbb{R}_+^4 \times \mathbb{R}$, and the partition consists of the sets:

$$\{[C_t, \Lambda_t, K_t, N_t, z_t]' \in \mathbb{R}_+^4 \times \mathbb{R} \mid \Lambda_t > \left(e^{z_t} K_t^{\alpha} N_t^{1-\alpha}\right)^{-\tau}\}$$
$$\{[C_t, \Lambda_t, K_t, N_t, z_t]' \in \mathbb{R}_+^4 \times \mathbb{R} \mid \Lambda_t \leq \left(e^{z_t} K_t^{\alpha} N_t^{1-\alpha}\right)^{-\tau}\}$$

The partition can always be characterized using the known model structure.

Many highly non-linear general equilibrium models which arise in macroeconomic theory satisfy assumption 2.3.1. The assumption includes business cycle models incorporating the ZLB (Fernandez-Villaverde et al. (2012); Aruoba and Schorfheide (2012); Nakata (2012)), models of sudden stop crises (Mendoza and Smith (2006); Mendoza (2010); Korinek and Mendoza (2013))), and models of breakdown in financial

[10] Restricting to a piecewise smooth structure in the contemporaneous control variables is without loss of generality because any non-differentiabilities in terms of state variables can be absorbed into the control variables by augmentation. The augmentation requires a new control variable which replaces the state or shock variable. Then an additional equation is introduced which imposes that the new control variable equals the state or shock which it replaced.

intermediation (Boissay et al. (2013); Gertler and Kiyotaki (2013)). Discrete time versions of continuous time macro-finance models (He and Krishnamurthy (2012); Sannikov and Brunnermeier (2012); He and Krishnamurthy (2013); Sannikov and Brunnermeier (2013)) also have a piecewise smooth structure and satisfy this assumption.

Note that typical everywhere differentiable DSGE models arise as the case when $S = 1$. Additionally, the class of piecewise smooth models includes Markov-switching rational expectations models, as in Foerster et al. (2013), and endogenous-switching rational expectations models, as in Davig and Leeper (2006). Both these types of models arise as special cases, and appendix 2.E shows this connection in detail. The class of piecewise smooth models satisfying assumption 2.3.1 encompasses many discrete time macroeconomic models.

2.3.2 A Regime-Switching Characterization of Equilibrium

Regime-switching perturbation solves for regime-switching representations of rational expectations equilibria. These regime-switching equilibria use the partition $\{\mathcal{Y}_s\}_{s=1}^{S}$ in assumption 2.3.1 to define regimes such that the active piece of the partition identifies the present regime. This construction re-interprets the a piece-wise smooth model as an endogenous-regime model. Regime-switching equilibrium and the standard concept of a rational expectations equilibrium are equivalent since they make identical predictions. However, this representation is useful because it effectively absorbs all non-smooth model features into the regime variable, which facilitates a perturbation approach.

To formalize these ideas, I first fix equilibrium concepts:

Definition 2.3.1 *A minimum-state-variable rational expectations equilibrium (REE) is a pair of measurable functions g and h such that from any initial point $X_1 \in \mathcal{X}$ the* *stochastic process generated recursively according to*

$$Y_t = g(X_t), \tilde{X}_t = h(X_t)$$
$$X_{t+1} = \tilde{X}_t + \Sigma\varepsilon_{t+1},\ \varepsilon_{t+1} \overset{iid}{\sim} F$$

satisfies model (2.13). *Namely, for each* $x \in \mathcal{X}$

$$0 = \int f(g(h(x)+\Sigma\varepsilon'), g(x), h(x), x)dF(\varepsilon')$$

This definition of equilibrium is commonly used as the basis of global solution methods for non-linear rational expectations models. The success of global methods in approximating equilibrium functions depends both on the dimension of the state space and how well the chosen numerical scheme can handle the non-linear structure of the model[11].

Next, I need an equilibrium concept that incorporates the notion of an endogenous regime. Generally, I allow for an arbitrary regime variable r_t which takes values of $1,\dots,R$ and is not identified with the partition $\{\mathcal{Y}_s\}_{s=1}^{S}$:

Definition 2.3.2 *A minimum-state-variable order-R regime-switching equilibrium (RSE) is a collection of measurable functions* $\{g_r, h_r\}_{r=1}^{R}$ *and a conditional distribution function* Π *such that*

1. *Rational Expectations: For any initial point* $X_1 \in \mathcal{X}$ *the outcomes generated*

[11]The first consideration limits the number of pre-determined and shock variables which can be incorporated into the model. Leveraging Taylor approximations reduces this computational bottle neck. The second often requires clever partitioning of the model state space to account for kinks and discontinuities. In contrast, I use the known partition of the control-space.

recursively according to

$$\begin{aligned} r_t \mid X_t &\sim \Pi_r(X_t) \\ Y_t = g_{r_t}(X_t), \tilde{X}_t &= h_{r_t}(X_t) \\ X_{t+1} = \tilde{X}_t + \Sigma\varepsilon_{t+1},\ \varepsilon_{t+1} &\overset{iid}{\sim} F \end{aligned}$$

satisfy model (2.13). *Specifically, for each* $(r,x) \in \{1,\ldots,R\} \times \mathcal{X}$*, if* $\Pi_r(x) > 0$ *then*

$$0 = \int \sum_{r'=1}^{R} f(g_{r'}(h_r(x) + \Sigma\varepsilon'), g_r(x), h_r(x), x)\Pi_{r'}(h_r(x) + \Sigma\varepsilon')dF(\varepsilon')$$

2. *Fundamental Volatility: Either* $\Pi_r(x) = 0$ *or* $\Pi_r(x) = 1$.

The first requirement states that the self-fulfilling beliefs of agents in the model can only put positive support on regimes which also satisfy the model equations. Regimes whose solution functions do not satisfy the model at the value of the state variables must occur with probability zero. The second condition requires that these self-fulfilling beliefs only place support on a single regime at each point of the state space. This imposes that the only source of volatility comes from fundamental shocks[12].

The REE concept and the RSE concept are closely related because they make identical predictions. This idea can be stated precisely by defining a notion of equivalence:

Definition 2.3.3 *A REE* (g,h) *and an order-R RSE are* equivalent equilibria *if for each* $x \in \mathcal{X}$ *and each* $r \in \{1,\ldots,R\}$ *if* $\Pi_r(x) > 0$ *then* $g_r(x) = g(x)$ *and* $h_r(x) = h(x)$.

[12]More generally, I could consider equilibria where sunspot shocks select between regimes so that multiple outcomes can occur in-equilibrium for a given point of the state space. From a solution strategy point of view, sunspot equilibria simply condition on an additional state variable, which can be appended to the model. For an example of this approach see Lubik and Schorfheide (2003) and Farmer and Khramov (2013).

An REE and an RSE are equivalent when they agree on contemporaneous outcomes across all points of the state space.

The formal link between these equilibrium concepts is summarized in the following proposition:

Proposition 2.3.1 *The collection* $(\{g_r, h_r\}_{r=1}^R, \Pi)$ *is an order-R RSE if and only there exists an equivalent REE.*

The proof of this proposition is constructive and shows precisely how the two concepts are related.

Proof. $(\Rightarrow)$: We can construct an REE from a given RSE by splicing its solution functions together. Given $(\{g_r, h_r\}_{r=1}^R, \Pi)$ choose $g(x) = \sum_{r=1}^R g_r(x)\Pi_r(x)$ and $h(x) = \sum_{r=1}^R h_r(x)\Pi_r(x)$. By construction, (g, h) is equivalent, provided that it is an REE. Due to the fundamental volatility requirement, Π is a degenerate distribution and so

$$\begin{aligned}
&\int f(g(h(x) + \Sigma\varepsilon'), g(x), h(x), x)\mathrm{d}F(\varepsilon') \\
&= \int f(g_{r'}(h_r(x) + \Sigma\varepsilon'), g_r(x), h_r(x), x)\Pi_{r'}(h_r(x) + \Sigma\varepsilon')\Pi_r(x)\mathrm{d}F(\varepsilon')
\end{aligned}$$

When $\Pi_r(x) = 0$ the last term must equal zero. However, if $\Pi_r(x) > 0$ this term will also equal zero due to the rational expectations requirement of an order-R RSE. Therefore (g, h) is an REE.

$(\Leftarrow)$: Necessity follows from realizing that for a given REE, (g, h), there are many equivalent RSE, and they differ according to their specification for out-of-equilibrium outcomes. Let Π be any conditional distribution over $\{1, \ldots, R\}$ which is degenerate (satisfies the fundamental volatility requirement). For each $x \in \mathcal{X}$, if $\Pi_r(x) > 0$ set $g_r(x) = g(x)$ and $h_r(x) = h(x)$. If $\Pi_r(x) = 0$, arbitrarily choose values for $g_r(x)$ and $h_r(x)$. These values are the out-of-equilibrium segments of the solution functions and may be chosen in any way without influencing equilibrium predictions. If the collection

$(\{g_r, h_r\}_{r=1}^R, \Pi)$ is an RSE, it is equivalent to (g, h) by construction.

Now, check the rational expectations requirement. For each in-equilibrium state-regime pair (so that $\Pi_r(x) > 0$) we have:

$$\begin{aligned}
&\int \sum_{r'=1}^{R} f(g_{r'}(h_r(x)+\Sigma\varepsilon'), g_r(x), h_r(x), x)\Pi_{r'}(h_r(x)+\Sigma\varepsilon')\mathrm{d}F(\varepsilon') \\
&= \int \sum_{r'=1}^{R} f(g(h(x)+\Sigma\varepsilon'), g(x), h(x), x)\Pi_{r'}(h_r(x)+\Sigma\varepsilon')\mathrm{d}F(\varepsilon') \\
&= \int f(g(h(x)+\Sigma\varepsilon'), g(x), h(x), x)\mathrm{d}F(\varepsilon')
\end{aligned}$$

Since (g, h) is an REE, the last term is equal to zero, and so $(\{g_r, h_r\}_{r=1}^R, \Pi)$ is an order-R RSE. ■

This proposition shows the relationship between the two equilibrium concepts. The in-equilibrium parts of an RSE's solution functions always match with the solution functions of some REE. The out-of-equilibrium parts of the functions g_r and h_r can never occur and do not influence equilibrium predictions. They are completely unrestricted by the definition of an RSE, and we have many possible RSE's simply by varying out-of-equilibrium outcomes.

Since out-of-equilibrium events are unrestricted, we can focus attention on sub-classes of RSE's by placing restrictions on out-of-equilibrium outcomes. In the previous section, we used this idea to work exclusively with linear solution functions. In the general case, we can focus on smooth solution functions due to assumption 2.3.1.

This is accomplished by identifying each of the R regimes with a segment of the piece-wise smooth model. This choice allows the regime variable to absorb all of the non-differentiable features in the model, effectively converting it into a smooth model with an endogenous regime variable.

This idea can be formalized by defining a concept of partition consistency:

Definition 2.3.4 *Given a control-space partition $\{\mathcal{Y}_s\}_{s=1}^{S}$, an order-R RSE is* partition consistent *if for each $r \in \{1,\ldots,R\}$ there exists some $s \in \{1,\ldots,S\}$ so that for every $x \in \mathcal{X}$ if $\Pi_r(x) > 0$ then $g_r(x) \in \mathcal{Y}_s$.*

Partition consistency implies that each regime is tied to a unique piece of the natural partition arising from the model's structural equations. The outcomes generated by each regime must lie in a single piece of the model's partition, and so knowing the current regime implies knowing the active piece of the model partition.

But since the model is smooth on the interior of each piece $\mathcal{Y}_s$, this implies that each solution function can be smoothly pinned down by the model structure while the regime stays fixed. To see this, fix the regime at r and use assumption 2.3.1 and partition consistency of the RSE to write:

$$
\begin{aligned}
&\Pi_r(x) > 0 \implies 0 \\
&= \int \sum_{r'=1}^{R} f(g_{r'}(h_r(x)+\Sigma\varepsilon'), g_r(x), h_r(x), x)\Pi_{r'}(h_r(x)+\Sigma\varepsilon')\mathrm{d}F(\varepsilon') \\
&= \int \sum_{r'=1}^{R}\sum_{s=1}^{S} f_s(g_{r'}(h_r(x)+\Sigma\varepsilon'), g_r(x), h_r(x), x)\mathbf{1}\{g_r(x)\in\mathcal{Y}_s\}\Pi_{r'}(h_r(x)+\Sigma\varepsilon')\mathrm{d}F(\varepsilon') \\
&= \int \sum_{r'=1}^{R} f_{s(r)}(g_{r'}(h_r(x)+\Sigma\varepsilon'), g_r(x), h_r(x), x)\Pi_{r'}(h_r(x)+\Sigma\varepsilon')\mathrm{d}F(\varepsilon')
\end{aligned}
$$

where $s : \{1,\ldots,r\} \to \{1,\ldots,S\}$ is the map from the current regime to the currently active partition. This mapping is well-defined because of partition consistency. The model is now entirely smooth, apart from any non-smoothness arising from a perfectly forecastable jump in the future regime that could be enduced by varying x. Provided that future regime changes are not extremely sensitive to current choices which determine the state, the model is completely smooth, and it is reasonable to look for smooth regime-conditional solution functions.

Remark 2.3.1 *A partition consistent RSE implies a completely smooth regime-switching* *model of the form*

$$0 = \int \sum_{r'=1}^{R} f_{s(r)}(g_{r'}(h_r(x) + \Sigma\varepsilon'), g_r(x), h_r(x), x)\Pi_{r'}(h_r(x) + \Sigma\varepsilon')dF(\varepsilon')$$

Intuitively, if regimes are based on the pieces of the partition $\{\mathcal{Y}_s\}_{s=1}^{S}$, then the model equations are completely smooth once cast into a regime-switching form. Focusing on smooth solution functions amounts to using analytic continuation to extend how the smooth structure of the model determines in-equilibrium outcomes to pin down out-of-equilibrium outcomes.

Because any RSE implies a unique equivalent REE, solving for RSE is without loss of generality and completely equivalent to solving for REE. This result suggests that we focus on solving for smooth solution functions in a regime-switching equilibrium, provided the regime variable is consistent with the partition $\{\mathcal{Y}_s\}_{s=1}^{S}$.

2.3.3 Regime-Switching Perturbation

The perturbation method approximates each regime specific solution function with a polynomial generated by a Taylor expansion. This procedure requires points for calculating derivatives of the structural equations. Traditionally, the deterministic steady state is used, but here reference points must move with the regime variable.

Instead of using a single steady state, the method uses multiple reference points which can be chosen a-priori. To apply the implicit function theorem, I setup an augmented model which nests model (2.13) and a slack model whose solution involves the chosen points. This procedure breaks perturbation's dependence on the model's steady state.

The implicit function theorem, applied to the augment model, gives a second order matrix equation in the derivatives of the model solution. Solving this matrix equation ends up being analogous to solving a Markov-switching model as in Foerster et al. (2013). This result connects non-linear models with endogenous regimes (via the RSE concept) to solution methods developed for Markov-switching rational expectations models.

Perturbation delivers the collection of approximate solution functions $\{\hat{g}_s, \hat{h}_s\}_{s=1}^S$ associated with a fed regime distribution, Π. To solve for the regime distribution, I use a backward induction approach based on the requirement that Π is a partition consistent regime distribution. In order for Π to be partition consistent, the outcomes generated by the system

$$\begin{aligned} Y_t &= \hat{g}_{s_t}(X_t) \\ X_{t+1} &= \hat{h}_{s_t}(X_t) + \Sigma\varepsilon_{t+1} \\ \varepsilon_{t+1} &\overset{iid}{\sim} F,\, s_{t+1} \mid X_{t+1} \sim \Pi_s(X_{t+1}) \end{aligned}$$

must satisfy $\mathbb{P}[Y_t \in \mathcal{Y}_s \mid X_t = x] = \Pi_s(x)$. This requirement is a fixed point condition which imposes rational expectations and implies that $(\{\hat{g}_s, \hat{h}_s\}_{s=1}^S, \Pi)$ is an approximate partition consistent RSE. At each iteration of the algorithm, I update the regime distribution based on the known partition $\{\mathcal{Y}_s\}_{s=1}^S$ so that it is consistent with solution function approximations. Using this distribution update step after each perturbation step, gives a complete iteration of the solution algorithm.

This procedure amounts to performing backward induction on the conditional distribution of the regime. Let an exponent of (n) denote an equilibrium object associated with the n-th iteration of the algorithm. Then each iteration n consists of a perturbation step to find estimates $\{\hat{g}_s^{(n)}, \hat{h}_s^{(n)}\}_{s=1}^S$ given $\Pi^{(n-1)}$, and then a distribution update step which calculates $\Pi^{(n)}$ given $(\{\hat{g}_s^{(n)}, \hat{h}_s^{(n)}\}_{s=1}^S, \Pi^{(n-1)})$. This generates a sequence $\{\Pi^{(n)}\}$. If this sequence converges, then the limit is an approximate equilibrium distribution.

The Perturbation Step

The augmented model which facilitates the perturbation step uses a nesting parameter $\eta \in [0,1]$ to combine model (2.13) with a slack model:

$$\begin{aligned} 0 &= \mathbb{E}_t\left[f(Y_{t+1},Y_t,\tilde{X}_t,X_t)-(1-\eta)\Delta_{s_{t+1},s_t}\right] \\ X_{t+1} &= \tilde{X}_t+(1-\eta)(\bar{x}_{s_{t+1}}-\tilde{x}_{s_t})+\eta\Sigma\varepsilon_{t+1} \\ \varepsilon_{t+1} &\overset{iid}{\sim} F, \text{ and } s_{t+1}\mid\tilde{X}_t,\varepsilon_{t+1}\sim\Pi_s(\tilde{X}_t+\Sigma\varepsilon_{t+1}) \end{aligned} \tag{2.15}$$

where the residual term $\Delta_{s,s'}$ is given by:

$$\Delta_{s,s'} = f(\tilde{y}_{s'},\tilde{y}_s,\tilde{x}_s,\bar{x}_s)$$

and the points $\{\tilde{y}_s,\tilde{x}_s,\bar{x}_s\}_{s=1}^{S}$ are any points such that each $\tilde{y}_s$ is in the interior of $\mathcal{Y}_s$. The region of the outcome space which is well approximated via perturbation is controlled by this collection of points.

The additive slack terms incorporated into this augmented model relax the original model (2.13). Including these slack terms modifies the evolution of state variables and adds a residual to the model equations. The distortion of the state variable evolution shifts the beliefs of agents so that the future state is close to the reference points at which the implicit function theorem is applied, while the residual term ensures that the slack model (the case when $\eta = 0$) always has the outcome $(\tilde{y}_s,\tilde{x}_s)$ as a solution when the state is $\bar{x}_s$. Setting $\eta = 1$ reduces this augmented model to the original model given in system (2.13).

An approximation of the underlying REE arises from using perturbation to solve for an approximate RSE. Since the solution functions of the RSE must solve (2.15) when $\eta = 1$, approximating a solution to the augmented model delivers a solution to the original model.

Holding the collection $\{\Pi_s, \tilde{y}_s, \tilde{x}_s, \bar{x}_s\}_{s=1}^S$ fixed, consider finding functions $\{g_s, h_s\}_{s=1}^S$ with domain of $\mathcal{X} \times [0,1]$ which solve (2.15). Then for every point $(s,x,\eta) \in \{1,\ldots,S\} \times \mathcal{X} \times [0,1]$ these solution functions must satisfy:

$$0 = \int \sum_{s'=1}^{S} \left[\begin{array}{r} f(g_{s'}(h_s(x,\eta) + (1-\eta)(\bar{x}_{s'} - \tilde{x}_s) \\ +\eta\Sigma\varepsilon', \eta), g_s(x,\eta), h_s(x,\eta), x) \\ -(1-\eta)\Delta_{s,s'} \end{array} \right] \Pi_{s'}(h_s(x) + \Sigma\varepsilon')\mathrm{d}F(\varepsilon') \tag{2.16}$$

Due to the residual term $\Delta_{s,s'}$, the solution functions associated with regime s must map the point $(x,\eta) = (\bar{x}_s, 0)$ to the outcome $(\tilde{y}_s, \tilde{x}_s)$:

$$\begin{aligned} \tilde{y}_s &= g_s(\bar{x}_s, 0) \\ \tilde{x}_s &= h_s(\bar{x}_s, 0) \end{aligned}$$

Given this known solution point, we can use the implicit function theorem to calculate the derivatives in the Taylor approximation:

$$\begin{aligned} g_s(x,\eta) &\approx \tilde{y}_s + \frac{\partial}{\partial x'} g_s(\bar{x}_s,0)(x-\bar{x}_s) + \frac{\partial}{\partial \eta} g_s(\bar{x}_s,0)\eta \\ h_s(x,\eta) &\approx \tilde{x}_s + \frac{\partial}{\partial x'} h_s(\bar{x}_s,0)(x-\bar{x}_s) + \frac{\partial}{\partial \eta} h_s(\bar{x}_s,0)\eta \end{aligned}$$

Implicitly differentiate equation (2.16) in x and evaluate at $(x,\eta) = (\bar{x}_s, 0)$ to get

$$0 = \sum_{s'=1}^{S} \left[F_{s,s'}^{(1,0,0,0)} G_{s'}^{(1,0)} H_s^{(1,0)} + F_{s,s'}^{(0,1,0,0)} G_s^{(1,0)} + F_{s,s'}^{(0,0,1,0)} H_s^{(1,0)} + F_{s,s'}^{(0,0,0,1)} \right] \bar{\Pi}_{s,s'} \tag{2.17}$$

where

$$\begin{aligned}
F_{s,s'}^{(i,j,k,l)} &\equiv \left.\frac{\partial f(y',y,x',x)}{\partial y'^i \partial y^j \partial x'^k \partial x^l}\right|_{(y',y,x',x)=(\tilde{y}'_s,\tilde{y}_s,\tilde{x}_s,\bar{x}_s)} \\
G_s^{(i,j)} &\equiv \left.\frac{\partial g_s(x,\eta)}{\partial x^i \partial \eta^j}\right|_{(x,\eta)=(\bar{x}_s,0)} \\
H_s^{(i,j)} &\equiv \left.\frac{\partial h_s(x,\eta)}{\partial x^i \partial \eta^j}\right|_{(x,\eta)=(\bar{x}_s,0)} \\
\bar{\Pi}_{s,s'} &\equiv \int \Pi_{s'}(\tilde{x}_s + \Sigma\varepsilon')\mathrm{d}F(\varepsilon')
\end{aligned}$$

denotes evaluated derivatives and average transition probabilities.

This second order matrix equation implicitly defines the collection of partial derivative matrices $\{G_s^{(1,0)}, H_s^{(1,0)}\}_{s=1}^{S}$. Note that the regime distribution influences the solution since the transition probabilities enter this condition. Foerster et al. (2013) arrive at an analogous second order system when working with Markov-switching models and use Groebner basis methods to calculate solutions to this problem. In this way, the first order approximation of the model can be calculated using any algorithm designed to solve Markov-switching rational expectations models[13].

I focus on solutions whose eigenvalues are as small as possible. I use a Bernoulli iteration method which is equivalent to Cho's (2013) Markov-switching generalization of the forward method introduced by Cho and Moreno (2008). This particular solution method selects an equilibrium with a number of theoretically desirable properties, as demonstrated by Cho (2013). It also can be understood as finding an equilibrium in which beliefs comes from a process of backwards induction, which makes it conceptually consistent with the backwards induction iteration that I use to solve for Π.

In particular, I iterate on the non-linear map M defined by solving for $[G_s^{(1,0)}, H_s^{(1,0)}]'$

[13]Linear or local approximation based solution approaches include Farmer et al. (2011), Bianchi and Melosi (2012), Foerster et al. (2013), and Cho (2013). The numerical linear algebra techniques in Dreesen et al. (2012) provide another alternative approach to solving this matrix equation.

in equation (2.17) :

$$\begin{bmatrix} H_s^{(1,0)} \\ G_s^{(1,0)} \end{bmatrix} = -\left(\sum_{s'=1}^{S}\left[F_{s,s'}^{(1,0,0,0)}G_{s'}^{(1,0)} + F_{s,s'}^{(0,0,1,0)} \quad F_{s,s'}^{(0,1,0,0)}\right]\bar{\Pi}_{s,s'}\right)^{+}\left(\sum_{s'=1}^{S}F_{s,s'}^{(0,0,0,1)}\bar{\Pi}_{s,s'}\right)$$
$$\equiv M_s(\{G_{s'}^{(1,0)}\}_{s'=1}^{S}), \ \forall s$$

where $+$ denotes the Moore-Penrose pseudo-inverse. This recursion is equivalent to the Markov-switching forward method of Cho (2013). In practice, it converges quickly. If convergence fails, the techniques in Foerster et al. (2013) can be used to calculate the full set of solutions to see if any exist.

Given the matrices $\{G_s^{(1,0)}, H_s^{(1,0)}\}_{s=1}^{S}$, next implicitly differentiate in η and evaluate at $(x,\eta) = (\bar{x}_s, 0)$ to get

$$0 = \sum_{s'=1}^{S}\left\{\begin{array}{c} F_{s,s'}^{(1,0,0,0)}\left[G_{s'}^{(1,0)}(H_s^{(0,1)} + \mu_{s,s'}) + G_{s'}^{(0,1)}\right] \\ +F_{s,s'}^{(0,1,0,0)}G_s^{(0,1)} + F_{s,s'}^{(0,0,1,0)}H_s^{(0,1)} \end{array}\right\}\bar{\Pi}_{s,s'}$$

where $\mu_{s,s'} \equiv \frac{\int(\Sigma\varepsilon' - \bar{x}_{s'} + \tilde{x}_s)\Pi_{s'}(\tilde{x}_s + \Sigma\varepsilon')\mathrm{d}F(\varepsilon')}{\int \Pi_{s'}(\tilde{x}_s + \Sigma\varepsilon')\mathrm{d}F(\varepsilon')}$ is a correction term to account for how the slack model was constructed and the regime-conditional mean of the fundamental innovation. Just like in Foerster et al. (2013), since the regime-conditional mean of ε_t doesn't drop out – as it does in standard perturbation – certainty equivalence does not hold[14].

This equation is linear in the collection of matrices $\{G_s^{(0,1)}, H_s^{(0,1)}\}_{s=1}^{S}$ and is straightforward to solve. As is the case in standard perturbation, from this point onwards, any n-th order approximation ($n \geq 2$) can be calculated using the $(n-1)$-th order approximation by implicit differentiation in x and η and evaluation at $(x,\eta) = (\bar{x}_s, 0)$ to get additional linear systems of equations[15].

[14] In standard perturbation this term is exactly zero since there is a single regime and the model is solved relative to the deterministic steady state (so that $\Pi_1 = 1$ and $\bar{x}_1 = \tilde{x}_1$).

[15] For instance, see Schmitt-Grohe and Uribe (2004) and Kim et al. (2008)

Focusing on the first order approximation for simplicity, set $\eta = 1$ to get an approximation to the original model's solution functions:

$$\begin{aligned}\hat{g}_s(x) &\equiv \left[\tilde{y}_s + G_s^{(0,1)}\right] + G_s^{(1,0)}(\bar{x}_s,0)(x-\bar{x}_s) \\ \hat{h}_s(x) &\equiv \left[\tilde{x}_s + H_s^{(0,1)}\right] + H_s^{(1,0)}(\bar{x}_s,0)(x-\bar{x}_s)\end{aligned} \tag{2.18}$$

Notice that the derivative in the nesting parameter η adjusts the constant term of the solution to correct for the additive slack terms used to specify the augmented model (2.15). Provided that Π is an equilibrium distribution, these functions provide approximations to the solution functions in an RSE of model (2.13).

This procedure takes a choice of $\{\Pi_s, \tilde{y}_s, \tilde{x}_s, \bar{x}_s\}_{s=1}^{S}$ and returns the approximate solution functions $\{\hat{g}_s, \hat{h}_s\}_{s=1}^{S}$. Within the overall algorithm, this perturbation step is used at each iteration to find approximate solution functions consistent with a given regime distribution.

Updating the Regime Distribution

Together, the conditional distribution and its implied solution function approximations define a regime-switching state-space system:

$$\begin{aligned}Y_t &= \hat{g}_{s_t}(X_t) \\ X_{t+1} &= \hat{h}_{s_t}(X_t) + \Sigma\varepsilon_{t+1} \\ \varepsilon_{t+1} &\overset{iid}{\sim} F,\ s_{t+1} \mid X_{t+1} \sim \Pi_s(X_{t+1})\end{aligned} \tag{2.19}$$

Solving for a partition consistent equilibrium requires finding a regime distribution that is self-fulfilling. Specifically, $\Pi_s(x)$ must be the conditional law $\mathbb{P}[Y_t \in \mathcal{Y}_s \mid X_t = x]$ of the process $\{Y_t, X_t, s_t\}$ that is generated by this system.

I propose using backwards induction to solve for the regime distribution. Given that outcomes evolve according to system (2.19), the probability that Y_t lies in the region

$\mathcal{Y}_s$ given that $X_t = x$ is precisely

$$\Pi'_s(x) = \sum_{\tilde{s}=1}^{S} \mathbf{1}\{\hat{g}_{\tilde{s}}(x) \in \mathcal{Y}_s\}\Pi_{\tilde{s}}(x) \tag{2.20}$$

This updating equation can be calculated exactly at any point of the state space for given $(\{\hat{g}_s, \hat{h}_s\}_{s=1}^S, \Pi)$. Recall that the fundamental volatility condition of an RSE requires that the regime distribution is degenerate. In this updating condition, if Π is a degenerate distribution, then Π' must also be a degenerate distribution. Therefore, by initializing the algorithm with a guess which is degenerate, we can always ensure that the fundamental volatility requirement is satisfied.

The Taylor approximation of the true solution functions is accurate only if regime-conditional outcomes tend to be close to the associated regime-conditional reference point. Therefore, updating the reference points $\{\tilde{y}_s, \tilde{x}_s, \bar{x}_s\}_{s=1}^S$ improves accuracy. However, to estimate the regime-conditional means generally requires simulation of the model, which is costly. I propose simulating the model only occasionally to update these reference points. I simulate data from system (2.19) to update the reference points with sample averages:

$$\begin{aligned} \tilde{y}'_s &= \frac{\sum_{t=1}^T Y_t \mathbf{1}\{Y_t \in \mathcal{Y}_s\}}{\sum_{t=1}^T \mathbf{1}\{Y_t \in \mathcal{Y}_s\}} \\ \tilde{x}'_s &= \frac{\sum_{t=1}^T (X_{t+1} - \Sigma\varepsilon_{t+1}) \mathbf{1}\{Y_t \in \mathcal{Y}_s\}}{\sum_{t=1}^T \mathbf{1}\{Y_t \in \mathcal{Y}_s\}} \\ \bar{x}'_s &= \frac{\sum_{t=1}^T X_t \mathbf{1}\{Y_t \in \mathcal{Y}_s\}}{\sum_{t=1}^T \mathbf{1}\{Y_t \in \mathcal{Y}_s\}} \end{aligned} \tag{2.21}$$

The simulated data can also be used to define the collection of points in the state space at which the distribution update is calculated.

2.3.4 The Algorithm

The most general statement of the proposed solution approach is given as algorithm 2.3.1. In practice, it is useful to initialize based on a first order solution from a linear approximation at a locally determinate steady state.

This is accomplished by solving for the steady state, solving for the first order approximation, simulating the model based on this initial guess, and then using the cloud of simulated data to form approximation nodes by averaging according to (2.21) and to choose state-space points at which to track the regime distribution. The initial guess for the regime distribution then follows from an initial updating step (as in equation (2.20)) based on the linear approximation.

Algorithm 2.3.1

1. *Choose initial* $(\{\tilde{y}^{(0)}, \tilde{x}^{(0)}, \bar{x}^{(0)}\}_{s=1}^{S}, \hat{\Pi}^{(0)})$. *And set* $n = 1$.

2. *Use implicit differentiation of the augmented model* (2.15) *with*

$$(\{\tilde{y}_s, \tilde{x}_s, \bar{x}_s\}_{s=1}^{S}, \Pi) = (\{\tilde{y}_s^{(n-1)}, \tilde{x}_s^{(n-1)}, \bar{x}_s^{(n-1)}\}_{s=1}^{S}, \hat{\Pi}^{(n-1)})$$

   to calculate the solution function approximations $\{\hat{g}_s^{(n)}, \hat{h}_s^{(n)}\}_{s=1}^{S}$ *according to equation* (2.18).

3. *Update to* $\Pi^{(n)}$ *using equation* (2.20).

4. *If updating reference points at this iteration, then simulate the state-space system* (2.19) *to generate data* $\{Y_t, X_t, \varepsilon_t\}_{t=1}^{T}$. *Then update to* $\{\tilde{y}_s^{(n)}, \tilde{x}_s^{(n)}, \bar{x}_s^{(n)}\}_{s=1}^{S}$ *according to system* (2.21). *Otherwise, set* $\{\tilde{y}_s^{(n)}, \tilde{x}_s^{(n)}, \bar{x}_s^{(n)}\}_{s=1}^{S} = \{\tilde{y}_s^{(n-1)}, \tilde{x}_s^{(n-1)}, \bar{x}_s^{(n-1)}\}_{s=1}^{S}$.

5. *If* $\hat{\Pi}^{(n)} \approx \hat{\Pi}^{(n-1)}$ *stop, else set* $n = n+1$ *and return to step 2.*

2.4 Simple ZLB Model: Numerical Results

This section presents numerical results for the simple Fisherian model of inflation determination presented in section 2.2. Instead of assuming that the real interest rate is *iid* in each period, I allow for persistence in the real interest rate. In this case, the non-linearity introduced by the ZLB implies equilibria are no longer piece-wise linear.

Although the model has no closed-form solution in this case, policy function iteration solves the model arbitrarily well. Using this global solution as a benchmark, I examine the accuracy of regime-switching perturbation based solutions of the model. I show that both first and second order regime-switching perturbation solutions of the model match the global solution very well.

For this exercise, I calibrate the model so that persistence enters both through the monetary policy rule and the real rate shock. I set the mean of the real rate shock to $\bar{r} = 0.03$, its persistence to $\rho_r = 0.5$, and its standard deviation to $\sigma = 0.01$. The persistence of the desired rate is set to $\rho_i = 0.5$ and the policy reaction coefficient is set to $\theta = 0.75$. These policy parameters correspond to an inflation reaction coefficient of $\phi = 1.5$ in a rule of the form

$$i_t^* = \rho_i i_{t-1}^* + (1-\rho_i)(r + \phi\pi_t)$$

Figure 2.4 compares a first order regime-switching perturbation based solution to the model with a solution using policy function iteration. The two solutions are very close to each other apart for very negative values of the real rate. They differ for $r_t < -0.02$. Since the mean of the shock of the real rate is 0.03 and the standard deviation of its innovation is 0.01 and persistence is 0.5, this region is visited very infrequently. In particular the standard deviation of the real rate is approximately 0.012 so this region is about 5 standard deviations away from the real rate's mean.

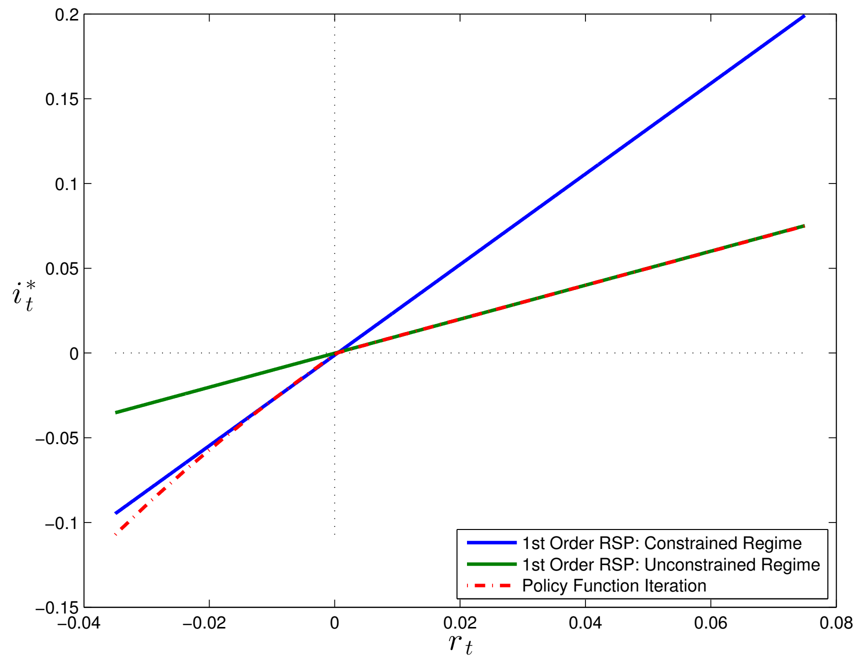

Figure 2.4: First Order Regime-Switching Perturbation Vs. Policy Function Iteration

Figure 2.5 shows the second order regime-switching perturbation solution of the model. Going to second order very slightly changes the solution by adding a small amount of curvature to the solution functions. This shows how the non-linearity in the true solution of the model is almost entirely absorbed into the regime variable. Higher orders of perturbation will do little to improve the accuracy of the solution.

This example model is not sufficiently complicated to demonstrate the computational gains from using perturbation. The model is solved in milliseconds. Further numerical work will demonstrate the computational scaling properties of the method.

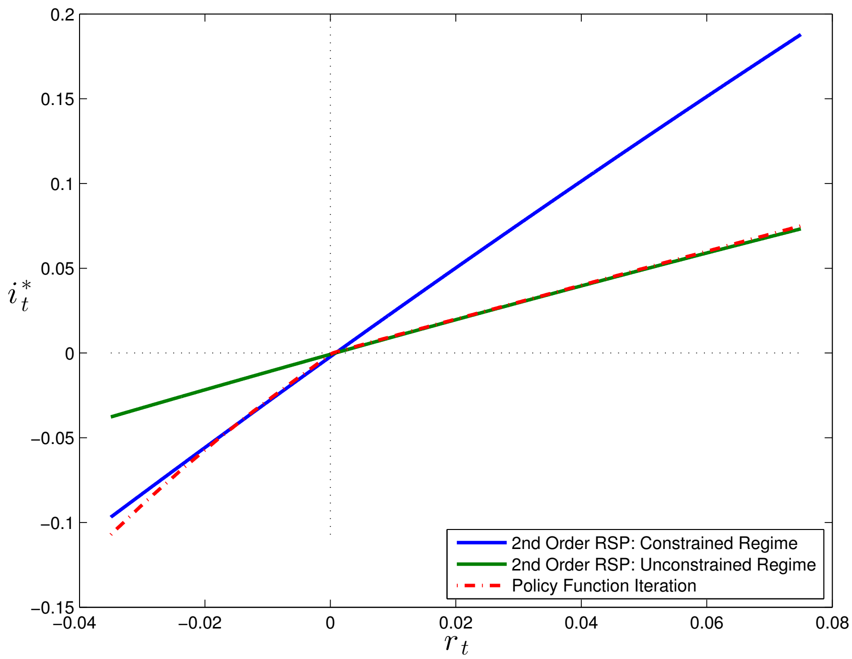

Figure 2.5: Second Order Regime-Switching Perturbation Vs. Policy Function Iteration

2.5 Conclusion

This paper introduces a generalization of perturbation which can be applied to highly non-linear equilibrium models. Using this method, economists can easily and systematically solve DSGE models that incorporate highly non-linear structural features – such as occasionally binding constraints and discontinuous equilibrium conditions. Previous solution methods either struggle to capture strong non-linearities without sacrificing model size or struggle to capture expectational effects associated with regime changes. Because the approach introduced in this paper uses perturbation but accounts for regime transition probabilities during solution, it overcomes these limitations.

The regime-switching equilibrium concept which underlies the perturbation approach provides a formal framework for modeling economic regimes. This concept allows economists to model the underlying determinants of sudden and large shifts in an economy. Markov-switching frameworks rely on assuming that these breaks occur exogenously, and previous endogenous-switching approaches require a known relationship between economic fundamentals and the regime. The regime-switching equilibrium concept allows the notion of an endogenous regime to be defined rigorously, and allows the relationship between fundamentals and the regime to be determined during solution. This approach provides a framework for building theoretical models to explain crises, policy changes, and other significant breaks in economic structure.

In addition, casting equilibria into a regime-switching form allows regime-switching econometric models to have a direct link to economic theories. In particular, the class of piecewise smooth rational expectations models can be identified equivalently as the class of endogenous regime models. The equilibria in these models can be represented using linear regime-switching state-space systems, and this connects theory to common empirical models. This connection opens the possibility of using two step estimation approaches based on reduced-form regime-switching econometric models to estimate highly non-linear DSGE models.

Appendix

2.A Proof of Lemma 2.2.1:

Substituting the hypothesized solution functions into the model gives (for each $s = u, c$):

$$
\begin{aligned}
0 \;\; = \;\; & \mathbb{E}_{r' \sim \mathcal{N}(\bar{r},\sigma^2)} \left\{ \begin{array}{r} \left[g_c((1-\eta)\bar{r}_c + \eta r', \eta) - (1-\eta)\Delta_{s,c} \right] \Pi_c(r') \\ + \left[g_u((1-\eta)\bar{r}_u + \eta r', \eta) - (1-\eta)\Delta_{s,u} \right] \Pi_u(r') \end{array} \right\} \\
& -(1-\rho_i)\bar{r} - \max\left\{ \rho_i g_s(r,\eta), (\rho_i + \theta) g_s(r,\eta) \right\} + \theta r
\end{aligned}
$$

Implicitly differentiate with respect to r at points (r, η) such that $g_s(r, \eta) \neq 0$ to get

$$
0 = -[\rho_i + \mathbf{1}\{g_s(r,\eta) > 0\}\theta] \frac{\partial}{\partial r} g_s(r,\eta) + \theta
$$

Note that the higher-order derivatives of this equation are all zero, and so there are no gains from higher order approximations. This is expected since the true solution is linear in r. Evaluate at $(r, \eta) = (\bar{r}_s, 0)$ to get

$$
\frac{\partial}{\partial r} g_s(\bar{r}_s, 0) = \frac{\theta}{\rho_i + \mathbf{1}\{s = u\}\theta} \equiv b_s
$$

In this calculation I replace the term $\mathbf{1}\{g_s(\bar{r}_s, 0) > 0\}$ with $\mathbf{1}\{s = u\}$ by using $g_s(\bar{r}_s, 0) = \bar{i}^*_s$, $\bar{i}^*_u > 0$, and $\bar{i}^*_c < 0$. This step reveals the importance of choosing the points $\bar{i}^*_s$ on both sides of the kink induced by the ZLB constraint. Choosing these points this way is necessary to get local information for both regimes from which to construct the regime-conditional solution functions.

The derivative with respect to η provides an adjustment term to the guess for the intercept of $\bar{i}_s^*$. Implicitly differentiating gives:

$$
\begin{aligned}
0 \quad = \quad & \mathbb{E}_{r' \sim \mathcal{N}(\bar{r},\sigma^2)} \left\{ \begin{array}{c} \left[\begin{array}{c} \frac{\partial}{\partial r} g_c((1-\eta)\bar{r}_c + \eta r', \eta)(r' - \bar{r}_c) \\ + \frac{\partial}{\partial \eta} g_c((1-\eta)\bar{r}_c + \eta r', \eta) + \Delta_{s,c} \end{array} \right] \Pi_c(r') \\ + \left[\begin{array}{c} \frac{\partial}{\partial r} g_u((1-\eta)\bar{r}_u + \eta r', \eta)(r' - \bar{r}_u) \\ + \frac{\partial}{\partial \eta} g_u((1-\eta)\bar{r}_u + \eta r', \eta) + \Delta_{s,u} \end{array} \right] \Pi_u(r') \end{array} \right\} \\
& -[\rho_i + \mathbf{1}\{g_s(r,\eta) > 0\}\theta] \frac{\partial}{\partial \eta} g_s(r,\eta)
\end{aligned}
$$

Note that any higher order derivatives in η must be zero because $\frac{\partial^2}{\partial \eta \partial r} g_s(r,\eta) = 0$. Evaluate at $(r,\eta) = (\bar{r}_s, 0)$ to get

$$
\begin{aligned}
0 \quad = \quad & \mathbb{E}_{r' \sim \mathcal{N}(\bar{r},\sigma^2)} \left[\left(b_c(r' - \bar{r}_c) + \frac{\partial}{\partial \eta} g_c(\bar{r}_c, 0) + \Delta_{s,c} \right) \Pi_c(r') \right] \\
& + \mathbb{E}_{r' \sim \mathcal{N}(\bar{r},\sigma^2)} \left[\left(b_u(r' - \bar{r}_u) + \frac{\partial}{\partial \eta} g_u(\bar{r}_u, 0) + \Delta_{s,u} \right) \Pi_u(r') \right] \\
& -[\rho_i + \mathbf{1}\{s = u\}\theta] \frac{\partial}{\partial \eta} g_s(\bar{r}_s, 0)
\end{aligned}
$$

where I have used $\mathbf{1}\{g_s(\bar{r}_s,0) > 0\} = \mathbf{1}\{s = u\}$ and $\frac{\partial}{\partial r} g_s(\bar{r}_s, 0) = b_s$. This expression defines a system of two linear equations in the two unknown derivatives $\frac{\partial}{\partial \eta} g_c(\bar{r}_c, 0)$ and $\frac{\partial}{\partial \eta} g_u(\bar{r}_u, 0)$.

Rewrite the system as a linear system in the coefficient

$$
a_s(\Pi) \equiv \left[\bar{i}_s^* + \frac{\partial}{\partial \eta} g_s(\bar{r}_s, 0) - b_s \bar{r}_s \right]
$$

by using the definition of the slackness variable $\Delta_{s,s'}$ and the fact that $\theta\bar{r}_s = [\rho_i + \mathbf{1}\{s = u\}\theta]b_s\bar{r}_s$:

$$[\rho_i + \mathbf{1}\{s = u\}\theta]a_s(\Pi) = \mathbb{E}_{r' \sim \mathcal{N}(\bar{r},\sigma^2)}\left[\left(a_c(\Pi) + b_c r'\right)\Pi_c(r') + \left(a_u(\Pi) + b_u r'\right)\Pi_u(r')\right] \\ - (1-\rho_i)\bar{r}$$

The right hand side of this equation does not depend on s and so $\rho_i a_c(\Pi) = (\rho_i + \theta)a_u(\Pi)$. Now, re-write this equation for $s = c$ as:

$$\rho_i a_c(\Pi) = \mathbb{E}_{r' \sim \mathcal{N}(\bar{r},\sigma^2)}\left[\frac{\rho_i a_c(\Pi) + \theta r'}{\rho_i}\Pi_c(r') + \frac{\rho_i a_c(\Pi) + \theta r'}{\rho_i + \theta}\Pi_u(r')\right] - (1-\rho_i)\bar{r}$$

which leads to

$$\rho_i a_c(\Pi) = \frac{\mathbb{E}_{r' \sim \mathcal{N}(\bar{r},\sigma^2)}\left[\frac{\theta}{\rho_i} r'\Pi_c(r') + \frac{\theta}{\rho_i+\theta} r'\Pi_u(r')\right] - (1-\rho_i)r}{1 - \mathbb{E}_{r' \sim \mathcal{N}(\bar{r},\sigma^2)}\left[\frac{1}{\rho_i}\Pi_c(r') + \frac{1}{\rho_i+\theta}\Pi_u(r')\right]} \equiv -\theta r^*(\Pi)$$

and

$$a_c(\Pi) = -b_c r^*(\Pi) \qquad a_u(\Pi) = -b_u r^*(\Pi)$$

2.B Proof of Proposition 2.2.1:

Suppose that Π is an equilibrium distribution. Using lemma 2.2.1, the equilibrium restrictions on Π are that $a_c(\Pi) + b_c r > 0 \Rightarrow \Pi_c(r) = 0$ and $a_u(\Pi) + b_u r \leq 0 \Rightarrow \Pi_u(r) = 0$. But $a_c(\Pi) + b_c r \leq 0 \Leftrightarrow r \leq r^*(\Pi)$ and $a_u(\Pi) + b_u r > 0 \Leftrightarrow r > r^*(\Pi)$, so these conditions imply that

$$\Pi_s(r) = \begin{cases} 1\{r \leq r^*(\Pi)\} & \text{if } s_t = c \\ 1\{r > r^*(\Pi)\} & \text{if } s_t = u \end{cases} \tag{2.22}$$

The threshold $r^*(\Pi)$ is a summary statistic for the equilibrium distribution since it is the threshold in the state space at which regime change occurs.

Define the quantity $i^{*E}(\Pi)$ (not necessarily equal to i^{*e}) as

$$i^{*E}(\Pi) \equiv \mathbb{E}_{r' \sim \mathcal{N}(\bar{r},\sigma^2)}\left[\left(a_c(\Pi)+b_c r'\right)\Pi_c(r')+\left(a_u(\Pi)+b_u r'\right)\Pi_u(r')\right]$$

Then from the definition of $r^*(\Pi)$, the constant term in the approximation is:

$$a_s(\Pi) = \frac{i^{*E}(\Pi)-(1-\rho_i)r}{\rho_i+\mathbf{1}\{s=u\}\theta}$$

Substitute the terms $a_c(\Pi)$ and $a_u(\Pi)$ into the definition of $i^{*E}(\Pi)$ to get:

$$\begin{aligned}
i^{*E}(\Pi) &= \int\left[\left(a_c(\Pi)+b_c r'\right)\Pi_c(r')+\left(a_u(\Pi)+b_u r'\right)\Pi_u(r')\right]\frac{1}{\sigma}\phi\left(\frac{r'-\bar{r}}{\sigma}\right)\mathrm{d}r' \\
&= \int\left[\frac{i^{*E}(\Pi)-(1-\rho_i)\bar{r}+\theta r'}{\rho_i}\Pi_c(r')+\frac{i^{*E}(\Pi)-(1-\rho_i)\bar{r}+\theta r'}{\rho_i+\theta}\Pi_u(r')\right] \\
&\qquad \times\frac{1}{\sigma}\phi\left(\frac{r'-\bar{r}}{\sigma}\right)\mathrm{d}r' \\
&= \int \max\left\{\frac{i^{*E}(\Pi)-(1-\rho_i)\bar{r}+\theta r'}{\rho_i},\frac{i^{*E}(\Pi)-(1-\rho_i)\bar{r}+\theta r'}{\rho_i+\theta}\right\}\frac{1}{\sigma}\phi\left(\frac{r'-\bar{r}}{\sigma}\right)\mathrm{d}r'
\end{aligned}$$

where the last line follows since Π is an equilibrium distribution and must be of the form (2.22). This condition is the same fixed point condition in (2.4). Therefore $i^{*E}(\Pi) = i^{*e}$ where i^{*e} is some solution of condition (2.4). Therefore, the constant terms must be identical to the constant terms in the true solution.

2.C Proof of Proposition 2.2.2:

The update equation (2.12) implies that after a single iteration of the algorithm, the threshold $r^*(\hat{\Pi}^{(n-1)})$ is a sufficient statistic for the distribution $\hat{\Pi}^{(n)}$. In particular:

$$\hat{\Pi}_s^{(n)}(r) = \begin{cases} \mathbf{1}\left\{r \leq r^*(\hat{\Pi}^{(n-1)})\right\} & \text{if } s = c \\ \mathbf{1}\left\{r > r^*(\hat{\Pi}^{(n-1)})\right\} & \text{if } s = u \end{cases}$$

This recursion can be re-expressed as a recursive condition on the threshold using the definition of r^*:

$$\begin{aligned} r^*(\hat{\Pi}^{(n)}) &= \frac{(1-\rho_i)\bar{r} - \int \left[\hat{\Pi}_c^{(n)} \frac{\theta}{\rho_i}(r' - r^*(\hat{\Pi}^{(n)})) + \hat{\Pi}_u^{(n)} \frac{\theta}{\rho_i+\theta}(r' - r^*(\hat{\Pi}^{(n)}))\right] \frac{1}{\sigma}\phi\left(\frac{r'-\bar{r}}{\sigma}\right) \mathrm{d}r'}{\theta} \\ &= \frac{(1-\rho_i)\bar{r}}{\theta} - \int \max\left\{\frac{1}{\rho_i}(r' - r^*(\hat{\Pi}^{(n-1)})), \frac{1}{\rho_i+\theta}(r' - r^*(\hat{\Pi}^{(n-1)}))\right\} \\ &\qquad \frac{1}{\sigma}\phi\left(\frac{r'-\bar{r}}{\sigma}\right) \mathrm{d}r' \\ &\qquad - \Phi\left(\frac{r^*(\hat{\Pi}^{(n-1)}) - \bar{r}}{\sigma}\right) \frac{1}{\rho_i}[r^*(\hat{\Pi}^{(n-1)}) - r^*(\hat{\Pi}^{(n)})] \\ &\qquad - \left[1 - \Phi\left(\frac{r^*(\hat{\Pi}^{(n-1)}) - \bar{r}}{\sigma}\right)\right] \frac{1}{\rho_i+\theta}[r^*(\hat{\Pi}^{(n-1)}) - r^*(\hat{\Pi}^{(n)})] \end{aligned}$$

where Φ is the cdf of the standard normal distribution. Let $r_n^* \equiv r^*(\Pi^{(n)})$. This threshold update equation can be written as:

$$\begin{aligned} r_n^* &= r_{n-1}^* + \frac{\frac{(1-\rho_i)\bar{r}}{\theta} + \int \max\left\{\frac{1-\rho_i}{\rho_i} r_{n-1}^* - \frac{1}{\rho_i} r', \frac{1-\rho_i-\theta}{\rho_i+\theta} r_{n-1}^* - \frac{1}{\rho_i+\theta} r'\right\} \frac{1}{\sigma}\phi\left(\frac{r'-\bar{r}}{\sigma}\right) \mathrm{d}r'}{1 - \left\{\Phi\left(\frac{r_{n-1}^*-\bar{r}}{\sigma}\right)\frac{1}{\rho_i} + \left[1 - \Phi\left(\frac{r_{n-1}^*-\bar{r}}{\sigma}\right)\right]\frac{1}{\rho_i+\theta}\right\}} \\ &\equiv m(r_{n-1}^*) \end{aligned} \tag{2.23}$$

The map m is continuously differentiable everywhere except for a single singularity at the point $\sigma\Phi^{-1}\left(\frac{\rho_i(1-\rho_i-\theta)}{\theta}\right)$. Moreover, it is strictly decreasing every except at this singularity. There are then precisely two points r^* such that $r^* = m(r^*)$. For sufficiently small σ, both of these points are dynamically stable. Therefore $r_n^* \to r^*$ for one of these two points.

It remains to show that any point r^* such that $r^* = m(r^*)$ must be associated with an equilibrium distribution. However, this result follows directly from the definition of the map m. Indeed, the above equalities can be reversed to conclude that

$$\Pi_s(r) = \begin{cases} \mathbf{1}\{r \leq r^*\} & \text{if } s = c \\ \mathbf{1}\{r > r^*\} & \text{if } s = u \end{cases}$$

is an equilibrium distribution for either solution r^*.

2.D RSP Solution of Model (2.1) with $\rho_r > 0$

With solution functions substituted in, the model reads:

$$0 = \int \sum_{s'=c,u} \begin{bmatrix} g_{s'}((1-\rho_r)\bar{r} + \rho_r r \\ +(1-\eta)(\bar{r}_{s'} - \tilde{r}_s) \\ +\eta\varepsilon', \eta) - (1-\rho_i)\bar{r} \\ -\max\{\rho_i g_s(r,\eta), (\rho_i + \theta) g_s(r,\eta)\} \\ +\theta r - (1-\eta)\Delta_{s,s'} \end{bmatrix} \Pi_{s'}((1-\rho_r)\bar{r} + \rho_r r + \varepsilon') \frac{1}{\sigma}\phi\left(\frac{\varepsilon'}{\sigma}\right) d\varepsilon' \tag{2.24}$$

where $\tilde{r}_s = (1-\rho_r)\bar{r} + \rho_r \bar{r}_s$. Note that

$$\Delta_{s,s'} \equiv \tilde{i}_{s'} - (1-\rho_i)\bar{r} - \max\left\{\rho_i \tilde{i}_s, (\rho_i + \theta)\tilde{i}_s\right\} + \theta \bar{r}_s$$

Differentiate with respect to r and η to get (suppressing function arguments):

$$0 = \int \sum_{s'=c,u} \left\{ \begin{array}{r} \left[\frac{\partial g_{s'}}{\partial r'}\rho_r - (\rho_i + \mathbf{1}\{g_s > 0\}\theta)\frac{\partial g_s}{\partial r} + \theta\right]\Pi_{s'} \\ + \left[g_{s'} - (1-\rho_i)\bar{r} - \max\{\rho_i g_s, (\rho_i+\theta)g_s\} + \theta r - (1-\eta)\Delta_{s,s'}\right]\frac{\partial \Pi_{s'}}{\partial r'}\rho_r \end{array} \right\} \frac{1}{\sigma}\phi\left(\frac{\varepsilon'}{\sigma}\right)d\varepsilon'$$

$$0 = \int \sum_{s'=c,u} \left\{ \left[\frac{\partial g_{s'}}{\partial r'}(\varepsilon' - \bar{r}_{s'} + \tilde{r}_s) + \frac{\partial g_{s'}}{\partial \eta} - (\rho_i + \mathbf{1}\{g_s > 0\}\theta)\frac{\partial g_s}{\partial \eta} + \Delta_{s,s'}\right]\Pi_{s'}\right\} \frac{1}{\sigma}\phi\left(\frac{\varepsilon'}{\sigma}\right)d\varepsilon'$$

In this short hand notation:

$$g_s = g_s(r,\eta)$$

$$g_{s'} = g_{s'}((1-\rho_r)\bar{r} + \rho_r r + (1-\eta)(\bar{r}_{s'} - \tilde{r}_s) + \eta\varepsilon', \eta)$$

$$\Pi_{s'} = \Pi_{s'}((1-\rho_r)\bar{r} + \rho_r r + \varepsilon')$$

Note that the derivative of the distribution is the weak derivative.

Evaluate these equations at $(r,\eta) = (\bar{r}_s, 0)$ using the fact that $\tilde{i}_s = g_s(\bar{r}_s, 0)$ and $\tilde{i}_c < 0 < \tilde{i}_u$ to get a system of linear equations in the unknown derivatives:

$$0 = \sum_{s'=c,u} G_{s'}^{(1,0)}\rho_r \bar{\Pi}_{s,s'} + \theta - (\rho_i + \mathbf{1}\{s=u\}\theta)G_s^{(1,0)}$$

$$0 = \int \sum_{s'=c,u} \left[G_{s'}^{(1,0)}\mu_{s,s'}^{(1)} + G_{s'}^{(0,1)} + \Delta_{s,s'}\right]\bar{\Pi}_{s,s'} - (\rho_i + \mathbf{1}\{s=u\}\theta)G_s^{(0,1)}$$

with the notation:

$$G_s^{(i,j)} \equiv \left.\frac{\partial^{i+j} g_s(r,\eta)}{\partial r^i \partial \eta^j}\right|_{(r,\eta)=(\bar{r}_s,0)}$$

$$\mu_{s,s'}^{(i)} \equiv \frac{\int (\varepsilon' - \bar{r}_{s'} + \tilde{r}_s)^i \Pi_{s'}(\tilde{r}_s + \varepsilon')\sigma^{-1}\phi(\varepsilon'/\sigma)d\varepsilon'}{\int \Pi_{s'}(\tilde{r}_s + \varepsilon')\sigma^{-1}\phi(\varepsilon'/\sigma)d\varepsilon'}$$

$$\bar{\Pi}_{s,s'} \equiv \int \Pi_{s'}(\tilde{r}_s + \varepsilon')\sigma^{-1}\phi(\varepsilon'/\sigma)d\varepsilon'$$

To get this expression, note that the term $g_{s'} - (1-\rho_i)\bar{r} - \max\{\rho_i g_s, (\rho_i+\theta)g_s\} + \theta r -$

$(1-\eta)\Delta_{s,s'}$ evaluates to zero by the construction of $\Delta_{s,s'}$. The solution of this system gives the first order approximation:

$$g_s(r,\eta) \approx \tilde{i}_s + G_s^{(1,0)}(r-\bar{r}_s) + G_s^{(0,1)}\eta$$

Next, take the second order derivatives of (2.24):

$$0 = \int \sum_{s'=c,u} \left\{ \begin{array}{r} \left[\frac{\partial^2 g_{s'}}{\partial r'^2}\rho_r^2 - (\rho_i + \mathbf{1}\{g_s>0\}\theta)\frac{\partial^2 g_s}{\partial r^2}\right]\Pi_{s'} \\ +2\left[\frac{\partial g_{s'}}{\partial r'}\rho_r - (\rho_i+\mathbf{1}\{g_s>0\}\theta)\frac{\partial g_s}{\partial r}+\theta\right]\frac{\partial \Pi_{s'}}{\partial r'}\rho_r \\ +\left[g_{s'} - (1-\rho_i)\bar{r} - \max\{\rho_i g_s, (\rho_i+\theta)g_s\} + \theta r - (1-\eta)\Delta_{s,s'}\right]\frac{\partial^2 \Pi_{s'}}{\partial r'^2}\rho_r^2 \end{array} \right\} \frac{1}{\sigma}\phi\left(\frac{\varepsilon'}{\sigma}\right) d\varepsilon'$$

$$0 = \int \sum_{s'=c,u} \left\{ \left[\begin{array}{r} \frac{\partial^2 g_{s'}}{\partial r'^2}\left(\varepsilon' - \bar{r}_{s'} + \tilde{r}_s\right)^2 + 2\frac{\partial g_{s'}}{\partial r'\partial\eta}\left(\varepsilon' - \bar{r}_{s'} + \tilde{r}_s\right) + \frac{\partial^2 g_{s'}}{\partial \eta^2} \\ -(\rho_i+\mathbf{1}\{g_s>0\}\theta)\frac{\partial^2 g_s}{\partial \eta^2} \end{array} \right] \Pi_{s'} \right\} \frac{1}{\sigma}\phi\left(\frac{\varepsilon'}{\sigma}\right) d\varepsilon'$$

$$0 = \int \sum_{s'=c,u} \left\{ \begin{array}{r} \left[\frac{\partial^2 g_{s'}}{\partial r'^2}\rho_r\left(\varepsilon' - \bar{r}_{s'} + \tilde{r}_s\right) + \frac{\partial^2 g_{s'}}{\partial r'\partial\eta}\rho_r - (\rho_i+\mathbf{1}\{g_s>0\}\theta)\frac{\partial^2 g_s}{\partial r\partial\eta}\right]\Pi_{s'} \\ +\left[\frac{\partial g_{s'}}{\partial r'}\left(\varepsilon' - \bar{r}_{s'} + \tilde{r}_s\right) + \frac{\partial g_{s'}}{\partial \eta} - (\rho_i+\mathbf{1}\{g_s>0\}\theta)\frac{\partial g_s}{\partial \eta} + \Delta_{s,s'}\right]\frac{\partial \Pi_{s'}}{\partial r'}\rho_r \end{array} \right\} \frac{1}{\sigma}\phi\left(\frac{\varepsilon'}{\sigma}\right) d\varepsilon'$$

Evaluate at $(r,\eta) = (\bar{r}_s, 0)$ to get a another linear system of equations in the second order derivatives given the first order derivatives

$$0 = \sum_{s'=c,u} \left\{ \begin{array}{r} \left[G_{s'}^{(2,0)}\rho_r^2 - (\rho_i+\mathbf{1}\{g_s>0\}\theta)G_s^{(2,0)}\right]\bar{\Pi}_{s,s'} \\ +2\left[G_{s'}^{(1,0)}\rho_r - (\rho_i+\mathbf{1}\{g_s>0\}\theta)G_s^{(1,0)}+\theta\right]\bar{\Pi}_{s,s'}^{(1)}\rho_r \end{array} \right\}$$

$$0 = \int \sum_{s'=c,u} \left\{ \left[\begin{array}{r} G_{s'}^{(2,0)}\mu_{s,s'}^{(2)} + 2G_{s'}^{(1,1)}\mu_{s,s'}^{(1)} + G_{s'}^{(0,2)} \\ -(\rho_i+\mathbf{1}\{g_s>0\}\theta)G_s^{(0,2)} \end{array} \right] \bar{\Pi}_{s,s'} \right\}$$

$$0 = \int \sum_{s'=c,u} \left\{ \begin{array}{r} \left[G_{s'}^{(2,0)}\rho_r\mu_{s,s'}^{(1)} + G_{s'}^{(1,1)}\rho_r - (\rho_i+\mathbf{1}\{g_s>0\}\theta)G_s^{(1,1)}\right]\bar{\Pi}_{s,s'} \\ +\left[G_{s'}^{(1,0)}\mu_{s,s'}^{(1/1)} + G_{s'}^{(0,1)} - (\rho_i+\mathbf{1}\{g_s>0\}\theta)G_s^{(0,1)} + \Delta_{s,s'}\right]\bar{\Pi}_{s,s'}^{(1)}\rho_r \end{array} \right\}$$

where now

$$\begin{aligned}
\mu_{s,s'}^{(i/j)} &\equiv \frac{\int (\varepsilon' - \bar{r}_{s'} + \tilde{r}_s)^i \frac{\partial^j \Pi_{s'}(\tilde{r}_s+\varepsilon')}{\partial r'^j} \sigma^{-1}\phi(\varepsilon'/\sigma)d\varepsilon'}{\int \frac{\partial^j \Pi_{s'}(\tilde{r}_s+\varepsilon')}{\partial r'^j} \sigma^{-1}\phi(\varepsilon'/\sigma)d\varepsilon'} \\
&= \frac{\int (\varepsilon' - \bar{r}_{s'} + \tilde{r}_s)^i \Pi_{s'}(\tilde{r}_s + \varepsilon')\sigma^{-1-j}\phi^{(j)}(\varepsilon'/\sigma)d\varepsilon'}{\int \Pi_{s'}(\tilde{r}_s + \varepsilon')\sigma^{-1-j}\phi^{(j)}(\varepsilon'/\sigma)d\varepsilon'} \\
\bar{\Pi}_{s,s'}^{(i)} &\equiv \int \Pi_{s'}(\tilde{r}_s + \varepsilon')\sigma^{-1-j}\phi^{(j)}(\varepsilon'/\sigma)d\varepsilon'
\end{aligned}$$

where $\phi^{(j)}$ denotes the j-th derivative of the standard normal pdf and the second equality in both lines follows from the definition of the weak derivative.

Solving for these second order derivatives leads to the second order Taylor approximation:

$$\begin{aligned}
g_s(r,\eta) \approx\ & \tilde{i}_s + G_s^{(1,0)}(r - \bar{r}_s) + G_s^{(0,1)}\eta \\
& + \frac{1}{2}G_s^{(2,0)}(r - \bar{r}_s)^2 + G_s^{(1,1)}(r - \bar{r}_s)\eta + \frac{1}{2}G_s^{(0,2)}\eta^2
\end{aligned}$$

2.E Exogenous and Threshold Regime-Switching as a Special Case

To construct a Markov-switching model, factor the control vector as $Y_t = (Y_t^0, Y_t^s)$ where $Y_t^s = s_t$ denotes the present regime. Then let one of the state variables keep track of the past regime, so factor the state vector in a similar way as $X_t = (X_t^0, X_t^s)$ and include an equation which specifies that the next period value of this state variable is the current value of the regime: $X_{t+1}^s = Y_t^s$. Assume that there is a component of the exogenous shock vector υ_t which (in-equilibrium) governs the regime given the past regime (which is equal to the value of X_t^s). Partition the shocks as: $\varepsilon_t = (\varepsilon_t^0, \upsilon_t)$. Assume that $\upsilon_t = (\upsilon_t^1, \ldots, \upsilon_t^S)$ with υ_t^s an IID extreme value type 1 random variable. Let $\{\beta_{s,s'}\}_{s,s'=1}^S$ be coefficients which will control the switching probabilities and specify the present regime as:

$$X_{t+1}^{s} = Y_t^{s} = s_t = \sum_{r=1}^{S} s 1 \left\{ s = \arg\max_{s} \left[\beta_{s_0,s} 1 \left\{ X_t^{s} \in R_{s_0} \right\} + \upsilon_t^{s} \right] \right\} \tag{2.25}$$

Here $\{R_s\}_{s=1}^{N_S}$ defines a partition of $\mathbb{R}$ so that for each $s = 1, \ldots, S$ the integer s is in the interior of R_s. As a result, the equation is well defined on the entire real line for any $X_t^s \in \mathbb{R}$. Moreover, the equilibrium value of the regime variable is $s_t \in \{1, \ldots, S\}$. Since this equation takes the form of a multinomial logit random utility model, the regime transition probabilities are:

$$\mathbb{P}\left[s_{t+1} = s' \mid s_t = s\right] = \mathbb{P}\left[s_{t+1} = s' \mid x_t^s \in R_s\right] = \frac{\exp\left[\beta_{s,s'}\right]}{\sum_{s'=1}^{S} \exp\left[\beta_{s,s'}\right]}$$

The matrix P with entries of $P_{s,s'} = \exp\left[\beta_{s,s'}\right] / \sum_{s'=1}^{S} \exp\left[\beta_{s,s'}\right]$ is the regime transition matrix. In this way, the present methodology generalizes Foerster et al. (2013). The generalization to threshold-switching (called endogenous switching in Davig and Leeper (2006)) arises from letting the coefficients in specification (2.25) depend on the state vector (X_t^0). Then the transition probabilities become

$$\mathbb{P}\left[s_{t+1} = s' \mid X_{t+1} = (x^0, s)\right] = \frac{\exp\left[\beta_{s,s'}(x^0)\right]}{\sum_{s'=1}^{S} \exp\left[\beta_{s,s'}(x^0)\right]}$$

2.F Derivation of Model (2.1)

A representative consumer has a stochastic endowment of Y_t which follows an AR(1) process in logs: $\ln Y_t = \rho_y \ln Y_{t-1} + \varepsilon_t^y$ where $\varepsilon_t^y \overset{iid}{\sim} \mathcal{N}(\sigma^2/2, \sigma^2)$. This consumer has unlimited access to risk-free nominal bonds, and derives utility from consumption and real balances. Its inter-temporal optimization problem is:

$$\begin{aligned}\max_{\{C_t, v_{t+1}, M_{t+1}\}_{t=0}^{\infty}} \quad & \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t [\ln C_t + v(M_{t+1}/P_t)] \\ \text{s.t.} \quad & P_t C_t + Q_t B_{t+1} + M_{t+1} \leq P_t Y_t + B_t + M_t + T_t\end{aligned}$$

where P_t is the nominal price of the consumption good, Q_t denotes the price of a one period risk-less nominal bond, B_{t+1} denotes bond holdings, M_{t+1} denotes end of period cash-on-hand, and T_t denotes monetary transfers to the representative household conducted by the central bank. The function v is differentiable, concave, increasing, and has a satiation point μ such that $v'(m) = 0$ for all levels of real balances greater than μ.

Since output has no alternative use to consumption, the goods market clearing condition states that $C_t = Y_t$. Using this condition in the first order conditions of this problem gives:

$$\begin{aligned}Q_t Y_t^{-1} &= \beta \mathbb{E}_t Y_{t+1}^{-1} \frac{P_t}{P_{t+1}} = \beta Y_t^{-\rho_y} \mathbb{E}_t \Pi_{t+1}^{-1} \\ v'(M_{t+1}/P_t) &= 1 - Q_t\end{aligned}$$

where $\Pi_{t+1} \equiv P_{t+1}/P_t$ is the gross inflation rate.

Due to the satiation point in the function v, any choice of money supply greater than $P_t\mu$ will imply that the bond price is 1. For levels of money supply below this satiation level, it must be the case that nominal bonds sell at a discount.

Assuming that monetary policy uses monetary transfers T_t to control the yield on the nominal bond, we can abstract from the level of money supply and assume that the bond price is set directly by the central bank, so long as the implied nominal interest rate $i_t \equiv \ln Q_t^{-1}$ satisfies $i_t \geq 0$. That is, policy must respect the zero-lower-bound on interest implied by to coexistence of cash and nominal bonds.

The policy rule used by the central bank is assumed to take a shadow-rate Taylor-rule form with smoothing:

$$
\begin{aligned}
i_t &= \max\{0, i_t^*\} \\
i_t^* &= \rho_i i_{t-1}^* + (1-\rho_i)\bar{r} + \theta\pi_t
\end{aligned}
$$

where i_t^* denotes the central bank's desired (shadow) rate, $\pi_t \equiv \ln(P_t/P_{t-1})$ is the inflation rate, and $\bar{r} \equiv \ln(\beta^{-1} - \sigma^2/2)$. These two equations are identical to the final two equations of model (2.1). Note that this policy rule respects the zero-lower-bound by construction.

Finally, using the approximation that $\Pi_{t+1} \approx 1$, and defining the real interest rate as $r_t = \ln\beta^{-1} - (1-\rho_y)\ln Y_t$ we can write the household's Euler equation and the law of motion for the real interest rate as:

$$
\begin{aligned}
r_t &= i_t - \mathbb{E}_t\pi_{t+1} \\
r_t &= \rho_r r_{t-1} + (1-\rho_r)\bar{r} + \varepsilon_{t+1},\, \varepsilon_{t+1} \overset{iid}{\sim} \mathcal{N}(\bar{r}, \sigma^2)
\end{aligned}
$$

These two equations are identical to the first two equations of model (2.1).

Chapter 3

Credit Regimes and the Seeds of Crisis

Beginning in 2003, the US mortgage market underwent an unprecedented pivot in lending standards — between 2003 and 2004, the share of non-prime mortgage originations increased from 10% to 28%, peaking at 39% in 2006. Simultaneously, the labor market began to recover from the recession of 2001 and the housing market boomed. Employment growth jumped from 0% in 2003 to just below 2% in 2004, and growth in national home price indices doubled from 6% to 12%. This boom ended in the housing crisis of 2006-2007, the financial crisis of 2007-2008, and, ultimately, the Great Recession. A large and growing literature points to the initial change in mortgage markets as central to understanding this period. To assess policies meant to protect against similar episodes, economists need a theory explaining this shift in lending standards.

This paper develops a theory of rational non-prime lending and housing boom-bust cycles. In addition to matching existing microeconomic evidence, the theory allows us to understand the role of changing lending standards in fueling a housing boom, and how rising home prices influence the allocation of credit across different borrowers.

The theory is built on rational and competitive behavior by lenders who face asymmetric information about borrower income risk. Borrowers are either safe, having stable income, or risky, with uncertain income. Lenders can only observe borrower outcomes such as their credit history and employment status, and based on this limited information they offer mortgage contracts to individual borrowers.

The market is endogenously split into a pooling segment and a screening segment. Lenders pool together borrowers who appear unlikely to be risky, and offer them a low interest rate. Safe borrowers in this segment have little gain from revealing their type. In contrast, the screening segment consists of those borrowers who appear likely to be risky. Safe borrowers differentiate themselves by either making a high downpayment or undergoing income verification, and benefit from a low interest rate after revealing their type. As a result, risky borrowers in the screening segment can only borrow at a high interest rate.

The pooling segment can be interpreted as the market for Alt-A mortgages — where loans require high credit scores but do not require income verification. In contrast, the screening segment corresponds to prime and subprime mortgage markets. Prime mortgages typically have low interest rates but require income documentation or a sizable downpayment, while subprime mortgages have high interest rates and may have either limited income documentation or a low downpayment. Together, the pooling segment and the subprime portion of the screening segment represent non-prime mortgages. By endogenizing the share of the overall mortgage market where pooling occurs, the theory can speak to the rise of non-prime lending during the mid 2000's.

In doing so, we also get a theory of housing boom-bust cycles driven by shifts in lending standards. Due to how segmentation in the mortgage market changes with economic fundamentals and home prices, the economy may operate in two different credit regimes — a "pooling regime" and a "screening regime" — corresponding to to whether the marginal home owner is offered a pooling contract or is screened out of the prime market.

When the marginal homeowner is screened, they only have access to a subprime mortgage with a high interest rate. Given this elevated interest rate, they have a low home valuation. Since the marginal home buyer must be indifferent between owning and renting, home prices must be low in the screening regime.

However, if risky borrowers become less likely to default — say during an economic expansion — then safe borrowers have less reason to differentiate themselves and the pooling segment of the mortgage market expands. If lenders begin to offer a pooling contract to the marginal home buyer, then this expansion triggers the pooling regime.

During this regime change, interest rates discretely fall, home prices rapidly rise, and lenders substitute away from using income verification toward using downpayment requirements to ration credit. These results come from joint mortgage and housing market clearing, and how the screening in the mortgage market varies with home prices. When a fall in income risk increases the extent of pooling, it generates an excess demand for housing. Risky borrowers gain access to low interest rate loans, which attracts former renters into the housing market. They begin bidding up home prices. For the housing market to clear, some of these former renters must be rationed back out of the mortgage market as home prices rise.

The market clearing mechanism comes from how incentive constraints link downpayment requirements to home prices. Rising home prices make the outside option of renting more attractive, which relaxes incentive constraints. As a result, if lenders are using high downpayments to screen out risky borrowers, then they can relax their downpayment requirements. This change in lending standards make screening more attractive, which reduces the size of the pooling segment. With less pooling and more screening, more risky borrowers get rationed out of the market, which dampens housing demand.

However, if lenders do not use downpayment-based screening, then the housing market cannot clear through this mechanism. For the housing market to clear, lenders must first stop using income verification to screen borrowers. They will do so only if they can attract safe borrowers by offering a small downpayment. A large enough appreciation of home prices will relax downpayment requirements enough to incentivize a switch away from income verification. As a result, joint equilibrium between housing and mortgage markets implies a rapid home price appreciation, a relaxation of downpayment requirements, and an abandonment of income verification. This general equilibrium mechanism provides an explanation for the sudden fall in income documentation that occurred during the housing boom (Jiang et al. (2014)).

The credit regime captures strong non-linear feedback between housing and mortgage markets and leads to housing boom-bust cycles[1]. If the underlying fundamentals of the economy are mean reverting, then a boom triggered by falling income risk must eventually end. When the economy returns to the screening regime, credit tightens for risky borrowers, demand for housing falls, and home values collapse. As home prices fall, borrowers go underwater on their homes, which generates a foreclosure crisis. Note that both safe and risky borrowers go into default, which lines up with the microeconomic evidence on foreclosures during the housing crisis (Ferreira and Gyourko (2015), Adelino et al. (2016)). In this theory, a foreclosure crisis is the ultimate unwinding of a credit-fueled housing boom.

Non-linearity is central to this result. The regime summarizes a discontinuity in equilibrium outcomes, which makes standard perturbation solution methods inappropriate.

[1]The mechanism in this theory operates through the terms of mortgage contracts and lets us match the shift in lending standards during this period. Boom-bust cycles may arise through other mechanisms such as adaptive or extrapolative expectations and speculative and social dynamics — see Lo (2004), Barberis et al. (2015), DeFusco et al. (2017), and Burnside et al. (2016).

To solve the infinite horizon economy, I apply the regime-switching perturbation method in Chapter 2. This solution method leverages the speed of perturbation and absorbs non-linearities into an endogenous regime variable. As a result, despite having a highly non-linear mechanism and many state variables, we can study equilibrium dynamics driven by endogenous credit regimes.

By highlighting the importance of state-dependent amplification in driving boom-bust cycles[2], this paper builds on and complements the results of Gorton and Ordoñez (2012) and Gorton and Ordoñez (2016). In these papers, financial crises arise in short-term collateralized debt markets when market participants have an ex-post incentive to secretly produce information about collateral (mortgage-backed security) quality. Similarly, my model is built on asymmetric information and generates highly non-linear dynamics[3]. It differs in that I study the joint determination of mortgage contracts and home prices, and the crisis does not come from an ex-post incentive problem. Rather, the key mechanism comes from general equilibrium feedback between the mortgage market and the housing market. Since the foreclosure crisis led investors to worry about the quality of mortgage-backed securities, this paper complements the contributions of Gorton and Ordoñez (2012) and Gorton and Ordoñez (2016). It endogenizes the housing crisis that precipitated the financial crisis.

To diagnose the crisis, economists have turned to credit bureau data to ask whether or not there was reallocation of credit towards risky borrowers during the boom. Different studies use different measures of borrower riskiness and find seemingly conflicting results. Mian and Sufi (2015) use a measure of *historical* riskiness based on fixed groups of borrowers. They find that borrowers with low FICO scores in 1997 had a larger increase

[2]For recent work on non-linear dynamics that generate limit cycles see Beaudry et al. (2015).

[3]Guerrieri and Uhlig (2016) also examine how adverse selection in credit markets influences the housing market. Their model generates housing cycles through switching between multiple equilibria. In contrast, the theory in this paper focuses on a unique stable equilibrium and generates boom-bust cycles in response to shocks to fundamentals through a non-linear feedback mechanism between housing and mortgage markets.

in debt than borrowers with high FICO scores in 1997. In contrast, Albanesi et al. (2016) use a measure of *observable* riskiness. They examine the distribution of debt across borrowers by FICO score two years before origination and show that debt by FICO score grew more for high FICO score borrowers.

The theory proposed by this paper can simultaneously match the evidence from both papers. The difference in their results highlights the importance of distinguishing between fundamental borrower types (risky versus safe) and lender information about borrower types. To the extent that types are persistent, I argue that the Mian and Sufi (2015) approach is more likely to measure reallocation across borrower types. In my model, the average loan-to-value ratio of safe borrowers falls during the boom, while risky borrower loan-to-value ratios are stable — matching the evidence on loan-to-value ratios in Mian and Sufi (2015). Given this difference in loan-to-value ratios, safe borrowers accumulate less debt than risky borrowers during a housing boom.

However, these results reverse if we condition on lender beliefs about borrower types rather than fundamental types. Those safe borrowers who appear risky to lenders end up using high-downpayment mortgages to signal their type during the boom. Those safe borrowers who are obviously safe do not signal their type and have stable loan-to-value ratios in parallel to risky borrowers. As a result, loan-to-values are more stable among borrowers who seem safe from the perspective of lenders at the time of origination. Therefore, conditioning on lender information reverses the pattern of debt accumulation. This predicted reversal is consistent with the evidence of Albanesi et al. (2016), who study debt growth conditional on lender information close to the time of origination. By matching both sets of results, the theory may help explain why these authors get seemingly conflicting results from the same data — the measure of riskiness that each study uses captures a distinct theoretical concept. Patterns of debt accumulation can differ significantly depending on whether or not we condition on the information available to lenders at the time of mortgage origination.

The theory also provides an explanation for the mortgage rate conundrum of Justiniano et al. (2016). These authors document that from mid-2003 to 2006 mortgage rates were low relative to their historical relationship with US treasury yields. In the summer of 2003 the yield curve began to rise following the Federal Reserve's June 2003 announcement that it was ending its easing cycle. Following similar past annoucements, mortgage rates rose with the yield curve. This time, they did not.

The theory explains this puzzle through a switch from the screening regime to the pooling regime. The interest rates of risky borrowers fall when they get pooled with safe borrowers — pooling dilutes their default risk. The interest rates of safe borrowers also fall. Those safe borrowers who switch to high-downpayment based screening have reduced loan-to-value ratios. With a larger equity cushion, they are less likely to default and have a reduced risk premium. The switch away from income verification toward downpayment-based rationing implies that mortgage rates fall for safe borrowers on average. As a result, when a fall in income risk triggers a shift in the credit regime, mortgage rates fall relative to the risk free rate for all borrowers. The endogenous response of monetary policy to improving labor market conditions accounts for the rise in the yield curve, and an endogenous shift in the credit regime accounts for the low level of mortgages rates relative to treasury yields.

This paper also speaks to the macroeconomic literature on the housing boom by endogenizing the shocks that explain the boom-bust cycle. Previous research points to credit supply shocks as an explanation of the housing boom (for instance, Mian and Sufi (2009), Justiniano et al. (2015a), and Mian and Sufi (2016)). The model can explain these shocks through switches in the credit regime. Even a small change in income risk can trigger a regime change, generating a large shift away from high documentation lending and a rapid appreciation in home prices. When this occurs, credit rationing decreases and it is as-if a credit supply shock hit the economy.

Kaplan et al. (2016) replicate the boom-bust episode in a quantitative heterogenous agent macroeconomic model calibrated to cross-sectional facts. They find that three shocks are necessary to match the time series data: an income shock, a loan-to-value shock, and a home price expectation shock[4]. The theory I propose explains the loan-to-value ratio shock through the endogenous increase in loan-to-value ratios that occurs when rising home prices relax incentive constraints. The shock to home price expectations is explained through the shift in the credit regime. From the perspective of a traditional asset pricing condition for homes, a discrete switch in the contract faced by the marginal home owner appears as a wedge shock. With rational expectations and some persistence in the credit regime, it generates an endogenous anticipation of high home prices while lenders continue to offer pooling contracts to the marginal home buyer. The mechanism developed in this paper may enable economists to reduce the number of shocks that are required to explain the housing boom and bust.

Because the mechanism generates endogenous boom-bust cycles by amplifying underlying business cycle fundamentals, this theory points to possible non-prime lending booms during future macroeconomic expansions — indeed, the re-emergence of subprime lending in both mortgages and car loans has made headlines recently[5]. This prediction is consistent with the historical evidence, which suggests that credit-driven housing booms and busts are rare but systemic events. For instance, Jordà et al. (2015) show that, across countries, an expansion of credit during a housing boom predicts a subsequent bust. The theory proposed in this paper formalizes this perspective.

More broadly, this model provides a framework for studying credit rationing in business cycle analysis. One lesson is that the price of assets that serve as collateral can

[4]Similarly, Gelain et al. (2015) find that shocks to lending standards can explain the boom-bust episode in a representative agent framework.

[5]For example, see "Subprime Bonds Are Back With Different Name Seven Years After U.S. Crisis" Bloomberg News, January 27, 2015.

change rapidly and even discontinuously in fundamentals when there is adverse selection in credit markets[6]. This strong form of amplification is distinct from workhorse models of collateralized debt used when studying business cycles — for example, Kiyotaki and Moore (1997) and Bernanke et al. (1999). This theory provides a new approach to modeling asset price booms and financial frictions in business cycle analysis.

The paper is organized as follows. Section 3.1 introduces a partial equilibrium model of endogenous segmentation in the mortgage market. Section 3.2 examines joint mortgage market and housing market equilibrium and shows how small changes to income risk can sometimes generate big changes in mortgage contracts and home prices. In Section 3.3, I extend the model to examine equilibrium dynamics. I show that the accumulation of small shocks to income risk can trigger a housing boom which then leads to a housing crisis. Section 3.4 shows that the theory can simultaneously match the seemingly conflicting results in the empirical literature on debt reallocation and also accounts for the mortgage rate conundrum. Finally, Section 3.5 concludes.

3.1 Endogenous Segmentation in the Mortgage Market

This section develops a model of the mortgage market where borrowers are endogenously segmented by the type of equilibrium that prevails — either a screening, separating, or pooling equilibrium. This segmentation arises from how lenders tailor mortgage offers to individual borrowers based on their observable characteristics.

I first setup the economic environment, and then I examine equilibrium segmentation. Note that for this partial equilibrium analysis home prices are exogenous. In Section 3.2 I endogenize home prices.

[6] See Bigio (2015) for a model where adverse selection plays a key role in business cycles via the market for liquidity.

3.1.1 Setup

The economy consists of a unit mass of borrowers and at least two lenders. There are two periods, $t = 0, 1$, with mortgage contracts signed at $t = 0$ before uncertainty is resolved at $t = 1$.

Each borrower $i \in [0, 1]$ has the same linear preferences:

$$u(C_{i,0}, C_{i,1}) = C_{i,0} + \beta \mathbb{E} C_{i,1}$$

where $C_{i,0}$ denotes period 0 consumption of numeraire, $C_{i,1}$ denotes period 1 consumption, and $\beta > 0$. Borrowers come in two types, safe and risky. The variable τ_i indicates whether or not individual i is risky.

$$\tau_i = \begin{cases} 0 & \text{if } i \text{ is safe} \\ 1 & \text{if } i \text{ is risky} \end{cases}$$

Safe borrowers have income of $Y_i = Y$ for sure at $t = 1$, while risky borrowers have income of $Y_i = Y$ with chance $1 - \phi$ and $Y_i = 0$ with chance ϕ. The parameter

$$\phi \equiv \mathbf{P}[Y_i = 0 \mid \tau_i = 1]$$

measures the *income risk* of risky borrowers. Since income risk falls in expansions and increases in recessions[7], ϕ captures the state of the business cycle.

Lenders are risk neutral and maximize expected profits given an opportunity cost of funds of $1 + r < 1/\beta$. Because lenders are patient relative to borrowers, there are potential gains from trade.

Several frictions limit the scope for trade. First, limited liability and limited commitment implies that borrowers may default either by choice or if they have insufficient

[7] See Busch et al. (2016).

income to cover their debts. Second, borrower income realizations are private information, which makes state-contingent contracting impossible. Third, there is asymmetric information between borrowers and lenders about borrower types.

Lenders have access to a technology that reveals borrower types — which I will refer to as income verification. This technology requires borrowers to supply income documentation and requires lenders to pay a fixed cost. Denote the lender's cost of income verification by κ. I assume that providing income documentation imposes some infinitesimal cost on borrowers. This cost will keep risky borrowers from applying when they anticipate being rejected after their type is revealed.

If a lender does not verify a borrower's income, then they must form a belief based on freely available information (for example, a borrower's credit history and employment status). Let X_i be a vector of all variables that a lender can costlessly observe related to individual i's type. Denote the density of X_i conditional on borrower type by $f_X(x \mid \tau_i = \tau)$.

Lenders care about an individual borrower's likelihood of being the risky type. By Bayes' law, the share of risky borrowers among borrowers with $X_i = x$ is

$$\mathbf{P}[\tau_i = 1 \mid X_i = x] = \frac{f_X(x \mid \tau_i = 1)\mathbf{P}[\tau_i = 1]}{f_X(x \mid \tau_i = 1)\mathbf{P}[\tau_i = 1] + f_X(x \mid \tau_i = 0)\mathbf{P}[\tau_i = 0]} \equiv \rho(x).$$

The function ρ maps borrower observables to lender beliefs about borrower types. Instead of working directly with X_i, it will be convenient to work with the following *risk index*:

$$\rho_i \equiv \rho(X_i).$$

This index is the proportion of risky borrowers among borrowers with observables of X_i. Individual borrowers with a large value of ρ_i are perceived by lenders as likely to be the risky type.

Given that there are no state-contingent contracts, mortgage contracts specify period 0 and period 1 transfers. They also specify whether or not lenders require borrowers to provide income documentation.

Definition 3.1.1 *A* mortgage contract *is a triple,* (T, L, Doc)*, where*

1. *T is a period 0 transfer from the lender to the borrower.*

2. *L is a period 1 transfer from the borrower to the lender.*

3. *$Doc = 0$ if income documentation is not required, and $Doc = 1$ if it is required.*

I will refer to contracts with $Doc = 0$ as *low-documentation contracts* and contracts with $Doc = 1$ as *high-documentation contracts*.

Each borrower lives in a single unit of housing which they either own or rent. The purchase price of a house at $t = 0$ is P units of numeraire and the rental rate is R. The ex-post value of individual i's home (at $t = 1$) is idiosyncratic:

$$\text{Period 1 home value} = Z_i \overset{iid}{\sim} F.$$

Z_i is strictly positive, independent of borrower income, and continuously distributed with density F'.

Let $h(z) \equiv F'(z)/(1 - F(z))$ be the hazard rate of Z_i. I assume that the hazard rate satisfies the following monotonicity condition.

Assumption 3.1.1 (Monotone Scaled Hazard Rate) *$zh(z)$ is strictly increasing and continuous in z.*

This condition is a weak regularity condition that is common in the literature[8] and is satisfied by many standard distributions[9].

[8] Bernanke et al. (1999) assume the same condition on the idiosyncratic return on capital.

[9] For example, any distribution with a monotone hazard rate, like the log-normal distribution, satisfies this condition.

Due to limited liability and limited commitment, unsecured borrowing is impossible. However, a foreclosure technology exists which enables lenders to recover contracted debts through the liquidation of a borrower's home following default at the start of $t = 1$.

Foreclosure is costly[10]. When a lender initiates foreclosure, this reduces the sale value of the home to $(1-\lambda)Z_i$. The parameter λ captures losses from foreclosure[11]. The foreclosure technology then transfers funds from the home sale to the lender to cover any debt outstanding against the home. Given debt of L, the payoffs from foreclosure are

$$\text{Borrower Foreclosure Payoff} = \max\{(1-\lambda)Z_i - L, 0\}$$
$$\text{Lender Foreclosure Payoff} = \min\{L, (1-\lambda)Z_i\}.$$

Figure 3.1 depicts these payoffs. This technology provides an ex-post incentive for borrowers to re-pay contracted debts, allowing homes to serve as collateral for borrowing.

I make the following assumption to rule out contracts which imply certain default.

Assumption 3.1.2 (Certain Default is Inefficient) *The scaled hazard rate gets sufficiently large:*

$$\lim_{z\to\infty} zh(z) > \frac{1-(1+r)\beta}{(1-\phi)\lambda}.$$

This condition implies that arbitrarily large loans — which all borrowers will default on for sure — are inefficient. Borrowers will strategically default when underwater on their mortgage and as their loan size increases their probability of default increases. This condition states that at large loan sizes, the likelihood of strategic default increases too much to be worth increasing any borrowers loan size. Together with Assumption 3.1.1,

[10] As a result, there are possible gains from ex-post renegotiation. I assume that there are sufficiently large renegotiation costs that rule out ex-post renegotiation.

[11] For a discussion of foreclosure costs see Hayre and Saraf (2008). For empirical estimates of losses due to forced sales of homes see Campbell et al. (2011).

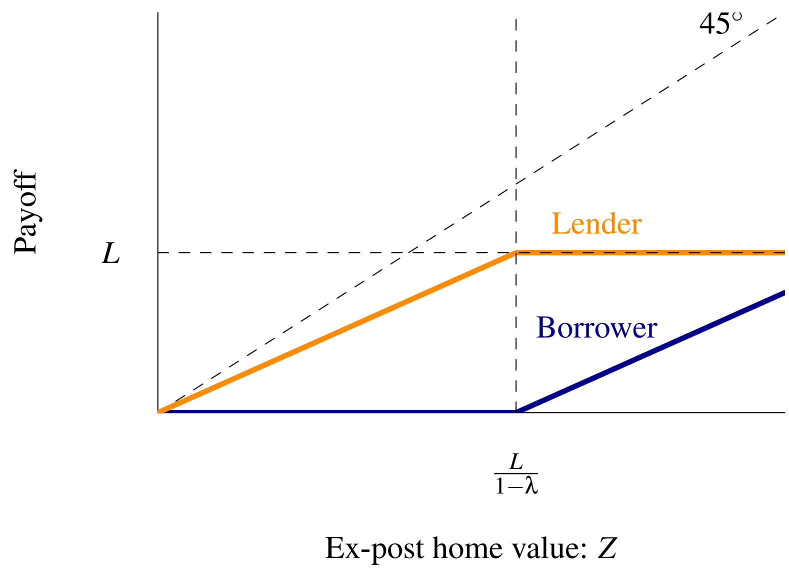

Figure 3.1: Foreclosure payoffs given debt level of L.

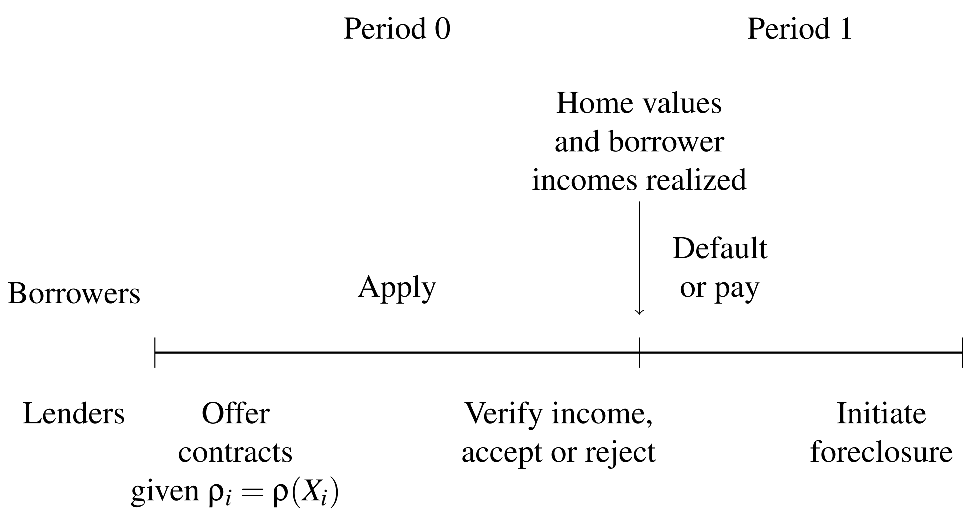

Figure 3.2: Within-period timing.

this condition ensures that equilibrium contract offers are finite and unique.

Timing within the two periods is as follows (see Figure 3.2). During period 1, lenders offer mortgages to individual borrowers. Then, borrowers apply to an individual lender. Finally, lenders perform income verification, and accept or reject applicants. At the end of $t = 0$, transfers occur and borrowers consume numeraire, which is not storable. Between period 0 and period 1, borrowers learn the value of their home and whether or not they have income. They then choose to make their mortgage payment, or not, and lenders initiate foreclosure. Finally, homes are liquidated and borrowers consume numeraire at the end of $t = 1$.

3.1.2 Equilibrium

Models of asymmetric information typically have a large number of equilibria. I focus on the unique efficient equilibrium. It is stable in the spirit of Kohlberg and Mertens (1986) and satisfies a number of selection criteria used in the literature on signaling games[12]. This selection is similar to the approach taken by Dell'Ariccia and Marquez (2006), who examine how equilibrium in a credit market changes as the composition of types in the market varies[13].

Up to the identity of the lenders who make contract offers and the application behavior of borrowers when indifferent, there is a unique stable Nash equilibrium.

Proposition 3.1.1 (Existence and Uniqueness) *Let Assumption 3.1.1 and let Assumption 3.1.2 hold. Suppose that risky borrowers do not apply for screening contracts whenever they anticipate rejection. Then for almost every* $\rho \in [0,1]$*, there exists a unique stable Nash equilibrium up to deletion of non-traded mortgage offers and the identity of the lender making each offer. In this equilibrium lenders make zero profits and safe borrowers only use contracts that maximize their value.*

Proof. See Appendix 3.A. ■

The assumption that risky borrowers do not apply for high-documentation contracts can be formalized by a small cost to supply income documentation. This selects

[12] See also Cho and Kreps (1987), Banks and Sobel (1987), Cho and Sobel (1990), and Ramey (1996)
[13] Also see Wilson (1977), Stiglitz and Weiss (1981), Bester (1985), and Hellwig (1987).

the equilibrium where risky borrowers do not apply for high-documentation contracts.

In Appendix 3.A I provide proofs and characterize equilibrium. The following sub-sections summarize equilibrium outcomes.

Period 1: Foreclosure, and Default

Since homes have strictly positive value, initiating foreclosure following a borrower default is a strictly dominant strategy. Borrowers then anticipate that foreclosure will always occur following any default.

Borrowers prefer to default if the foreclosure value of $\max\{(1-\lambda)Z_i - L, 0\}$ is greater than their equity of $Z_i - L$. This condition holds if any only if they have negative equity: $Z_i < L$. Underwater borrowers strategically default on their mortgage.

Even if they are willing to repay, they may be unable to pay. Borrowers are unable to make any payment which is larger than their potential income of Y. For simplicity, I assume that Y is large enough that repayment is feasible whenever income is available[14]. This assumption means that the model highlights the role of income risk in influencing mortgage contract terms and abstracts from the level of borrower income[15].

Given debt of L and ex-post home value of $Z_i = z$, the ex-post default rate among borrowers with risk index of ρ is

$$D(L, \rho\phi, z) = \mathbf{1}\{z < L\} + \mathbf{1}\{z \geq L\}\rho\phi. \tag{3.1}$$

If the borrower's home is underwater, then they default for sure. Otherwise, they only default if their income is unavailable, which occurs with chance of 0 for safe borrowers

[14] More generally, if a lender made a loan with $Y < L$, the loan will be defaulted on with probability 1. In the presence of foreclosure costs, these loans are inefficient and will not be made. Lenders will then impose a debt-to-income constraint: $L \leq Y$.

[15] See Greenwald (2016) for a model in which the level of borrower income and debt-to-income constraints are key determinants of monetary transmission and macroeconomic dynamics.

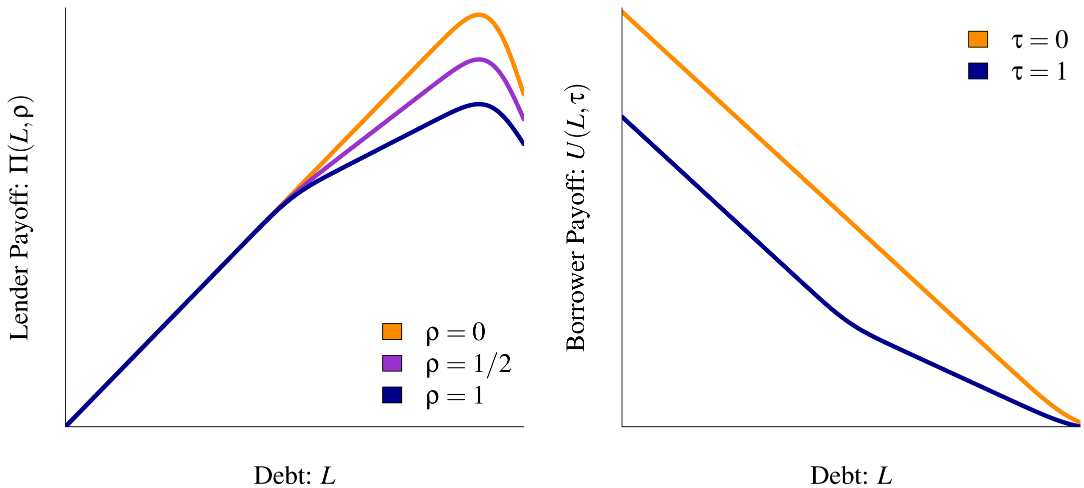

Figure 3.3: Anticipated contract payoffs varying ρ and τ when $\phi = 1/2$ and Z is log-normal.

and chance of ϕ for risky borrowers. Among borrowers with $\rho_i = \rho$, a fraction ρ are risky, and a fraction ϕ of these risky borrowers are forced into default.

Period 0: Anticipated Contract Payoffs

Lenders anticipate defaults when originating mortgages. Let $\Pi(L,\rho\phi)$ denote a lender's expected payoff on a loan of size L to a borrower with risk index of ρ given income risk of ϕ:

$$\begin{aligned}\Pi(L,\rho\phi) &= \mathbb{E}\left[D(L,\rho\phi,Z_i)\min\{(1-\lambda)Z_i,L\} + (1-D(L,\rho\phi,Z_i))L\right] \\ &= L - \int_0^\infty D(L,\rho\phi,z)\max\{L-(1-\lambda)z,0\}F'(z)\mathrm{d}z \end{aligned} \tag{3.2}$$

If default occurs, the lender gets their foreclosure payoff, and if the household doesn't default then the lender receives the face value of the debt. The second line expresses the payoff as the loan's face value net of anticipated losses from foreclosure. The left panel of Figure 3.3 plots Π for $\rho = 0, 1/2, 1$ when $\phi = 1/2$ and Z_i is log-normal.

Let $U(L,\tau\phi)$ denote the expected payoffs of a borrower with type τ given income risk of ϕ:

$$U(L,\tau\phi) = \mathbb{E}\left[D(L,\tau\phi,Z_i)\max\{(1-\lambda)Z_i - L, 0\} + (1-D(L,\tau\phi,Z_i))(Z_i - L)\right]$$
$$= \bar{z} - L + \int_0^\infty D(L,\tau\phi,z)\max\{-\lambda z, L-z\}F'(z)\mathrm{d}z. \tag{3.3}$$

where $\bar{z} = \int_0^\infty zF'(z)\mathrm{d}z$. If the borrower defaults, they get the payoff from foreclosure. If they do not default, they keep their home equity. The second line expresses this valuation in terms of expected home equity plus a term capturing the net benefit to the household of forced and strategic defaults. Borrowers benefit from default whenever they are underwater. Otherwise, default is costly. As a result, safe borrowers always benefit from their default option, while risky borrowers occasionally default on positive equity. The right panel of Figure 3.3 plots U for $\tau = 0, 1$ when $\phi = 1/2$ and Z_i is log-normal.

The following lemma establishes a single crossing condition which ensures that safe borrowers can signal their type by making a large down payment.

Lemma 3.1.1 (Single Crossing Condition) *For any debt level L such that* $\mathbf{P}[(1-\lambda)Z_i < L < Z_i] > 0$*, risky borrowers value a marginal increase in debt more than safe borrowers. In particular:*

$$\frac{\partial}{\partial L}\left[U(L,\phi) - U(L,0)\right] = \phi\left[F\left(\frac{L}{1-\lambda}\right) - F(L)\right].$$

Proof. See Appendix 3.A. ■

Although risky borrowers value mortgages less than safe borrowers (U is decreasing in its second argument), they value debt more *on the margin* than safe households. This occurs because of the difference in default behavior between the two types. Safe borrowers only default strategically. They never face foreclosure costs because they only default when they have negative equity, passing the costs onto their lender. Similarly, risky borrowers always default on negative equity, but will also default on positive equity with chance ϕ. Whenever $\mathbf{P}[(1-\lambda)Z_i < L < Z_i] > 0$, there is a chance that a risky borrower will default on positive equity but the cost of foreclosure destroys that equity by reducing the value of the home. In this case, risky borrowers do not internalize the cost of holding more debt, but safe borrowers do internalize this cost. As a result, risky borrowers value increases in debt more than safe borrowers.

Period 1: Income Verification, Rejections, and Applications

Lenders will reject any application which they anticipate a loss on. For instance, a high-documentation contract might be profitable only if the lender rejects all risky applicants. Due to zero profits, lenders will have an ex-post incentive to verify income and reject risky borrowers[16].

Anticipating lender income verification and rejection choices, borrowers will apply to whichever contract gives them the best payoff among all offers that are profitable for the lender. Given an infinitesimal application cost, risky borrowers will not apply for high-documentation contracts since they anticipate rejection.

Period 1: Offers

When making offers, lenders observe individual observables of X_i and infer that individual i is risky with chance $\rho_i = \rho(X_i)$. They then make an offer based on this belief.

To characterize equilibrium contracts, it is useful to first consider the full information case where lenders can observe borrower types. With full information, the safe

[16]Note that there exist unstable equilibria where lenders randomly reject additional borrowers. These equilibria arise because lenders make zero profits ex-post and therefore are indifferent between accepting and rejecting applications of borrowers that remain after rejecting borrowers who are unprofitable to keep in the pool. However, these equilibria are unstable because borrowers have a strong incentive to accept slightly worse terms in order to give lenders a small positive profit and avoid these arbitrary rejections.

borrower's most preferred contract solves the following program.

$$V^{safe} = \max_{T,L} \quad T + \beta U(L,0)$$
$$\text{s.t.} \quad -T + \frac{1}{1+r}\Pi(L,0) \geq 0 \tag{PC}$$

The constraint PC is the lender's zero profit (participation) constraint. Competition drives lender profits to zero so that this constraint binds. We can solve the problem by substituting in the participation constraint, finding the optimal loan size, and then backing out the initial transfer consistent with zero profits. Denote the loan size for this safe contract by L^{safe}.

The full information contract for risky borrowers gives them value of

$$V^{risky} = \max_{L} \quad \frac{1}{1+r}\Pi(L,\phi) + \beta U(L,\phi).$$

Here, I have used the zero profit condition for lenders to eliminate the period 0 transfer from the program.

Given these full information contracts, safe borrowers get more value from being a homeowner than risky borrowers.

Remark 3.1.1 *Safe borrowers value their full information contract more than risky borrowers value their full information contract:*

$$V^{safe} > V^{risky}.$$

Proof. Define $V(x) = \max_L \frac{1}{1+r}\Pi(L,x) + \beta U(L,x)$. We have $V(0) = V^{safe}$ and $V(\phi) = V^{risky}$. By the envelope theorem $V' < 0$ because Π and U are both decreasing in their second arguments. Therefore, $V(0) > V(\phi)$ and $V^{safe} > V^{risky}$. ■

Relative to risky borrowers, safe borrowers are natural homeowners. Even for relatively high home values, safe borrowers will prefer to purchase a house. In contrast, risky borrowers will only be willing to use their full information contract when home values are relatively low.

Now consider a safe borrower's most preferred low-documentation contract which signals their type through a high downpayment — the competitive high-downpayment contract. The outside option for a risky borrower is to either apply for their full information contract, or to rent. The value of their full information contract net of the home price is $-P+V^{risky}$, and the value of renting is $-R$. Risky borrowers prefer their outside option to a low-documentation contract $(T, L, 0)$ whenever the following incentive compatibility constraint holds.

$$T - P + \beta U(L, \phi) \leq \max\{-R, -P + V^{risky}\}. \tag{IC}$$

The competitive high-downpayment contract maximizes value for safe borrowers subject to this incentive constraint.

$$\begin{aligned} V^{down}(P-R) = \max_{L} \quad & \frac{1}{1+r}\Pi(L,0) + \beta U(L,0) \\ \text{s.t.} \quad & \frac{1}{1+r}\Pi(L,0) + \beta U(L,\phi) \leq \max\{P-R, V^{risky}\} \qquad \text{(IC)} \end{aligned}$$

Here, I used the zero profit condition to eliminate T, and simplified the IC constraint by adding P to both sides. I call the quantity $P-R$ the *net home price*. It captures how the relative cost of owning versus renting influences this contract.

Note that the value of the high-downpayment contract does not depend on borrower observables — in a separating equilibrium, risky borrowers and safe borrowers reveal their types and so a lender's initial information is irrelevant. However, this value may depend on the net home price, $P-R$. Changes in home prices relative to rental rates change the incentive for risky borrowers to enter the mortgage market. When home prices are high, risky borrowers find owning less attractive, and it takes a smaller downpayment for safe borrowers to signal their type (high $P-R$ relaxes the IC constraint). As a result, the value of separation for safe borrowers will vary with the net home price. Once we consider joint housing and mortgage market equilibrium in Section 3.2, this link will lead to general equilibrium feedback and amplification of shocks to income risk.

When is the IC constraint slack? When risky borrowers do not wish to use the safe borrower's full information contract. If a risky borrower were to gain access to this contract, they'd get value of

$$V^* = \frac{1}{1+r}\Pi(L^{safe},0) + \beta U(L^{safe},\phi).$$

Whenever the net home price is above this threshold, risky borrowers prefer to rent over using the loan intended for safe borrowers. As a result, the IC constraint will not bind if $P-R \geq V^*$. In this case, there is no issue of adverse selection. Safe borrowers use their full information contract and become homeowners, risky borrowers become renters, and types are revealed in equilibrium.

But for lower home prices — when $P-R < V^*$ — risky borrowers will enter the market and introduce adverse selection (the IC constraint will bind). In this case, safe borrowers may prefer the allocation under income-verification based screening or under pooling.

Consider income verification. Since only safe borrowers apply, the value of a high-documentation contract is constant across safe borrowers:

$$V^{doc} = \max_L \quad \frac{1}{1+r}\Pi(L,0) + \beta U(L,0) - \kappa = V^{safe} - \kappa.$$

They get the value of their full information contract net of the income verification cost. Note that, like with downpayment-based screening, the value of screening through income verification is constant across borrowers. Income verification removes all risky borrowers from the pool of funded loans. The chance that a risky borrower has no income is irrelevant for the terms a safe borrower receives on a high-documentation contract.

Define the overall value of screening as

$$V^{screen}(P-R) = \max\{V^{down}(P-R), V^{doc}\}.$$

Lenders will use whichever method of screening is most preferred by safe borrowers. If they didn't another lender could offer a contract using the alternative form of screening and capture the market. The implication is that whether or not the value of screening changes with the net home price depends on whether or not lenders are actively using income verification or high downpayment requirements to ration credit. When they use income verification, the value of screening does not depend on home prices. At high enough home prices, lenders will substitute away from income verification toward downpayment requirements and the value of screening will vary with the net home price.

Finally, consider pooling. In a pooling equilibrium, lenders offer a single contract and accept all applicants. Due to zero profits, the value of the mortgage to safe borrowers depends on how many risky borrowers there are in the pool. In the group of borrowers with risk index of ρ, a fraction ρ are risky and a fraction $1-\rho$ are safe. The value of the contract to safe borrowers will depend on how many of these risky borrowers apply for the contract.

In equilibrium, risky borrowers will always apply whenever they are offered a contract that leads to pooling. This outcome comes from the fact that whenever risky borrowers are willing to opt out of a pooling contract, lenders have an incentive to offer the high-downpayment contract. Intuitively, a contract which fails to attract risky borrowers away from renting effectively acts like a separating contract. Any contract other than the high-downpayment contract is relatively inefficient at inducing separation.

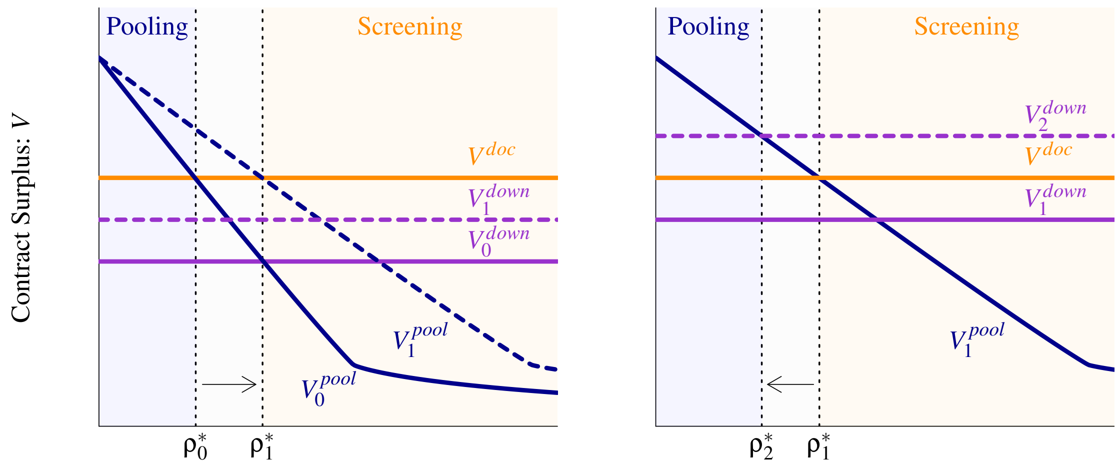

Figure 3.4: Comparative statics of market segmentation: fall in ϕ (left) followed by increase in $P-R$ (right).

Because of this result, we can focus on a pooling equilibrium in which all risky borrowers apply. In this case, the lender zero profit condition implies a transfer of $T = \frac{1}{1+r}\Pi(L,\rho\phi)$. A safe borrower with risk-index of ρ then prefers a low-documentation contract that solves

$$V^{pool}(\rho) = \max_{L} \quad \frac{1}{1+r}\Pi(L,\rho\phi) + \beta U(L,0).$$

Note that the value of pooling across ρ-values will change as income risk varies. In particular, a fall in income risk will increase the value of pooling across all borrowers:

$$\frac{\partial}{\partial\phi}V^{pool}(\rho) = \frac{1}{1+r}\Pi_2(L^{pool}(\rho),\rho\phi)\rho = -\rho\int_{L^{pool}(\rho)}^{\frac{L^{pool}(\rho)}{1-\lambda}}(L^{pool}(\rho)-(1-\lambda)z)F'(z)\mathrm{d}z < 0.$$

Intuitively, safe borrowers face a higher interest rate when they are pooled with risky borrowers, but as risky borrowers become less likely to default, loan terms improve. As a result, the pooling value function for safe borrowers will increase when the income risk of risky borrowers falls.

In equilibrium, lenders will make offers that maximize value for safe borrowers. We can determine the pattern of segmentation by comparing the pooling and screening value functions. The following proposition characterizes the pattern of market segmentation.

Proposition 3.1.2 (Equilibrium Segmentation) *Let Assumption 3.1.1 and let Assumption 3.1.2 hold. When $P-R \geq V^*$ there is no segmentation. Lenders offer the full information safe contract to all borrowers. Safe borrowers apply and are accepted to this contract. Risky borrowers do not apply and instead rent their home. When $P-R < V^*$, the market is divided into a pooling segment and a screening segment. Let ρ^* satisfy $V^{screen}(P-R) = V^{pool}(\rho^*)$. Borrowers with $\rho < \rho^*$ are in a pooling equilibrium and those with $\rho > \rho^*$ are in a screening equilibrium. Lenders use downpayment requirements to screen borrowers when $V^{doc} < V^{down}(P-R)$ and use income verification when $V^{doc} > V^{down}(P-R)$.*

Proof. See Appendix 3.A. ■

Lenders offer contracts that maximize value for safe borrowers, as otherwise another a competing lender could attract away safe borrowers. As a result, lenders offer those borrowers with $\rho_i = \rho$ whichever contract gives the most value to safe borrowers and the pattern of equilibrium segmentation corresponds to the upper envelope of the pooling, high-documentation, and high-downpayment value functions.

Note that the screening segment consists of the borrowers who are likely to be risky from the perspective of lenders. The value of screening is constant in ρ, while the value of pooling is decreasing in ρ. In high risk parts of the market, safe borrowers face high interest rates when using pooling contracts, and find it worthwhile to differentiate themselves by either undergoing income verification or making a large downpayment.

There are two key comparative statics for the pattern of market segmentation which lead to the key general equilibrium results in Section 3.2. First, as income risk in the economy falls (increases), the pooling segment expands (contracts). The boundary ρ-value between the pooling and screening segments, ρ^*, quantifies the size of the pooling segment. Formally, if $P-R<V^*$ so that $\rho^*>0$ then the pooling segment is decreasing in income risk:

$$\frac{\partial\rho^*}{\partial\phi}<0.$$

Intuitively, as risky borrowers become more likely to default safe borrowers are less willing to be pooled and a smaller fraction of the market is in a pooling equilibrium.

Second, so long as lenders use downpayment requirement to ration credit, then the pooling segment is decreasing in the net home price. If lenders instead use income verification, then the size of the pooling segment is constant in $P-R$. Formally, if $P-R<V^*$ and $V^{down}(P-R)>V^{doc}$ then the pooling segment is decreasing in the net home price:

$$\frac{\partial\rho^*}{\partial(P-R)}<0,$$

while if $V^{down}(P-R)<V^{doc}$ then it is constant:

$$\frac{\partial\rho^*}{\partial(P-R)}=0.$$

Figure 3.4 depicts these comparative statics. The left diagram shows a fall in income risk which expands the pooling segment. The right diagram then shows an increase in the net home price that offsets the expansion by reducing the size of the pooling segment. Note that a smaller increase in the net home price — which didn't shift the value of downpayment-based screening above the value of income-verification based screening — would not reduce the size of the pooling segment.

Note that the homeownership rate in the economy is the number of borrowers who gain access to credit and purchase a home. Since credit access varies with the net home price through the incentive constraint channel, the comparative static for the net home price implies a housing demand curve and the shape of this housing demand curve will depend on the nature of screening — whether lenders use downpayment requirements or income verification to ration credit.

Also, for a given home price, changes in the mortgage market due to variation in income risk will change credit access to risky borrowers and influence homeownership. For example, if income risk falls, then a portion of the market that previously was in a screening equilibrium shifts to pooling, this expands credit access to risky borrowers, and increases housing demand.

Together, these two observations suggest an important interaction between mortgage markets and housing markets. An analysis of joint housing and mortgage market equilibrium is necessary to fully characterize how the share of non-prime loans varies with economic fundamentals.

3.2 Joint Mortgage and Housing Market Equilibrium

This section examines joint housing market and mortgage market equilibrium. I show that general equilibrium feedback between mortgage and housing markets leads to a strong form of state-dependent amplification. In particular, small shocks to income risk can sometimes lead to sudden and large changes in equilibrium outcomes. We can interpret this result as coming from endogenous switching between two distinct credit regimes: a screening regime where the marginal home buyer is credit rationed, and a pooling regime where they instead are pooled with safe borrowers. The key result is that, during a transition from screening to pooling, home prices boom and lenders both rapidly relax down-payment requirements and stop verifying borrower incomes.

3.2.1 Housing Demand

Housing demand equals the number of borrowers in the economy who choose to enter the mortgage market and become homeowners. We can trace out a housing demand curve by considering how equilibrium in the mortgage market changes as home prices vary.

First, consider when home prices are sufficiently high so that all risky borrowers prefer to rent — which occurs when $P-R > V^{*}$. In this case, lenders are able to offer safe borrowers their full information contract without worrying about risky borrowers applying. So long as the home price is not too high, safe borrowers will become homeowners. They are indifferent between owning and renting when $P-R = V^{safe}$. They prefer to rent when $P-R > V^{safe}$ and prefer to own when $P-R < V^{safe}$.

Once $P-R \leq V^{*}$, risky borrowers begin to enter the mortgage market. Their entry leads to adverse selection and the introduction of rationing. The market is split into a pooling segment, where risky borrowers gain access to credit, and a screening segment, where they are rationed out of the prime market.

First consider when the home price is high enough that high-documentation contracts are not introduced. By Proposition 3.1.2, this outcome occurs when $V^{doc} < V^{down}(P-R)$. In this case, lenders use downpayment requirements to ration credit in the screening segment of the market. As the net home price falls, the size of the pooling segment increases, more risky borrowers become homeowners, and we have a downward sloping housing demand curve. Note that in this case — which I refer to as the pooling regime — the marginal home buyer is a risky borrower who has access to a pooling contract.

Once home prices are low enough, the market enters the alternative screening regime. This occurs once lenders begin requiring income documentation and verifying borrower incomes — once $V^{doc} > V^{down}(P-R)$. Now, further reductions in the net

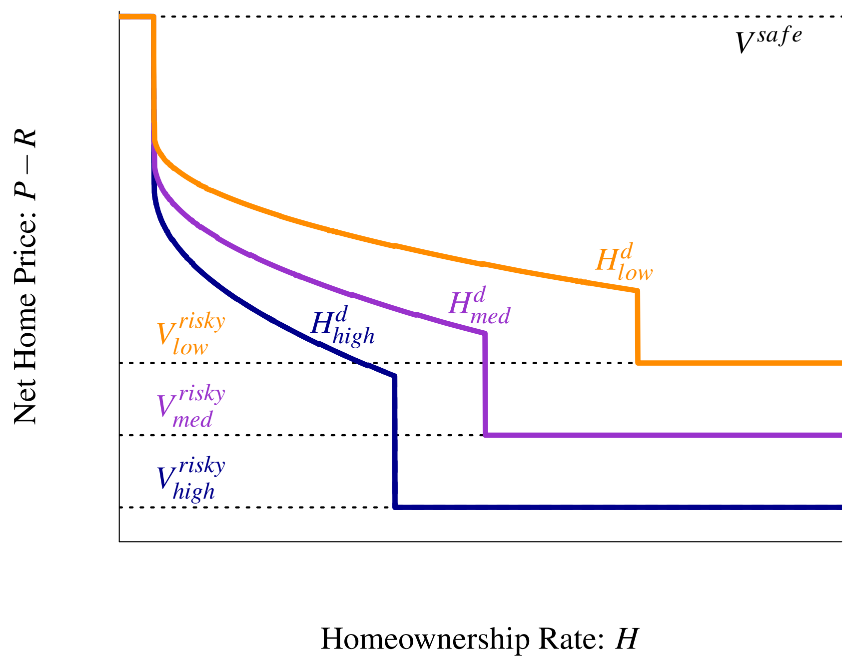

Figure 3.5: Housing demand as income risk varies. The orange corresponds to low income risk, purple to medium income risk, and blue to high income risk.

home price no longer increase credit access to risky borrowers and the size of the pooling segment does not change with the net home price.

As a result, the homeownership rate becomes inelastic. Expansion in homeownership requires more risky borrowers to become homeowners, but the introduction of income verification rations the marginal borrower out of the market. Additional risky borrowers will become homeowners only if they find it worthwhile to use their full information contract, which occurs once $P-R \leq V^{risky}$.

Figure 3.5 shows this housing demand curve as income risk varies. Housing demand is 0 when $P-R$ is above V^{safe}. Once $P-R < V^{safe}$ all safe borrowers become homeowners. Then, if $P-R < V^{*}$, the mortgage market enters the pooling regime and housing demand increases as $P-R$ falls. For a sufficiently low home price, the mortgage market enters the screening regime, and housing demand becomes is inelastic. Once

$P - R = V^{risky}$, risky borrowers are willing to use their full information contract to buy a home, and if $P - R < V^{risky}$ then all risky borrowers become homeowners. The shape of the housing demand curve summarizes how adverse selection and changes in credit rationing impact homeownership.

This demand curve shifts with income risk. A fall in income risk increases V^{risky} (because risky borrowers value homes more when they are less likely to default), and also makes pooling relatively attractive for safe borrowers. The expansion in the pooling segment of the mortgage market increases credit access to risky borrowers, which shifts the inelastic portion of the housing demand curve rightward.

3.2.2 General Equilibrium

I now consider joint housing market and mortgage market equilibrium. For simplicity, I focus on the case of a fixed stock of housing. The intuition developed here will carry forward once I consider dynamics and the accumulation of housing in Section 3.3.

With an inelastic supply of housing, if a fall in income risk expands the pooling segment of the mortgage market enough (if the inelastic portion of housing demand shifts far enough to the right), then it will trigger a rapid increase in home prices, a full shift away from income verification (a switch to the pooling regime), and a rapid relaxation of down-payment requirements.

Let H^s denote the stock of homes. As income risk falls, housing demand increases, as in Figure 3.5. If the level of housing demand attributable to risky borrowers gaining access to pooling contracts remains below H^s, then for the housing market to clear some risky borrowers must be willing to use their full information contract and become homeowners. Their participation in the market requires that $P - R \leq V^{risky}$. All risky borrowers would want to become homeowners if $P - R < V^{risky}$ and the market would

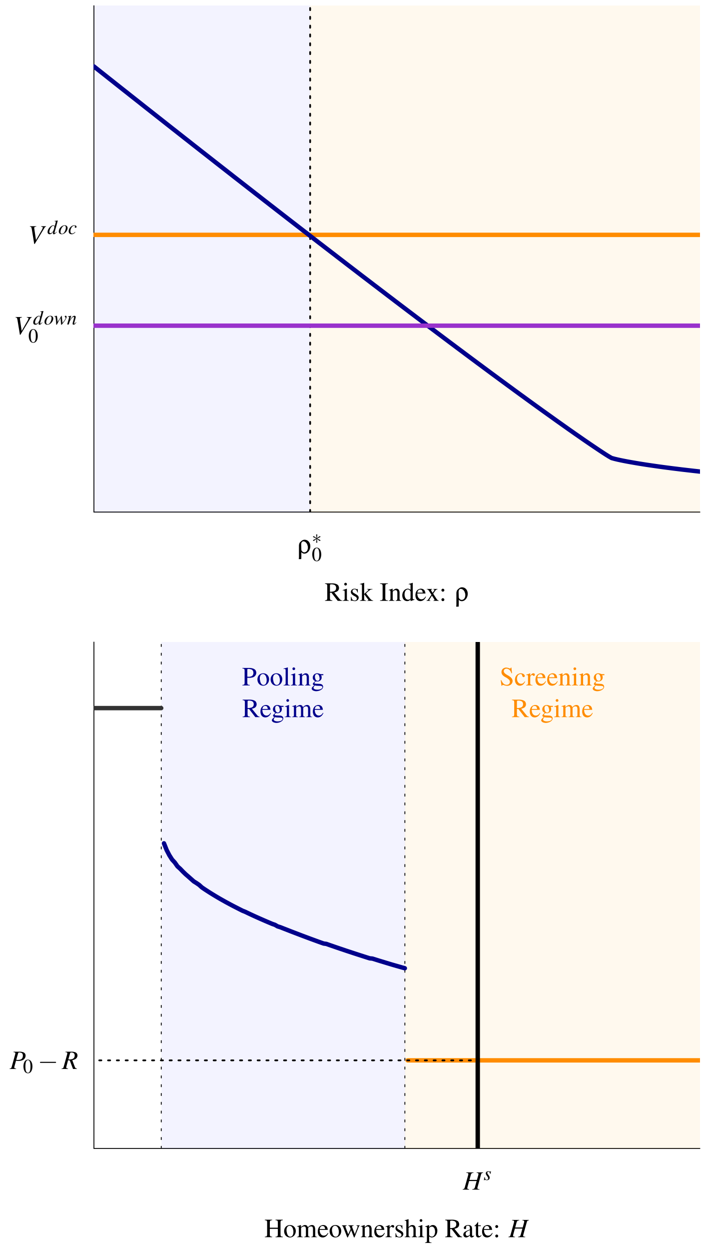

Figure 3.6: Joint equilibrium in the mortgage market (top) and housing market (bottom) when the screening regime occurs (high income risk).

not clear. We must have $P - R = V^{risky}$. The marginal homeowner is a risky borrower using a high interest rate loan who is indifferent between owning and renting. Note that in this case, lenders must use income verification to ration credit access, so the screening regime prevails.

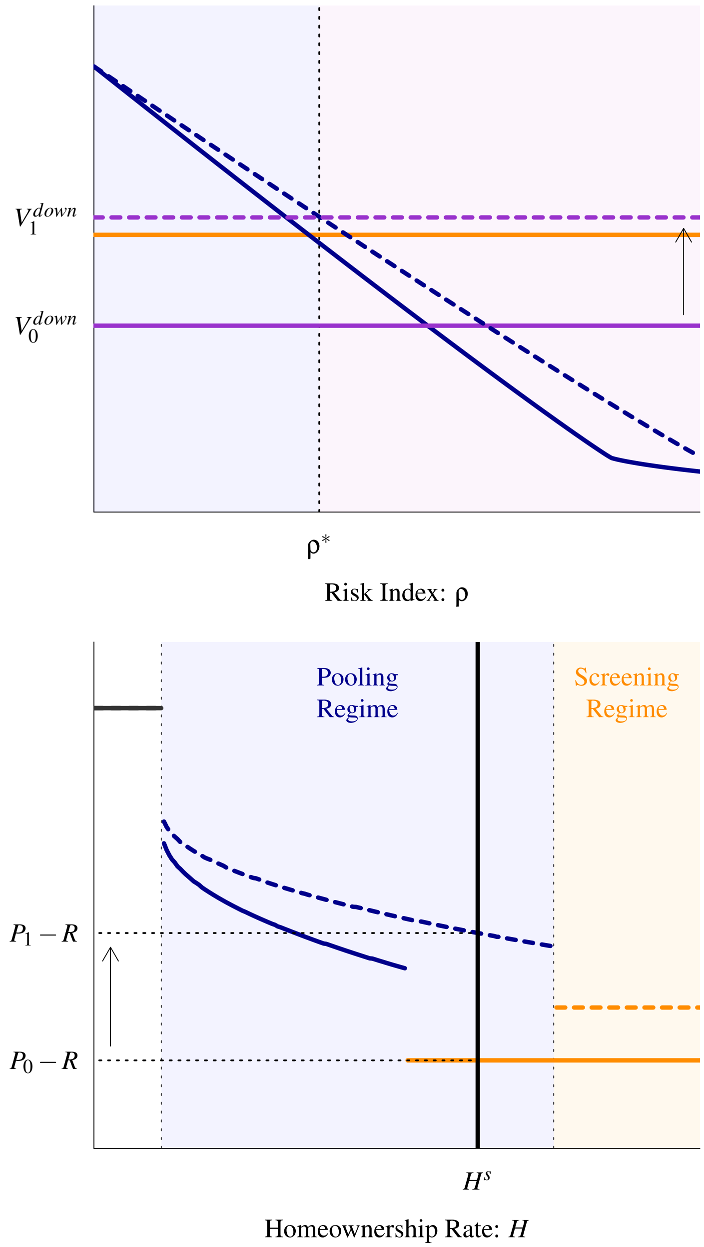

Figure 3.7: Joint equilibrium in mortgage market (top) and housing market (bottom) during a shift from the screening regime to the pooling regime (fall in income risk).

Figure 3.6 depicts this outcome. The top panel corresponds to the mortgage market (as in Figure 3.4), depicting the pattern of segmentation across risk indices. The bottom panel corresponds to the housing market, with a demand curve as in Figure 3.5 and an inelastic housing supply.

Now consider a fall in income risk and suppose that the inelastic portion of the housing demand curve surpasses H^s. This increase beyond H^s represents risky borrowers who previously rented gaining access to pooling loans. Because these risky borrowers are pooled with safe borrowers, they face a relatively low interest rate, get more value from the pooling contract than they get from their full information contract, and therefore strictly prefer to enter the housing market. Given the fixed housing stock, this increase in demand cannot be met. Since all market participants value homeownership more than the prevailing home price of $P - R = V^{risky}$, borrowers begin to bid up home prices.

As home prices rise, homeownership becomes less attractive relative to renting, which reduces downpayment requirements (the IC constraint relaxes). Figure 3.7 depicts this outcome. The net home price rises from $P_0 - R$ to $P_1 - R$, relaxing the IC constraint, and reducing downpayment requirements. This improvement in loan terms increases the value of the downpayment contract for safe borrowers from V_0^{down} to V_1^{down}, and shifts the mortgage market toward using downpayment requirements to ration credit instead of using income verification.

It is through the switch to downpayment-based credit rationing and the incentive constraint channel that we get housing market clearing. Home prices must increase enough to generate a shift away from income-verification based screening toward downpayment-based screening. With the incentive constraint channel active, further increases in the net home price will reduce the size of the pooling segment, reducing credit access to risky borrowers, and dampening housing demand.

Before the shock to income risk, the marginal home buyer was a risky borrower who was screened out of the market and was indifferent between owning and renting. Now, that same marginal home buyer (because the homeownership rate has not changed) instead has access to a pooling contract. The economy has switched from the original screening regime to the pooling regime.

Note that outcomes change discontinuously during the regime change. Joint housing and mortgage market equilibrium implies that a small fall in income risk can generate a sudden collapse in income documentation requirements, a rapid appreciation in home prices, and a relaxation of down-payment requirements.

We can interpret the rise in the home price as coming from a switch in the indifference condition pricing homes. In the screening regime, risky borrowers had to be indifferent between owning and renting so that they would be willing to purchase homes. But in the pooling regime, risky borrowers gain access to low interest rate loans and now strictly prefer to become homeowners. Instead of through their indifference condition, homes are now priced implicitly through the indifference condition of those safe borrowers who are on the margin between using a pooling contract and a high-downpayment contract. It is their indifference between contract types which determines the extent of credit access to risky borrowers and the level of housing demand. In effect, homes are priced implicitly through the IC constraint.

This result has two interpretations which tie to previous results in the literature on the origins of the housing boom. Justiniano et al. (2015b) show that shocks to an exogenous limit on credit supply can generate a housing boom and bust. In this model, the screening regime captures such a credit supply constraint, and the shift to the pooling regime is an endogenous relaxation of this constraint. While income risk is high, adverse selection in the mortgage market leads to credit rationing through income documentation requirements. When the credit regime switches from screening to pooling, it is as-if a credit limit ceases to bind.

The results can also be interpreted as endogenizing the shock to loan-to-value ratios that Kaplan et al. (2016) find are necessary to explain the boom and bust. This interpretation comes from how the regime shift changes the role of the IC constraint. This constraint can be interpreted as a loan-to-value constraint because it specifies (given prevailing home prices), how large of a down payment a safe borrower must make to signal their type. When home prices rise, this constraint relaxes, allowing safe borrowers who use high-downpayment contracts to increase their leverage as home prices appreciate. This mechanism generates a shift in loan-to-value ratios as-if there were a shock to a loan-to-value constraint.

Note that the mechanism works both to generate booms and busts. Starting from the pooling regime, an increase in income risk can lead to a tightening of lending standards and a collapse in home prices. This result implies that mean reversion in income risk can explain a bust following a housing boom. The next section extends the model to a dynamic setting and formalizes this intuition.

3.3 Housing Boom-Bust Cycles

To examine the dynamics of home prices, I extend the model to an infinite horizon setting and introduce simple house production and rental unit sectors. The resulting rational expectations equilibrium is characterized by housing cycles due to shifts in the credit regime. These cycles are driven by small underlying shocks to borrower income risk which get amplified through the general equilibrium mechanism from Section 3.2.

The non-linear mechanism at the heart of this theory generates boom-bust cycles. When there is a switch to pooling followed by a return to screening, it is as-if a large positive credit supply shock hits the economy and then a large negative credit supply shock follows up to generate a housing crisis. Rather than relying on a series of negatively correlated shocks, the theory generates the entire episode through joint housing and mortgage market clearing. As income risk falls, the economy reaches the threshold triggering the pooling regime and an endogenous credit-boom occurs. Then, as income risk smoothly returns to steady state, the economy reverts to the screening regime. Credit contracts and a housing crisis occurs. In this model, a housing crisis is the ultimate unwinding of the initial credit-fueled boom.

3.3.1 Dynamic Extension

Time is indexed by $t \in \{0, 1, \dots\}$. I re-interpret period 0 of the previous static model as the end of each time period t, and period 1 as the beginning of $t+1$. All notation carries forward, but with a time subscript t appended.

For instance, I denote income risk by ϕ_t, which I assume evolves as a second order autoregressive process in logs:

$$\ln \phi_t = (1 - \rho_1^\phi - \rho_2^\phi)\mu^\phi + \rho_1^\phi \ln \phi_{t-1} + \rho_2^\phi \ln \phi_{t-2} + \sigma^\phi \varepsilon_t^\phi.$$

This setup generates momentum, which will lead to persistent rather than transitory shifts in the credit regime. I view this specification for income risk as a reduced form for a model of the labor market in which individuals have a higher risk of losing income during an economic downturn[17]. For expositional simplicity, I take income risk as exogenous. This exogeneity assumption shuts down potential general equilibrium feedback from home prices to the state of the business cycle which may be important for understanding the housing boom and bust (Mian and Sufi (2011) Berger et al. (2015)). However, abstracting from this additional source of amplification helps to clarify the mechanism introduced in this paper.

I assume that the resale value of a home purchased in $t-1$ is proportional to the equilibrium price at time t, so that $P_t Z_{it}$ is the value of individual i's home[18]. The value

[17] Appendix 3.C presents a model of labor market search where fundamental shocks to firm productivity generate fluctuations in employment risk and therefore generate fluctuations in income risk.

[18] This assumption can be formalized by a setup where homeowners have idiosyncratic home maintenance and investment opportunities. The amount of effective housing that they own after maintenance and investment is Z_{it}.

of the home in foreclosure is then $P_t(1-\lambda)Z_{it}$.

Given this proportionality, anticipated contract payoffs in $t+1$ for loan of size L_{it} to borrower i at time t are

$$\mathbb{E}_t\Pi\left(\frac{L_{it}}{P_{t+1}},\rho_{it}\phi_{t+1}\right)P_{t+1}, \quad \text{and} \quad \mathbb{E}_tU\left(\frac{L_{it}}{P_{t+1}},\tau_i\phi_{t+1}\right)P_{t+1}.$$

Ex-post payoffs are homogenous of degree one in the realized home price and the face value of debt, and so we can re-use the functions Π and U after normalizing by the next period home price.

I assume that borrower observables take three values: $\rho_{it} \in \{0, \rho_L, \rho_H\}$ where $0 < \rho_L < \rho_H < 1$. Borrowers with $\rho_{it} = 0$ are *known safe borrowers*. In this segment of the market, there is no problem of adverse selection. All borrowers will receive the full information contract intended for safe borrowers and so they face neither an income documentation requirement or a high downpayment requirement. We can interpret this segment of the market as representing the traditional Alt-A segment of the mortgage market — low documentation loans to borrowers with very high credit scores.

In contrast, adverse selection bites in the other two portions of of the market. Borrowers with $\rho_{it} = \rho_L$ are *low risk borrowers* while those with $\rho_{it} = \rho_H$ are *high risk borrowers*. These two portions of the market both represent borrowers with less-than-perfect credit. A portion ρ_L of low risk borrowers are risky and a portion ρ_H of high risk borrowers are risky. Let n_L denote the fraction of the population with $\rho_{it} = \rho_L$ and n_H denote the fraction with $\rho_{it} = \rho_H$. The remaining $1-n_L-n_H$ borrowers are known safe borrowers with $\rho_{it} = 0$.

Because I no longer assume a continuous distribution for ρ_{it}, equilibrium now requires that lenders occasionally use mixed strategies. In particular, in the parameterization I will consider, it is the middle segment of the market which will determine the prevailing credit regime. The segment of known safe borrowers always receives their full information contract and the high risk segment will always be in a screening equilibrium. The middle segment will either be in a screening equilibrium where all risky borrowers have their type revealed, or lenders will randomly mix between offering pooling contracts and high-downpayment contracts. Those risky borrowers who are offered the pooling contract will gain access to credit, while those risky borrowers who are offered the high-downpayment contract will opt to rent their home.

The homeownership rate in the pooling regime equals the number of risky borrowers who gain access to pooling contracts (since they strictly prefer pooling), plus the number of safe borrowers (who all become homeowners). Housing market clearing then pins down the fraction of risky borrowers in segment $\rho_{it} = \rho_L$ who are offered pooling contracts. Safe borrowers in this group are happy to receive a random offer so long as they are indifferent between pooling and high-downpayment contracts — which is precisely the indifference condition pricing housing in the pooling regime.

Let H_t denote the homeownership rate and let μ_t denote the probability that a lender offers the pooling contract to low risk borrowers in the pooling regime. Housing market clearing requires that

$$H_t = \underbrace{(1-n_L-n_H)}_{\text{Known Safe}} + \underbrace{n_L(1-\rho_L)}_{\text{Safe Among Low Risk}} + \underbrace{n_L\rho_L}_{\text{Risky Among Low Risk}} \mu_t + \underbrace{n_H(1-\rho_H)}_{\text{Safe Among High Risk}} .$$

The equilibrium homeownership rate must equal the demand for housing. All safe borrowers own a home — captured by the first, second, and fourth terms on the right hand side. The third term is the number of risky borrowers who become homeowners. These borrowers only come from the low risk segment because the high risk segment is in a screening equilibrium. A fraction μ_t of risky borrowers in the low risk segment get offered a pooling contract.

Solving for μ_t gives an expression for the chance that a low risk borrower is offered a pooling contract.

$$\mu_t = \frac{H_t - (1 - n_L\rho_L - n_H\rho_H)}{n_L\rho_L}.$$

The numerator is the portion of homeowners who are risky borrowers. The number of safe borrowers in the whole population is $1 - n_L\rho_L - n_H\rho_H$ and any housing that is not purchased by a safe borrower must be purchased by a risky borrower. The denominator is the portion of borrowers in the low risk segment who are risky.

Together, mortgages that neither require income documentation nor a high down-payment, and mortgages intended for risky borrowers are *non-prime*. The non-prime share is

$$\text{Non-Prime Share}_t = \begin{cases} \frac{(1-n_L-n_H)+H_t-(1-\rho_L n_L-\rho_H n_H)}{H_t} & \text{if } S_t = \text{Screening} \\ \frac{(1-n_L-n_H)+n_L\mu_t}{H_t} & \text{if } S_t = \text{Pooling}. \end{cases}$$

The non-prime share suddenly increased between 2003 and 2004, and then collapsed in the housing crisis. In the screening regime, non-prime lending consists of mortgages to known safe borrowers and mortgages to risky borrowers. In the pooling regime, it consists of mortgages to known safe borrowers and pooling contracts in the low risk segment.

The credit regime depends on which kind of contract safe borrowers in the low risk segment prefer. The high-documentation contract has a loan size satisfying

$$0 = \frac{1}{1+r}\mathbb{E}_t\Pi_1\left(\frac{L_t^{doc}}{P_{t+1}}, 0\right) + \beta\mathbb{E}_t U_1\left(\frac{L_t^{doc}}{P_{t+1}}, 0\right).$$

The value of this contract to a safe borrower in the low risk segment is

$$V_t^{doc} = \frac{1}{1+r}\mathbb{E}_t\Pi\left(\frac{L_t^{doc}}{P_{t+1}},0\right)P_{t+1} + \beta\mathbb{E}_t U\left(\frac{L_t^{doc}}{P_{t+1}},0\right)P_{t+1} - \kappa.$$

The pooling contract in the low risk segment has a loan size satisfying

$$0 = \frac{1}{1+r}\mathbb{E}_t\Pi_1\left(\frac{L_t^{pool}}{P_{t+1}},\rho_L\phi_{t+1}\right) + \beta\mathbb{E}_t U_1\left(\frac{L_t^{pool}}{P_{t+1}},0\right).$$

Safe borrowers get value of

$$V_t^{pool} = \frac{1}{1+r}\mathbb{E}_t\Pi\left(\frac{L_t^{pool}}{P_{t+1}},\rho_L\phi_{t+1}\right)P_{t+1} + \beta\mathbb{E}_t U\left(\frac{L_t^{pool}}{P_{t+1}},0\right)P_{t+1}$$

from this contract.

The prevailing credit regime, indicated by S_t, is determined by these two values.

$$S_t = \begin{cases} \text{Screening} & \text{if } V_t^{doc} \geq V_t^{pool} \\ \text{Pooling} & \text{if } V_t^{doc} < V_t^{pool} \end{cases}$$

If $V_t^{doc} \geq V_t^{pool}$ then safe borrowers in the low risk segment prefer screening, and otherwise they prefer pooling.

In the screening regime, risky borrowers are indifferent between their full information contract and renting. The first order condition determining the loan size of their full information contract, L_t^{risky}, is

$$0 = \frac{1}{1+r}\mathbb{E}_t\Pi_1\left(\frac{L_t^{risky}}{P_{t+1}},\phi_{t+1}\right) + \beta\mathbb{E}_t U_1\left(\frac{L_t^{risky}}{P_{t+1}},\phi_{t+1}\right).$$

They get value of

$$V_t^{risky} = \frac{1}{1+r}\mathbb{E}_t\Pi\left(\frac{L_t^{risky}}{P_{t+1}},\phi_{t+1}\right)P_{t+1} + \beta\mathbb{E}_t U\left(\frac{L_t^{risky}}{P_{t+1}},\phi_{t+1}\right)P_{t+1}$$

from this contract. The pricing condition in the screening regime is then

$$P_t = R_t + V_t^{risky} \quad \text{if} \quad S_t = \text{Screening}.$$

In the pooling regime, safe borrower indifference between pooling and high-downpayment contracts pins down the home price. The binding IC constraint implies that the high-downpayment contract's loan size must satisfy

$$\frac{1}{1+r}\mathbb{E}_t\Pi\left(\frac{L_t^{down}}{P_{t+1}},0\right)P_{t+1} + \beta\mathbb{E}_t U\left(\frac{L_t^{down}}{P_{t+1}},\phi_{t+1}\right)P_{t+1} = P_t - R_t$$

giving safe borrowers value of

$$V_t^{down} = \frac{1}{1+r}\mathbb{E}_t\Pi\left(\frac{L_t^{down}}{P_{t+1}},0\right)P_{t+1} + \beta\mathbb{E}_t U\left(\frac{L_t^{down}}{P_{t+1}},0\right)P_{t+1}.$$

Their indifference condition pricing the home is then

$$V_t^{pool} = V_t^{down} \quad \text{if} \quad S_t = \text{Pooling}.$$

The last contract we need is the full information safe contract. It has loan size satisfying

$$0 = \frac{1}{1+r}\mathbb{E}_t\Pi_1\left(\frac{L_t^{safe}}{P_{t+1}},0\right) + \beta\mathbb{E}_t U_1\left(\frac{L_t^{safe}}{P_{t+1}},0\right).$$

Safe borrowers get value of

$$V_t^{safe} = \frac{1}{1+r}\mathbb{E}_t\Pi\left(\frac{L_t^{safe}}{P_{t+1}},0\right)P_{t+1} + \beta\mathbb{E}_t U\left(\frac{L_t^{safe}}{P_{t+1}},0\right)P_{t+1}.$$

from this contract.

Given these contracts, we can calculate the ex-post depreciation of the housing stock due to foreclosure losses. On a contract with loan size of L, the ex-post loan-to-value ratio is L/P_t and the ex-post default rate among borrowers with type τ is $D(L/P_t, \tau\phi_t, z)$. As a result, the average home quality after foreclosure losses given loan size of L and type τ is

$$\Gamma(L/P_t, \tau\phi_t) = \int_0^\infty \left[1 - \lambda D\left(L/P_t, \tau\phi_t, z\right)\right] zF'(z)\mathrm{d}z$$

We can then calculate the stock of housing after idiosyncratic value shocks and foreclosure losses. This value depends on the allocation of contracts from the prior period, and therefore depends on the prior credit regime. Let $\tilde{H}_t$ denote this remaining stock of housing.

If the screening regime prevailed in the previous period, then this quantity satisfies

$$\begin{aligned}\tilde{H}_t = {} & \Gamma\left(\frac{L_{t-1}^{safe}}{P_t},0\right)(1-n_L-n_H) + \Gamma\left(\frac{L_{t-1}^{doc}}{P_t},0\right)\left[n_L(1-\rho_L)+n_H(1-\rho_H)\right] \\ & + \Gamma\left(\frac{L_{t-1}^{risky}}{P_t},\phi_t\right)\left(H_{t-1} - (1-n_L\rho_L - n_H\rho_H\right) \quad \text{if} \quad S_{t-1} = \text{Screening}\end{aligned}$$

Known safe borrowers used the safe contract. The remaining safe borrowers used the high-documentation contract. All remaining housing demand must come from risky borrowers who opted to use the risky contract.

On the other hand, if the pooling regime prevailed in the prior period, then the remaining stock of housing is

$$\tilde{H}_t = \Gamma\left(\frac{L_{t-1}^{safe}}{P_t}, 0\right)(1 - n_L - n_H) + \Gamma\left(\frac{L_{t-1}^{pool}}{P_t}, 0\right) n_L(1-\rho_L)\mu_{t-1}$$
$$+ \Gamma\left(\frac{L_{t-1}^{down}}{P_t}, 0\right) n_L(1-\rho_L)(1-\mu_{t-1}) + \Gamma\left(\frac{L_{t-1}^{pool}}{P_t}, \phi_t\right) n_L \rho_L \mu_{t-1}$$
$$+ \Gamma\left(\frac{L_{t-1}^{down}}{P_t}, 0\right) n_H(1-\rho_H) \quad \text{if} \quad S_{t-1} = \text{Pooling}.$$

This expression differs from the previous because risky borrowers no longer use the risky contract, safe borrowers in the low risk segment use both the pooling and high-downpayment contracts, safe borrowers in the high risk segment use high-downpayment contracts, and risky borrowers in the low risk segment use the pooling contract.

3.3.2 Rental and Housing Supply

I assume that absentee landlords supply rental units using apartment buildings. Apartment buildings are produced from numeraire using a constant returns to scale technology. Given the cost of funds of r faced by landlords, the rental rate is equal to the user cost of apartments:

$$R_t = \Xi - \frac{1-\delta_A}{1+r}\Xi$$

where Ξ is the landlord's conversion rate between apartments and numeraire and δ_A is the depreciation rate on apartments. Under this assumption, the supply of rental units is perfectly elastic. This setup simplifies the analysis and lets the model match the stability of rental rates over the housing boom. Although a stark simplification, the main results go through if landlords can produce rental services by combining housing and apartments — see Appendix 3.B for an example extension.

Housing units are produced by a competitive housing sector. New homes are built using a Cobb-Douglas production technology in land and numeraire. In each period, a quantity $\ell > 0$ of new land is auctioned to home builders by the government. The supply of new housing then solves:

$$\max_{H_t^{new}, I_t^H} \quad P_t H_t^{new} - I_t^H$$
$$\text{s.t.} \quad H_t^{new} \leq \ell^{\alpha} \left(\frac{I_t^H}{1-\alpha} \right)^{1-\alpha}$$

where I_t^H denotes investment in new housing. Since land is purchased from the government and the production of housing has constant returns to scale, developers make zero profits and the government receives land rents. The implied stock of new housing is

$$H_t^{new} = P_t^{\frac{1-\alpha}{\alpha}} \ell.$$

The quantity $(1-\alpha)/\alpha$ is the elasticity of new home construction to the home price.

The supply of housing is then the amount of housing from the prior period after depreciation and foreclosure costs, plus new housing. This quantity is

$$H_t = \tilde{H}_t + H_t^{new}$$

Note that both $\tilde{H}_t$ and H_t^{new} are increasing in the home price so housing supply is upward sloping. It is no longer inelastic, as in Section 3.2.

3.3.3 Numerical Example: Parameterization

I parameterize the model to benchmark values from the macroeconomic literature on housing. For the new parameters, I target historical mortgage rates and spreads, as well as the historical foreclosure rate. For the exogenous process for income risk, I choose parameter values that generate momentum. The qualitative results only require that income risk exhibits hump-shaped impulse responses. Table 3.1 summarizes the model parameterization.

Mortgage contract payoffs depend on the foreclosure loss parameter, λ, and the distribution of Z_{it}. I choose $\lambda = .27$ to match the estimate of the sale discount on a foreclosed homes from Campbell et al. (2011). I assume that Z_{it} has a log-normal distribution with mean equal to $1-\delta_o$ and with standard deviation parameter of σ_z. The parameter δ_o can be interpreted as the depreciation rate on owner-occupied housing. I choose this depreciation rate and the standard deviation parameter based on Piazzesi and Schneider (2016). Specifically, I choose δ_o equal to the lower bound of the range that they suggest, and use the standard deviation parameter to match idiosyncratic volatility in home prices.

In order to be consistent with the literature, I set the discount factor for borrowers to match the discount factor for impatient households in Iacoviello (2005). Similarly, I choose the lender cost of funds parameter, r, equal to the discount rate for patient households in Iacoviello (2005).

I parameterize land supply to match homeownership — given the steady state home price, we can use ℓ to match a 65% homeownership rate — the average homeownership rate in the three decades prior to 2003. I choose the expenditure share on land, α, to imply an housing supply elasticity, $\frac{1-\alpha}{\alpha}$, equal to the median of 1.5 in Saiz (2010).

In steady state, the full information contract for risky borrowers depends crucially on the income risk parameter μ^{ϕ}. Since this contract represents a high interest rate loan used by borrowers who are known to have elevated default risk, I choose this parameter to imply a 6% spread of the contract's interest rate above the cost of funds for lenders. This choice implies a steady state mortgage rate of about 10% faced by risky borrowers in the subprime market, which is in the typical range for interest rates faced by subprime

Table 3.1: Baseline parameterization (quarterly)

Parameter	Value	Target/Source
λ	0.27	Campbell et al. (2011)
δ_o	.0038	1.5% ann. depreciation (Piazzesi and Schneider (2016))
σ_z	.07	.14 ann. s.d. of home prices (Piazzesi and Schneider (2016))
β	.95	Iacoviello (2005)
r	$\ln(1/.99)$	Iacoviello (2005)
R	1	normalize rental rate to 1
ℓ	3.52×10^{-5}	65% s.s. homeownership rate
$\frac{1-\alpha}{\alpha}$	1.5	median supply elasticity in Saiz (2010)
$\exp(\mu^{\phi})$	11.1%	6% s.s. spread on risky contract
κ	0.0119	0.75% s.s. spread on high-doc contract in segment L
ρ_L	8.35%	0.75% s.s. spread on pooling contract in segment L
ρ_H	47.4%	0.36% foreclosure rate in steady state (Corbae and Quintin (2013))
n_L, n_H	29.3%, 70.1%	1% s.s. share pooling and 1% s.s. share risky borrowers
ρ_1^{ϕ}	1.6	
ρ_2^{ϕ}	$-.63$	
σ^{ϕ}	.05	

borrowers prior to the housing boom.

Next, I choose values for ρ_L and κ to imply that both the pooling contract in the low risk segment and high-documentation contracts have a 0.75% spread over the lender's cost of funds. In the model, when these two contracts have the same interest rate, safe borrowers in the low risk segment will prefer the high-documentation contract because it gives them their optimal loan size. However, keeping the spread on the two contracts close to each other ensures that there is some chance that shocks to income risk can trigger a switch to the pooling regime. Since the high-documentation contract interest rate depends on both parameters, I first choose ρ_L to imply a 0.75% spread on the pooling contract, and then choose κ (given ρ_L) to get a 0.75% spread on the screening contract. This parameterization implies that the mortgage rate for the high-documentation contract is about 4.5%, in line with levels of prime mortgage rates (after accounting for inflation) during the early 2000's (see Justiniano et al. (2016)).

Finally, we can jointly pin down the share of risky borrowers in the high risk segment, the size of the low risk segment, and the size of the high risky segment to target the share of pooling contracts, the share of contracts going to borrowers who are revealed to be risky (subprime mortgages), and the steady state foreclosure rate. I set the former two targets to 1% to reflect that both forms of non-prime lending were rare prior to the housing boom. For the later moment, I match a 0.36% foreclosure rate following Corbae and Quintin (2013).

Finally, I choose parameter values for the exogenous income risk process to imply that income risk has some momentum. The exact parameter values are not important for the qualitative results. A hump-shaped response is necessary to get persistent changes in the credit regime from slightly different paths for income risk. I choose the standard deviation parameter σ^{ϕ} to be just large enough so that the solution algorithm (see Appendix 3.D) converges to an equilibrium where credit regime changes occasionally arise. If this parameter is too small, then equilibrium remains in a small enough neighborhood of the steady state that the algorithm converges to the usual local linear approximation coming from traditional perturbation methods. For a sufficiently large standard deviation, the algorithm will instead converge to a regime-switching linear approximation, capturing the non-linear equilibrium dynamics arising in this model.

3.3.4 Numerical Example: Simulations

Figure 3.8 shows equilibrium outcomes from two simulations of the model. In the first simulation (dashed lines), I start the economy at its stochastic steady state[19], and then shock the economy for three periods with a series of -0.5 standard deviation income risk shocks. These shocks push the economy to the brink of regime change, but do not quite trigger the pooling regime. In the second simulation, I perform the same exercise, but add a fourth -0.5 standard deviation shock, which pushes the economy past

[19]The stochastic steady state is the long-run outcome in the absence of shocks and is not necessarily equal to the deterministic steady state.

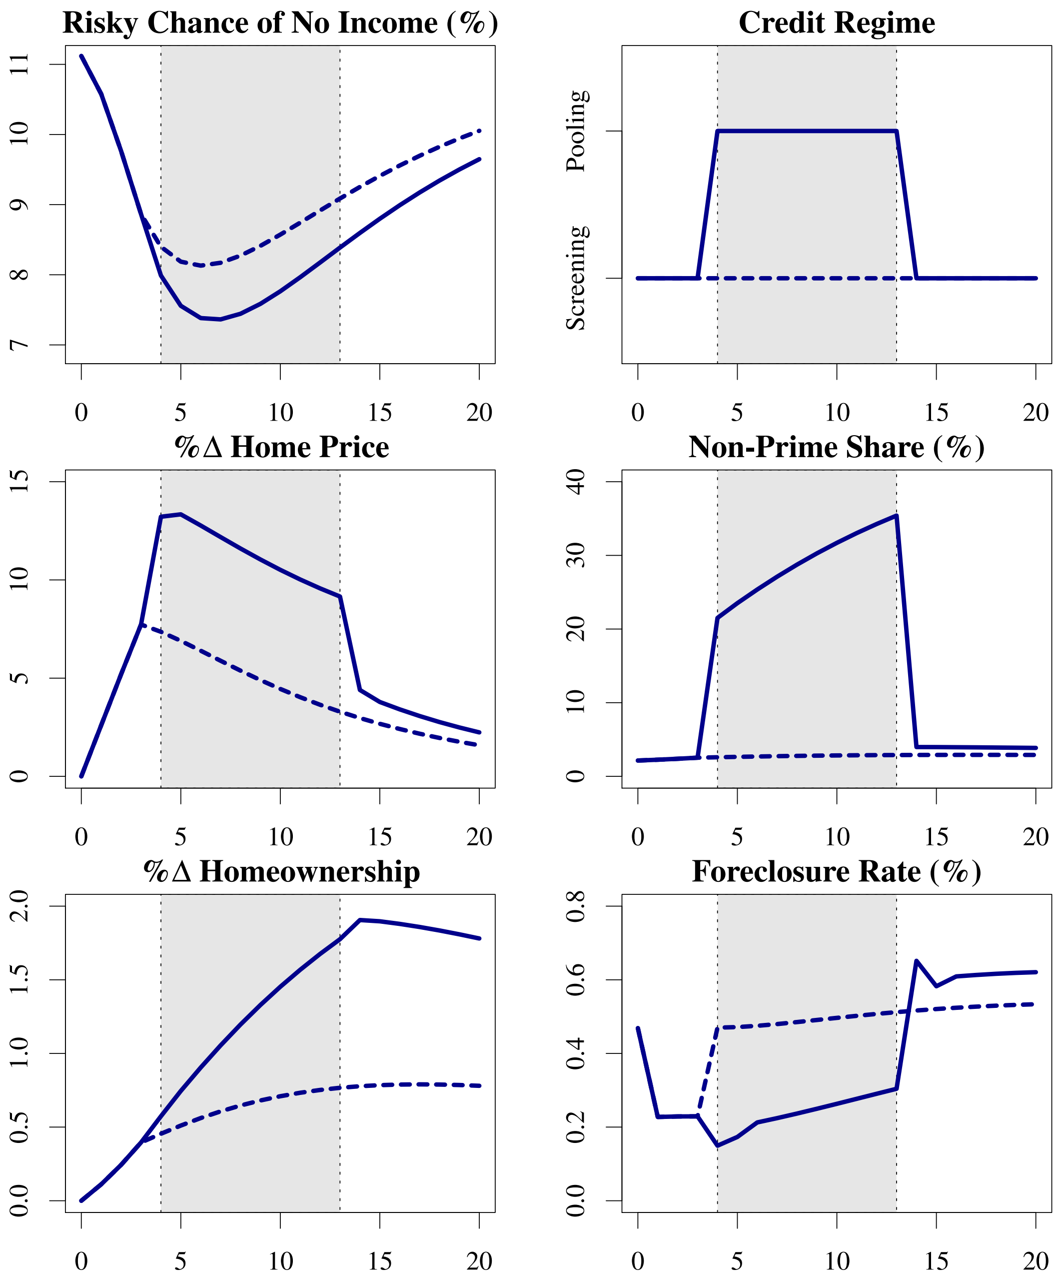

Figure 3.8: Response to 3 (dashed) and 4 (solid) -1/2 standard deviation shocks to income risk. Shaded region indicates pooling regime in second simulation.

the tipping point and triggers the pooling regime. The difference between the dashed and solid impulse responses lets us visualize the effect of the regime change.

The top left panel shows the implied paths for income risk in the two scenarios. Both have a smooth hump-shaped response to the shocks. The responses are identical for the first three periods, but then diverge in the fourth period. For both simulations, income risk steadily and smoothly reverts to steady state after peaking.

These two slightly different paths for income risk lead to very different equilibrium outcomes. The top right panel shows the equilibrium credit regime. In the three-shock simulation, the economy stays in the screening regime throughout and exhibits a smooth impulse response. In the four-shock simulation, the additional shock in period 4 triggers the pooling regime which persists until period 14.

The regime change leads to a sudden and rapid shift in equilibrium outcomes. Home prices rapidly appreciate (middle left panel), the share of non-prime lending in the economy jumps up and begins to grow (middle right panel), the growth rate of homeownership increases (bottom left panel), and the foreclosure rate slightly falls and then begins to grow (bottom right panel).

As income risk reverts smoothly toward steady state, the economy eventually reaches the regime-change threshold again. The resulting shift back to the screening regime leads to a credit contraction (collapse in non-prime share) and a sudden collapse in home prices. The collapse in home prices pushes homeowners underwater and generates a foreclosure crisis — which appears as a spike in the foreclosure rate in period 14.

Again, note that no shock hits the economy to generate the bust — the only driving force underlying these responses is the smoothly varying path of income risk. From the accumulated effect of shocks in the early periods of the simulation, we get a rich boom-bust pattern through the model's state-dependent amplification mechanism. Although it appears as-if a series of highly correlated shocks hit the economy in periods 4 and then again in period 14, the true underlying income risk shocks all occurred within the first 4 periods of the simulation. The turning points of the cycle can be entirely attributed to the propagation and amplification of these initial shocks.

It took a number of these initial shocks to push the economy to the tipping point

— specifically four -0.5 standard deviation shocks — so this housing boom-bust episode is a relatively rare event. The conditions that trigger the housing cycle develop steadily through the cumulative effect of small shocks to the aggregate state of the economy. This result formalizes the idea that developing conditions in an economy may lead up to a boom. This model provides micro-foundations to study how evolving economic conditions might signal an imminent credit boom.

This theory provides a framework for modeling credit-driven housing booms which end in housing crises. The predictions of the model line up with the historical evidence of Jordà et al. (2015), who show that a simultaneous boom in credit and home prices predicts a subsequent bust. This theory proposes adverse selection in the mortgage market and general equilibrium between housing and the mortgage market as an explanation for this type of boom-bust pattern. Because of the non-linear mechanism underlying this result, the model does not rely on shocks to drive the turning points in the cycle. Rather, small shocks accumulate into the state of the economy to trigger the boom, then mean reversion in fundamentals triggers the bust.

3.4 Discussion of Empirical Evidence

This section examines the theory's predictions for the evolution of debt across borrowers, and for mortgage rates. I argue that the model matches existing empirical evidence from individual-level mortgage outcomes during the housing boom and bust.

3.4.1 The Allocation of Debt During the Boom

The microeconomic literature has recently turned to credit bureau data to examine how mortgage markets allocated debt during the housing boom. Mian and Sufi (2015) examine the growth in debt of individuals borrowers, binned according to their FICO score in 1997. They find that debt grew more for borrowers with low FICO scores in 1997 than for borrowers with high FICO scores in 1997 and show that the difference across the credit score distribution comes from the dynamics of loan-to-value ratios. High FICO score borrowers reduces their loan-to-value ratios during the boom while low FICO score borrowers did not. In contrast, Albanesi et al. (2016) look at the allocation of debt within each year by FICO score. They find that the percent difference in debt between low FICO score borrowers in 2005 and low FICO score borrowers in 2000 was smaller than the percent difference in debt between high FICO score borrowers in 2005 and high FICO score borrowers in 2000. Based on different conditional moments calculated from similar data, these authors come to different conclusions[20].

I argue that the theory presented in this paper may help to reconcile these results. When comparing the theory to this empirical evidence, it is crucial to distinguish between borrower time-invariant characteristics (captured by τ_i), and borrower observable characteristics relevant for credit decisions (captured by ρ_i). A switch from the screening regime to the pooling regime leads to a change in how lenders ration credit. They substitute down-payment based rationing for income-verification based rationing. The implication is that safe borrowers in the prime market will reduce their loan-to-value ratios after the regime change. In contrast, risky borrower loan-to-value ratios are relatively stable since they find high-downpayment loans undesirable. As a result, risky borrowers have debt which moves in proportion to home prices while safe borrower debt moves less than one-for-one with home prices. This pattern of loan-to-value ratios over time by borrower type implies that debt grows more for risky borrowers during the boom.

However, fundamental types are not directly measurable. The variable, ρ_i, measures lender initial beliefs about the riskiness of borrowers before they make mortgage offers — which is likely to be correlated with a borrowers credit score close to the time of

[20] Albanesi et al. (2016) replicate the result of Mian and Sufi (2015) so the difference is unlikely due to a discrepancy in their samples.

mortgage origination. Since the mortgage market is segmented according to this variable, the pattern of debt growth will vary by ρ_i. Average debt growth in each market segment depends on both the composition of funded loans (related to the share of risky borrowers in the segment) and the terms of those loans. In low risk parts of the market, lenders pool borrowers which leads to a relatively stable loan-to-value ratio regardless of the credit regime. In contrast, lenders screen in the high risk segment, and loan-to-value ratios will change depending on whether or not they use downpayment requirements or income verification to ration credit. The implication is that loan-to-value ratios will fall during a switch to the pooling regime because lenders begin to rely on downpayment requirements to separate out borrowers. As a result, debt growth will be larger for borrowers that are *observably low risk* (low ρ_i) because they have more stable loan-to-value ratios.

As a result, the distinction between fundamental types and lender beliefs about fundamental types can reverse the pattern of debt growth predicted by the theory. Whether a given measure of borrower riskiness is more related to borrower types or lender information about types changes how it relates to the theory.

To the extent that the measure used by Mian and Sufi (2015) captures time-invariant dimensions of borrower riskiness, I argue that the theory is consistent with their evidence. They track fixed groups of individuals based on their FICO scores in 1997 and report how debt changed for each group during the boom. If borrower fundamental types are persistent, we would expect the share of risky borrowers to be larger among those borrowers with low FICO scores in 1997 relative to those borrowers with high FICO scores in 1997. As a result, I argue that their measure of borrower riskiness relates to τ_i in the theory. Under this interpretation, the model's predictions for debt growth and loan-to-value ratios are consistent with the evidence of Mian and Sufi (2015).

Simultaneously, I argue that the model also matches the evidence of Albanesi et al. (2016). Their measure of borrower riskiness reflects observables close to the time

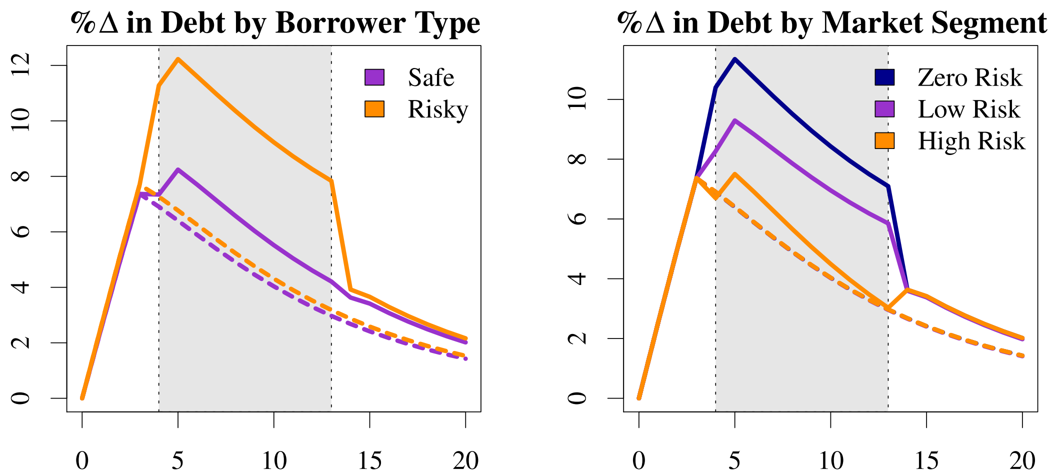

Figure 3.9: Debt accumulation by fundamental borrower types (left) and borrower observables at origination (right). Response to 3 (dashed) and 4 (solid) -1/2 standard deviation income risk shocks. Shaded region indicates pooling regime in second simulation.

of mortgage origination, which relates to the risk index, ρ_i. Along this dimension, the model predicts that borrowers that appeared risky at the time or mortgage origination had less debt growth during the boom — consistent with their findings.

Figure 3.9 shows the model's predictions for debt growth by borrower fundamental types (τ_i) and by borrower observables (ρ_i). I show the cumulative change in debt in the two simulations considered in Section 3.3. In the first simulation (dashed lines) three -1/2 standard deviation shocks hit the economy. In the second (solid lines), an additional fourth shock hits the economy and triggers the pooling regime.

In the three-shock simulation, debt grows symmetrically across fundamental borrower types and across borrower observables. In the absence of a change in the credit regime, lenders use income verification to screen risky borrowers which leads to stable loan-to-value ratios within individual borrowers. As a result, debt moves in proportion to home prices and the pattern of debt accumulation reflects the response of home prices to the fall in income risk. There are no differential patterns of debt growth.

In contrast, we do get differences in debt growth by borrower type and borrower observables when a fourth shock triggers the pooling regime. First consider the right panel — debt growth by borrower observables (ρ_i). In the zero-risk segment all borrowers are known to be safe and are offered pooling contracts in both regimes. They have relatively stable loan-to-value ratios and so their cumulative debt growth (blue line) moves one-for-one with home prices. On the other extreme, the high risk segment of the market consists entirely of safe borrowers who switch from income-verification based screening to downpayment based screening. These borrowers see a fall in their loan-to-value ratios because of this shift in the nature of credit rationing. As a result, their cumulative debt growth (orange line) moves less than one-for-one with the home price during the boom. In the low risk segment of the market, safe borrowers are indifferent between pooling and making a large downpayment. Lenders mix between offering the two types of contracts. As a result, a portion of borrowers switches from income-verification based screening to pooling while the remaining portion switches from income verification based screening to downpayment-based screening. As a mixture of the two kinds of contracts, their response for debt accumulation is between the extremes of the zero risk and the high risk segments. This pattern — with more debt growth among observably safe borrowers — is qualitatively consistent with the findings of Albanesi et al. (2016) if we assume that their measure of borrower riskiness captures ρ_i in the theory.

Moving to the left graph, we get a reverse pattern that I argue is consistent with the evidence of Mian and Sufi (2015). The debt growth of risky borrowers (orange line) follows the path of home prices because they have relatively stable loan-to-value ratios. In the screening regime they use their full information contract and in the pooling regime they use a pooling contract. Neither require a particularly large downpayment, and so their debt moves in proportion to the home price. Safe borrowers use a mixture of contracts because they get loans from multiple segments of the overall market. However, because a large portion of safe borrowers in the prime market switch from income-verification based screening to downpayment based screening, there is a fall in the average loan-to-value ratio among safe borrowers. The shift in the nature of credit rationing implies that safe borrower debt moves less than one-for-one with home prices. In contrast, risky borrowers use their full information contract in the screening regime and switch to using pooling contracts in the pooling regime. Neither has a large downpayment requirement, and so loan-to-value ratios do not change very much for risky borrowers. Their pattern for debt growth reflects the movement of home prices. As a result, risky borrower debt grows more than safe borrower debt during the boom. These patterns for debt growth and loan-to-value ratios are consistent with Mian and Sufi (2015) if we assume that their measure of borrower riskiness proxies for borrower fundamental types. This assumption is reasonable so long as borrower fundamental types are relatively persistent.

Note that the above discussion abstracts from possible compositional effects from more risky borrowers entering the market. In the theory, the homeownership rate is relatively constant over the simulation (increasing by two percentage points). If home prices didn't rise, then more risky borrowers would enter the market and drive up the homeownership rate. But, in general equilibrium, rising home prices lead to relaxed incentive constraints, and a shift to downpayment based credit rationing. The rationing effect of higher home prices keeps risky borrowers from driving up the homeownership rate and ensures that the housing market clears. Consistent with this result, Foote et al. (2016) show that, indeed, the boom was a bad time to become a new homeowner.

In summary, I argue that the theory is consistent with both the evidence in Mian and Sufi (2015) and the evidence in Albanesi et al. (2016). Rather than being conflicting, I argue that these studies complement each other by giving two distinct sets of moments which any theory of the boom ought to match. By following fixed individuals over time, Mian and Sufi (2015) let us think about reallocation of debt by persistent borrower characteristics. By examining how debt was allocated in each year across borrower observables, Albanesi et al. (2016) show how the market responded to available information. Each approach captures a distinct theoretical concept — borrower fundamental types versus observable characteristics. The theory offered by this paper is consistent with these moments, and highlights the importance of distinguishing between borrower types and borrower observables.

3.4.2 The Mortgage Rate Conundrum

Justiniano et al. (2016) examine individual level loan data and document a puzzle — during the boom, mortgage rates were low relative to the level we would expect given prevailing treasury yields. In particular, they show that after composition adjusting the sample of loans, average mortgage rates fell relative to treasury yields by about 50 basis points during the later half of 2003, and this persisted for the duration of the housing boom.

Figures 3.10 and 3.11 show a stylized replication of their findings using aggregate data. Figure 3.10 shows the average spread on 15 year and 30 year fixed rate mortgages over the ten year treasury yield from January of 2000 to January of 2008. Between the middle of 2003 and the end of 2005, mortgage rate spreads were low relative to the overall sample.

We should expect these spreads to move endogenously with various aggregate factors. I follow Justiniano et al. (2016) and calculate a measure of residual variation in mortgage rates that is unexplained by the historical relationship of mortgage rates to the yield curve and macroeconomic factors.

Specifically, on a sample from January of 1990 to January of 2008 I run the regression

$$i_t = \alpha_m + \beta' X_t + u_t \tag{3.4}$$

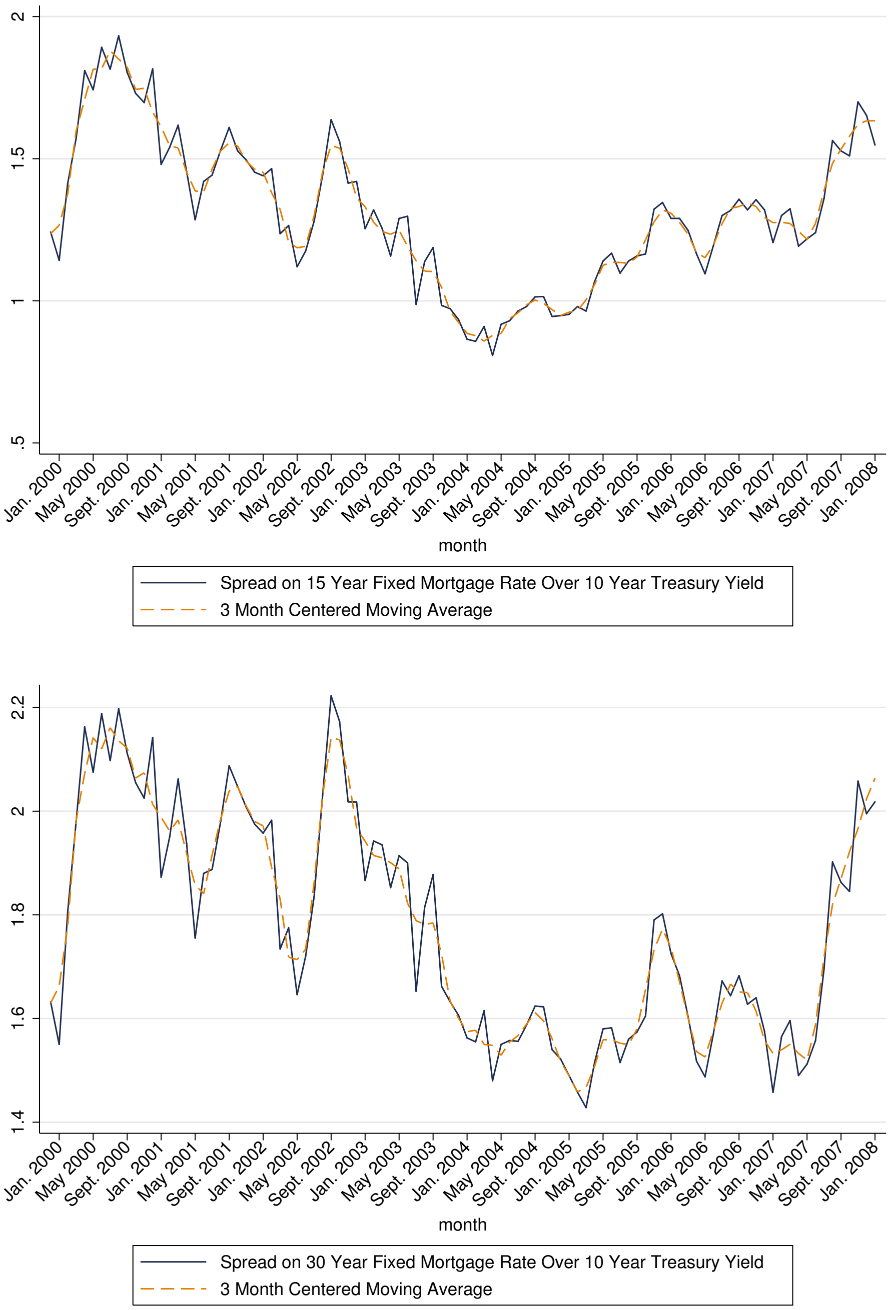

Figure 3.10: Average spread on fixed rate mortgages over ten year treasuries

where i_t is either the average interest rate on 15 or 30 year fixed rate mortgages, α_m is a dummy for the month in the year, and X_t is a vector containing the contemporaneous and lagged 1, 5 ,and 10 year treasury yields, the lagged unemployment rate, lagged growth in

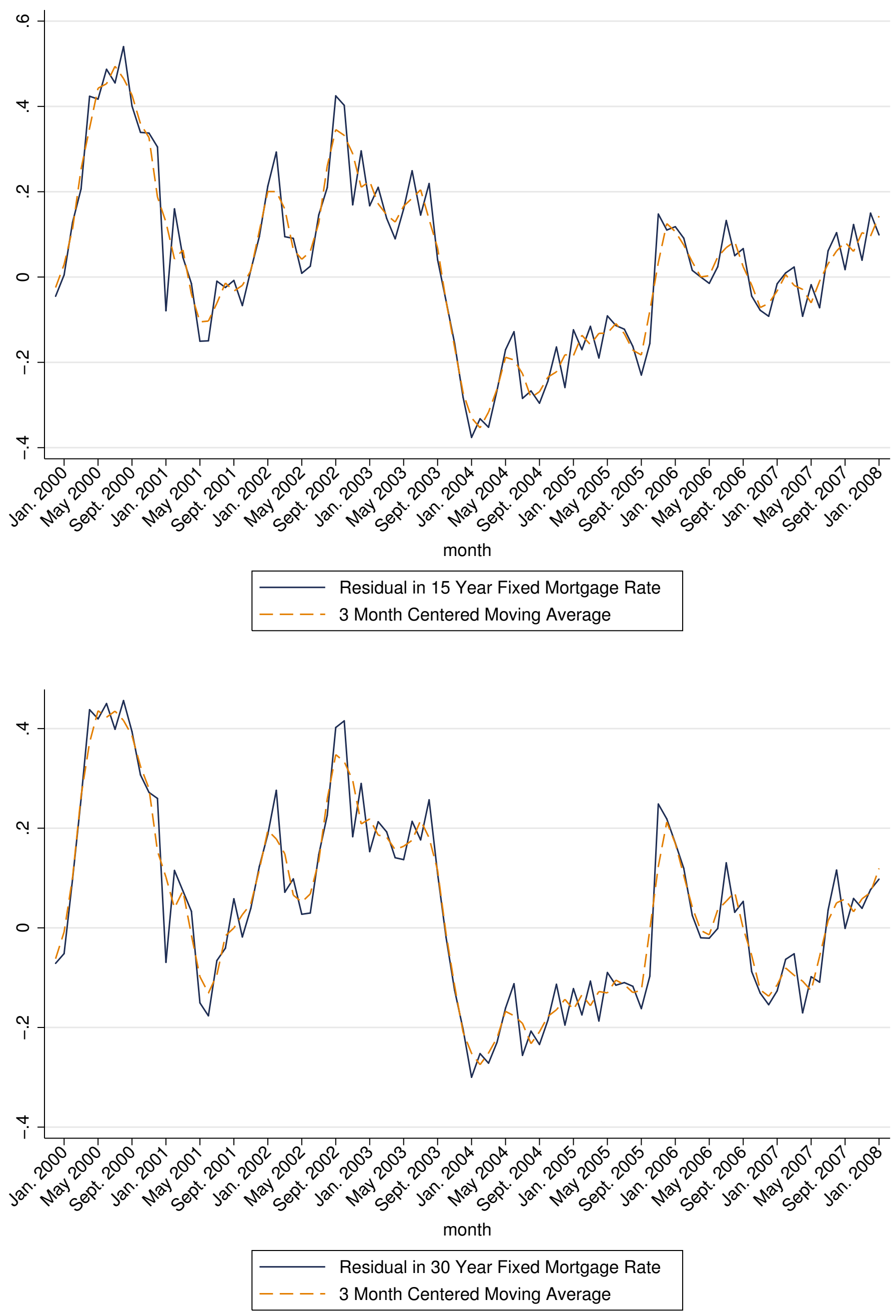

Figure 3.11: Residual variation in fixed rate mortgages (OLS residuals from (3.4)).

industrial production, and lagged inflation. The residuals from this regression represent movements in mortgage rates that are not explained by historical co-movement with these aggregate factors. Figure 3.11 shows these residuals from January of 2000 to January of

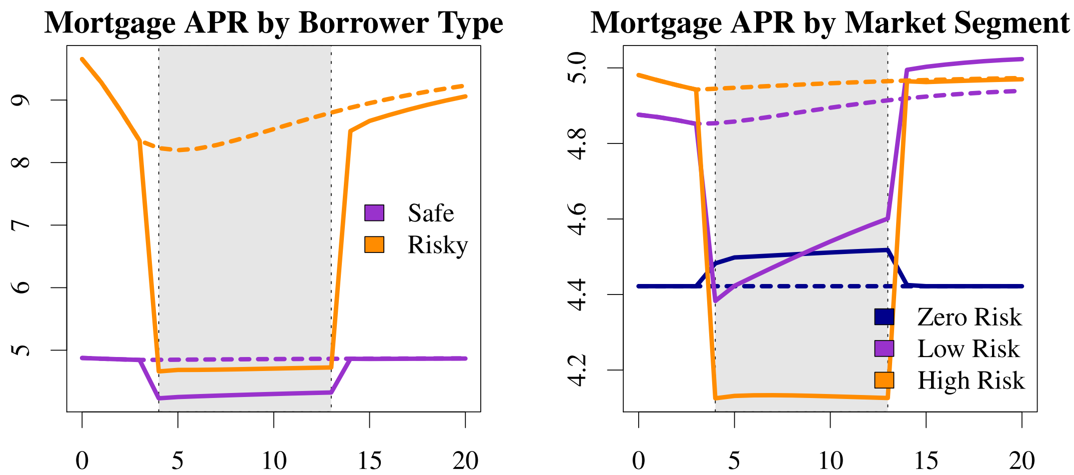

Figure 3.12: Mortgage rates by fundamental borrower types (left) and borrower observables at origination (right). Response to 3 (dashed) and 4 (solid) -1/2 standard deviation income risk shocks. Shaded region indicates pooling regime in second simulation.

2008.

The residual variation in mortgage rates moves closely together for both series. Moreover, the pattern follows precisely the results of Justiniano et al. (2016). In the later half of 2003, both series show a sudden and rapid fall in mortgage rates relative to what we would expect given historical co-movements with aggregate factors. This effect persists during the housing boom, reverting to zero around 2005-2006.

The theory offered by this paper can explain the mortgage rate conundrum documented by Justiniano et al. (2016) as coming from a compression of risk spreads on mortgages, and as well as a reduction in screening costs that get priced into contracts. In particular, it explains how a sudden fall in mortgage rates relative to aggregate factors can occur — through a shift from a screening regime to a pooling regime in the mortgage market.

Figure 3.12 shows the paths for mortgage rates in a simulation where three -1/2 standard deviation shocks to income risk hit the economy and a simulation where four -1/2 standard deviation shocks hit the economy. It shows responses of mortgage rates disaggregated by borrower type and by borrower observables at origination. In the first simulation the economy remains in the screening regime throughout, while the economy enters and exits the pooling regime in the second simulation. The grey shaded region shows when the pooling regime occurs.

A switch to the pooling leads to a collapse in mortgage rates — especially for risky borrowers and borrowers in the low risk and high risk market segments. This result is consistent with the findings of Justiniano et al. (2016).

Why do mortgage rates fall? Consider a risky borrower who gains access to a pooling contract. Initially, they used a full information contract which fully priced in their elevated default risk. When they gain access to a pooling contract, this default risk gets diluted through pooling with safe borrowers. As a result, they see a fall in the risk premium on their mortgage due to an implicit subsidy from safe borrowers.

For safe borrowers who switch from high-documentation to high-downpayment contracts, the fall in the interest rate comes from reduced screening costs for lenders. When screening, lenders sink resources into producing information about individual borrowers and the screening contract prices these costs into the interest rate. When a safe borrower moves from a high-documentation contract to a high-downpayment contract, these cost are no longer paid by lenders. A high down payment additionally reduces risk of strategic default. As a result, these contracts have relatively low interest rates and the shift implies a fall in mortgage spreads.

The non-linear mechanism in this theory, where small shocks to income risk can trigger large changes in equilibrium outcomes, provides an explanation for the mortgage rate conundrum as an endogenous outcome driven by macroeconomic fundamentals. As emphasized by Justiniano et al. (2016), the timing of the relative fall in mortgage rates coincided with the end of the Federal Reserve's easing cycle. Specifically, in the summer of 2003, the Fed signaled that, due to improving conditions in labor markets, they would no longer continue to lower their policy rate. At this point, long-term yields began to rise in anticipation of a higher future policy rate, yet mortgage rates didn't rise with treasury yields. I argue that we can understand the mortgage rate puzzle as coming from the simultaneous endogenous response of both the mortgage market and policy makers to improvements in the macroeconomy.

Justiniano et al. (2016) also show that the size of the mortgage origination industry didn't follow its historical pattern of shrinking at the end of a Fed easing cycle. Similar to past business cycles, there was a mortgage refinance boom as monetary policy lowered interest rates after the recession of 2001. During this period, employment in the mortgage industry grew steadily. But this time it continued to grow when the Fed signaled higher future interest rates, likely due to the boom in non-prime lending that began at the end of 2003. The theory of this paper explains the non-prime lending boom as the rational profit maximizing response of lenders to falling income risk, and therefore can rationalize the continued expansion of the mortgage industry at the end of the Fed's policy cycle.

The expansion in non-prime lending explains both low mortgage rates relative to the yield curve, and the continued expansion of the mortgage industry. This pattern was at odds with the historical pattern for precisely the same reasons that the emergence of non-prime lending was at odds with historical patterns. This narrative explains the mortgage rate conundrum by explaining the simultaneity of the non-prime lending boom, the housing boom, rising yields, and improving labor market conditions.

3.5 Conclusion

In this paper, I develop a theory of mortgage credit which provides an explanation for the rise of non-prime lending during the early 2000's, the associated housing boom, and the subsequent housing crisis. This theory is built on the assumption that competitive lenders face asymmetric information about borrower income risk during mortgage origination, which leads to endogenous segmentation of the mortgage market. As a result, the model makes rich predictions for the allocation of credit and mortgage rates across borrowers. I show that it can match and may to help reconcile the existing microeconomic evidence on debt allocation during the boom, as well as explain the mortgage rate conundrum. The dynamics in the theory rely solely on variation in income risk over the business cycle, and the model generates boom-bust cycles through the endogenous and state-dependent amplification of these fundamentals. This theory allows us to explain housing boom-bust episodes without appealing to either new or large shocks.

The key non-linear mechanism operates through general equilibrium feedback between housing and mortgage markets. During an expansion, income risk falls, which increases the share of borrowers offered pooling mortgage contracts. If this expansion in credit access attracts former renters into the housing market, then housing demand increases. This increase in demand triggers a regime change — lenders stop offering mortgages with high documentation requirements, loan-to-value ratios suddenly increase, and home prices rapidly appreciate. As the cycle peaks and income risk begins to revert, there is a reversal in the credit regime. Credit access falls, home prices collapse, and, as a result, there is a foreclosure crisis as borrowers go underwater on their homes. Through the non-linear mechanism implied by joint housing and mortgage market clearing, we can explain housing boom-bust cycles resulting from small shocks to economic fundamentals.

The dynamic model has a structure that can be embedded into standard business cycle models. An important next step is to introduce the mechanism into a production economy that includes frictions in financial intermediation. Feedback from home prices to consumption and output may amplify housing boom-bust episodes and influence business cycle dynamics generally. Also, because the housing crisis precipitated a financial panic in markets that relied on mortgage-backed securities as collateral, frictions in the intermediation of funds between lenders and savers may be particularly important for understanding recessions which follow after housing crises. This theory provides a novel mechanism for understanding the origins of housing crises and provides micro-foundations for studying policies intended to reduce the likelihood of future crises and recessions.

Appendix

3.A Equilibrium Characterization and Proofs

Lemma 3.A.1 (Sub-game Perfect Period 1 Expected Payoffs) *Consider expected payoffs from the sub-game beginning in period 1. Suppose that an individual borrower i has type τ_i and has been accepted into a contract with transfer of T_i and loan size of $L_i > 0$. Then the expected payoff from the contract for the borrower is $T_i + \beta U(L_i, \tau_i \phi)$ and the expected payoff to the lender is $-T_i + \frac{1}{1+r}\Pi(L_i, \rho_i \phi)$ where ρ_i is the lender's subjective belief that i is the risky type.*

Proof. The result follows from backward induction in period 1. At the end of period 1, lenders choose whether or not to initiate foreclosure following a borrower default. Their payoff from initiating foreclosure is $-T_i + \frac{1}{1+r}\min\{L_i, (1-\lambda)Z_i\}$. If they do not initiate foreclosure, their payoff is $-T_i$. As a result, since $\mathbf{P}[Z_i > 0] > 0$, they have a strictly dominant strategy to initiate foreclosure.

Borrowers anticipate that foreclosure will occur following any default. If $Y_i = 0$, then they cannot make their payment of L_i, and must default. If they have income and do not default, they get $-T_i + \beta(Y_i + Z_i - L_i)$. If they default they get $-T_i + \beta(Y_i + \max\{(1-\lambda)Z_i - L_i, 0\})$. The gain from default is $\beta \max\{-\lambda Z_i, L_i - Z_i\}$. The gain is strictly positive when $L_i > Z_i$, is zero when $L_i = Z_i$, and is strictly negative when $L_i < Z_i$. Since Z_i is continuously distributed, the borrower has a strictly dominant strategy with probability 1.

Therefore, the payoffs from this sub-game are

$$\mathbf{1}\{Z_i < L_i \text{ or } Y_i = 0\}\,[T_i + \beta(Y_i + \max\{(1-\lambda)Z_i - L_i, 0\})] \\ + \mathbf{1}\{Z_i \geq L_i \text{ and } Y_i = Y\}\,[Y_i + Z_i - L_i]$$

for the borrower, and

$$\mathbf{1}\{Z_i < L_i \text{ or } Y_i = 0\}\left[-T_i + \frac{1}{1+r}\min\{L_i, (1-\lambda)Z_i\}\right] \\ + \mathbf{1}\{Z_i \geq L_i \text{ and } Y_i = Y\}\left[-T_i + \frac{1}{1+r}L_i\right]$$

for the lender. Given type τ_i and the lender's belief about i's type of ρ_i, the result follows from calculating expected values of these payoffs. ■

These payoff functions are the continuation payoffs for the sub-game following the borrower's application decision and the lender's screening and accept/reject decision. For any finite set of contracts that are offered to an individual borrower, this sub-game is a finite signaling game. The message space of the borrower (sender) is the set of contracts and the action space of the lender (receiver) is the choice of screening (if the contract applied to is a screening contract) and the choice of whether to accepting or reject the contract. Note that because this sub-game is a standard signaling game, there can be many equilibria.

However, competition among lenders over contract offers will select among these possible continuation outcomes, and I can show that, up to the application behavior of risky borrowers when offered screening contracts, the sub-game perfect equilibrium of the overall game is unique, despite there existing many equilibria of the sub-game that begins with the borrower application decision.

To do so, I define competitive contracts that may emerge as equilibrium offers, and show that the contracts satisfying these definitions are unique.

First, I calculate the derivatives of Π and U in L. We have

$$\begin{aligned}\Pi(L,\phi) &= L - \int_0^\infty D(L,\rho\phi,z)\max\{L-(1-\lambda)z,0\}F'(z)\mathrm{d}z \\ &= L - \int_0^L [L-(1-\lambda)z]F'(z)\mathrm{d}z - \phi\int_L^{\frac{L}{1-\lambda}}[L-(1-\lambda)z]F'(z)\mathrm{d}z\end{aligned}$$

and so, by Leibniz's rule,

$$\begin{aligned}\Pi_1(L,\phi) &= 1-\lambda LF'(L) - F(L) - \phi\left[-\lambda LF'(L) + F\left(\frac{L}{1-\lambda}\right) - F(L)\right] \\ &= 1-(1-\phi)(\lambda LF'(L)+F(L)) - \phi F\left(\frac{L}{1-\lambda}\right).\end{aligned}$$

Similarly,

$$\begin{aligned}U(L,\phi) &= \bar{z} - L + \int_0^\infty D(L,\tau\phi,z)\max\{-\lambda z, L-z\}F'(z)\mathrm{d}z \\ &= \bar{z} - L + \int_0^L (L-z)F'(z)\mathrm{d}z + \phi\int_L^{\frac{L}{1-\lambda}}(L-z)F'(z)\mathrm{d}z + \phi\int_{\frac{L}{1-\lambda}}^\infty(-\lambda z)F'(z)\mathrm{d}z\end{aligned}$$

and

$$\begin{aligned}U_1(L,\phi) &= -1 + F(L) + \phi\frac{1}{1-\lambda}\left(L - \frac{L}{1-\lambda}\right)F'\left(\frac{L}{1-\lambda}\right) \\ &\qquad + \phi\left[F\left(\frac{L}{1-\lambda}\right) - F(L)\right] + \phi\frac{\lambda}{1-\lambda}\frac{L}{1-\lambda}F'\left(\frac{L}{1-\lambda}\right) \\ &= -1 + (1-\phi)F(L) + \phi F\left(\frac{L}{1-\lambda}\right).\end{aligned}$$

From this result, we can prove Lemma 3.1.1.

Proof of Lemma 3.1.1. We have

$$\frac{\partial}{\partial L}\left[U(L,\phi)-U(L,0)\right]=U_1(L,\phi)-U_1(L,0)=\phi\left[F\left(\frac{L}{1-\lambda}\right)-F(L)\right].$$

Since $\phi > 0$, this quantity is strictly positive if and only if $\mathbf{P}[L < Z < L/(1-\lambda)] = \mathbf{P}[(1-\lambda)Z < L < Z] > 0$. ■

The following lemma establishes an intermediate result for proving the uniqueness of equilibrium contracts.

Lemma 3.A.2 *Let Assumption 3.1.1 and Assumption 3.1.2 hold. Fix some* $\rho \in [0,1]$ *and* $\tau \in [0,\rho]$. *Then the program*

$$\begin{aligned} \max_{T,L} \quad & T+\beta U(L,\tau\phi) \\ s.t. \quad & T \le \frac{1}{1+r}\Pi(L,\rho\phi) \end{aligned} \tag{PC}$$

has a unique and finite solution.

Proof. Since the objective is strictly increasing in T, the lender's zero profit (participation) constraint must bind, pinning down T for any given L. The first order necessary condition for the loan size is then

$$\begin{aligned} 0 &= \frac{1}{1+r}\Pi_1(L,\rho\phi)+\beta U_1(L,\tau\phi) \\ &= \frac{1}{1+r}\left[1-(1-\rho\phi)(\lambda LF'(L)+F(L))-\rho\phi F\left(\frac{L}{1-\lambda}\right)\right] \\ &\qquad\qquad +\beta\left[-1+(1-\tau\phi)F(L)+\tau\phi F\left(\frac{L}{1-\lambda}\right)\right] \end{aligned}$$

Re-arrange to get

$$
\begin{aligned}
(1-(1+r)\beta)(1-F(L)) &= (1-\rho\phi)\lambda L F'(L) + (\rho-(1+r)\beta\tau)\phi\left[F\left(\frac{L}{1-\lambda}\right)-F(L)\right] \\
&= (1-\rho\phi)\lambda L F'(L) + (\rho-(1+r)\beta\tau)\phi\int_1^{\frac{1}{1-\lambda}} LF'(Lx)\mathrm{d}x
\end{aligned}
$$

For any L such that $F(L/(1-\lambda))=0$, this condition cannot hold because the left-hand-side is strictly positive, while the right hand side must equal zero. For any L such that $F(L)=1$, the condition will hold — this case corresponds to when all borrowers default for sure on the loan. To see if such a large loan is optimal, we need to compare to the case where $F(L)<1$.

In this case, we can re-write in terms of the hazard function $h(z)=F'(z)/(1-F(z))$.

$$
\begin{aligned}
1-(1+r)\beta &= (1-\rho\phi)\lambda L\frac{F'(L)}{1-F(L)} + (\rho-(1+r)\beta\tau)\phi\int_1^{\frac{1}{1-\lambda}} L\frac{F'(Lx)}{1-F(L)}\mathrm{d}x \\
&= (1-\rho\phi)\lambda L h(L) + (\rho-(1+r)\beta\tau)\phi\int_1^{\frac{1}{1-\lambda}} Lh(Lx)\frac{1-F(Lx)}{1-F(L)}\mathrm{d}x \qquad (3.5)
\end{aligned}
$$

Note that because $zh(z)$ is strictly increasing in z

$$
\begin{aligned}
\frac{\partial}{\partial L}\ln\frac{1-F(Lx)}{1-F(L)} &= \frac{xF'(Lx)}{1-F(Lx)} - \frac{F'(L)}{1-F(L)} = \frac{1}{L}Lxh(Lx) - \frac{F'(L)}{1-F(L)} \\
&> \frac{1}{L}h(L) - \frac{F'(L)}{1-F(L)} = \frac{1}{L}\frac{LF'(L)}{1-F(L)} - \frac{F'(L)}{1-F(L)} = 0
\end{aligned}
$$

which implies $Lh(Lx)\frac{1-F(Lx)}{1-F(L)}$ is increasing and continuous in L for each x. As a result, the right-hand side of Equation (3.5) is strictly increasing and continuous in L. Since the left-hand side is constant in L, there can be at most one L solving this equation.

The right-hand side is equal to 0 at $L=0$. The limit of the left-hand side of

Equation (3.5) as $L \to \infty$ is

$$\begin{aligned}
\text{Limit} &\equiv \lim_{L\to\infty} (1-\rho\phi)\lambda L h(L) + (\rho-(1+r)\beta\tau)\phi \int_1^{\frac{1}{1-\lambda}} Lh(Lx)\frac{1-F(Lx)}{1-F(L)}\mathrm{d}x \\
&= (1-\rho\phi)\lambda \lim_{L\to\infty} Lh(L) + (\rho-(1+r)\beta\tau)\phi \int_1^{\frac{1}{1-\lambda}} \lim_{L\to\infty} Lh(Lx)\frac{1-F(Lx)}{1-F(L)}\mathrm{d}x \\
&= (1-\rho\phi)\lambda \lim_{L\to\infty} Lh(L) + (\rho-(1+r)\beta\tau)\phi \int_1^{\frac{1}{1-\lambda}} \frac{1}{x}\lim_{L\to\infty} Lh(L)\mathrm{d}x \\
&= (1-\rho\phi)\lambda \lim_{L\to\infty} Lh(L) + (\rho-(1+r)\beta\tau)\phi \ln\frac{1}{1-\lambda}\lim_{L\to\infty} Lh(L) \\
&= \left[(1-\rho\phi)\lambda + (\rho-(1+r)\beta\tau)\phi\ln\frac{1}{1-\lambda}\right]\lim_{L\to\infty} Lh(L)
\end{aligned}$$

We then have

$$\begin{aligned}
\text{Limit} &= \left[(1-\rho\phi)\lambda + (\rho-(1+r)\beta\tau)\phi\ln\frac{1}{1-\lambda}\right]\lim_{L\to\infty} Lh(L) \\
&> \left[(1-\rho\phi)\lambda + (\rho-\tau)\phi\ln\frac{1}{1-\lambda}\right]\lim_{L\to\infty} Lh(L) \\
&> (1-\rho\phi)\lambda \lim_{L\to\infty} Lh(L) \\
&> (1-\phi)\lambda \lim_{L\to\infty} Lh(L) > 1-(1+r)\beta
\end{aligned}$$

where the final inequality follows from Assumption 3.1.2. Since the last expression is equal to the left-hand side of Equation (3.5), a solution to Equation (3.5) exists. Finally, since the objective is decreasing in L as $L \to \infty$, we can rule out the case that $F(L) = 1$. Therefore, the program has a unique and finite solution. ■

We can apply this result to establish the uniqueness of contracts satisfying various requirements. For example, consider full information contracts:

Definition 3.A.1 (Competitive Full Information Contracts) *A* competitive safe contract *is a non-screening contract* $(T, L, 0)$ *whose transfer and loan size solve*

$$V^{safe} = \max_{T,L} \quad T + \beta U(L, 0)$$
$$s.t. \quad T \leq \frac{1}{1+r}\Pi(L, 0) \qquad \text{(PC)}$$

Similarly, a competitive risky contract *is a non-screening contract* $(T, L, 0)$ *whose transfer and loan size solve*

$$V^{risky} = \max_{T,L} \quad T + \beta U(L, \phi)$$
$$s.t. \quad T \leq \frac{1}{1+r}\Pi(L, \phi) \qquad \text{(PC)}$$

Corollary 3.A.1 (Unique Competitive Full Information Contracts) *Let both of Assumption 3.1.1 and Assumption 3.1.2 hold. Then each full information contract is unique. Denote the competitive safe contract by* $(T^{safe}, L^{safe}, 0)$ *and the competitive risky contract by* $(T^{risky}, L^{risky}, 0)$.

Proof. The result follows from Lemma 3.A.2 for the case where $\rho = 0$ and $\tau = 0$, and the case where $\rho = 1$ and $\tau = 1$. ■

Similarly, there are unique competitive pooling contracts across risk indices.

Definition 3.A.2 (Competitive Pooling Contract) *A* competitive pooling contract *among borrowers with risk index of* ρ *is a non-screening contract* $(T, L, 0)$ *whose transfer and loan size solve*

$$V^{pool}(\rho) = \max_{T,L} \quad T + \beta U(L, 0)$$
$$s.t. \quad T \leq \frac{1}{1+r}\Pi(L, \rho\phi) \qquad \text{(PC)}$$

Corollary 3.A.2 (Unique Competitive Pooling Contract) *Let Assumption 3.1.1 and*

Assumption 3.1.2 hold. Then there is a unique competitive pooling contract for each $\rho \in [0,1]$. *Denote its transfer by* $T^{pool}(\rho)$ *and its loan size by* $L^{pool}(\rho)$. *Moreover,* $T^{pool}(0) = T^{safe}$, $L^{pool}(0) = L^{safe}$, *and* $V^{pool}(\rho)$ *is strictly decreasing, continuous, and convex in* ρ.

Proof. For each ρ, there is a unique competitive pooling contract due to Lemma 3.A.2 for the case where $\tau = 0$. The pooling contract for $\rho = 0$ is identical to the full information competitive safe contract because both contracts correspond to the case when $\rho = 0$ and $\tau = 0$. Finally, note that

$$V^{pool}(\rho) = \max_L \quad \frac{1}{1+r}\Pi(L, \rho\phi) + \beta U(L, 0)$$

and Π is affine in its second argument. As a result, by the envelope theorem $V^{pool}(\rho)$ is continuous and differentiable in ρ with

$$\frac{\partial}{\partial\rho}V^{pool}(\rho) = \frac{1}{1+r}\Pi_2(L^{pool}(\rho), \rho\phi)\phi = -\phi\int_{L^{pool}(\rho)}^{\frac{L^{pool}(\rho)}{1-\lambda}}(L^{pool}(\rho) - (1-\lambda)z)F'(z)\mathrm{d}z < 0.$$

Finally, as the maximum of a collection of affine functions, $V^{pool}(\rho)$ is convex. Specifically, for any $x \in [0,1]$ and $\rho, \rho' \in [0,1]$ we have

$$\begin{aligned}
&xV^{pool}(\rho') + (1-x)V^{pool}(\rho) \\
&= x\max_L\left\{\frac{1}{1+r}\Pi(L, \rho'\phi) + \beta U(L,0)\right\} \\
&\qquad\qquad + (1-x)\max_L\left\{\frac{1}{1+r}\Pi(L, \rho\phi) + \beta U(L,0)\right\} \\
&\geq \max_L \quad x\frac{1}{1+r}\Pi(L, \rho'\phi) + (1-x)\frac{1}{1+r}\Pi(L, \rho\phi) + \beta U(L,0) \\
&= \max_L \quad \frac{1}{1+r}\Pi(L, (x\rho' + (1-x)\rho)\phi) + \beta U(L,0) \\
&= V^{pool}(x\rho' + (1-x)\rho).
\end{aligned}$$

■

The competitive screening contracts for each screening technology are also unique.

Definition 3.A.3 (Competitive High-Documentation Contract) *A high-documentation contract* $(T, L, 1)$ *whose transfer and loan size solve*

$$
\begin{aligned}
V^{doc} = \max_{T,L} \quad & T + \beta U(L, 0) \\
s.t. \quad & T \le \frac{1}{1+r}\Pi(L, 0) - \kappa_m
\end{aligned}
\tag{PC}
$$

is a competitive high-documentation contract.

Corollary 3.A.3 (Unique Competitive Screening Contract) *Let Assumption 3.1.1 and Assumption 3.1.2 hold. Then there is a unique competitive high-documentation contract. Denote this contracts by* $(T^{doc}, T^{doc}, 1)$. *These quantities are*

$$T^{doc} = T^{safe} - \kappa, \quad and \quad L^{doc}) = L^{safe}.$$

Proof. Fix $\rho \in [0, 1]$. Let $\tilde{T} = T + \kappa$. The program is equivalent to solving

$$
\begin{aligned}
\max_{\tilde{T},L} \quad & \tilde{T} + \beta U(L, 0) \\
\text{s.t.} \quad & \tilde{T} \le \frac{1}{1+r}\Pi(L, 0)
\end{aligned}
\tag{PC}
$$

We can therefore apply Lemma 3.A.2 for the case where $\rho = 0$ and $\tau = 0$, implying that the unique solution to this program must have $\tilde{T} = T^{safe}$ and $L = L^{safe}$. Therefore, $T = \tilde{T} - \kappa = T^{safe} - \kappa$. ■

Definition 3.A.4 (Competitive High-Downpayment Contract) *Given net home price of* $P - R$, *a* competitive high-downpayment contract *is a low-documentation contract* *whose transfer and loan size solve the following program*

$$V^{down}(P-R) = \max_{T,L} \quad T + \beta U(L,0)$$

$$s.t. \quad T + \beta U(L,\phi) \leq \max\{P-R, V^{risky}\} \tag{IC}$$

$$T \leq \frac{1}{1+r}\Pi(L,0) \tag{PC}$$

Lemma 3.A.3 (Unique Competitive Separating Contracts) *Let Assumption 3.1.1 and Assumption 3.1.2 hold. Then there is a unique high-downpayment contract. Denote its transfer by $T^{down}(P-R)$ and its loan size by $L^{down}(P-R)$.*

Proof. The PC constraint must bind at any solution. Suppose that the IC constraint is slack. Then L must solve

$$\max_{L} \quad \frac{1}{1+r}\Pi(L,0) + \beta U(L,0)$$

this program has a unique solution by Lemma 3.A.2 for the case where $\rho = 0$ and $\tau = 0$. Denote it by L^{safe}. Define $V^* = \frac{1}{1+r}\Pi(L^{safe},0) + \beta U(L^{safe},\phi)$. We can see that if $V^* < P-R$, then we must indeed have a slack IC constraint in the original program, and if $V^* = P-R$, then L^{safe} remains feasible in the original program. In either case L^{safe} must solve the original program. Otherwise, when $V^* > P-R$, L^{safe} is not in the feasible set of the original program. In this case, the IC constraint must bind and the program reduces to

$$\max_{L \in \{L | \frac{1}{1+r}\Pi(L,0) + \beta U(L,\phi) = P-R\}} \quad \max\{P-R, V^{risky}\} + \beta U(L,0) - \beta U(L,\phi)$$

Due to Lemma 3.1.1, the objective of this program is strictly decreasing in L. The solution must be the smallest value of L such that the IC constraint binds. Note that

$$\begin{aligned}\frac{1}{1+r}\Pi_1(L,0)+\beta U_1(L,\phi) &= \frac{1}{1+r}\left[1-\lambda LF'(L)-F(L)\right]\\ &\qquad +\beta\left[-1+(1-\phi)F(L)+\phi F\left(\frac{L}{1-\lambda}\right)\right]\\ &> \frac{1}{1+r}\left[1-\lambda LF'(L)-F(L)\right]+\beta\left[-1+F(L)\right]\\ &= \frac{1}{1+r}\Pi_1(L,0)+\beta U_1(L,0)\end{aligned}$$

and we have $\frac{1}{1+r}\Pi_1(L,0)+\beta U_1(L,\phi)>0$ for $L<L^{safe}$. Since $\frac{1}{1+r}\Pi(0,0)+\beta U(0,\phi)<\max\{P-R,V^{risky}\}$ and $\frac{1}{1+r}\Pi(L^{safe},0)+\beta U(L^{safe},\phi)>\max\{P-R,V^{risky}\}$, there is then a unique solution for $L<L^{safe}$. Therefore, there is a unique solution to the original program. Denote this solution by $L^{down}(P-R)$ and define $T^{down}(P-R)=\frac{1}{1+r}\Pi(L^{down}(P-R),0)$. ■

We can now use these definitions and results to simultaneously prove Proposition 3.1.1 and Proposition 3.1.2. The proof proceeds by showing that only these competitive contracts can be offered in any equilibrium, and any equilibrium where lenders randomly reject borrowers must be unstable.

Proof of Proposition 3.1.1 and Proposition 3.1.2.

Let Assumption 3.1.1 and Assumption 3.1.2 hold. Consider when $P-R>V^{safe}$. Since V^{safe} is an upper bound on both safe borrower and risky borrower valuations for any contract which gives lenders non-negative profits, all borrowers have a strictly dominant strategy to become renters for any set of contract offers. As a result, the set of contracts offered is indeterminate, and all borrowers become renters.

Now suppose that $P-R\leq V^{safe}$. I first show that there is a unique strictly dominant competitive contract for almost all risk indices, and characterize the sets of ρ-values where each competitive contract strictly dominates. Then, I show that equilibrium is unique for ρ-values in each of these sets.

Consider whether or not the IC constraint for the competitive high-downpayment contract binds.

1. Suppose that $V^* \leq P - R < V^{safe}$. In this case the IC constraint for the competitive separating contract doesn't bind and so $T^{down}(P-R) = T^{safe}$, $L^{down}(P-R) = L^{safe}$, and $V^{down}(P-R) = V^{safe}$. Since $V^{safe} > V^{pool}(\rho)$ for all $\rho \in (0,1]$ and $V^{safe} > V^{doc}$, the competitive high-downpayment contract strictly dominates both the competitive pooling and competitive high-documentation contracts for almost all borrowers.

2. Suppose that $P - R < V^*$. Now, the IC constraint binds and so $V^{down}(P-R) < V^{safe}$. Let $V^{screen}(P-R) = \max\{V^{down}(P-R), V^{doc}\} \leq V^{safe}$. Since $V^{pool}(0) = V^{safe}$, we may have a $\rho > 0$ such that $V^{pool}(\rho) = V^{screen}(P-R)$. If $V^{pool}(1) < V^{screen}(P-R)$ then there is a unique solution for $\rho \in (0,1)$. Otherwise, we have $V^{pool}(\rho) > V^{screen}(P-R)$ for $\rho \in [0,1)$. Take ρ^* to be the solution in the former case, and otherwise take $\rho^* = 1$. We then have that the competitive pooling contract strictly dominates for $\rho \in [0,\rho^*)$. For $\rho \in (\rho^*,1]$ either the competitive high-downpayment contract or the competitive high-documentation contract will dominate depending on whether $V^{down}(P-R) > V^{doc}$, $V^{down}(P-R) = V^{doc}$, or $V^{down}(P-R) < V^{doc}$.

These results establish that for almost all ρ-values from 0 to 1, one of the competitive contracts strictly dominates. I next show that within each region where a specific type of competitive contract dominates, that contract must be offered in equilibrium.

First consider ρ such that the competitive high-downpayment contract dominates. Consider the sub-game beginning with a borrower's decision to apply for contracts. Let C_i denote the set of contract offers made to individual i. Suppose C_i does not include the competitive high-downpayment contract. I will show that lenders can make a profitable deviation by introducing a contract of the form $(T^{down}(P-R)-\varepsilon, T^{down}(P-R), 0)$ for $\varepsilon > 0$ sufficiently small along side the full information contract for risky borrowers.

Let $\varepsilon > 0$ and consider equilibrium in the sub-game starting with borrower application choices when the contract offer set is

$$\{(T^{down}(P-R)-\varepsilon, T^{down}(P-R), 0), (T^{risky}, L^{risky}, 0)\} \cup C_i.$$

There exist two types of equilibrium. In one type, enough risky borrowers apply for the new contract that it becomes unprofitable, and lenders reject all applications. I show that these equilibria are unstable. In the second type of equilibrium, risky borrowers never apply for the new contract, and lenders make a strictly positive profit on the new contract. This equilibrium is unique and stable. Selecting this stable equilibrium to the sub-game, we conclude that lenders can make a profitable deviation by introducing the new contract.

Let θ denote the anticipated rejection rate when i applies to the new contract. If $\tau_i = 1$, then their expected value from applying to contract $(T^{down}(P-R)-\varepsilon, T^{down}(P-R), 0)$ is

$$\begin{aligned}(1-\theta)(T^{down}(P-R) - \varepsilon + \beta U(T^{down}(P-R), \phi)) + \theta(P-R) \\ = (1-\theta)(\max\{V^{risky}, P-R\} - \varepsilon) + \theta(P-R)\end{aligned}$$

The left hand side is strictly less than $\max\{V^{risky}, P-R\}$ whenever $\theta < 1$ or $\theta = 1$ and $V^{risky} > P-R$. In either case, we have that applying for contract $(T^{down}(P-R)-\varepsilon, T^{down}(P-R), 0)$ is strictly dominated. If $\theta = 1$ and $P-R \geq V^{risky}$, then risky borrowers may apply. However, any equilibrium with $\theta = 1$ is unstable because a small reduction of θ will imply that applying is strictly dominated. We can therefore focus on the case where $\theta < 1$.

In this case, risky borrowers will never apply, and any applicants must be safe borrowers. Lenders then have expected profit of

$$-(T^{down}(P-R)-\varepsilon)+\frac{1}{1+r}\Pi(T^{down}(P-R),0)=\varepsilon>0$$

on any accepted application. Since they have a strictly positive profit, they must accept all applications so that $\theta=0$. Anticipating this acceptance, borrowers with $\tau_i=0$ get value from applying of

$$T^{down}(P-R)-\varepsilon+\beta U(T^{down}(P-R),0))=V^{down}(P-R)-\varepsilon$$

Since $V^{down}(P-R)$ is strictly larger than $V^{pool}(\rho)$ and V^{doc}, and both these values represent the best case for pooling and high-documentation contracts, by taking ε sufficiently small, we can conclude that the new contract strictly dominates any contract in C_i. As a result, we can conclude that after introducing the new contract, there is a unique stable equilibrium in the following sub-game and in this new equilibrium, risky borrowers do apply for the new contract while all safe borrowers do apply for the new contract, all applications are accepted, and lenders make a strictly positive expected profit. Given this profitable deviation, we can conclude that whenever $P-R\leq V^{safe}$, and ρ is in the region where the competitive high-downpayment contract dominates, the competitive separating contract must be offered.

Now, suppose that the competitive high-downpayment contract is in C_i. Zero profits requires that no risky borrowers apply for the contract, but also implies that lenders are indifferent between accepting and rejecting applicants. This indifference might lead to multiple equilibria because lenders may randomly reject applications. Suppose that they reject a fraction $\theta>0$. I show that another lender could introduce a contract of the form $(T^{down}(P-R)-\varepsilon, L^{down}(P-R),0)$, capture the market for safe borrowers, and make a strictly positive profit. Safe borrowers get expected value from the competitive high-downpayment contract of

$$(1-\theta)V^{down}(P-R)+\theta(P-R)<V^{safe}$$

We can then choose an $\varepsilon>0$ such that $\varepsilon<V^{safe}-[(1-\theta)V^{down}(P-R)+\theta(P-R)]$. Given that risky borrowers previously didn't use the high-downpayment contract, the new contract must be strictly dominated from their perspective since it has value reduced by ε. As a result, in the equilibrium following the introduction of the new contract, lenders must believe that only safe borrowers apply. They then get expected profit of

$$-(T^{down}(P-R)-\varepsilon)+\frac{1}{1+r}\Pi(L^{down}(P-R),0)=\varepsilon>0$$

and so safe borrowers anticipate that they will be accepted if they apply for the new contract. Although they get a lower transfer, they no longer face a chance of rejection and strictly prefer the new contract because

$$V^{down}(P-R)-\varepsilon>(1-\theta)V^{down}(P-R)+\theta(P-R).$$

Therefore, all safe borrowers apply for the new contract, get accepted with probability 1, and lenders make a strictly positive profit. Therefore, we can conclude that in any equilibrium, lenders must accept all applicants for the competitive high-downpayment contract (so that $\theta=0$).

Using the same line of reasoning for high-documentation and pooling contracts, we get that for almost all ρ-values and in any stable equilibrium, the strictly dominating competitive contract must be offered and used by all safe borrowers. Risky borrowers use either pooling contracts or their full-information contract. Since the identity of the lender offering specific contracts is irrelevant, this stable equilibrium is unique up to the identity of which lender makes each contract offer. It is also unique up to the deletion of un-traded contracts.

■

3.B Extension: Conversion of Housing to Rental Units

To capture conversion of housing into rental units, I extend the absentee landlord's production of rental services to incorporate both houses as well as apartment buildings. I assume that the production of rental services uses both housing and apartments. The production function is

$$S_t^R = \mu A_{t-1}^{\gamma} (H_t^R)^{1-\gamma}.$$

where S_t^R denotes rental services, A_{t-1} denotes the stock of apartments from the previous period, and H_t^R denotes the current stock of homes owned by landlords. These stocks evolve as

$$H_t^R = (1-\delta_H)H_{t-1}^R + I_t^R, \quad \text{and} \quad A_t = (1-\delta_A)A_{t-1} + I_t^A$$

where I_t^A is investment in new apartments. This specification implies that landlords produce apartment buildings one-for-one from numeraire, which normalizes the units of the stock of apartments.

Given an opportunity cost of funds of $1+r$, the landlord's first order condition for investment in housing is

$$P_t = R_t(1-\gamma)\mu \left(\frac{A_{t-1}}{H_t^R}\right)^{\gamma} + \frac{1-\delta_H}{1+r}\mathbb{E}_t P_{t+1}.$$

This condition states that the converted-house-to-apartment ratio is an increasing function of the rental rate relative to the user cost of housing:

$$H_t^R = A_{t-1}\left((1-\gamma)\mu\frac{R_t}{P_t - \frac{1}{1+r}\mathbb{E}_t P_{t+1}}\right)^{\frac{1}{\gamma}}$$

Their first order condition for investment in apartments is

$$r + \delta_A = \mathbb{E}_t R_{t+1}\gamma\mu\left(\frac{H_{t+1}^R}{A_t}\right)^{1-\gamma}$$

where δ_A is the depreciation rate on apartment buildings. This condition states that landlords equate the expected marginal revenue product of apartment investment to their cost of funds plus depreciation.

3.C Relating Income Risk to Labor Market Fundamentals

This appendix presents a simple extension of the baseline model in 3.3 in which borrower income risk comes from employment risk due to search in the labor market.

There is a large family of workers who are either employed or unemployed in each period. This family provides imperfect income insurance to its members. When an individual worker is unemployed, the family will help them to cover their debt payments with chance $1-\bar{\phi}$. As a result, an individual worker who is unemployed will be forced into default with chance $\bar{\phi}$. When they are employed, they are never forced to default on their mortgage.

Employed workers separate from their current employer with chance s at the end of each period. Individual workers learn whether or not they will keep their job between t and $t+1$ prior to signing mortgage contracts at time t. Lenders cannot observe their separation, and so there is asymmetric information between workers and lenders about individual worker's labor search status.

Both recently separated workers and currently unemployed workers search for a new jobs between periods, while workers who retain their job between periods do not need to search for work. The former have some risk of having no income, while the later have income for sure. As a result, we can interpret the group of workers searching for work as the risky borrowers in the baseline model, and the group with stable employment as the safe borrowers.

I assume that current employment status is observable to lenders. As a result, the population is split into a group of potentially risky borrowers — those who were employed today and might have separated from their current employer — and a group of known risky borrowers — those who didn't have work this period and therefore must be searching for work. Among borrowers with current employment, a fraction s are searching for work, and so we have $\rho_{it} = s$ among workers with current employment. The fraction searching among workers without current employment is 1, so $\rho_{it} = 1$. This models maps precisely into the setup in section 3.3 when we set $\rho = s$, and let the proportion of potentially risky borrowers in baseline model, n, be time varying and equal to the employment rate in this model, N_t.

The ex-post chance of having no income among workers searching for work is

$$\phi_{t+1} = \bar{\phi}\mathbf{P}_{t+1}[E_{i,t+1} = 0 \mid \text{Search}_{it}] = \bar{\phi}(1 - f_{t+1})$$

where f_{t+1} is the ex-post job-finding rate. Borrowers who do not need to search for work have income for sure.

I assume that competitive firms post job vacancies and are randomly matched to workers. Define θ_t as labor market tightness

$$\theta_t = \frac{V_t}{1-(1-s)N_{t-1}}$$

where V_t is total vacancies posted at time t and N_{t-1} is the employment rate in the prior period.

Let M denote the matching function, mapping total vacancy postings and total number of searching workers to a total number of matched workers to vacancies. Assume it is homogenous of degree one, so we can write the job-finding rate and hiring rate by

$$f_t = M(\theta_t, 1), \quad \text{and} \quad h_t = M(1, 1/\theta_t).$$

Assume that posting a vacancy costs a firm κ_v units of numeraire. Firms operate a constant returns to scale production function in labor. Let A_t denote the marginal product of labor in this technology, and let W_t denote the wage paid by firms. Then the job creation condition for firms is

$$\frac{\kappa_v}{h_t} = A_t - W_t + \frac{1-s}{1+r}\mathbb{E}_t \frac{\kappa_v}{h_{t+1}}$$

where I assume that firms face the same exogenous opportunity cost of funds as lenders. The quantity κ_v/h_t represents the effective cost of adding a worker to the labor force.

Following Haefke et al. (2013) and Michaillat (2012), I assume that wages are rigid and linked to the marginal product of labor as

$$W_t = A_t^{\gamma}$$

where $\gamma < 1$. As a result, we can solve for the current hiring rate as a function of expectations of future productivity

$$\frac{\kappa_v}{h_t} = \sum_{j=0}^{\infty} (1-s)^j \left(\frac{1}{1+r}\right)^j (A_{t+j} - A_{t+j}^{\gamma})$$

Since hiring rates pin down labor market tightness and therefore job finding rates, we get that the job finding rate depends on expectations of the path of productivity. In turn, the income risk of workers who are searching for work depends on the ex-post job-finding rate. Since productivity and the opportunity cost of funds are exogenous, the job finding rate and income risk are both effectively exogenous.

3.D Sketch of Solution Algorithm

The model of Section 3.3 is a infinite horizon rational expectations model with endogenous regime-switching. In particular, the distribution of the credit regime, conditional on the state of the economy, is an equilibrium object. I apply a solution algorithm based on the approach in Chapter 2 to jointly solve for this conditional regime distribution and the associated regime-conditional policy functions.

In general the approach applies to models whose equilibrium conditions can be expressed as

$$\begin{aligned} 0 &= \mathbb{E}_t f_{S_{t+1},S_t,S_{t-1}}(Y_{t+1}, Y_t, X_{t+1}, X_t) \\ X_{t+1} &= h_{S_t,S_{t-1}}(X_t, Y_t) + \Sigma \varepsilon_{t+1} \\ S_t &= \sigma_{S_{t-1}}(Y_t, X_t) \end{aligned}$$

where Y_t is an endogenous vector of control variables, X_t is an endogenous vector of state variables, S_t is an endogenous regime variable taking values in $\{1, \ldots, N_S\}$, ε_t is an exogenous vector of *iid* shocks, and Σ is a matrix mapping shocks into future states.

The function f captures the forward looking equilibrium condition of the model, the function h captures the endogenous evolution of the state variables, and the function σ assigns a regime to a particular equilibrium outcome at time t. Note that Y_t and S_t are simultaneously determined because S_t and expectations of S_{t+1} can influence Y_t through the forward looking equilibrium conditions and realizations of Y_t are linked to S_t through the regime assignment equation.

I focus on a minimum-state variable equilibrium concept that accounts for this simultaneity between the regime and contemporaneous outcomes.

Definition 3.D.1 *A* regime-switching equilibrium *is a collection of measurable functions* $\{g_{s,s_-}(x)\}_{s=1}^{N_S}$ *and a conditional regime distribution* $Q_{s',s}(x)$ *(measurable in x) such that*

1. *(**Rational Expectations**) For each* (s,s_-) *and* x

$$0 = \int \sum_{s_+=1}^{N_S} f_{s_+,s,s_-}(g_{s_+,s}(h_{s,s_-}(x, g_{s,s_-}(x)) + \Sigma\varepsilon), g_{s,s_-}(x), h_{s,s_-}(x, g_{s,s_-}(x)) + \Sigma\varepsilon, x)$$
$$\times Q_{s_+,s}(h_{s,s_-}(x, g_{s,s_-}(x)) + \Sigma\varepsilon) f_\varepsilon(\varepsilon) d\varepsilon$$

2. *(**Regime Consistency**) For each* (s,s_-) *and* x*, if* $Q_{s,s_-}(x) > 0$ *then* $s = \sigma(g_{s,s_-}(x), x)$.

3. *(**Fundamental Volatility**) For each* (s,s_-) *and* x*,* $Q_{s,s_-}(x) \in \{0,1\}$.

The first condition requires that the policy functions solve the model given rational expectations. Implicitly, this definition requires that out-of-equilibrium outcomes (in regimes that never occur given $S_{t-1} = s_-$ and $X_t = x$) also solve the model equations, though a more general definition need not make this particular equilibrium selection. In this context, the selection is inconsequential. The second condition requires that regimes can only occur if outcomes in that regime are consistent with the regime assignment equation. The third condition requires that all randomness in outcomes can be attributed to variation in the state variables (say due to fundamental shocks) — a requirement which rules out sunspots.

The solution algorithm uses backwards induction and perturbation based approximation to approximate such an equilibrium. At iteration n, let $g^{(n-1)}_{s,s_-}(x)$ and $Q^{(n-1)}_{s,s_-}(x)$ denote the policy functions and regime distribution from the previous iteration. Consider finding an updated policy function, $g^{(n)}_{s,s_-}$, solving

$$0 = \int \sum_{s_+=1}^{N_S} f_{s_+,s,s_-}(g^{(n-1)}_{s_+,s}(h_{s,s_-}(x, g^{(n)}_{s,s_-}(x)) + \Sigma\varepsilon), g^{(n)}_{s,s_-}(x), h_{s,s_-}(x, g^{(n)}_{s,s_-}(x)) + \Sigma\varepsilon, x)$$
$$\times Q^{(n-1)}_{s_+,s}(h_{s,s_-}(x, g^{(n)}_{s,s_-}(x)) + \Sigma\varepsilon) f_\varepsilon(\varepsilon) d\varepsilon.$$

Introduce a perturbation parameter of $\eta \in [0,1]$ as

$$0 = \int \sum_{s_+=1}^{N_S} f_{s_+,s,s_-}(g^{(n-1)}_{s_+,s}(h_{s,s_-}(x, g^{(n)}_{s,s_-}(x,\eta),1) + \eta\Sigma\varepsilon + (1-\eta)\bar{v}^{(n-1)}_{s_+,s,s_-}),$$
$$g^{(n)}_{s,s_-}(x,\eta), h_{s,s_-}(x, g^{(n)}_{s,s_-}(x,\eta)) + \eta\Sigma\varepsilon + (1-\eta)\bar{v}^{(n-1)}_{s_+,s,s_-}, x)$$
$$\times Q^{(n-1)}_{s_+,s}(h_{s,s_-}(x, g^{(n)}_{s,s_-}(x)) + \Sigma\varepsilon) f_\varepsilon(\varepsilon) d\varepsilon.$$

where $\bar{v}^{(n-1)}_{s_+,s,s_-}$ is a regime-transition conditional approximation node.

We can approximate the policy function $g^{(n)}_{s,s_-}$ in two steps. First, consider the case with $\eta = 0$, and solve for a nodes $\bar{y}^{(n)}_{s,s_-}$ such that the equation is satisfied for a given point in the state space of $x = \bar{x}^{(n)}_{s,s_-}$. Then, since the model is solved for $y = \bar{y}^{(n)}_{s,s_-}$ for $x = \bar{x}^{(n)}_{s,s_-}$ and $\eta = 0$, we can apply the implicit function theorem to get a first order approximation of $g^{(n)}_{s,s_-}$.

Next we can update the conditional distribution of the regime to $Q^{(n)}$ such that it is consistent with the assignment of regimes implied by the approximation to $g^{(n)}_{s,s_-}$ and the regime function σ_{s,s_-}. I do so by keeping track of conditional regime transitions on a grid of ε-values (transitions to s_+ conditional on $S_t = s$ and $S_{t-1} = s_-$). Also, I update the shock approximation nodes, $\bar{v}^{(n)}_{s_+,s,s_-}$, to equal the expected value of $\Sigma\varepsilon'$ condition on the regime transition. Finally, I allow the approximation nodes $\bar{x}^{(n)}_{s,s_-}$ to slowly adjust based on the average predicted outcomes for the state at each step of the algorithm. I iterate on these approximation and update steps until the algorithm converges.

Bibliography

K. Adam and M. Grill. Optimal sovereign default. Discussion Papers 09/2013, Deutsche Bundesbank, Research Centre, 2013.

M. Adelino, A. Schoar, and F. Severino. Loan originations and defaults in the mortgage crisis: The role of the middle class. *Review of Financial Studies*, page hhw018, 2016.

S. Adjemian, H. Bastani, F. Karamé, M. Juillard, J. Maih, F. Mihoubi, G. Perendia, J. Pfeifer, M. Ratto, and S. Villemot. Dynare: Reference manual version 4. Technical report, CEPREMAP, 2014.

S. Albanesi, G. De Giorgi, J. Nosal, and M. Ploenzke. Credit growth and the financial crisis: A new narrative. Technical report, Mimeo, Ohio State University, 2016.

C. Arellano. Default risk and income fluctuations in emerging economies. *American Economic Review*, 98(3):690–712, 2008.

S. B. Aruoba and F. Schorfheide. Macroeconomic dynamics near the zlb: A tale of two equilibria. Technical report, 2012.

J. S. Banks and J. Sobel. Equilibrium selection in signaling games. *Econometrica: Journal of the Econometric Society*, pages 647–661, 1987.

N. Barberis, R. Greenwood, L. Jin, and A. Shleifer. X-capm: An extrapolative capital asset pricing model. *Journal of Financial Economics*, 115(1):1–24, 2015.

J. Barthélemy and M. Marx. State-Dependent Probability Distributions in Non Linear Rational Expectations Models. Technical report, 2013.

P. Beaudry, D. Galizia, and F. Portier. Reviving the limit cycle view of macroeconomic fluctuations. 2015.

D. Berger, V. Guerrieri, G. Lorenzoni, and J. Vavra. House prices and consumer spending. Technical report, National Bureau of Economic Research, 2015.

B. S. Bernanke, M. Gertler, and S. Gilchrist. The financial accelerator in a quantitative business cycle framework. *Handbook of macroeconomics*, 1:1341–1393, 1999.

H. Bester. Screening vs. rationing in credit markets with imperfect information. *The American Economic Review*, 75(4):850–855, 1985.

F. Bianchi. Evolving monetary/fiscal policy mix in the united states. *The American Economic Review*, 102(3):167–172, 2012.

F. Bianchi. Regime switches, agents' beliefs, and post-world war ii us macroeconomic dynamics. *The Review of Economic Studies*, 80(2):463–490, 2013.

F. Bianchi and L. Melosi. Modeling the evolution of expectations and uncertainty in general equilibrium. Technical report, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania, 2012.

J. Bianchi and E. G. Mendoza. Optimal time-consistent macroprudential policy. NBER Working Papers 19704, National Bureau of Economic Research, Inc, Dec 2013.

S. Bigio. Endogenous liquidity and the business cycle. *The American Economic Review*, 105(6):1883–1927, 2015.

L. Bocola. The pass-through of sovereign risk. *Journal of Political Economy*, 124(4): 879–926, 2016.

F. Boissay, F. Collard, and F. Smets. Booms and banking crises. Technical report, Working Paper, 2013.

R. A. Braun, L. M. Körber, and Y. Waki. Some unpleasant properties of log-linearized solutions when the nominal rate is zero. Technical report, 2012.

M. Brzoza-Brzezina, M. Kolasa, and K. Makarski. A penalty function approach to occasionally binding credit constraints. Dynare Working Papers 27, CEPREMAP, June 2013.

C. Burnside, M. Eichenbaum, and S. Rebelo. Understanding booms and busts in housing markets. *Journal of Political Economy*, 124(4):1088–1147, 2016.

C. Busch, D. Domeij, R. Madera, and F. Guvenen. Asymmetric business cycle risk and government policy. Technical report, 2016.

J. Y. Campbell, S. Giglio, and P. Pathak. Forced sales and house prices. *The American Economic Review*, 101(5):2108–2131, 2011.

I.-K. Cho and D. M. Kreps. Signaling games and stable equilibria. *The Quarterly Journal of Economics*, pages 179–221, 1987.

I.-K. Cho and J. Sobel. Strategic stability and uniqueness in signaling games. *Journal of Economic Theory*, 50(2):381–413, 1990.

S. Cho. Characterizing markov-switching rational expectations models. Technical report, 2013.

S. Cho and A. Moreno. The Forward Solution for Linear Rational Expectations Models. Technical report, 2008.

D. Corbae and E. Quintin. Leverage and the foreclosure crisis. Technical report, National Bureau of Economic Research, 2013.

T. Davig and E. M. Leeper. Endogenous monetary policy regime change. NBER Working Papers 12405, National Bureau of Economic Research, Inc, Aug. 2006.

A. A. DeFusco, C. G. Nathanson, and E. Zwick. Speculative dynamics of prices and volume. Technical report, 2017.

G. Dell'Ariccia and R. Marquez. Lending booms and lending standards. *The Journal of Finance*, 61(5):2511–2546, 2006.

P. Dreesen, K. Batselier, and B. De Moor. Back to the roots: Polynomial system solving, linear algebra, systems theory. *Proc 16th IFAC Symposium on System Identification (SYSID 2012)*, pages 1203–1208, 2012.

R. Farmer and V. Khramov. Solving and Estimating Indeterminate DSGE Models. IMF Working Papers 13/200, International Monetary Fund, Oct. 2013.

R. E. Farmer, D. F. Waggoner, and T. Zha. Understanding markov-switching rational expectations models. *Journal of Economic Theory*, 144(5):1849–1867, 2009.

R. E. Farmer, D. F. Waggoner, and T. Zha. Minimal state variable solutions to markov-switching rational expectations models. *Journal of Economic Dynamics and Control*, 35(12):2150–2166, 2011.

J. Fernandez-Villaverde, G. Gordon, P. A. Guerron-Quintana, and J. Rubio-Ramirez. Nonlinear adventures at the zero lower bound. Working Paper 18058, National Bureau of Economic Research, May 2012.

F. Ferreira and J. Gyourko. A new look at the us foreclosure crisis: Panel data evidence of prime and subprime borrowers from 1997 to 2012. Technical report, National Bureau of Economic Research, 2015.

A. Foerster, J. Rubio-Ramirez, D. Waggoner, and T. Zha. Perturbation methods for markov-switching dsge model. Technical report, 2013.

A. T. Foerster. Monetary policy regime switches and macroeconomic dynamic. Technical report, 2013.

C. L. Foote, L. Loewenstein, and P. S. Willen. Cross-sectional patterns of mortgage debt during the housing boom: Evidence and implications. 2016.

W. T. Gavin, B. D. Keen, A. Richter, and N. Throckmorton. Global dynamics at the zero lower bound. *Federal Reserve Bank of St. Louis Working Paper Series*, (2013-007), 2013.

P. Gelain, K. J. Lansing, and G. J. Natvik. Explaining the boom-bust cycle in the us housing market: A reverse-engineering approach. 2015.

M. Gertler and P. Karadi. A model of unconventional monetary policy. *Journal of Monetary Economics*, 58(1):17–34, January 2011.

M. Gertler and N. Kiyotaki. Banking, liquidity and bank runs in an infinite horizon economy. 2013 Meeting Papers 59, Society for Economic Dynamics, 2013.

G. Gorton and G. Ordoñez. Good booms, bad booms. Technical report, National Bureau of Economic Research, 2016.

G. B. Gorton and G. Ordoñez. Collateral crises. Technical report, National Bureau of Economic Research, 2012.

D. L. Greenwald. The mortgage credit channel of macroeconomic transmission. *Available at SSRN 2735491*, 2016.

L. Guerrieri and M. Iacoviello. Occbin: A toolkit for solving dynamic models with occasionally binding constraints easily, 2014.

V. Guerrieri and H. Uhlig. Housing and credit markets: Booms and busts. Handbook of Macroeconomics, pages –. Elsevier, 2016.

C. Gust, J. D. López-Salido, and M. E. Smith. The empirical implications of the interest-rate lower bound. Technical report, CEPR Discussion Papers, 2012.

C. Haefke, M. Sonntag, and T. Van Rens. Wage rigidity and job creation. *Journal of monetary economics*, 60(8):887–899, 2013.

L. S. Hayre and M. Saraf. A loss severity model for residential mortgages. *The Journal of Fixed Income*, 18(2):5–31, 2008.

Z. He and A. Krishnamurthy. A macroeconomic framework for quantifying systemic risk. Working Paper Research 233, National Bank of Belgium, Oct 2012.

Z. He and A. Krishnamurthy. Intermediary asset pricing. *American Economic Review*, 103(2):732–70, 2013.

M. Hellwig. Some recent developments in the theory of competition in markets with adverse selection? *European Economic Review*, 31(1-2):319–325, 1987.

M. Iacoviello. House prices, borrowing constraints, and monetary policy in the business cycle. *American economic review*, pages 739–764, 2005.

W. Jiang, A. A. Nelson, and E. Vytlacil. Liar's loan? effects of origination channel and information falsification on mortgage delinquency. *Review of Economics and Statistics*, 96(1):1–18, 2014.

H. Jin and K. L. Judd. Perturbation methods for general dynamic stochastic models. *Manuscript, Stanford University*, 2002.

Ò. Jordà, M. Schularick, and A. M. Taylor. Leveraged bubbles. *Journal of Monetary Economics*, 76:S1–S20, 2015.

A. Justiniano, G. E. Primiceri, and A. Tambalotti. Credit supply and the housing boom. Technical report, 2015a.

A. Justiniano, G. E. Primiceri, and A. Tambalotti. Household leveraging and deleveraging. *Review of Economic Dynamics*, 18(1):3–20, 2015b.

A. Justiniano, G. E. Primiceri, and A. Tambalotti. The mortgage rate conundrum. Technical report, 2016.

G. Kaplan, K. Mitman, and G. Violante. Consumption and house prices in the great recession. Technical report, 2016.

J. Kim, S. Kim, E. Schaumburg, and C. A. Sims. Calculating and using second-order accurate solutions of discrete time dynamic equilibrium models. *Journal of Economic Dynamics and Control*, 32(11):3397 – 3414, 2008. ISSN 0165-1889.

N. Kiyotaki and J. Moore. Credit cycles. *The Journal of Political Economy*, 105(2): 211–248, 1997.

E. Kohlberg and J.-F. Mertens. On the strategic stability of equilibria. *Econometrica*, 54 (5):1003–1037, 1986.

A. Korinek and E. G. Mendoza. From sudden stops to fisherian deflation: Quantitative theory and policy implications. NBER Working Papers 19362, National Bureau of Economic Research, Inc, Aug. 2013.

A. W. Lo. The adaptive markets hypothesis. *The Journal of Portfolio Management*, 30 (5):15–29, 2004.

T. A. Lubik and F. Schorfheide. Computing sunspot equilibria in linear rational expectations models. *Journal of Economic Dynamics and Control*, 28(2):273–285, November 2003.

E. G. Mendoza. Sudden stops, financial crises, and leverage. *American Economic Review*, 100(5):1941–66, 2010.

E. G. Mendoza and K. A. Smith. Quantitative implications of a debt-deflation theory of sudden stops and asset prices. *Journal of International Economics*, 70(1):82 – 114, 2006. ISSN 0022-1996.

E. G. Mendoza and V. Z. Yue. A general equilibrium model of sovereign default and business cycles. *The Quarterly Journal of Economics*, 127(2):889–946, 2012.

A. Mian and A. Sufi. The consequence of mortgage credit expansion: Evidence from the us mortgage default crisis. *Quarterly Journal of Economics*, 2009.

A. Mian and A. Sufi. House prices, home equity–based borrowing, and the us household leverage crisis. *The American Economic Review*, 101(5):2132–2156, 2011.

A. Mian and A. Sufi. Household debt and defaults from 2000 to 2010: Facts from credit bureau data. Technical report, National Bureau of Economic Research, 2015.

A. R. Mian and A. Sufi. Household debt and defaults from 2000 to 2010: The credit supply view. 2016.

P. Michaillat. Do matching frictions explain unemployment? not in bad times. *The American Economic Review*, 102(4):1721–1750, 2012.

T. Nakata. Optimal fiscal and monetary policy with occasionally binding zero bound constraints. Technical report, 2012.

M. Piazzesi and M. Schneider. Housing and macroeconomics. *Handbook of Macroeconomics*, 2016.

A. Prestipino. Financial crises and policy. Technical report, 2014.

G. Ramey. D1 signaling equilibria with multiple signals and a continuum of types. *Journal of Economic Theory*, 69(2):508–531, 1996.

A. Saiz. The geographic determinants of housing supply. *The Quarterly Journal of Economics*, 125(3):1253–1296, 2010.

Y. Sannikov and M. Brunnermeier. A macroeconomic model with a financial sector. 2012 Meeting Papers 507, Society for Economic Dynamics, 2012.

Y. Sannikov and M. Brunnermeier. The i-theory of money. 2013 Meeting Papers 620, Society for Economic Dynamics, 2013.

S. Schmitt-Grohe and M. Uribe. Solving dynamic general equilibrium models using a second-order approximation to the policy function. *Journal of Economic Dynamics and Control*, 28(4):755–775, January 2004.

J. E. Stiglitz and A. Weiss. Credit rationing in markets with imperfect information. *The American economic review*, 71(3):393–410, 1981.

C. Wilson. A model of insurance markets with incomplete information. *Journal of Economic theory*, 16(2):167–207, 1977.