

Essays in Computational Macroeconomics and Finance

Author: Isaiah Hull

Persistent link: <http://hdl.handle.net/2345/bc-ir:104376>

This work is posted on [eScholarship@BC](#),
Boston College University Libraries.

Boston College Electronic Thesis or Dissertation, 2013

Copyright is held by the author, with all rights reserved, unless otherwise noted.

Boston College

The Graduate School of Arts and Sciences

Department of Economics

ESSAYS IN COMPUTATIONAL MACROECONOMICS AND FINANCE

a dissertation

by

ISAIAH HULL

submitted in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

May 2013

© copyright by ISAIAH JAMES HULL

2013

ESSAYS IN COMPUTATIONAL MACROECONOMICS AND FINANCE

Dissertation Abstract

by

ISAIAH HULL

Dissertation Committee:

PETER IRELAND (chair)

GEORG STRASSER

ROBERT MURPHY

This dissertation examines three topics in computational macroeconomics and finance. The first two chapters are closely linked; and the third chapter covers a separate topic in finance. Throughout the dissertation, I place a strong emphasis on constructing computational tools and modeling devices; and using them in appropriate applications.

The first chapter examines how a central bank's choice of interest rate rule impacts the rate of mortgage default and welfare. In this chapter, a quantitative equilibrium (QE) model is constructed that incorporates incomplete markets, aggregate uncertainty, overlapping generations, and realistic mortgage structure. Through a series of counterfactual simulations, five things are demonstrated: 1) nominal interest rate rules that exhibit cyclical behavior increase the average default rate and lower average welfare; 2) welfare can be substantially improved by adopting a modified Taylor rule that stabilizes house prices; 3) a decrease in the length of the interest rate cycle

will tend to increase the average default rate; 4) if the business and housing cycles are not aligned, then aggressive inflation targeting will tend to increase the mortgage default rate; and 5) placing a legal cap on loan-to-value ratios will lower the average default rate and lessen the intensity of extreme events. In addition to these findings, this paper also incorporates an important mechanism for default, which had not previously been included in the QE literature: default spikes happen when income falls and home equity is degraded at the same time. The paper concludes with a policy recommendation for central banks: if they wish to avoid crises where many households default simultaneously, they should either adopt a rule that generates interest rates with slow-moving cycles or use a modified Taylor rule that also targets house price growth.

The second chapter generalizes the solution method used in the first and compares it to more common techniques used in the computational macroeconomics literature, including the parameterized expectations approach (PEA), projection methods, and value function iteration. In particular, this chapter compares the speed and accuracy of the aforementioned modifications to an alternative method that was introduced separately by Judd (1998), Sutton and Barto (1998), and Van Roy et al. (1997), but was not developed into a general solution method until Powell (2007) introduced it to the Operations Research literature. This approach involves rewriting the Bellman equation in terms of the post-decision state variables, rather than the pre-decision state variables, as is done in standard dynamic programming applications in economics. I show that this approach yields considerable performance benefits over common global solution methods when the state space is large; and has the added benefit of not forcing modelers to assume a data generating process for shocks. In addition to this, I construct two new algorithms that take advantage of this approach to solve heterogeneous agent models.

Finally, the third chapter imports the SIR model from mathematical epidemiology; and uses it to construct a model of financial epidemics. In particular, the paper demonstrates how the SIR model can be microfounded in an economic context to make predictions about financial epidemics, such as the spread of asset-backed securities (ABS) and exchange-traded funds (ETFs), the proliferation of zombie financial institutions, and the expansion of financial bubbles and mean-reverting fads. The paper proceeds by developing the 1-host SIR model for economic and financial contexts; and then moves on to demonstrate how to work with the multi-host version of the model. In addition to showing how the SIR framework can be used to model economic interactions, it will also: 1) show how it can be simulated; 2) use it to develop and estimate a sufficient statistic for the spread of a financial epidemic; and 3) show how policymakers can impose the financial analog of herd immunity—that is, prevent the spread of a financial epidemic without completely banning the asset or behavior associated with the epidemic. Importantly, the paper will focus on developing a neutral framework to describe financial epidemics that can be either bad or good. That is, the general framework can be applied to epidemics that constitute a mean-reverting fad or an informational bubble, but ultimately yield little value and shrink in importance; or epidemics that are long-lasting and yield a new financial instrument that generates permanent efficiency gains or previously unrealized hedging opportunities.

Acknowledgements

I am very grateful to Professors Peter Ireland, Georg Strasser, and Robert Murphy for their generous advice and encouragement. They have not only helped me to refine my dissertation, but have also guided my development as a researcher and an economist. I must also thank Professor Andrew Beauchamp, who has been a constant source of helpful advice.

In addition to this, I would like to thank my classmates and friends—especially John O’Trakoun, Stacey Chan, and Mikhail Dmitriev. They have helped me to clarify and improve my ideas; and have been supportive throughout my time in graduate school.

Finally, I would like to thank my parents, James and Gale; my sisters, Tali, Noel, and Cherith; my mother-in-law, Kathy, and my sister-in-law, Susan, for their constant encouragement. I would especially like to thank my wife, Jamie, for her patience and unconditional support over the last five years. This dissertation is dedicated to her.

CHAPTER 1

INTEREST RATE RULES AND MORTGAGE DEFAULT

Isaiah Hull

ABSTRACT

This paper examines how a central bank's choice of interest rate rule impacts the rate of mortgage default and welfare. I do this by constructing a quantitative equilibrium (QE) model that incorporates incomplete markets, aggregate uncertainty, overlapping generations, and realistic mortgage structure. Through a series of counterfactual simulations, I demonstrate five things: 1) nominal interest rate rules that exhibit cyclical behavior increase the average default rate and lower average welfare; 2) welfare can be substantially improved by adopting a modified Taylor rule that stabilizes house prices; 3) a decrease in the length of the interest rate cycle will tend to increase the average default rate; 4) if the business and housing cycles are not aligned, then aggressive inflation targeting will tend to increase the mortgage default rate; and 5) placing a legal cap on loan-to-value ratios will lower the average default rate and lessen the intensity of extreme events. In addition to these findings, my model also incorporates an important mechanism for default, which had not previously been included in the QE literature: default spikes happen when income falls and home equity is degraded at the same time. Overall, my results suggest that the univariate time series properties of interest rates (i.e. wavelength, persistence, and variance) may play a substantial role in generating mass mortgage-default events. If a central bank wishes to avoid such crises, they should either adopt a rule that generates interest rates with slow-moving cycles or use a modified Taylor rule that also targets house price growth.

JEL Classification: E50, E52, C63, C68

Keywords: housing, monetary policy, Taylor rule, mortgage default, aggregate uncertainty, incomplete markets

1 Introduction

In July of 2000, the Federal Reserve initiated a series of rate cuts that lowered the effective federal funds rate (FFR) from a peak of 6.54% to 1% in late 2004, dropping it below 2% by the end of 2001. During this same period, house price growth jumped from an already high rate of 9.6% to 14.6%; and grew at an average of 6.8 percentage points per year faster than it did from 1988 to 2000. In late 2004 to mid-2007, interest rates shot up once again to 5.25%; and, with a lag, house price growth peaked at 15.9% and then dropped to -4.6% by mid-2007, falling further as the recession deepened. Not long after the house price drop, the default rate on first mortgages rose from a historically stable 1% to a peak of 5.39% in 2009.

After a casual glance at the data, we might conclude that the FFR cycle that spanned late 2000 to mid-2007 can safely be blamed for Great Recession—that is, it caused the twin housing and financial crises, bridged by a mortgage default spike. Indeed, when John Taylor addressed the Jackson Hole conference in 2007, that is exactly what he argued.¹ In this paper, I will not attempt to find the cause of the “Great Recession”; however, I will try to elucidate how monetary policy can play a role in preventing such crises. In particular, I ask: how does the central bank’s choice of interest rate rule impact the frequency and intensity of mortgage default?

To give the reader a sketch of the mechanisms that connect mortgage default crises and monetary policy, consider what happens during a typical interest rate cycle. First, assume that the real, risk-free interest rate falls sharply and remains low; and that these changes are at least weakly transmitted to mortgage interest rates. This will increase demand for housing by reducing the cost to borrow. If the increase in demand also pushes up house prices, then household equity positions will improve, allowing

¹See Figure 1 in the appendix for a comparison of the Taylor Rule-implied interest rate and the effective federal funds rate leading up to the crisis.

homeowners to borrow more (i.e. withdraw additional equity), which could increase the amount of debt that homeowners hold—and possibly increase household leverage. Finally, a drop in the risk free rate could divert marginal bank depositors to instead invest in capital, increasing output (income).

Now, assume that the first part of this cycle is followed by a sharp increase in interest rates—just as the prolonged period of low interest rates from 2001 to 2004 was followed by a sharp rise in the FFR. Suddenly, the effects of the housing boom will be reversed. Demand will collapse, pushing house prices down, and leaving households who extracted equity or bought homes near the trough of the cycle with negative equity. This will prevent households from withdrawing equity to smooth consumption or to refinance into lower interest rate loans should they become available. High interest rates will also divert investment away from capital, lowering income, and will push up adjustable-rate mortgage payments with a lag. The combination of low incomes, high mortgage payments, and negative equity positions will push up default rates, which will deteriorate financial intermediary balance sheets, and may cause a credit crunch.

Prior to the Great Recession, few DSGE models contained many of the aforementioned elements. Additionally, economists had only started to consider how to model multi-period mortgage structure realistically in general equilibrium.² And many of the macro housing models that did exist lacked heterogeneity, eliminating the effect of a prolonged house price rise completely (Jeske 2005).

In this paper, I attempt to contribute to the now-vibrant housing literature that has drawn inspiration from Aiygari (1993), and Krusell and Smith (1998). In particular, I construct a quantitative equilibrium model that incorporates incomplete markets, aggregate uncertainty, overlapping generations, price stickiness, credit-scoring,

²See Chambers, Garriga, and Schlagenhaut (2009) for a multi-period treatment of mortgages.

optimal default, inter-sectoral productivity correlation, and realistic mortgage structure. I then calibrate the model to match the cross-sectional and time series dimensions in the data; and run a series of counterfactual simulations under different interest rate rules to determine how they impact default rates and welfare. In addition to considering popular classes of interest rate rules (e.g. the Taylor Rule), I also look at more fundamental components of interest rate behavior by testing rules that generate cycles, but do not endogenously respond to other macrovariables, such as autoregressive rules and sine wave rules.

2 Related Literature

This paper contributes to three subliteratures. The first consists of papers that attempt to explain the Great Recession and financial crises in general. The second consists of macro papers that have attempted to integrate housing. And the final looks at the optimality of the Taylor rule and other monetary policy rules. I will review each of these literatures in order below.

2.1 Great Recession Causes Literature

Mian and Sufi (2010) find that household leverage yields considerable predictive power for the 2007-09 recession. They show that the household debt-to-disposable income ratio accurately predicts movements in aggregate variables, such as unemployment, consumer default, and house prices. They suggest that measures of household leverage could provide a well-grounded empirical basis for explaining macroeconomic fluctuations. This finding is consistent with my theme, which explores interest rate rules as the cause of household balance sheet deteriorations, which lead to increased default rates and related macroeconomic fluctuations.

In a separate paper, Mian and Sufi (2009) consider the importance of the subprime loans and their securitization in the financial crisis. They perform their analysis at the county-level; and find three important results. First, they determine that counties that experienced aggressive growth in subprime lending tended to also experience the greatest default intensities. Second, they find that the credit expansion in 2002-2005 was negatively correlated with income growth—an unexpected reversal of normal credit expansion patterns. And third, they demonstrate that subprime loans tend to be more common in areas with decreasing income. While my paper does not explicitly incorporate subprime borrowers, it does provide substantial heterogeneity and permits a changing, endogenous asset distribution (aggregate uncertainty), which allows us to consider whether borrower quality declines prior to crises.

Other plausible explanations for the Great Recession include Mayer (2009), who attempts to explain the crisis through a decline in underwriting standards and a large post-teaser rate jump on mortgage interest rates. Bucks and Pence (2008) argue that the large cluster of defaults was caused by a change in the composition of borrowers. That is, less financially literate individuals were induced to become homeowners. When they were hit with negative income shocks, they were less able to optimize budgeting correctly (often ignoring shocks entirely), which lead to a sharp rise in defaults.

In line with the theme of this paper, Leamer (2007) argues that recessions in the U.S. have largely been driven (or preceded) by downturns in housing investment and consumer durables. He claims that monetary policy should explicitly incorporate housing investment smoothing as an objective; and suggests that a modified Taylor Rule could incorporate such an objective. Similarly, Ahearne et. al (2005) draws the connection between monetary policy and house prices. In a study of 18 major industrial countries, they find that monetary easing typically precedes an increase in

housing investment and prices.

Hott and Jokipii (2012) further cement the relationship between interest rate movements and housing investment. Using a multi-country dataset, they show that deviations in interest rates from the Taylor Rule account for 50% of the housing overvaluation that occurred prior to the Great Recession. Assesnmacher-Wesche and Gerlach (2008) find a similar relationship between interest rates and house prices using a 17-country VAR that spans 1986-2006. They find that a 25 basis point increase in short term rates pushes down GDP by 0.125% and housing prices by 0.375% with a lag. However, in contrast to Hott and Jokipii, they argue against using interest rates to smooth house prices.

Finally, Foote, Geraradi, Goette, and Willen (2008) provide an additional empirical nuance to to the debate. They show that borrowers do not simply default if they have negative equity or if they receive a negative income shock, but rather, they default if both conditions are present. They call this the “double trigger” condition for mortgage default—a term that originated in finance.³ To this author’s knowledge, this paper is the only rigorously-microfounded model of default in general equilibrium that replicates this empirical regularity.

2.2 Macro-Housing Literature

This paper also contributes to the macro-housing literature, which has become increasingly focused on heterogeneity, incomplete markets, and aggregate uncertainty. Early papers in the literature, such as Yang (2006), demonstrated the importance of lifecycle elements in housing; and devised mechanisms for modeling them correctly. Li and Yao (2007) showed that the impact of house price changes depend on the degree of heterogeneity in the economy. Indeed, depending on how agents are modeled,

³See Abraham 1993.

price changes may have no effect on the macroeconomy, but can cause inter-group wealth transfers.

Ortalo-Magne and Rady (2005) contribute one of the foundational papers in this literature. They incorporate life-cycle elements, property ladders, and credit constraints into an equilibrium model with many heterogeneous agents. While they greatly simplify the life-cycle elements (by including only four periods) and build a highly stylized model with no aggregate uncertainty, they still introduce a housing market prototype that is reused elsewhere in the literature. In particular, they assume that property comes in two varieties: “flats” and “houses”; and explicitly model the utility-based differences between the two. As for credit constraints, they assume that wealth cannot fall below some fraction of the value of the property. Finally, they assume that the supply of both flats and housing is fixed.

Chatterjee, Corbae, and Rios-Rull (2009) expand the literature further by incorporating default and credit scoring into a large scale macro context. They model agents who take out loans and repay them with an unknown probability. This leads to a system of credit-scoring, where lenders use a known credit-scoring function and pricing kernel to create loan terms. While this model provides a reasonable structure for incorporating default into a model with heterogeneous agents, it can say little about the timing and cause of default, since default probabilities are determined by an agent’s type (good or bad). In contrast, I allow agents to default optimally and use decision rules to compute the probability of default.

Other papers in this literature use the new class of rigorously-microfounded macro-housing models to run policy simulations. Jeske & Krueger (2005), for instance, find that mortgage interest rate subsidies tend to benefit and to increase homeownership rates among high-income and high-net worth individuals in general.

Another important part of modeling default choices is determining the structure

of mortgages, since this may play a large role in determining default behavior. While there are a number of continuous time finance models that permit variation in payment schedules and default timing in a partial equilibrium context, this is not true in general for DSGE models, which frequently do not even model housing. In the wake of the Great Recession, however, a number of authors have attempted to create a serious role for mortgage structure in DSGE models. Most notably, Chambers, Garriga, and Schlagenauf (2009) have constructed a general equilibrium model that permits individuals to choose between FRMs and ARMs—and even to obtain combo or piggyback loans. While I do not incorporate mortgage choice into my model, I do borrow modeling devices from the paper, and work to make them tractable in an environment with default.

Perhaps most closely related to this paper, Iacoviello and Pavan (2010) create a quantitative equilibrium model that accurately reproduces empirical co-movements in aggregate debt accumulation and housing investment. In their paper, they introduce (and borrow) several computational macro modeling devices that permit them to perform simulations in a realistic environment, but without making the computational stage intractable. In particular, they use deterministic life-cycle productivity profiles to generate income heterogeneity, small pension payments after retirement to generate lifecycle asset accumulation motives, lump-sum taxes to simplify interactions with government, housing transaction costs that are proportional to the change in housing position (to make changes larger and infrequent), and a simple no-arbitrage condition to determine the price of housing. My paper borrows many of the modeling devices from Iacoviello and Pavan (2010), and adds optimal default, individual credit-scoring, inter-sectoral productivity correlation, and mortgage structure—and applies them to a different research question.

Goodhart, Osario, and Tsomocos (2009) attempt to advance the literature by

creating a template for a new generation of macro-financial models. While they use only a handful of representative agents (unlike the other referenced housing papers), they model both default and complex linkages between the financial sector and the macroeconomy. They include corporate lending, interbank lending, and central bank lending into a model where most agents can “default.” However, since only a handful of representative heterogeneous agents are used, individual agents do not default in a strict sense—rather, they choose a repayment rate. Banks respond to this by using rational expectations to predict repayment rates and to penalize default. Additionally, interbank and central bank lending is used to rescue banks in the event of mass defaults.

2.3 Optimal Monetary Policy

Finally, this paper contributes to the part of the optimal monetary policy literature that examines the welfare properties of interest rate rules. Woodford (2001), a landmark paper in the subliterature, considers whether the Taylor rule can be rationalized through a plausible central bank objective. He does this by constructing a simple model and testing its welfare properties. He finds that Taylor-style rules perform well, but suffer from two problems: 1) they rely on the output gap being measured correctly; and 2) they do not vary with the Wicksellian natural rate of interest, but instead assume a fixed natural interest rate. He suggests that future work should attempt to “analyze the consequences of inertial rules in the context of more detailed models.” This is one of the primary objectives of this paper.

Julliard et. al (2006) test the welfare properties of the Taylor rule by constructing a DSGE model, using Bayesian methods to estimate its parameters, and performing counterfactual simulations under different parameter values. They find that the standard Taylor rule performs reasonably well. Ahrend (2010), on the other hand,

performs an empirical investigation of the Taylor rule in practice; and finds that central bank departures from the rule can lead to substantial increases in asset prices. This paper explores both of these topics in a theoretical context.

Other papers, such as Giannoni (2012) find that simple Taylor-style rules may be inferior to Wicksellian rules, which permit some type of history-dependence. Giannoni (2012) claims that such rules are less prone to indeterminacy and are more robust to model misspecification; and argues that they are especially effective when coupled with a “high degree of interest rate inertia.” Forlati and Lambertini (2011) also consider interest rate inertia, but do it in the context of richer model that includes housing. They find that inertial interest rate rules lead to larger contractions in output.

In summary—I expand on the macro-housing literature by adding optimal default, credit scoring, inter-sectoral productivity correlation, and mortgage structure to a quantitative equilibrium model. I add to the largely empirical Great Recession and financial crisis literature by performing counterfactual simulations using a theoretically consistent framework and a rigorously-microfounded model. And, finally, I add to the optimal monetary policy literature by evaluating a variety of different interest rate rules to determine their impact on default and welfare in a detailed model.

3 The Model

In order to answer the questions above, I start by constructing a macro-housing model in which agents default optimally. That is, default is not forced, arbitrary, or determined by type, but emerges from optimization. In building this model, my primary goal is to attack the problem as simply as possible, but with enough detail to capture the important features of default and housing choices. For this reason, I construct a rigorously-microfounded model in the style of Aiyagari (1993), and Krusell

and Smith (1998).

The initial formulation will focus on a real model with capital and endogenously-determined house prices. However, I will later modify the model by removing capital and the endogenous component of house prices. I will also explain how sticky and flexible prices are added to the model.

In the model, there are infinitely many heterogeneous households of measure 1, who consume non-durables, invest in capital, purchase housing, and choose whether or not to default on mortgage debt. Firms produce a non-durable good with Cobb-Douglas technology. There is a representative financial intermediary, which accepts deposits from households, and then uses those deposits to issue mortgages to households. The government collects lump-sum taxes, pays pension benefits, insures deposits at financial intermediaries, and collects housing from deceased agents. Finally, a central bank implicitly determines the risk-free return on deposits by setting the interbank lending rate.

3.1 Firms

The firm side of the economy consists of 1) a consumption good producer who rents capital and labor services; and 2) a technology that permits all households to transform the consumption good into housing units.

3.1.1 Consumption Goods

The consumption goods are produced using Cobb-Douglas technology:

$$Y_t = e^{A_t} K_t^\alpha N_t^{1-\alpha}, \tag{1}$$

where $A_t = c_A + \rho A_{t-1} + \epsilon_{At}$, $\epsilon_A \sim N(0, \sigma_\epsilon)$, and $\mu_A = \frac{c_A}{1-\rho}$. Firms maximize profits, yielding the familiar first order conditions:

$$w_t = (1 - \alpha)e^{A_t} \left(\frac{K_t}{N_t}\right)^\alpha \quad (2)$$

$$R_t = \alpha e^{A_t} \left(\frac{N_t}{K_t}\right)^{1-\alpha} - \delta_K, \quad (3)$$

Note that capital is assumed to depreciate at a constant rate, δ_K . Furthermore, N_t is the mass of employed workers. The mass (or fraction) of workers employed in any given period is derived from the conditional Markov process for employment and the assumption that households supply labor inelastically. For simplicity, I have assumed that this process does not depend on housing investment.

For computational purposes, I discretize A_t using the Rouwenhorst (1995) method,⁴ which Kopecky and Suen (2009) argue is more accurate than other approximation algorithms for highly persistent AR(1) processes. In particular, this method allows me to generate a discrete approximation of an AR(1) process by setting four parameters and the desired number of states: ρ_A , q_A , σ_A , and μ_A , where ρ_A is the probability of the highest state, q_A is the probability of the lowest state, σ_A is the desired standard deviation of the process, and μ_A is the desired mean of the process. The algorithm generates the Markov chain and the associated transition probability matrix.

To pin down the mass of employed workers, I use a conditional Markov process that depends on the technology shock. That is, roughly speaking, when the technology shock is high, the probability of transitioning into employment will be high (and vice versa). More specifically, I will calibrate the conditional Markov process to generate unemployment statistics that match U.S. business cycle data.

I use $Pr(\epsilon'_E | \epsilon_E, \epsilon_A)$ to denote the probability of transitioning to employment state

⁴See Appendix 1 for an explanation of the Rouwenhorst method.

ϵ'_E , given last period's employment state, ϵ_E , and this period's technology shock state, ϵ_A . Note that the Markov chain will remain the same in all periods; however, the transition matrix will change, depending on the state of the technology shock, which means that we will have a different transition matrix for each state of ϵ_A .

3.1.2 Housing Investment

The housing investment specification is similar to Glover, Heathcote, Krueger, and Rios-Rull (2011). In particular, I assume that households have access to a linear technology that transforms the consumption good into housing. That is, if the household builds with c_{it}^h units of the consumption good, it will yield h_{it}^n new units of housing:

$$h_{it}^n = \delta(IH_t)c_{it}^h e^{U_t}, \quad (4)$$

where $U_t = u_H + \rho_H + \epsilon_H$, where $\epsilon_H \sim N(0, \sigma_H)$, IH_t is aggregate housing investment, $\delta(IH_t)$ is an analogy to capacity utilization,⁵ and $\delta'(IH_t) < 0$.

Notice that the housing specification implies a relative price limit. No one will pay more than $\frac{1}{\delta(IH_t)e^{U_t}}$ for a unit of housing, since it is possible to generate one using that many units of the numeraire good. Additionally, I assume that investment is reversible,⁶ so no one will sell housing for less than that price. This implies that house prices will rise when the economy is hit by a negative housing sector productivity shock or when housing investment demand increases.

⁵That is, when housing investment is high, the capacity of the sector to produce an additional unit is increasingly strained, making it more costly.

⁶While this is not a particularly realistic assumption, it should not have a qualitative impact on welfare and default results, since house prices in the model fall when investment is reversed.

Note that total housing investment can be written as follows:

$$IH_t = \sum_{i \in I} \delta(IH_t) c_{it}^h \mu_i e^{U_t} = \delta(IH_t) C_t^h e^{U_t}, \quad (5)$$

where μ_i is agent i 's mass and where IH_t and C_t^h are used to denote aggregates. Additionally, the housing stock evolves as follows:

$$H_{t+1} = H_t + IH_t - \delta_H H_t, \quad (6)$$

where δ_H is housing stock depreciation.

3.2 Households

Households are born at $a=1$ and work until $a=T$. After retirement, households receive a pension, x_t , for T^R periods, and then perish with certainty at age $T+T^R$. Note that this generates a hump-shaped paper asset profile for households, since they must accumulate assets in order to smooth post-retirement consumption.

At any point in time, heterogeneity across households is driven by two mechanisms: 1) employment status ($\epsilon_{it}^E = 1$ or 0); and 2) age-specific productivity, η_a . Following Heer and Maussner (2008), I assume that the former is generated by a conditional Markov process that depends on non-durables shocks (as given in the firms section). For the latter, I follow Iacoviello and Pavan (2010) in adopting a single, deterministic profile for age-specific productivity, which is computed using CPS data. These two mechanisms for heterogeneity will drive differences in asset holdings and default decisions.

Unemployed agents receive per period unemployment benefits, x_t^U , for the duration of their jobless spells.⁷ Similarly, retired agents will receive pension benefits of

⁷In the baseline specification, unemployed agents will receive benefits and will not be taxed. In

x_t^R . Employed household i , which is age a at time t receives a wage, $w_t \eta_a$, where $\sum_{a=1}^T \eta_a \mu_a = 1$, where μ_a is the density of age cohort a .

In the baseline model, households consume non-durable goods, c_{it} , and housing service flows, which are assumed to be directly proportional to their housing stock, h_{it} .⁸ They choose how much to save in bank deposits, d_{it} , how much capital to hold, k_{it} , how much collateralized debt to borrow in the form of mortgages, b_{it} , and whether or not to default on the mortgage they hold.

Investment in housing is lumpy; that is, households tend to make large and infrequent changes in housing size (i.e. by moving), rather than changing housing size frequently and in small increments. Here, I follow the standard assumption the literature that lumpiness is generated by housing stock adjustment costs, $\phi(h_{it}, h_{it-1})$, which depend on the size of the new and old housing stock.

Furthermore, I adopt Iacoviello and Pavan’s (2010) assumption that houses have a minimum size, \underline{h} . While they principally use this assumption to match the empirical fact that younger households tend to be renters (and that, in fact, households cannot purchase very small houses), I use this device largely to generate household leverage. When young households enter the model, they must take out a large mortgage in order to purchase a house. The high degree of leverage will translate into interest rate risk exposure—and, thus, a higher probability of default.

Additionally, I assume that all households have access to a small, fixed amount of non-housing shelter. This includes both defaulters and young households who have not yet purchased a home. This non-housing shelter can be interpreted as living with friends or relatives—or staying in a low-quality, but free apartment.⁹

the original set of simulations, I considered a version of the model without unemployment benefits. My findings did not differ qualitatively for the two versions of the model.

⁸The model is later extended to incorporate apartment rentals.

⁹For a discussion of the impact that non-housing shelter has on the model, see the appendix.

For simplicity, households are assumed to supply labor inelastically. Additionally, we may write the household's instantaneous utility from non-durables consumption and housing service flows as follows:

$$u(c_{it}, h_{it}) = \gamma \frac{c_{it}^{1-\sigma_c}}{1-\sigma_c} + (1-\gamma) \frac{(h_{it})^{1-\sigma_h}}{1-\sigma_h} \quad (7)$$

This specification is compatible with Jeske's (2005) finding that the ratio of housing to consumption tends to be hump-shaped over the life-cycle. Following Chambers, Garriga, and Schlafenhaus (2009), we set the curvature parameters for the utility function to $\sigma_c = 1$ (log utility) and $\sigma_h = 3$, which will give us a hump-shaped profile for $\frac{h}{c}$ over the lifecycle that matches the data.

Individuals have two sources of income: wages, net of taxes ($\Gamma_{a,t}$), from working at the consumption goods firm and pensions (or unemployment benefits):

$$y_{it} = \begin{cases} w_t n_a - \Gamma_{a,t} & \text{if employed} \\ x_t & \text{if unemployed or retired} \end{cases} \quad (8)$$

Household i faces the following budget constraint:

$$c_{it} + \phi(h_{it}, h_{it-1}) + d_{it} + p_t^h h_{it} + m_{it}^h + k_{it} = y_{it} + (1+r_{t-1})d_{it-1} + p_t \quad (9)$$

$$h_{it-1}(1-\delta_H) + (1+R_t)k_{it-1} + b_{it}$$

Note that p_t^h is the relative price of housing, m_{it} is the mortgage payment, b_{it} is the unpaid balance on the mortgage, R_t is the return to capital, and k_{it} denotes i 's capital holdings.

I depart from the current literature in applying a novel constraint that makes

holders of one-period mortgages behave as if they held long term debt instruments:

$$b_{it}^H \leq \begin{cases} \lambda p_t^h h_{it} & \text{if } h_{it} - h_{it-1} > 0 \\ \lambda p_t^h h_{it} & \text{if } b_{it-1} < \lambda p_t h_{it} \text{ \& } h_{it} = h_{it-1} \\ b_{it-1} & \text{otherwise,} \end{cases} \quad (10)$$

where $\lambda \in (0, 1)$ denotes the collateral constraint—that is, the maximum loan-to-value ratio.

According to the equation above, if homeowners move (adjust their housing stock), then they face a collateral constraint in that period, since they must obtain a new mortgage. Similarly, if they do not move, but carry forward less debt from the previous period than would be allowed by the collateral constraint in this period, then they have the option to borrow up to the constraint. Finally, if they did not move and have exceeded the collateral constraint, then they cannot borrow more, but do not have to reduce the size of their mortgage. The intent of these constraints is to achieve the following with one-period mortgages:

1. Avoid “forcing” default. In many endogenous default models, default is ultimately forced by a collateral constraint that is repeatedly applied to one-period mortgages. Empirically, this constraint only applies at origination—and not to existing loans. If it is forcefully applied to existing loans (i.e. by modeling them as repeated one-period loans), households will default whenever they are unable to borrow enough to repay last period’s debt (i.e. if house prices fall steeply). This generates a spurious channel for default (i.e. negative equity immediately triggers default), which is simply an artifact of one-period loan financing.

2. Allow mortgage equity withdrawal (MEW). Since we are interested in understanding how the path of interest rates impacts default, allowing for MEW may provide a critical channel for default. If households can borrow against the values of their homes, this may cause certain interest rate rules to generate a higher default rate.¹⁰

3. Allow negative equity. In one-period loan models with collateral constraints, households typically cannot have negative equity, since it will violate the collateral constraint. However, in reality, if the price of housing falls, but individuals choose to remain in their homes, there is no constraint that forces them to maintain positive equity. Since negative equity is an important part of most default crises, this constraint is maintained to generate more realistic household balance sheets.

Overall, this constraint will make it possible to maintain the simplicity of a one-period loan framework, while simultaneously generating household balance sheets and default behavior that more closely approximate what we would observe if 30-period debt instruments were available.

In addition to this, I borrow a constraint from Iacoviello and Pavan (2010) that limits borrowing to a fraction, γ , of discounted, remaining lifetime earnings:

$$b_{it}^I = \gamma E_t \sum_{j=t}^{T-a+j} \beta^{T-a+t} y_{ij} \quad (11)$$

The purpose behind this constraint is to impose a feasibility condition on repayment. If households cannot reasonably be expected to repay a mortgage with their remaining income flows, then a financial intermediary will not be willing to originate

¹⁰Certain classes of interest rate rules may make it favorable to borrow against the value of your home immediately prior to a hike in interest rates, which will lead to a high degree of interest rate risk.

it. Empirically, this is similar to the income and debt-servicing ratios that banks require borrowers to meet; however, it is tied to expected, discounted future income, rather than just current-period income and assets.¹¹ This constraint, coupled with the previous one, yields the final borrowing constraint:

$$b_{it} \leq \min\{b_{it}^H, b_{it}^L\} \quad (12)$$

That is, the maximum amount a household can borrow is the minimum implied by the two borrowing constraints. Furthermore, mortgage interest rates are adjustable and are given as follows:

$$r_{it}^* = \left(\frac{r_t}{1 - q_{it}} + \xi_t \right) \quad (13)$$

Note that r_t is the rate earned on deposits, q_{it} is household i 's probability of default (computed from household decision rules),¹² and ξ_t is the mortgage premium.

Default, ψ_{it}^d , is captured by a binary variable. Defaulters are not able to re-enter the mortgage market for a period of time, f_p . A separate binary variable, ψ_{it}^c , denotes whether an individual has defaulted recently enough to be excluded from the mortgage market.¹³

In contrast to Chatterjee (2009) and other papers in this literature, I assume that the housing market exclusion period is fixed, rather than random, for the sake of tractability. Furthermore, when a household defaults, the model requires that $h_{it} = 0$ and $b_{it} = 0$. That is, the housing stock (which serves as collateral) is transferred to the

¹¹An alternative—and arguably more realistic—specification of this constraint might incorporate current asset holdings. The impact of including this modification would be to permit more borrowing later in the life cycle, since there is no within-cohort heterogeneity. For the purposes of this paper, I do not consider this constraint explicitly, but may add it as a future extension.

¹²See the appendix for a shortcut for computing default probabilities without using the decision rules.

¹³Note that defaulters have access to non-housing shelter, which means that they will not receive negative infinity utility from defaulting.

financial intermediary and the mortgage debt is eliminated.¹⁴

Finally, note that some households in the model will perish with housing remaining. In the literature, it is common for the dying generation to either bequest the housing to the incoming generation or turn it over to the government. For simplicity, we assume that the government takes the housing when the outgoing cohort perishes.¹⁵

Now that all of the pieces of the model in place, we may collect the individual-level state variables, $z_{it} = \{d_{it-1}, \psi_{it-1}^c, \psi_{it}^d, h_{it-1}, b_{it-1}, k_{it-1}, \epsilon_{it}, a\}$, the aggregate-level state variables, $Z_t = \{K_t, A_{t-1}, U_{t-1}, IH_t, \epsilon_{At}, \epsilon_{Ut}, c_A\}$, and the parameters $\Omega = \{\alpha, \sigma_c, \sigma_h, \gamma, \lambda, \rho_A, \rho_U, \sigma_U, \sigma_A, \delta_K, \delta_H, c_A, c_U, \sigma_{A,U}\}$ to simplify notation. The dynamic programming problem (DPP) for the household may now be written as follows:

$$V_{it}(z_{it}, Z_t; \Omega) = \max_{\{c_{it}, d_{it}, k_{it}, h_{it}, b_{it}, \psi_{it}^d\}} u(c_{it}, h_{it}) + \quad (14)$$

$$\beta \sum_{A', U', \epsilon^{E'} \in \{1, 0\}} Pr(A') Pr(U') Pr(\epsilon^{E'} | \epsilon^E, A') V_{it+1}(z_{it+1}, Z_{t+1}; \Omega)$$

s.t.

$$c_{it} + \phi(h_{it}, h_{it-1}) + d_{it} + p_t^h h_{it} + m_{it}^h + k_{it} = y_{it} + (1 + r_{t-1})d_{it-1} + \quad (15)$$

$$p_t^h h_{it-1} (1 - \delta_H) + (1 + R_t)k_{it-1} + b_{it}$$

$$b_{it} \leq \min\{b_{it}^H, b_{it}^I\} \quad (16)$$

$$\text{If } \psi_{it}^c > 0, \text{ then } b_{it}, h_{it} = 0. \quad (17)$$

¹⁴It is important to note that 1) the foreclosure process is costly—and, thus, the amount recovered will be less than the value of the house prior to the foreclosure; and 2) the value of housing at default will not exceed the size of the mortgage. If it did, the household would simply sell it, rather than defaulting.

¹⁵Note that this will have a fairly insignificant impact on government's budget constraint, since each perishing cohort accounts for 1/60th of the population and will tend to draw down its housing position near the end of the lifecycle.

$$m_{it} = \begin{cases} (1 + r_{it-1}^*)b_{it-1} & \text{if } \psi_{it}^d = 0 \\ 0 & \text{if } \psi_{it}^d = 1 \end{cases} \quad (18)$$

This problem is solved using a custom approximate dynamic programming (ADP) algorithm, which is described in the appendix.

3.3 The Financial Intermediary

I adopt a largely novel specification for the financial intermediary that generates a number of desirable results related to mortgage pricing and solvency. In particular, I place structure on the financial intermediary's objective in order to obtain simple decision rules without solving a dynamic programming problem. I assume the following:

1. Deposits made at financial intermediaries yield the risk free rate, r_{t-1} .
2. Households may obtain competitively-priced, one-period mortgages from financial intermediaries, which are subject to the constraints given in the housing section.
3. Financial intermediaries are risk-neutral.
4. There are infinitely many financial intermediaries, which are represented by a single financial intermediary with zero net cashflows.
5. Financial intermediaries use rational expectations (i.e. a household's decision rules) to determine a household's probability of defaulting, q_{it} .
6. Financial intermediaries add a state-contingent premium, ξ_t , to the mortgage interest rate in order to generate zero net cash flows.

7. The foreclosure process is costly, leaving financial intermediaries with only a fraction, Λ of the housing, which they liquidate in the same period at the market rate, p_t^H .

Using these assumptions, the intermediary sets the mortgage payment for household i , who obtained a loan in period t as follows:

$$m_{it+1} = (1 + r_{it}^*) b_{it}, \quad (19)$$

where b_{it} is the size of the mortgage. Notice that the interest rate on the mortgage contains two components: 1) $(1 + r_{it}^*)$, which is specific to the individual and accounts for idiosyncratic default risk; and 2) a spread component, ξ_t , which is identical for all borrowers and clears the market.

Furthermore, note that the intermediary will receive $m_{it}\mu_i$ when a household repays a loan originated at time $t-1$ and $\Lambda p_t^H h_{it-1}$, when it does not. Thus, in order to obtain zero net cashflows from period t loans and deposits, it must set ξ_t to solve the following equation:

$$\underbrace{(1 + r_t) \sum_i d_{it-1} + \sum_i b_{it}}_{\text{Outflows}} = \underbrace{\sum_i \mathbf{1}(\psi_{it}^d = 0) m_{it} + p_t^H \sum_i \mathbf{1}(\psi_{it}^d = 1) h_{it-1} (1 - \delta_H) + \sum_i d_{it}}_{\text{Inflows}} \quad (20)$$

Notice that d_{it-1} , q_{it-1} , h_{it-1} are all predetermined at time t . Thus, the intermediary sets ξ_t to reduce or increase mortgage volume until net cashflows are zero.¹⁶

¹⁶There are two important things to note. First, in practice, we use a state-contingent function to set ξ_t , rather than setting it in all periods. This adds tractability; and is discussed more in the appendix. Second, ξ_t will also have an impact on the default rate and deposits at time t , but the effects will be substantially smaller than those on mortgage volume.

3.4 The Government

The government has one function in the model: to make transfer payments to retired and unemployed individuals using taxes collected. For simplicity, the government is assumed to use a constant replacement ratio, ζ . That is, it transfers ζw_t to retired and unemployed individuals.

3.4.1 Transfers

In order to cover payments to the unemployed and retired, the government must allocate the following amount to outgoing transfer payments:

$$\tau_t = \zeta w_t \left(\sum_{a=T^R+1}^{T^R+T} \mu_a + \sum_{a=1}^{T^R+T} (1 - \epsilon_{at}^E) \mu_a \right) \quad (21)$$

Note that $\sum_{a=T^R+1}^{T^R+T} \mu_a$, the mass of retired individuals, is constant in this model, so it may be rewritten as μ^R . The mass of unemployed, $\sum_{a=1}^{T^R+T} (1 - \epsilon_{at}^E) \mu_a$, changes over time, so it is denoted by μ_t^U . This yields:

$$\tau_t = \zeta w_t (\mu^R + \mu_t^U) \quad (22)$$

That is, the government must collect enough in taxes to pay the mass of retired, μ^R , and the mass of unemployed, μ_t^U , ζw_t in transfers.

3.4.2 Revenue

For simplicity, I assume the following about taxes: rates scale with productivity and unemployed agents do not pay taxes. With these assumptions, the tax for employed

households in cohort a can be written as follows:

$$\Gamma_{at} = \zeta w_t \eta_a \epsilon_{it}^E \frac{\mu_R + \mu_t^U}{(1 - \mu_t^U)} - \frac{h_t^D}{(1 - \mu_t^U)} \quad (23)$$

Note that h_t^D denotes the housing stock that households turn over to the government in period t after perishing. Additionally, notice that this specification for taxes will require the government to maintain a balanced budget at all times. That is, aggregate incoming transfer payments are equal to aggregate outgoing transfer payments:

$$\Gamma_t = \sum_{a=1}^{T+T^R} \left(\zeta w_t \epsilon_{it}^E \frac{\mu_R + \mu_t^U}{(1 - \mu_t^U)} - \frac{h_t^D}{(1 - \mu_t^U)} \right) \eta_a \mu_a + h_t^D = \zeta w_t (\mu_R + \mu_t^U) = \tau_t \quad (24)$$

To see why this is the case, recall that $\sum_{a=1}^T \eta_a \mu_a = 1$. Since $\eta_a = 0$ and $\epsilon_{it}^E = 0$ when $a > T$ (i.e. individuals are retired), it will also be the case that $\sum_{a=1}^{T+T^R} \eta_a \mu_a = 1$. Thus, $\sum_{a=1}^{T+T^R} \eta_a \epsilon_{it}^E \mu_a = (1 - \mu_t^U)$, which gives us the above equation.

Note that three components of the tax collected from households vary: 1) the wage; 2) the age-specific productivity component; and 3) the unemployment rate, μ_t^U . When an individual's productivity component is higher (i.e. the individual is earning more), she will pay more in taxes. This is also true if the wage is higher, which results in a generally progressive tax. However, an increase in unemployment will still tend to increase taxes, since the tax base will decline. However, if the wage also declines,¹⁷ this effect may be limited.

In order to test the magnitude of these effects, I used data from the autoregressive rule simulation with capital and endogenous house prices, which is described later in the paper. I found that the average tax rate was positively correlated with both

¹⁷This might happen if the consumer goods sector is hit with a negative productivity shock, since this will have a negative effect on wage through the productivity decline, but a positive effect through the increase in unemployment.

the consumption goods sector shock and aggregate income, which suggests that any countercyclical tax behavior generated by a shrinking tax base (as described above) is dominated by the magnitude of wage changes. That is, in the model, aggregate tax revenue tends to fall when output falls and rise when output rises.

3.5 The Central Bank

I assume that the central bank sets the interbank lending rate. Since each financial intermediary is indifferent between borrowing from other banks and from households, the interbank lending rate will also determine the risk-free rate earned on deposits, r_t .

Since the model employs a representative financial intermediary, net interbank lending is zero and is excluded. In the first set of exogenous interest rate rule simulations, I will assume that the central bank adopts an autoregressive interest rate rule. I will then test the properties of this rule by varying the autoregressive coefficient to determine how interest rate persistence impacts default and welfare.

In addition to autoregressive rules, the central bank will also employ other exogenous interest rate rules, such as a sine wave rule, in different simulations. The purpose of this exercise will be to capture the impact of certain interest rate behaviors on default.

The central bank will also employ a number of different rules that endogenously respond to the economy's dynamics. In different simulations, these rules will incorporate house price level targeting, inflation targeting, house price inflation targeting, and output gap targeting. The Taylor rule and modifications of it will also be tested to determine the optimal coefficients for reducing default and maximizing welfare.

Finally, it is important to note that inflation will have real effects in my model, even in the absence of sticky prices. Since agents are heterogenous and the model

incorporates realistic mortgage structure, an increase in inflation will reduce the real value of mortgage debt. Furthermore, inflation will tend to increase the market value of homes relative to the size of the mortgage contracts that were used to purchase those homes. This will improve homeowner balance sheets by increasing home equity positions. Ultimately, these effects (and others) make it possible to perform all of the analysis in a flexible price—rather than sticky price—framework. However, I will defer further discussion of sticky prices to later sections and the appendix.

3.6 Aggregate Consistency Conditions

In addition to satisfying individual-level constraints, the economy is also subject to aggregate consistency conditions. Each constraint requires an aggregate-level variable to be equal to the weighted sum of the individual-level variables:

$$K_t = \sum_{a=1}^{T^{R+T}} k_{at} \mu_a \quad (25)$$

$$N_t = \sum_{a=1}^{T^{R+T}} \epsilon_{at}^E \mu_a \quad (26)$$

$$C_t = \sum_{a=1}^{T^{R+T}} c_{at} \mu_a \quad (27)$$

$$\Gamma_t = \sum_{a=1}^{T^{R+T}} \Gamma_{at} \mu_a \quad (28)$$

$$B_t = \sum_{a=1}^{T^{R+T}} b_{at} \mu_a \quad (29)$$

$$D_t = \sum_{a=1}^{T^{R+T}} d_{at} \mu_a \quad (30)$$

$$H_t = \sum_{a=1}^{T^{R+T}} h_{at} \mu_a \quad (31)$$

$$C_t^h = \sum_{a=1}^{T^{R+T}} c_{at}^h \mu_a \quad (32)$$

$$\Phi_t = \sum_{a=1}^{T^{R+T}} \phi(h_{at}, h_{at-1}) \mu_a \quad (33)$$

$$Y_t = C_t + IH_t + K_t - (1 - \delta_K)K_{t-1} + \Phi_t \quad (34)$$

3.7 Price Stickiness

The model was originally solved and simulated with sticky prices. After further evaluation, it became clear that sticky prices were not an essential feature of my model—and, thus, were removed to highlight the importance of other mechanisms. Below, I briefly describe the version of sticky prices that were originally incorporated into the model. It is important to note that their inclusion in the model did not result in qualitatively different results. Other than that, I will restrict further discussion of the sticky price case to the appendix.

In order to generate sticky prices, I adopt an approach similar to the one taken in Chari, Kehoe, and McGrattan (2000) and Taylor (1979, 1980), but in a cashless economy. That is, there is a final goods producer who assembles intermediates with the following production function:

$$Y_t = \left[\int Y_t(j) di \right]^{\frac{1}{q}}, 0 < q \leq 1 \quad (35)$$

Additionally, there is a continuum of intermediate goods firms who use Cobb-Douglas technology to produce individual varieties:

$$Y_t(j) = e^{A_t} K_t(j)^\alpha N_t(j)^{1-\alpha}, \quad (36)$$

Intermediate good firm j will choose $P(j)$, $K(j)$, and $N(j)$ to maximize profits:

$$\pi_t(j) = P_t(j)Y_t(j) - RK_t(j) - W_tN_t(j) \quad (37)$$

I assume that only half of intermediate goods firms are able to set prices in each period.¹⁸ For instance, in period t , all j , such that $j \in [0, 0.5)$ will set prices for periods t and $t+1$; and in period $t+1$, all j , such that $j \in [0.5, 1]$ will set prices for periods $t+1$ and $t+2$. The price set by intermediate goods firms in period t is denoted by \bar{P}_t . This yields the following optimal price for intermediate goods firms:

$$\bar{P}_t = \frac{E_t(P_t^\theta V_t Y_t + (1+r_t)^{-1} P_{t+1}^\theta V_{t+1} Y_{t+1})}{q E_t(P_t^{\frac{1}{1-q}} + (1+r_t)^{-1} P_{t+1}^{\frac{1}{1-q}} Y_{t+1})}, \quad (38)$$

where $\theta = \frac{2-q}{1-q}$ and V_t is the minimized unit cost of production. Finally, the zero profit condition yields the following equation for the price index:

$$P_t = \left[\frac{1}{2} P_{t-1}^{\frac{q}{1-q}} + \frac{1}{2} \bar{P}_t^{\frac{q}{1-q}} \right]^{\frac{1}{q}} \quad (39)$$

In each period, I solve for P_t using an approximation scheme, which is outlined in the appendix.

4 Model Properties

Now that all of the pieces of the model are in place, we may take a more careful look at the economy, starting with the characteristics of the individuals who default.

¹⁸In the flexible price specification, all firms can set prices in all periods.

Figure 2 shows cumulative density function (CDF) plots of the age, income, home equity, and capital holding for simulated households. The plots show defaulters in red and non-defaulters in blue.

There are four useful things to take away from these plots. First, all defaulting households have negative equity positions, but not all non-defaulters have positive home equity positions. As outlined earlier, negative equity is a necessary—but not sufficient—condition for default. Second, the income of defaulters is low relative to non-defaulters. In fact, 47.86% of defaulters are unemployed; whereas, only 5.7% of non-defaulters are unemployed. This should not be surprising, since income shocks are the second trigger. Third, defaulters tend to be substantially younger. In fact, all defaulters in this particular simulation were under age 50. Individuals over age 50 typically have sufficient home equity and capital holdings to endure shocks to income and equity. Finally, the capital stocks of non-defaulters tend to be considerably higher. There are two reasons for this: 1) having more capital helps agents to smooth shocks that affect wages and house prices, which makes them less likely to default; and 2) the types of agents who are unlikely to default (i.e. older agents with substantial home equity) have had more time to accumulate capital to smooth consumption during retirement.

In addition to looking at default, I also consider other properties of individual-level variables in the model. The household shown in Figure 3 is drawn from the baseline model's simulation and is used to demonstrate the degree of heterogeneity in the model. Note that the household shown experiences multiple unemployment spells, defaults at an early age, holds a mortgage until age 60, and never manages to accumulate a substantial amount of assets. As a result, they experience unusually low consumption in retirement. Figure 4 shows the average household lifecycle profile for comparison. Note that Figure 4 replicates the stylized facts for U.S. homeowners

according to Jeske (2005). That is, they accumulate housing from age 25 until age 50-59, but then stay in the same house or move into a smaller one thereafter. Additionally, non-housing consumption increases until about age 40, but then declines thereafter. Net paper assets are initially negative, but then become positive and grow as the individual pays off debt and accumulates savings for retirement.

Additionally, note that the age-income profile, as shown in Figure 3 and 4 is calibrated to match Consumer Population Survey (CPS) data. That is, income increases until the individual reaches her early 40s. It then levels off and begins to decline. Furthermore, the individual in Figure 3 receives a random unemployment shock at age 62, which lasts a year. And, at age 65, the individual retires and accepts transfers from the government at the replacement ratio, ζ .

Some of the basic features of the aggregate economy are illustrated in Figure 5, which shows 100 periods of simulated data. One thing to note is that aggregate housing is substantially more volatile than consumption, as is also true in the historical data. Additionally, the simulated economy exhibits leverage cycles, as measured by the debt-to-income ratio. Finally, over the course of the 100 years of simulated data, we see significant changes in default rates, which range from 0% to 7%. This is roughly consistent with the U.S. housing market over the last 50 years.

As far as the calibration of the unemployment rate is concerned—the simulated rate has a mean of 6.49% and a standard deviation of 1.47% over the full simulation period (not just the 50 years shown in Figure 6). In the actual data for the U.S. from 1957 to 2012, the mean unemployment rate was 6.03% and the standard deviation was 1.58%. Additionally, it is important to note that idiosyncratic unemployment shocks are based off of shocks to consumption good production, which makes the unemployment rate move with the business cycle.

In addition to examining default in the model, I also consider how macrovariables

in the model respond to shocks by constructing impulse responses. I assume that the central bank uses the following standard version of the Taylor rule with $\alpha_1 = 1.5$ and $\alpha_2 = .5$:¹⁹

$$r_t = r^* + \alpha_1(\pi_t - \pi_t^*) + \alpha_2(\log(Y_t) - \log(Y_t^*)) + \epsilon_t \quad (40)$$

Note that the impulse responses do not come from a linearized version of the model, but instead are constructed by solving the nonlinear version, simulating without shocks initially, and then introducing a single, one-standard deviation shock. All responses shown are from the full-employment steady state, rather than the ergodic rate of unemployment steady state.²⁰ Finally, notice that the non-smoothness and the fluctuations in the impulse responses arise from two things: 1) non-convexities in the model; and 2) the non-linearity of the solution method. Occasionally-binding constraints, house-size minimums, mortgage structure, and post-default housing market exclusion result in discrete adjustments of individual-level (and aggregate) variables in response to shocks. Fluctuations may arise from both non-convexities and second (or higher) order effects, which are captured by the nonlinear solution method.

Figure 7 shows the impulse responses for a positive, one-standard deviation technology shock. The graphs for output, consumption, capital investment and the housing stock²¹ show percentage deviations from steady state. The graphs for the interest rate and the debt-to-GDP ratio are shown in deviations from their initial values. A one-standard deviation technology shock increases the productivity of workers and

¹⁹In the model, $\pi_t^* = 0$ and Y_t^* is equal to its ergodic mean in the full employment steady state.

²⁰For technical reasons, the ergodic rate of unemployment steady state presents several problems. Most importantly, however, even if the unemployment rate is fixed at its ergodic mean, macrovariables will continue to fluctuate slightly, since individuals with different characteristics will become unemployed in different periods.

²¹Note that there is very little housing investment in the steady state, since depreciation is zero. The only investment in the steady state is generated by replacing housing from agents who perish. For this reason, impulse responses are given for the housing stock, rather than housing investment.

capital, leading to an initial 4.1% increase in output, which decays over the course of 13 years. Consumption also rises by 3.5%, but does so more slowly and remains higher for 30 years. The interest rate increases by 200 basis points to respond to the jump in output; and remains high for 15 years. Investment in physical capital and the size of the housing stock increases in response to the shock—and remains above their respective ergodic means longer than output.

Figure 8 shows the impulse responses for a positive, one-standard deviation shock to housing productivity. Note that this shock reduces the price of houses by 6.3%, which fuels housing investment, resulting in an increase in the stock of housing by about 2.5% after 7 years. Physical capital investment, however, drops by 13% over 7 years as investment is diverted to housing. The drop in physical capital depresses output by 0.55% over 7 years; and reduces consumption with a similar magnitude, but lower and over a longer period of time.

Figure 9 shows the impulse responses for a positive, 100 basis point shock to the interest rate. First, notice that the shock increases the return to saving, diverting investment away from capital to bank deposits, which increase by 8% over the course of 20 years. The reduction in capital investment depresses output by .5% and also puts downward pressure on consumption. Note, however, that the reduction in consumption (3%) is much larger than the drop in output; and is most likely caused by the increase in interest rates, which pushes up mortgage payments, constraining households to consume less. The increase in interest rates puts downward pressure on housing investment, reducing the housing stock by 3% after 10 years. Average household leverage rises by 6 percentage points over the course of 20 years, suggesting that households may be withdrawing equity from their homes, even if they are not buying more housing.

Table 1 shows business cycle moments for 1000 periods of simulated data. This

is compared to the to the actual business cycle moments for the U.S. using annual data over the 1957-2011 period.²² In general, the magnitudes of the relative standard deviations and the signs of the covariances accurately represent the actual data; however, simulated investment tends to be less volatile and more procyclical than its empirical counterpart.

Table 2 shows the parameter values used in the autoregressive rule simulations. Note that some parameter values and all of the simulated business cycle moments will vary from simulation to simulation. For this reason, Tables 1 and 2 should not be used to interpret all of the simulated results.

Figure 10 compares the Lorenz Curve for the simulated economy (shown in blue) with an empirical Lorenz Curve for the United States (show in black). The straight, red line represents perfect equality. In general, the simulated economy reproduces the empirical income distribution well with two minor departures: 1) wealth is more evenly distributed within the top 50% of earners in the simulated data; and 2) the bottom 50% is relatively poorer in the simulated data.

5 Simulation Results

In the subsections below, I detail the simulation results, starting with what I will refer to as endogenous interest rate rules. This is intended to refer to any rule that specifies how interest rates should respond to other macrovariables. Alternatively, we might refer to them as policy rules. Later on, we will look at exogenous interest rate rules (or rules that do not specify comovement with other macrovariables) to get a better sense of how certain types of univariate time series properties of interest rates affect default and welfare. We will also use these rules to examine shock amplification

²²Note that the starting year is 1957 because it is the first year that I have observations for the housing investment data.

channels in the model.

Note that the results given in the subsections below are robust to several model specification changes. In particular, the results do not change qualitatively if we make any of the following modifications: 1) add sticky prices; 2) remove unemployment benefits; or 3) remove non-housing shelter. Sticky prices were included in the original set of simulations, but were removed because they played only a small role in the model and obscured important mechanisms. The other two items tended to increase the intensity of welfare results, but did not qualitatively change any findings. All of these modifications will be discussed briefly in the appendix; however, unless I state otherwise, all results shown below are for the flexible price case.

5.1 Endogenous Rules

In this section, I will consider endogenous (or policy) rules for setting interest rates. I will do this by examining the properties of inertial Taylor rules, Taylor rules with different coefficients, price-level targeting rules, output-targeting rules, and Taylor rules with a house price targeting component.

5.1.1 Inertial Taylor Rules

I will start this section by considering the class of inertial Taylor rules. This consists of standard Taylor rules that are augmented by autoregressive components. In the specification given below, r_t is the interest rate set by the central bank, α_ρ is a constant term, ρ_R is the autoregressive parameter, π_t is the rate of inflation, π_t^* is the target inflation rate, Y_t is aggregate demand, and Y_t^* is an “efficient” or target level of aggregate demand. I use the standard Taylor rule coefficients of $\alpha_1 = 1.5$ and $\alpha_2 = .5$.

$$r_t = \alpha_\rho + \rho_r r_{t-1} + \alpha_1(\pi_t - \pi_t^*) + \alpha_2(\log(Y_t) - \log(Y_t^*)) + \epsilon_t \quad (41)$$

Note that everything in the above equation will be mean zero other than the autoregressive component and intercept. Thus, the mean of the interest rate process will be determined by setting the intercept term. In order to be consistent with later simulations, I set $\alpha_\rho = (1 - \rho_r)\bar{r}$, which will generate an interest rate process with a mean equal to the ideal rate, \bar{r} .

In order to determine the impact of a higher degree of inertia, I solve and simulate the model 150 times and at 10 different parameter values. I then fit a curve to the results. Figure 11 and Figure 12 in the appendix show the curves for default and average normalized utility respectively. Note that each of these values is plotted against the value of α_ρ used in the simulation.

The results suggest that a high degree of interest rate inertia will increase the default rate and reduce normalized average utility. In particular, increasing the autoregressive parameter from 0 (a standard Taylor rule with no inertia) to .9 (a highly persistent Taylor rule) will increase the default rate from .55% to 2%. Furthermore, the increase in the default rate does not seem to be offset by utility-increasing benefits that might come from a slow-moving interest rate. Rather, an individual living in the economy with no interest rate inertia who is consuming the average amount would need to be given a 32% increase in consumption in order to be willing to live in the economy with $\alpha_\rho = .9$ —a non-trivial difference in living standards.

These results appear to confirm the findings in papers like Forlati and Lambertini (2011), which suggest that inertial interest rate rules lead to deeper output contractions—and, thus, may harm welfare. However, as we’ll see later, “inertia” itself may not be driving all of the reductions in welfare we see in this model. Rather, the

way in which we generate inertia (i.e. by using an autoregressive process) may create other undesirable univariate time series properties in interest rates—and these may be responsible for a substantial part of the default rate increase and welfare decline.

5.1.2 Standard Taylor Rules

Next, I’ll consider the set of standard Taylor rules. In the specification given below, r_t is the interest rate set by the central bank, \bar{r} is the ideal interest rate, π_t is the rate of inflation, π_t^* is the target inflation rate, Y_t is aggregate demand, and Y_t^* is an “efficient” or target level of aggregate demand.

$$r_t = \bar{r} + \alpha_1(\pi_t - \pi_t^*) + \alpha_2(\log(Y_t) - \log(Y_t^*)) + \epsilon_t \quad (42)$$

For the purposes of this first simulation exercise, I set $\pi_t^*=0$ and set Y_t^* equal to the ergodic mean of Y_t . I then simulate over $\alpha_1 \in (0, 2)$ with $\alpha_2 = .5$. Additionally, for the first simulation, productivity shocks in the housing sector and consumption goods sector are negatively correlated, which makes the price level and housing prices positively correlated.

My results suggest that default is not minimized at the standard Taylor Rule coefficient of $\alpha_1 = 1.5$. In fact, the maximum default rate occurs at $\alpha_1 = 1.3$, with a local minimum at $\alpha_1 = 2$ and a global minimum at $\alpha_1 = 1$.²³ Similarly, utility is maximized at $\alpha_1 = 1$.

Next, I consider the same simulation, but with positively correlated housing and consumption sector productivities (and negatively correlated prices). Here, the default rate is monotonically increasing in the coefficient on inflation. Furthermore,

²³In follow-up set of simulations, $\alpha_2 = .6$ was found to be the default-minimizing and utility-maximizing value of the output gap parameter when $\alpha_1 = 1$. While this isn’t necessarily a global maximum over the parameter space, it suggests that a Taylor rule with stronger output gap targeting, but weaker inflation targeting might be an improvement when housing and default are explicitly taken into consideration.

welfare is monotonically decreasing as α_1 increases, suggesting that the welfare benefits of lower inflation may be outweighed by the higher default rate.

Why is there such an important qualitative difference when the sign of the correlation is changed? The answer lies in house prices. Recall that $p_t^h = \frac{1}{\delta_{IH_t} e^{U_t}}$. If productivity is positively correlated across sectors, then A_t will then tend to be low when U_t is also low. Since the unemployment rate will tend to be high when A_t is low, output will also be low, which will depress the price level.²⁴ This will prompt the central bank to respond by reducing interest rates, which will stimulate housing investment. Since positive housing investment pushes up house prices, the central bank's actions will actually tend to destabilize the housing market. That is, when house prices are high, then central bank will often make them higher. And when house prices are low, the central bank will tend to push them lower. Additionally, as the central bank pursues inflation more aggressively, it will become an increasingly destabilizing force in the housing market.

To summarize—the empirical relationship between housing and consumption sector productivities is a critical determinant of the efficacy of monetary policy because it partially determines the comovement of the price level and house prices. If housing and consumption sector productivities tend to move in opposite directions, then aggressive inflation targeting will destabilize house prices, leading to higher rates of default. Conversely, if they move in the same direction, then aggressive inflation targeting will tend to stabilize house prices, reducing default.

²⁴It is important to note that an adverse technology shock does not increase the price level in my model. While it will push up the unit costs of production, it will simultaneously increase the rate of unemployment, which will have a strong, countervailing impact.

5.1.3 House Price Targeting

Next, I consider Taylor rules that incorporate house prices. The functional form for the rule is given below:

$$r_t = \bar{r} + \alpha_1(\pi_t - \pi_t^*) + \alpha_2(\log(Y_t) - \log(Y_t^*)) + \alpha_3(\log(p_t^h) - \log(p^{h*})) + \epsilon_t \quad (43)$$

Note that p_t^h is the price of housing and p^{h*} is the ergodic mean of house prices. For the purposes of this simulation exercise, I set $\alpha_1 = 1.5$ and $\alpha_2 = .5$; and then varied α_3 from 0 to 1. The primary impact from this rule is to smooth house prices by affecting the housing investment component of price. That is, when investment increases or housing productivity is falls, p_t^h will rise. If such a rule were in place, the central bank would respond by raising interest rates to push housing investment (and prices) down.

One important thing to note is that the central bank is still placing normal weights on inflation and the output gap; however, it has simply added a competing objective: house price smoothing. While this new term might occasionally lead to undesirable effects (i.e. high interest rates when output is low), the effects will be tempered by the original components of the rule. Thus, a recession (if accompanied by fast house price growth) will be met with a smaller reduction in the interest rate than would otherwise happen.

In my simulations, this rule substantially outperforms the standard version of the Taylor rule with respect to default rate minimization and welfare maximization. When $\alpha_3 = 0$, this rule is equivalent to a Taylor rule with typical coefficients. As α_3 increases, the house price targeting component becomes increasingly important. As can be seen in Figures 13 and 14, when $\alpha_3 = 1$, the default rate is 35% lower than it would be under a standard Taylor rule. Furthermore, in order to be indifferent

between the standard Taylor rule and the same rule with a house price targeting component, a household who is consuming the average amount would need to receive a consumption increase of 23%.

This suggests that there could be substantial welfare gains to house price level targeting. However, it is important to emphasize two nuances here: 1) even when $\alpha_3 = 1$, the central bank is not targeting only house price level deviations. Rather, it is adding a house price targeting component to a typical Taylor rule. And 2) the ergodic mean of house prices is known in the model, but not in reality. That is, unlike price level targeting, there is no obvious or reasonable range of levels to target. Furthermore, unlike inflation targeting, it is not clear what the appropriate rate of growth of house prices is.

The second nuance raises an important question: if the central bank does not know what the correct house price level or rate target is, should it just ignore house prices entirely and use a standard Taylor rule, inflation-target, or price-level target instead? I would suggest that ignoring housing altogether is, in fact, taking a position on it; and is no safer than explicitly including housing in the rule with an imperfect target. One possible strategy for getting around this problem might be to only activate the house price component if house price inflation passes out of some safe band of growth:

$$r_t = \bar{r} + \pi_t + \alpha_1(\pi_t - \pi_t^*) + \alpha_2(\log(Y_t) - \log(Y_t^*)) + \alpha_3 f(\log(p_t^h) - \log(p_{t-1}^h)) + \epsilon_t \quad (44)$$

$$f(\log(p_t^h) - \log(p_{t-1}^h)) = \begin{cases} \log(p_t^h) - \log(p_{t-1}^h) & \text{if } (\log(p_t^h) - \log(p_{t-1}^h)) > \nu^U \\ \log(p_t^h) - \log(p_{t-1}^h) & \text{if } (\log(p_t^h) - \log(p_{t-1}^h)) < \nu^L \\ 0 & \text{otherwise,} \end{cases}$$

That is, if the growth rate of house prices is greater than some upper bound, ν^U , then activate the housing component of the rule. Alternatively, if the growth rate of house

prices falls below some lower bound, ν^L , then activate the housing component. This will permit the central bank to conduct monetary policy as usual in normal times, but still respond to house price growth that is highly abnormal relative to its historical trend when it occurs.

5.1.4 House Price Inflation Targeting

For the sake of completeness, I perform simulations for house price inflation targeting in addition to house price level targeting. Here, instead of targeting the ergodic house price level, the central bank attempts to make house prices move slowly, whether they are returning to the ergodic house price level or departing from it:

$$r_t = \bar{r} + \pi_t + \alpha_1(\pi_t - \pi_t^*) + \alpha_2(\log(Y_t) - \log(Y_t^*)) + \alpha_3(\log(p_t^h) - \log(p_{t-1}^h)) + \epsilon_t \quad (45)$$

Note that this approach has interesting implications for housing booms: if house prices start growing rapidly, but taper off as the boom reaches its height, then the central bank will respond by sharply raising the interest rate initially, but then lowering it as the boom continues. This is very different from a house price level targeting rule, which will tighten monetary policy as the boom continues; and will keep it tight even after the peak, when house price growth becomes negative.

Additionally, recall that there is no obvious “ideal” rate of house price growth in the actual economy. Rather, there is perhaps some rate of growth that is consistent with fundamentals and is unlikely to result in a large price growth reversal in the future; however, it is unknown and probably changes over time. In contrast, the model presented in this paper is stationary, which renders a positive or negative house price growth target meaningless. It is also not clear whether fast house price growth ever deteriorates welfare; although, based on the results in this paper, fast drops in house

prices appear to reduce welfare because they outpace equity position improvements from mortgage payments—and result in negative equity, which is a necessary condition for default.

For all of the reasons above, this section will be considered for the sake of thoroughness, but will not be treated as a rule that could plausibly be used for policy-making purposes—at least not without a substantial amount of additional work and thought.

The results for house price inflation targeting suggest that default is minimized and welfare is maximized when $\alpha_3 = 1.2$ —that is, when the central bank changes the interest rate faster than the rate of house price growth. Note that there are two plausible explanations for this result: first, prolonged periods of house price increases are mitigated by aggressive interest rate hikes. And second (and perhaps more importantly), rapid house price declines are slowed by aggressive interest rate reductions.

5.1.5 Price Level Targeting

Next, I consider the class of price-level targeting rules, using the following specification:

$$r_t = \bar{r} + \alpha_p(\log(P_t) - \log(P^*)) + \epsilon_t \quad (46)$$

This rule will increase the interest rate when the price level exceeds its ergodic mean, P^* ; and reduce it otherwise. I performed simulations for $\alpha_p \in (0, 1)$ with positively correlated productivities, and found that default is minimized and welfare is maximized at $\alpha_p = 0.5$. Beyond $\alpha_p = 0.5$, welfare declines and default increases. This should not be surprising, since strong price-level targeting can generate high interest rates when income and house prices are low. That is, if income is low and house prices are low, but the price level is high, then the central bank will respond by pushing up

the interest rate. In the model, this would have particularly severe consequences for default. Conversely, under an inflation-targeting or output-gap targeting regime, this is much less likely to happen.

Similar to the inflation-targeting results, I find that price-level targeting performs poorly when productivities are positively correlated across sectors. In particular, default is minimized and welfare is maximized when $\alpha_p = 0$. That is—an interest rate that follows a white noise process around the ideal rate is preferable to price level targeting if house prices and the price level tend to move in opposite directions.

5.1.6 Output Smoothing Rules

Next, I'll consider rules that smooth the output gap. In the specification given below, r_t is the interest rate set by the central bank, \bar{r} is the ideal interest rate, Y_t is aggregate demand, and Y_t^* is an “efficient” or target level of aggregate demand.

$$r_t = \bar{r} + \alpha_Y (\log(Y_t) - \log(Y_t^*)) + \epsilon_t \quad (47)$$

For the purposes of this simulation exercise, I set Y_t^* equal to the ergodic mean of Y_t . I then simulate over $\alpha_Y \in (0, 1.5)$. I find the following three things: 1) increasing α_Y from 0 to 1.5 more than doubles the default rate; 2) positive correlation between sectoral productivities amplifies the increase in default caused by aggressive output smoothing; and 3) a moderate amount of output smoothing ($\alpha_Y = .5$) maximizes welfare, even if it increases the rate of default, as long as sectoral productivities are negatively correlated (i.e. house prices and the price level are positively correlated).

5.1.7 Regulating Mortgage Contracts

A final strategy that might be employed is to place legal limits on the availability of certain mortgage contracts.²⁵ If, for instance, banks were limited to offering fixed rate mortgages (FRMs) or low LTV ratio mortgages to individuals with poor credit histories or low incomes, this might help to eliminate one of the necessary conditions for default for many high-risk borrowers. Even if those individuals paid higher risk premia on FRMs or were excluded from the housing market temporarily (i.e. until they could make a sufficiently large downpayment), the impact might still be welfare-improving if it mitigated the aggregate risk that this group of borrowers imposes on the housing and mortgage markets.

Since mortgage choice is outside of the scope of this paper, I examine this problem by looking specifically at the LTV issue. In particular, I simulate the economy for an LTV ratio cap of .7 and of 1.1. I find that increasing the LTV ratio to 1.1 (i.e. allowing individuals to take out mortgages greater than the value of their homes) increases the average default rate by only 22.3%; however, it makes the economy more susceptible to catastrophes. When the cap is .7, the default rate never exceeds 10%. However, when the cap is 1.1, it exceeds 10% multiple times—and has a max default rate of 16.67%.

This suggests that regulating LTV maximums might be a way to create stability. Even though the impact of such a regulation will not be highly visible (i.e. it will not dramatically slash the default rate), it may prevent a crisis in which many households default at the same time; and house prices fall for a prolonged period of time.

²⁵Note that an alternative policy might involve taxing LTV ratios or placing stricter limitations on mortgages purchased by Freddie Mac.

5.1.8 Measures of Mispricing and Instability

In the years since the Great Recession, several papers have looked at the use of macroprudential policy and its role in creating financial stability without monetary intervention. Alich, Ryoo, and Hong (2012), for instance, look at macroprudential policy in South Korea. In particular, they consider the taxation of financial intermediaries' key financial ratios, such as the assets to non-core liabilities ratio. They also consider limiting or targeting other measures of leverage.

While I do not show results for macroprudential policy simulations, my model can be extended to incorporate some of these policies. For instance, the government in the model could generate some of its revenue by placing a tax on originations that is proportional to the LTV ratio. However, some macroprudential policies, such as the choice to tax non-core liabilities, as outlined in Shin's (2010) memo, could not be evaluated in my framework without the introduction substantial changes.

Finally, Aydin and Volkan (2011) suggest that modified monetary policy rules can be used to promote stability and discourage mispricing. They focus on non-financial sector borrowing spreads, bank foreign exchange leverage, credit volume, and asset prices as potential targets. The set of interest rate rules they suggest can be represented by the following equation:

$$r_t = \rho_m r_{t-1} + (1 - \rho_m) \{ \rho_\pi \tilde{\pi}_t + \rho_y \tilde{Y}_{d,t} + \rho_\psi \tilde{\psi}_t \} + \epsilon_{m,t} \quad (48)$$

Note that r_{t-1} is the lagged interest rate, $\tilde{\pi}_t$ is the inflation rate's deviation from its target, \tilde{Y}_t is output's deviation from its target, and $\tilde{\psi}_t$ captures the target financial variable's deviation from its target. In my model, this variable might be a key financial ratio, such as the debt-to-income ratio, the mortgage volume-to-income ratio, or the house price-to-GDP ratio—all of which could provide a measure of mispricing or

instability. Alternatively, the average default risk premium from my model could be used to measure risk in the financial sector.

In short–macroprudential policies that either focus on taxation or are incorporated into interest rate rules may provide a reasonable alternative to house price level (or growth) targeting. Rather than picking an arbitrary threshold for “acceptable” growth, the central bank can make an attempt to prevent financial ratios from deviating from historically stable values. While this may sometimes lead to tighter policy during contractions or looser policy during expansions, it will help to prevent the creation of financial vulnerabilities that have the possibility of greatly amplifying crises.

5.1.9 OER and House Prices

One final consideration is whether there is a need to target house prices when owner–equivalent rent (OER) is already incorporated into price indices like the CPI. That is, if the price index in my model were closer to the CPI, it might already capture house price growth; and there would be no need to modify the Taylor Rule, since it would already respond to house price increases.

In fact, OER is not a reasonable proxy for house prices—at least not in the years leading up to the Great Recession. As Cecchetti (2007) points out in a VOX editorial, targeting an inflation measure that incorporates OER is not equivalent to targeting an inflation measure that incorporates house prices. Between 2000 and 2007, OER increased by a mere 3.17% per year; whereas, house prices increased by 8.48% each year. Since OER accounts for roughly 30% of core CPI, changing this definition (or targeting house prices separately) would lead to dramatically different policy decisions.

Thus, relying on the OER component of CPI to capture house price growth may

be misguided; and is unlikely to prevent future crises. For this reason, it may make sense to target house prices or house price growth separately. See the appendix for an extended discussion of this topic.

5.2 Exogenous Interest Rate Rules

In the next set of simulations, I consider exogenous interest rate rules that have different types of univariate time series properties. The purpose behind these simulations is to determine whether certain properties are more likely to increase the default rate and lower welfare. In addition to this, we will also look more closely at the model by removing several critical components; and then adding them back one at a time. We will start by looking at the class of autoregressive interest rate rules in a model without capital or endogenous house prices.

5.2.1 Autoregressive Interest Rate Rules

In order to get the simulated interest rate to approximate federal funds rate (FFR) movements, I estimated the univariate time series properties of the FFR. This entailed two steps: first, I performed an AR(1) regression on the FFR series from 1970-2010 to recover the autoregressive coefficient (.76). Next, I computed the standard deviation of the residual series (0.0156).

In the baseline autoregressive interest rate rule simulation, I used the aforementioned FFR properties to simulate the interest rate. For computational purposes, I discretized the AR(1) process using the Rouwenhorst method (1995), which is discussed further in the appendix. I also removed capital and the endogenous component of house prices from the baseline simulation in order to examine the simplest version before adding further complexity.²⁶

²⁶Parameter values for the baseline are shown in Table 2.

For the purposes of this simulation, I assume that the central bank adopts a variety of different interest rate rules by varying the autoregressive parameter on the interest rate lag (ρ_r) between -1 and 1.²⁷ In each case, I set the constant term, c_r to $\mu_r(1 - \rho_r)$, where μ_r is the average interest rate target. This ensures that each simulation has the same average interest rate. In each case, I solve and simulate the model for a different ρ_r and c_r 150 times; and plot the main result in Figure 14. Note that the blue, solid line plots the simulated data, and the red, dotted line marks the estimated autoregressive coefficient for the annual FFR series.

Figure 15 shows a curve fitted to the default rate-autoregressive parameter pairs from the 150 simulations. Notice that the average default rate in the economy is minimized when the autoregressive coefficient is zero—that is, when the interest rate is a white noise process. Making the interest rate more persistent (by increasing the coefficient) or making it more volatile (by making the coefficient negative) generates a higher average rate of default. For example, adopting an interest rate that has the same univariate time series properties as the FFR will generate a default rate of 1.57%; whereas, adopting a white noise interest rate will generate an average default rate of 1.4%.

Of course, this is not the full story. This analysis completely ignores the possibility of endogenous interest rate responses to other macrovariables, which eliminates the virtue of having an interest rate rule entirely. However, this finding demonstrates something important: even if critical default amplifiers that are affected by interest rates (i.e. endogenous house prices) are ignored, adopting an interest rate rule that generates as much persistence as we see in the FFR leads to a default rate that is 12.14% higher than a white noise interest rate in the simulated data. This suggests

²⁷Note that simulations were conducted close to—but not at—the boundary values in order to avoid introducing nonstationarity into the model.

that there could be welfare gains from making the FFR less persistent, but without introducing volatility, which we've also shown also generates default. Later on, I will demonstrate that the gains might be substantially larger when we add default amplifiers.

In addition to looking at default, it is important to consider how the rule fares along other dimensions. If, for instance, interest rate persistence leads to a higher rate of default, but does so by generating a higher rate of homeownership (and, thus, adding more marginal borrowers), then persistence may have an positive impact on welfare overall. In fact, as shown in Figure 15, the exact opposite is true.

Figure 16 plots normalized, average utility against the autoregressive coefficient. In the graph, average utility is maximized when the autoregressive coefficient close to zero. For comparison's sake, consider an individual who is in the economy with a white noise interest rate and is consuming the average amount of the consumption good. In order to offset the utility loss of moving from the white noise interest rate economy to the economy with an interest rate that has the same univariate time series properties as the FFR, this individual will need to have her consumption increased by 11%.

When looking at the remaining simulation results, I will omit the findings for negatively autocorrelated interest rate rules, since these are less likely to be relevant for policy purposes. However, it is fair to state that interest rate volatility in general tends to push up the default rate—and this is perhaps the most important reason why negatively autocorrelated rates generate more default than white noise. Indeed, in a separate set of simulations, I found that increasing the standard deviation of the interest rate shock from .01 to .05 increased default by 20.8%.

5.2.2 Adding Endogenous House Prices

Next, I expand on the baseline version of the model used in the autoregressive rule simulations by making house prices endogenous. This is done by adding the capacity utilization term, $\delta(IH_t)$, as discussed in the model section. This makes the price of a unit of housing equal to $\frac{1}{\delta(IH_t)e^{U_t}}$ units of consumption. The purpose here is to match another important facet of the relationship between interest rates and default: when interest rates rise, financing new homes with a mortgage becomes more expensive, which pushes down housing investment. In the model without endogenous house prices, this effect could actually reduce the default rate by preventing marginal borrowers from obtaining homes. However, in the case with endogenous prices, rising interest rates will deteriorate housing equity positions by pushing down prices. This will tend to push up the default rate and amplify the effects of interest rates on housing cycles.

There's one more important thing to notice about this particular form for house prices: when the default rate rises and financial intermediaries foreclose on homes, housing investment from earlier periods will be reversed when intermediaries liquidate it, pulling prices down further. This will help the model to match the observation that rising foreclosure rates often depress house prices further.

Figure 17 yields two additional insights. First, the rate of default is higher at all autoregressive coefficients. Even when the interest rate follows a white noise process, the default rate is roughly 276% higher when house prices are endogenous. This suggests that the impact of interest rates on house prices may be one of the most important channels for default.

Adding endogenous house prices not only increases the default rate, but also amplifies the effects of persistence and volatility. For example, a white noise process generates a default rate of 3.95%; whereas, an interest rate with the persistence of

the FFR series generates a default rate of 4.7%—a 23% increase. This is more than twice as large proportionally as the change without endogenous house prices.

In Figure 18, we can see that the impact of interest rate persistence on utility is more pronounced. An individual in the white noise interest rate economy who is consuming the average level of the consumption good will now need a 18.18% increase in consumption in order to be willing to switch to the economy with the FFR-like interest rate rule.

5.2.3 Adding Capital

Finally, we add capital, which completes the model outlined originally. Capital will play several important roles. First, it will add another source of aggregate uncertainty to the model. Second, it will give households another asset they can use to hedge against house price drops. And third, it will add an additional channel through which interest rates may affect default. That is, if the interest rate on deposits falls relative to the return to capital, marginal households will invest in capital instead, increasing output. Conversely, an increase in interest rates will divert investment away from capital into deposits. Applying the logic of the double trigger requirement for default, lower income will not cause default by itself, but will cause more households to satisfy a necessary condition of default.

Figure 19 shows the results for the default rate. One important thing to notice is that the average default rate is substantially lower in this economy than in the economy with endogenous house prices, but no capital. In particular, when the interest rate is a white noise process, the economy has a default rate that is only 44.21% of the rate for the economy with endogenous house prices, but no capital.

Importantly, however, the original results have not changed. Higher interest rate persistence still leads to higher default rates. In fact, moving from a white noise

interest rate rule to a rule with an autoregressive coefficient of .76 increases the default rate by 30.95%. This is similar to the economy without endogenous house prices, but is considerably smaller than the case with endogenous house prices, but no capital.

Figure 20 shows similar, but dampened results for welfare. Normalized average utility drops as the interest rate becomes more persistent. In this case, however, the decline is more gradual. Using the welfare metric invoked in the previous simulation exercises again, the average individual's consumption would have to increase by 12.66% in the economy with $\rho_F = .76$ in order to to make him as well off as he would be under a white noise interest rate.

5.2.4 Sine Wave Rules

In the section on autoregressive rules, I considered the persistence and volatility properties of interest rate rules. One major finding was that either increased persistence or increased volatility would generate a higher rate of default. This section is intended to add a small, but important, nuance to the earlier findings. Here, the importance of cyclicalitly will be considered explicitly, rather than using persistence as a proxy. I will do this by having the central bank use a time-based sine wave rule to set interest rates.

In particular, a rule will consist of setting the z , g , and n parameters in the equation below:

$$r_t = \frac{\sin(zt) + 1}{g} + n \tag{49}$$

Note that n is the minimum interest rate, g determines the maximum interest rate, and z will change the wavelength of the interest rate cycle. For the purposes

of this simulation exercise, I set g to 25 and n to .02, which creates an upper bound for the interest rate at .1 and a lower bound at .02. I then vary z to generate a different wavelength for each simulation. I found that default is minimized when the wavelength of interest rate cycle is 22 years and utility is maximized when the wavelength is 17 years. For comparison's sake, if the nominal FFR's low frequency trend is removed, it has maintained an approximate wavelength of 5-10 years since 1956. Furthermore, if the interest rate is simulated at this approximate wavelength (7.5 years), the default rate increases by 1.27 percentage points over the minimum rate, suggesting that there could be welfare gains to increasing the wavelength of interest rate cycles.

This finding adds an important nuance to the findings from the previous section: persistence generates default because it creates cycles—not because it causes interest rates to move slowly. In fact, within the class of interest rates that exhibit cycles, slow moving interest rates (i.e. ones with long cycle wavelengths) may be preferable to fast ones. If this point seems too subtle, consider the following: if the average wavelength of a cycle were 60 years, then many households would never hold a mortgage during the turning point of an interest rate cycle. Conversely, if the average wavelength of a cycle were 5 years, then almost every homeowner will hold a mortgage during the turning point of a cycle. If these turning points are critical times for default—as we have argued thus far—then the wavelength of an interest rate cycle may be a critical consideration when constructing an interest rate rule.

Furthermore, if the cycles are longer, then house price equity declines (insofar as they are related to interest rates) will happen much more slowly. That is, if the interest rate moves very slowly, then the partial, negative effect that an increase in interest rates will have on home equity will be smaller than the size of mortgage payments. That is, in most periods, households will tend to improve their equity

positions over time. This will make it easier for younger households to deal with income shocks without defaulting. It will also prevent them from experiencing sharp payment increases associated with holding an ARM during a rapid increase in interest rates.

Finally, it is important to note that OLG models with realistic life cycles and mortgage structure are uniquely qualified to solve this class of problem. Had I used infinitely-lived agents with one period mortgages, it is not clear that I could say anything at all about optimal wavelength, since it depends critically on lifecycle elements.

5.3 Conclusion

This paper examines the relationship between a central bank's choice of interest rate rule, the mortgage default rate, and welfare. I do this by constructing a quantitative equilibrium model with rigorous microfoundations; and incorporate aggregate uncertainty, incomplete markets, overlapping-generations, credit-scoring, optimal default, inter-sectoral productivity correlation, and realistic mortgage structure. I then use this model to perform counterfactual simulations with a variety of interest rate rules, including modified Taylor rules, price-level targeting rules, autoregressive rules, and sine wave rules.

I find that the univariate time series properties implied by an interest rate rule are a critical determinant of default rates. In particular, rules that generate short interest rate cycles (i.e. have a shorter wavelength) tend to generate higher rates of default. For instance, the FFR series has an average wavelength that is substantially shorter than the optimal length implied by the simulation exercises.²⁸

²⁸I find the optimal wavelength to be 22 years for default-minimization and 17 years for utility maximization.

In addition to looking at the univariate time series properties, I also considered the importance of inter-sectoral productivity correlation. For the class of Taylor rules, very aggressive inflation targeting increased the default rate in an economy with positively correlated housing and consumption goods sector productivities. In contrast, in an economy with negatively correlated productivities, aggressive inflation targeting pushed down the default rate after initially increasing it. Furthermore, I suggested this may be related to the properties of house prices in the model: if sectoral productivities are negatively correlated, then an interest rate that is strongly procyclical (for the consumption goods sector) will tend to stabilize house prices. Conversely, if sectoral productivities are positively correlated, then an interest rate that is strongly procyclical will destabilize house prices.

Next, I considered price-level targeting rules. In contrast to an inflation target or output-gap target, a price-level target will continue to push up interest rates if the price level is high, even if income is low and house prices are low. With that said, a price-level targeting coefficient of .5 outperformed both white noise rules and rules with higher coefficients with respect to both default-minimization and welfare-maximization. This suggests that the business cycle stabilization benefits that come from weak price level targeting may outweigh the costs (i.e. high interest rates at bad times).

In addition to price-level targeting rules, I also examined house price level targeting rules. I found that modifying the Taylor Rule to add a house price level targeting component substantially reduced the average default rate and increased average welfare; however, it might not be possible to implement such a rule, since it would require the central bank to identify a target house price level or growth rate. With this said, it may still be possible to adopt such a more limited form of this rule if central banks only activate the house price component when growth is abnormally high or low.

Finally, I explored who actually defaulted in the model: young households²⁹ with little to no capital, negative equity, and low incomes. No households that had positive equity defaulted, even if they lacked employment. Additionally, even if households had negative equity, they never defaulted unless they also had a sufficiently low income or were unemployed. This suggests that the “double trigger” mechanism for default as described by Foote, Gerardi, Goette, and Willen (2008) is an accurate way to capture the behavior of optimizing agents in a DSGE model with default. Furthermore, it indicates that central banks should be mindful of this condition when implementing policy; and should actively avoid creating situations where many households have negative equity and low incomes simultaneously. This may involve slowing down interest rate cycles (i.e. increasing the interest rate wavelength), so that new borrowers are unlikely to enter the market at a turning point—or to experience house price declines at a rate that exceeds equity gains from mortgage payments.

²⁹See Figure 21 for the age-default profile from the baseline simulation.

6 References

Abraham, J. 1993. "A Cash Flow Model of Property Performance." Unpublished manuscript. McLean, VA: Freddie Mac Corporation.

Aldrich, E. & Fernandez-Villaverde, J. & Gallant, R., & Rubio-Ramirez, J. 2011. "Tapping the Supercomputer Under Your Desk: Solving Dynamic Equilibrium Models with Graphics Processors." *Journal of Economic Dynamics and Control* Vol. 35(3), 386-393.

Aiyagari, S.R. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics* Vol.109, 311-331.

Ahearne, A. et. al. 2005. "Monetary Policy and House Prices: A Cross-Country Study." Working Papers Central Bank of Chile 344, Central Bank of Chile.

Ahrend, R. 2010. "Monetary Ease: A Factor behind Financial Crises? Some Evidence from OECD Countries." *Economics: The Open-Access, Open-Assessment E-Journal* Vol. 4.

Assenmacher-Wesche, K., & Gerlach, S. 2008. "Monetary Policy, Asset Prices and Macroeconomic Conditions: A Panel-VAR Study." *IMFS Working Paper* Institute for Monetary and Financial Stability.

Bernanke, B., & Gertler, M. 2001. "Should Central Banks Respond to Movements in Asset Prices?" *The American Economic Review* Vol. 91, 253-257.

Bucks, B., & Pence, K. 2008. "Do Borrowers Know their Mortgage Terms?" *Journal of Urban Economics* Vol. 64, 218-33.

Cecchetti, S. 2007. "Housing in Inflation Measurement." Vox Editorial. Retrieved from <http://voxeu.org/article/housing-inflation-measurement>.

Chambers, M., Garriga, C., & Schlagenhauf, D. 2009. "The Loan Structure and Housing Tenure Decisions in an Equilibrium Model of Mortgage Choice." Working Paper 2008-024B, Federal Reserve Bank of St. Louis.

Chatterjee, S. 2009. "Foreclosures and House Price Dynamics: Quantitative Analysis of the Mortgage Crisis and the Foreclosure Prevention Policy." Research Department, Federal Reserve Bank of Philadelphia, Working Paper NO. 09-22.

Chatterjee, S., Corbae, D., & Ros-Rull, J. 2008. "A Finite-Life Private-Information Theory of Unsecured Consumer Debt." *Journal of Economic Theory* Vol. 142(1), 149-177.

Foote, C., Gerardi, K., Goette, L., & Willen, P. 2008. "Subprime Facts: What (We Think) We Know about the Subprime Crisis and What We Dont." Public Policy Discussion Paper 08-2, Federal Reserve Bank of Boston.

Forlati, C., & Lambertini, L. 2011. "Risky Mortgages in a DSGE Model." *International Journal of Central Banking* Vol. 7(1), 285-335.

Giannoni, M. 2012. “Optimal Interest Rate Rules and Inflation Stabilization versus Price-Level Stabilization.” Federal Reserve Bank of New York Staff Report No. 546, Federal Reserve Bank of New York.

Glover, A., Heathcote, J., Krueger, D., & Rios-Rull, V. 2011. “Intergenerational Redistribution in the Great Recession.” NBER Working Papers 16924, National Bureau of Economic Research, Inc.

Goodhart, C., Osorio, C., & Tsomocos, D. 2009. “Analysis of Monetary Policy and Financial Stability: A New Paradigm.” *CESifo Working Paper Series 2885*, CESifo Group Munich.

Heer, B., & Maussner, A. 2008. “Dynamic General Equilibrium Modeling: Computational Methods and Applications.” 2nd Edition. Berlin: Springer.

Hott, C., & Jokipii, T. 2012. “Housing Bubbles and Interest Rates.” Mimeo, Swiss National Bank.

Iacoviello, M., & Pavan, M. 2009. “Housing and Debt Over the Life Cycle and Over the Business Cycle.” *Boston College Working Papers in Economics 723*, Boston College Department of Economics.

Jeske, K. 2005. “Macroeconomic Models with Heterogenous Agents and Housing.” *Economic Review* Vol. 90(4), 39-56.

Judd, K., Maliar, L. & Maliar, S. 2010. “A Cluster-Grid Projection Method: Solving

Problems with High Dimensionality.” *NBER Working Papers 15965*, National Bureau of Economic Research, Inc.

Juillard, M., Philippe, K., Laxton, D. & Paolo, P. 2006. “Welfare-Based Monetary Policy Rules in an Estimated DSGE Model of the US Economy.” International Research Forum on Monetary Policy Working Paper No. 613, European Central Bank.

Keane, M., & Wolpin, K. 1994. “The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence.” *The Review of Economics and Statistics* Vol. 76(4), 648-72.

Kopecky, K., & Suen, R. 2010. “Finite State Markov-Chain Approximations to Highly Persistent Processes.” *Review of Economic Dynamics* Vol. 13, 701-714.

Krusell, P., & Smith, A. 1998. “Income and Wealth Heterogeneity in the Macroeconomy.” *Journal of Political Economy* Vol. 106, 867-896.

Leamer, E. 2007. “Housing is the Business Cycle.” Proceedings, Federal Reserve Bank of Kansas City, 149-233.

Li, W., & Yao, R. 2007. “The Life-Cycle Effects of House Price Changes.” *Journal of Money, Credit, and Banking* Vol. 39(6), 1375-1409.

Mayer, C., Pence, K., & Sherlund, S. 2009. “The Rise in Mortgage Defaults.” *Journal of Economic Perspectives* Vol. 23(1), 2750.

Mian, A., & Sufi, A. 2009. "House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis." NBER Working Papers 15283, National Bureau of Economic Research, Inc.

Mian, A., & Sufi, A. 2010. "Household Leverage and the Recession of 2007 to 2009." NBER Working Paper 15896, National Bureau of Economic Research, Inc.

Ortalo-Magn, F., & Rady, S. 2005. "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints." Discussion Papers in Economics 494, University of Munich, Department of Economics.

Powell, W. 2007. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. John Wiley & Sons, Inc.

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

Taylor, J. 1993. "Discretion versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy* Vol. 39, 195-214.

Tauchen, G. 1986. "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions." *Economic letters* Vol. 20, 177-181.

Vandell, K. 1995. "How Ruthless Is Mortgage Default? A Review and Synthesis of the Evidence." *Journal of Housing Research* Vol. 6., 245-264.

Woodford, M. 2001. "The Taylor Rule and Optimal Monetary Policy." *The American Economic Review* Vol. 91(2), 232-237.

Yang, F. 2006. "Consumption Along the Life Cycle: How Different is Housing?" Federal Reserve Bank of Minneapolis Working Paper 635, Federal Reserve Bank of Minneapolis.

7 Appendix

7.1 Figures and Tables

Figure 1: Taylor Rule vs. Effective Federal Funds Rate




Figure 2: Characteristics of Defaulters and Non-Defaulters

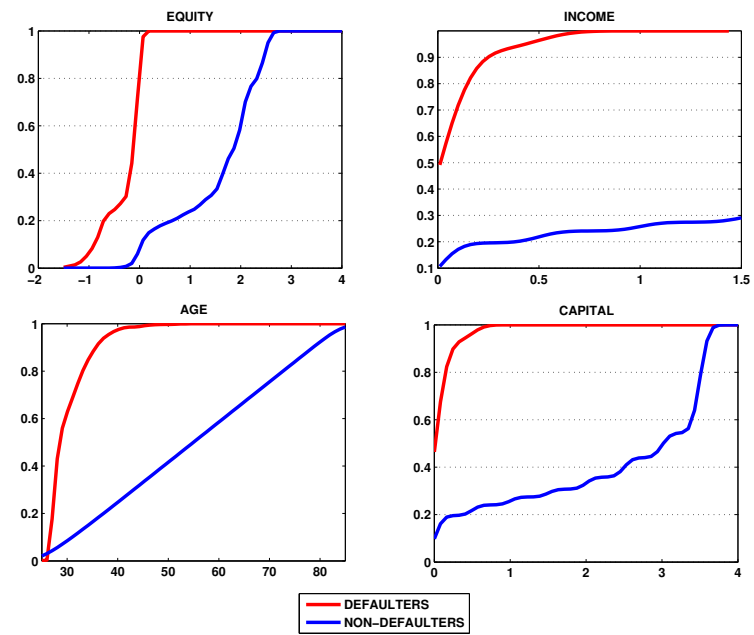


Figure 3: Individual-Level Variables for Random Household




Figure 4: Individual-Level Variables for Average Household

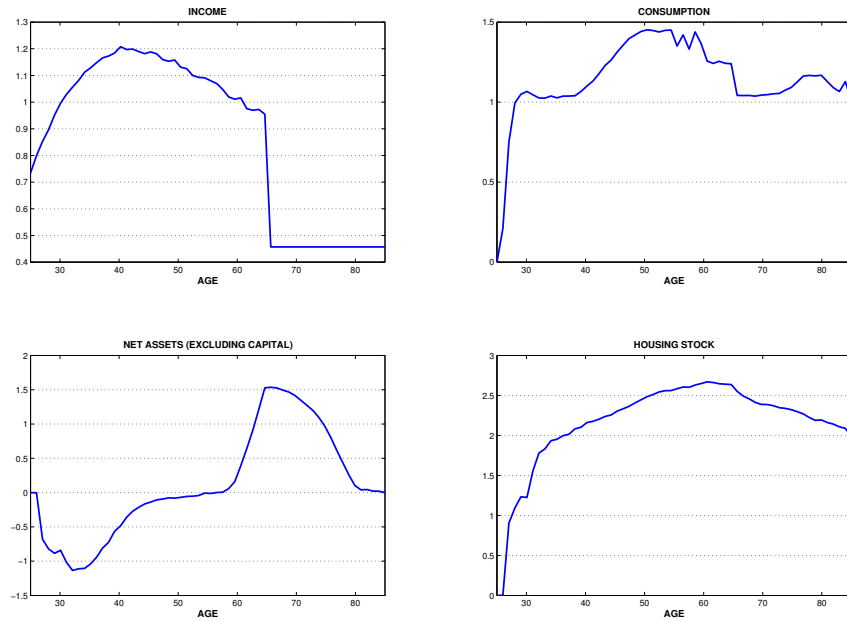


Figure 5: 100 Periods of Simulated Aggregate Data




Figure 6: 50 Periods of Simulated Unemployment Data




Figure 7: Technology Shock Impulse Responses

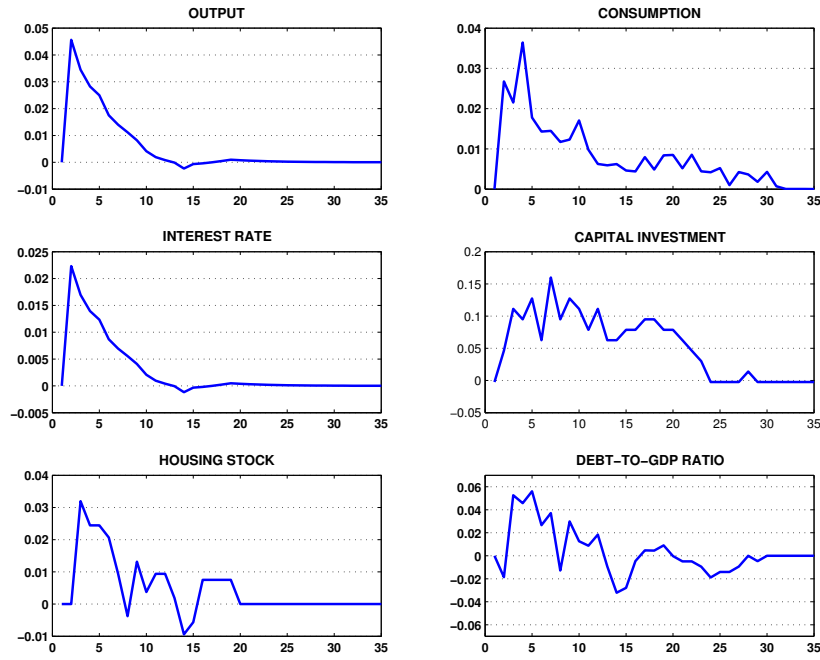


Figure 8: Housing Productivity Shock Impulse Responses

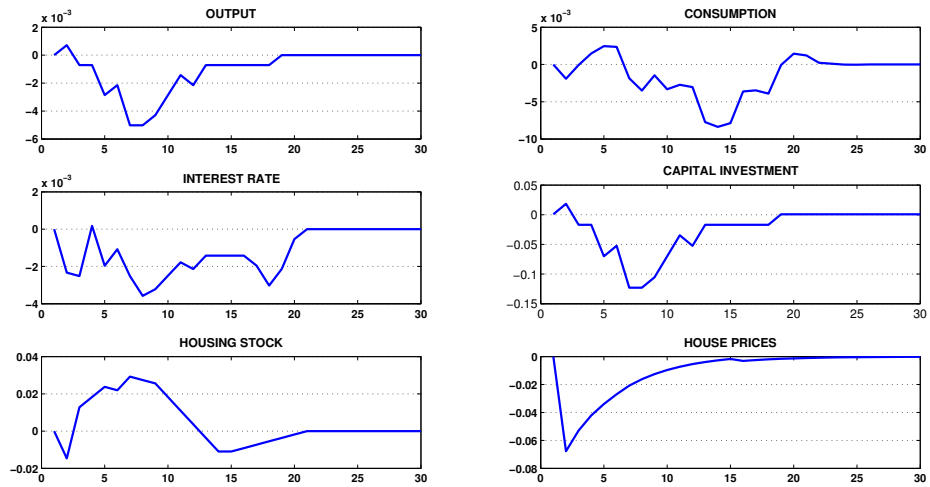


Figure 9: Monetary (Interest Rate) Shock Impulse Responses

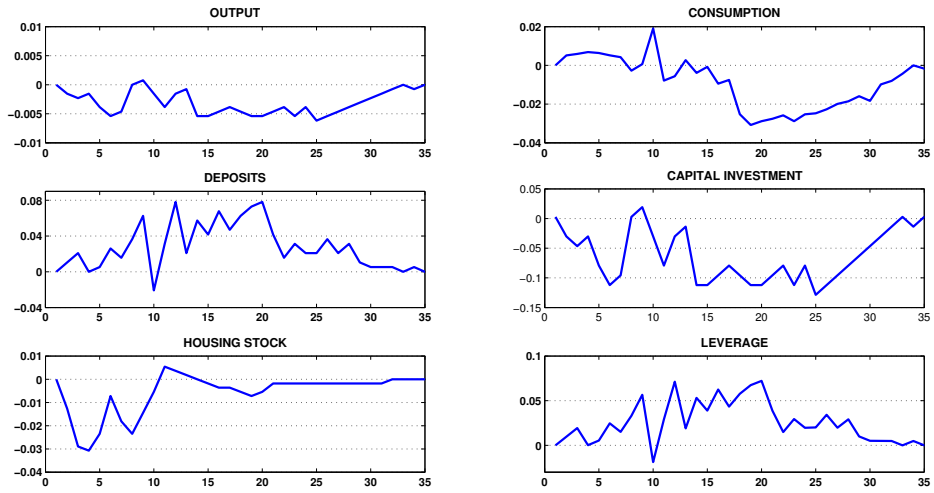


Figure 10: Empirical and Theoretical Lorenz Curves




Figure 11: Default with Inertial Taylor Rule




Figure 12: Welfare with Inertial Taylor Rule




Figure 13: Default with House Price Level Targeting




Figure 14: Welfare with House Price Level Targeting




Figure 15: Default without House Prices or Capital

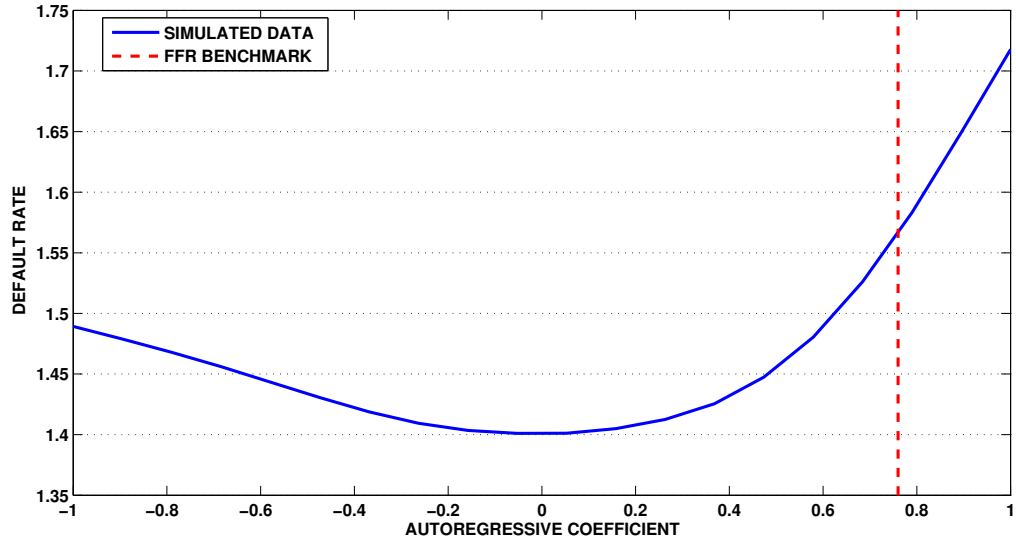


Figure 16: Welfare without House Prices or Capital




Figure 17: Default with Endogenous House Prices




Figure 18: Welfare with Endogenous House Prices




Figure 19: Default with Endogenous House Prices and Capital




Figure 20: Welfare with Endogenous House Prices and Capital




Figure 21: Age-Default Rate Profile




Figure 22: 150-Period AR(1) Series for Consumption Sector Productivity




Figure 23: 150-Period Rouwenhorst Approximation




Figure 24: Ergodic Employment Rate Convergence




Figure 25: 150-Period LoM-Generated Unemployment Rate




Table 1: Simulated vs. Actual (1957-2011) Moments of Log Detrended

Data ³⁰				
	Actual		Simulated	
X	$\frac{\sigma_X}{\sigma_Y}$	$\rho_{X,Y}$	$\frac{\sigma_X}{\sigma_Y}$	$\rho_{X,Y}$
C	0.91	0.94	0.94	0.78
I_H	3.94	0.38	2.51	0.68
I_K	2.43	0.53	1.31	0.64
U	0.38	-0.34	0.135	-0.1

Table 2: Baseline Parameter Values

Parameter	Description	Value
β	Discount Factor	.97
ζ	Replacement Ratio	.4
α	Capital's Share in production	.33
c_A	AR(1) Technology Process Constant Term	0
σ_A	Standard Deviation of Technology Shock	0.03
ρ_A	Technology Level AR(1) Coefficient	0.8
c_H	AR(1) Housing Productivity Process Constant Term	0
σ_H	Standard Deviation of Housing Productivity Shock	.1
ρ_H	Housing Productivity AR(1) Coefficient	0.8
δ_K	Physical Capital Depreciation Rate	.07
δ_H	Housing Depreciation Rate	0
λ	Max LTV Ratio	1
γ	Fraction of Lifetime Earnings Borrowable	.3
f_p	Default Exclusion Period	1

³⁰These simulations are performed with $\delta_H = .1$; however, for most simulations results shown, $\delta_H = 0$.

7.2 Extended Discussion of Topics

7.2.1 Sticky Prices

The original set of simulations was performed using the sticky price version of the model. However, after comparing the results to the flexible price version, it became clear that sticky prices play a limited role in my model. For this reason, they were replaced by the flexible price results in order to clarify the importance of critical mechanisms in the model. One reason why we might expect sticky prices to have a limited impact is because the unemployment rate in the model is pinned down by aggregate and idiosyncratic shocks. Thus, it is not possible to have a strong employment response after a monetary shock. Note that this differs substantially from the standard New Keynesian model, which features elastic labor supply, flexible wages, sticky prices, and demand-determined output in order to generate large output and employment responses to monetary shocks. At most, my model will generate mild output changes through the sticky price mechanism.

With this said, my results are “robust” to the inclusion of sticky prices. That is, when sticky prices are incorporated into the model in the way described in the paper, there are no substantial qualitative differences in my results. In some cases, the intensities of default and welfare results are slightly stronger when sticky prices are introduced; however, the magnitudes of the impacts are generally quite small.

7.2.2 Unemployment Benefits

Originally, simulations were performed without unemployment benefits in the model. This tended to push up default rates and to overstate the welfare impact of rules that caused default. Later, unemployment benefits were introduced for two reasons: 1) to ensure that consumption is always non-zero, even when a household is unem-

ployed (and, thus, avoid the need to set an arbitrary lower bound for utility from consumption); and 2) to avoid triggering defaults that might not immediately be caused by unemployment. No major qualitative results changed after introducing unemployment benefits into the model; however, the intensity of some results have been reduced by its inclusion.

7.2.3 Non-Housing Shelter

While incorporating non-housing shelter is not equivalent to adding a rental market to the model (since no one pays for the shelter), it provides young households and low-income households in the model with an attractive alternative to becoming highly leveraged. Instead, these groups have the option accumulate assets, live in non-housing shelter, and then purchase a home with a larger downpayment.³¹

The simulation results confirm the intuition above. When non-housing shelter is introduced into the model,³² the default rate drops by 51%. Furthermore, the homeownership rate drops from 93% to 74%, since households do not need to quickly purchase a home after entering the model. Relatedly, the median age of first time borrowers at the date of the mortgage origination rises from 27 to 31. While aggregate home equity in the economy drops by 26%, the impact is much less when we instead consider average equity among homeowners, which is only reduced by 6%.

Finally, while excluding non-housing shelter (or a rental market) will tend to magnify the intensity of my results by pushing up default rates and lowering the utility associated with not owning a house, it does not appear to have a qualitative impact on my results. At least in the case of autoregressive rules, the relationships were identical with and without non-housing shelter.

³¹Note that this will not be an attractive option in the case without non-housing shelter because housing's contribution to utility will be large and negative if an individual holds no housing.

³²Note that the comparison simulations were run using an autoregressive interest rate rule with a coefficient of .76

7.2.4 OER and House Prices

The Bank of England, which adopted an inflation target of 2% in 2003, recently debated whether or not to revise the CPI to incorporate house prices in order to deal with this issue.³³ Since the the Office of National Statistics' (ONS) measure of CPI excluded both housing prices and OER, the measure of inflation completely missed the run up in house prices prior to the Great Recession. In order to improve future results, the ONS plans to construct a separate measure of the CPI that will incorporate OER, but not house prices.

While this choice might not be as important in the U.S., where the FOMC is permitted to use discretion, it might be substantially more important for central banks that claim to target only inflation or are legally-required to keep inflation below a target threshold. In particular, re-defining the CPI to incorporate house prices, rather than OER, would have resulted in tighter monetary policy in many countries in the years leading up to the Great Recession.

It is important to note, however, that constructing an alternative version of the CPI that incorporates house prices, rather than OER, might not be a conceptually accurate measure of the “price level,” even if it yields better results for inflation-targeting central banks. Cecchetti (2007), for instance, points out that OER captures the opportunity cost of living in a house; whereas, a measure of the level of house prices confuses increases in cost with increases in the value of assets.

Thus, redefining the CPI in countries that have a legal mandate to target inflation could be a useful strategy for preventing the inflation of housing bubbles, even if doing so is not entirely conceptually accurate. It will permit the central bank to continue to communicate a clear target for a highly visible measurement of inflation; and it will make the choice to continue with rate hikes easier, since the public will see that the

³³See Osborne's (2007) letter to King.

rate of inflation is high, even if non-house price growth is contained. If house prices are not explicitly included in the index, adding OER could still provide a significant improvement for central banks that target an inflation measure that does not include the cost of housing at all.

7.3 Solution Algorithms

7.3.1 The Rouwenhorst Method

I use a number of routines to generate and calibrate the exogenous processes in the model. Central to most of these routines is the Rouwenhorst method (1995), which is used to discretize AR(1) processes into conditional Markov processes. Kopecky and Suen (2009) argue that this method is more accurate than other approximation algorithms for highly persistent AR(1) processes. Additionally, once the algorithm is constructed, we may generate a Markov process (i.e. the series of TFP shocks) by setting four parameters only: ρ , q , σ , and μ , where ρ is the probability of the highest state, q is the probability of the lowest state, σ is the desired standard deviation of the process, and μ is the mean of the process.

The Rouwenhorst method works by creating an initial Markov transition matrix for a 2-state approximation to an AR(1) process. It then uses that matrix as an input for the next step, where a 3-state approximation is created. This recursion continues until the desired level of discretization is achieved. In our case, exogenous processes are approximated with 7-state chains, which is standard in the literature. As outlined in Kopecky and Suen (2009), the process works as follows:

Step 1: Create the initial 2-state Markov transition matrix.

$$P_2 = \begin{bmatrix} p & 1-p \\ 1-q & q \end{bmatrix}$$

Step 2: Next, construct the 3-state Markov matrix as follows:

$$P_3 = \begin{bmatrix} p^2 & 2p(1-p) & (1-p)^2 \\ p(1-q) & pq + (1-p)(1-q) & q(1-p) \\ (1-q)^2 & 2q(1-q) & q^2 \end{bmatrix}$$

Step 3: Construct the nth-state Markov matrix in two steps. First, compute the following:

$$\begin{bmatrix} P_{n-1} & \mathbf{0} \\ \mathbf{0}' & 0 \end{bmatrix} + (1-p) \begin{bmatrix} \mathbf{0} & P_{n-1} \\ 0 & \mathbf{0}' \end{bmatrix} + (1-q) \begin{bmatrix} \mathbf{0}' & 0 \\ P_{n-1} & \mathbf{0} \end{bmatrix} + q \begin{bmatrix} 0 & \mathbf{0}' \\ \mathbf{0} & P_{n-1} \end{bmatrix}$$

Next, divide all rows 2 through (n-1) by 2 to complete the nth-state Markov matrix. Note that $\mathbf{0}'$ denotes a (n-1) row vector and $\mathbf{0}$ denotes a (n-1) column vector. This will yield an AR(1) process with an autoregressive coefficient of $p+q-1$.

Note that the Markov chain associated with the probability matrix is constructed by discretizing the interval $[\mu - v, \mu + v]$ into the same number of states as the matrix, where $v = \sqrt{\frac{n-1}{(p+q-1)^2-1}}\sigma$. This will generate a process with a standard deviation of σ . In our case, the mean is 1 the standard deviation is .05.

As an example of the Rouwenhorst Method, consider Figures 22 and 23. The first is a 150-period series with an autoregressive coefficient of .8 and a mean of 1. The second graph is a discrete, 7-state Rouwenhorst approximation.

Note also that the standard deviation of the approximated series over 2000 periods is .0491 and the mean is .9959, which suggests that the 7-state chain provides a

reasonably accurate approximation to an actual AR(1) process with the same mean, standard deviation, and autoregressive coefficient.

7.3.2 Unemployment Rate

The computation of individual unemployment spells and aggregate unemployment is done in the following steps:

Step 1: Draw a series of consumption productivity shocks.

Step 2: For each individual and time period, use the conditional Markov chain to compute the probability of becoming employed given the current consumption productivity shock and the previous period's employment state. Note that the probability of becoming employed will be higher if the agent was employed last period and if the consumption productivity shock is larger.

Step 3: For each individual and time period, draw a real number from a uniform distribution on the interval $[0,1]$. If the number drawn is greater than the conditional probability of employment, then the agent becomes unemployed. Otherwise, he becomes employed.

Step 4: Compute the unemployment rate as the fraction of agents who are employed in a given period.

Step 5: Adjust the conditional Markov transition matrix probabilities and re-simulate until individual-level unemployment spell profiles and the unemployment rate match calibration targets.

Note that this model is calibrated to generate an unemployment rate that fluctuates between 4% and 12%. Figure 25 shows a 50-period sample of the unemployment rate generated using this method.

7.3.3 Ergodic Rate of Unemployment

Since the ergodic rate of unemployment is used as an input in the solution algorithm, it is necessary to compute it prior to solving the model. With Markov processes, we can typically compute the steady state analytically; however, since there is a different Markov transition matrix for each TFP state, I instead solve for the ergodic rate of unemployment by simulating an individual employment path for over 50,000 periods and computing the fraction of periods unemployed.

I re-simulated and recalibrated to achieve a target ergodic unemployment rate of 6.3%. Figure 24 shows convergence in the employment rate after 5,000 simulation periods.

7.3.4 Krusell and Smith (1998) for Unemployment

In the model, the future unemployment rate affects factor prices, which means that agents must be able to form expectations about future unemployment rates. We use the Krusell and Smith algorithm in the following way to compute the law of motion for unemployment:

Step 1: Compute the unemployment rate in each period for a 5,000 period simulation of the economy.

Step 2: Update the law of motion coefficients. Construct a set of regressors by interacting the previous period unemployment rate and a consumption sector shock state-specific constant with a binary variable that equals one if the current shock state is j and 0 otherwise.

In this case, the estimating equation is the following, where $UR(t)$ denotes the unemployment rate at time t , $S(t)$ denotes the index of the consumption sector state at time t , $C(j)$ denotes a state-specific constant, $G(j)$ denotes the coefficient on capital, $\mathbf{1}\{.\}$ denotes an indicator function, and ns denotes the number of states:

$$UR(t) = \sum_{j=1}^{ns} C(j)\mathbf{1}\{S(t) = j\} + \sum_{j=1}^{ns} G(j)UR(t-1)\mathbf{1}\{S(t) = j\} \quad (50)$$

This law of motion is then checked for fit. If R^2 is sufficiently high, then the law of motion is adopted—and agents use it to form expectations about future factor prices. If the fit is poor, then a longer simulation is performed and the law is re-estimated.

Figure 25 shows the predicted unemployment rate for 150 periods.

7.4 Solution Method

Since the model contains many individual-level and aggregate-level continuous states, I employ three methods to make it computational feasible. First, I use GPU computing as outlined in Aldritch et. al (2010) to perform all matrix operations in my solution method.³⁴ Next, I use an approximate dynamic programming algorithm for simulations that would otherwise have prohibitively long run times. And, finally, I use a modification of Judd, Maliar, and Maliar (2010) to limit the household’s choice

³⁴These operations are performed using an Nvidia GeForce GTX 560 Ti 2b with 384 CUDA cores.

set to the ergodic set in the cases where I do not solve using backwards recursion.

The DDP is solved for households in one of two ways. First, in the cases where we have no capital or endogenous house prices, I solve the problem for the household using backwards recursion with cubic polynomial interpolation. For the more involved simulations, I use an approximate dynamic programming algorithm that is a modification of an algorithm in Powell (2007). It involves the following steps:

7.4.1 ADP Algorithm

Step 1: Draw a sequence of exogenous shocks for 5,000 periods and agents.

Step 2: Initialize the future value of all post-decision states and simulate the values of the aggregate variables using the initial laws of motion.

Step 3: For each household, step forward in time by choosing consumption, housing, capital, deposits, mortgage debt, and whether or not to default.

Step 4: After each household's choice, update the value of the post-decision state using the following equation:³⁵

$$V(s) = \alpha_j V_{new}(s) + (1 - \alpha_j) V_{old}(s) \quad (51)$$

Note that j denotes the number of times we have iterated through all households.

Step 5: After each household's problem is solved using the forward-pass portion of

³⁵Note that the new value of the post-decision state will always be different, since it will depend on the particular set of shocks drawn; however, $V(s)$ will start to approximate the expectation of the pre-decision state as we visit s more times and with different sets of shocks. This is why a variable step size α_j is used, which declines as the iteration number increases

the algorithm, we perform the back-propagation step. That is, we use the updated, post-decision state values to perform backwards recursion. After each household's step, we update the value of being in each post-decision state. Notice that—unlike the forward step—this gives us an unbiased estimate of being in each of these states.

Step 6: For all states, evaluate the post-decision value and compute the change from the previous iteration. If the maximum absolute change falls below the specified tolerance, terminate the algorithm and update the laws of motion for the aggregate states. If the tolerance criterion is not satisfied, then return to step 1.

7.4.2 Verification

In all cases where I use the ADP algorithm to solve the household's problem, I also solve the problem once for a baseline case using backwards recursion and cubic polynomial interpolation. This allows me to verify the quality of the approximation without solving and simulating 150 times using the more computationally-intensive approach.

7.4.3 Krusell & Smith (1998) Algorithm

After the optimal path for agents is determined in each step, the aggregate variables are updated, and their laws of motion are estimated using a modification the Krusell and Smith algorithm that uses neural networks.

7.4.4 State-Contingent Pricing Algorithm for Mortgage Market

In order to select a value of ξ_t that generates zero net cashflows for the financial intermediary, I use state-contingent pricing. That is, at each point in time, the financial intermediary observes Z_t and sets a corresponding $\xi(Z_t)$ to set cashflows equal to

zero.³⁶

The algorithm consists of the following steps:

Step 1: Set $\xi(Z_t)$ equal to the risk-free rate plus a fixed premium in all periods.

Step 2: Solve and simulate the model. Compute the average net cashflow for each set of aggregate states, Z_t .

Step 3: Modify the state-contingent mapping by setting $\xi(j) = \xi(j) + CO * NCF(j)$, where $CO > 0$ and $NCF(j)$ is the net cash-flow in state j . That is, if net cashflows are positive, then reduce $\xi(j)$ by a small number that is proportional to the net cashflow. If they are negative, then increase $\xi(j)$ by a small number that is proportional to the net cashflow.³⁷

Step 4: Repeat step 2. Check whether the maximum absolute net cashflow in each state falls below the tolerance value.

Step 5: Repeat steps 3 and 4 until all states have a maximum absolute net cashflow that falls under the tolerance value.

Note that this algorithm is nested within the main solution method; and will guarantee that deviations from market clearing in the mortgage market are small.

³⁶Note that cashflows deviate from zero in some periods, but deviations are at least two orders of magnitude smaller than the aggregate housing stock, which we use for comparison.

³⁷As mentioned in the section on financial intermediaries—the primary impact of ξ_t will be on borrowing at time t . Thus, increasing ξ_t will tend to increase net cashflows by reducing borrowing.

7.4.5 Price Level Approximation Using a Single Layer, 20-Node Neural Network

In order to compute the current price level, we must be able to solve for the expected value of a function of the next period price level, unit cost, and output. In order to simplify the computation, I apply approximation methods. In particular, I compute the price level as a state-contingent mapping using the following algorithm:

Step 1: Initialize the price level for all states.

Step 2: Solve the household's problem and perform the aggregation step using the state-contingent mapping for P_t .

Step 3: Compute \bar{P}_t and then P_t in each period, given the values of P_{t+1} , V_{t+1} , and Y_{t+1} from the previous iteration.

Step 4: Train a single-layer, 20-node neural network to approximate the nonlinear relationship between the aggregate states of the economy and P_t . Use the neural network to update the state-contingent mapping for P .

Step 5: Compute the maximum absolute difference between the new and old P in each state. If the maximum absolute difference is below the tolerance threshold, terminate this step and continue with the main algorithm. Otherwise, go back to Step 2.

Note that this part of the algorithm could be done by regressing P_t on basis functions.

I find that neural networks work particularly well because they have a low computational burden, provide a highly accurate nonlinear approximation, and do not require a substantial amount of restrictions the function being approximated.

CHAPTER 2

APPROXIMATE DYNAMIC PROGRAMMING WITH POST-DECISION STATES AS A SOLUTION METHOD FOR DSGE MODELS

Isaiah Hull

ABSTRACT

Dynamic programming (DP) is often used to solve DSGE models when a global solution is desired. DP using forward iteration (infinite horizon problems) or backwards recursion (finite horizon problems) will satisfy the Contraction Mapping Theorem (CMT) and its accuracy can be adjusted through the coarseness of the state space grid and the size of the convergence criterion. However, both of these common solution methods rely on the discretization of the state space and thus suffer from the curse of dimensionality, making them undesirable for problems with many continuous state variables. In macroeconomics, modifications of these approaches that parameterize expectations (PEA), parameterize some set of functional equations (projection methods), or interpolate between state space nodes are often employed to yield a higher degree of accuracy for a given program run time. In this paper, I compare the speed and accuracy of the aforementioned modifications to an alternative method that was introduced separately by Judd (1998), Sutton and Barto (1998), and Van Roy et al. (1997), but was not developed into a general solution method until Powell (2007) introduced it to the Operations Research literature. This approach involves rewriting the Bellman equation in terms of the post-decision state variables, rather than the pre-decision state variables, as is done in standard dynamic programming applications in economics. I show that this approach yields considerable performance benefits over common global solution methods when the state space is large; and has the added benefit of not forcing modelers to assume a data generating process for shocks. In addition to this, I construct two new algorithms that take advantage of this approach to solve heterogeneous agent models.

JEL Classification: C60, C61, C63

Keywords: Dynamic Programming, Projection Methods, Parameterized Expectation Approach, Value Function Iteration, Heterogeneous Agents, Dynamic Stochastic General Equilibrium Models, Solution Methods

1 Introduction

Dynamic programming (DP) was originally developed as a solution method for multi-period decision problems in Operations Research. The purpose of dynamic programming was to break down complex problems into a series of smaller, tractable problems. Prior to DP, solving multi-period decision problems required the complete enumeration and evaluation of all possible sequences of decisions. Bellman's innovation was to reduce the dimensionality of such decision problems through the Principle of Optimality (Bellman 1954):

PRINCIPLE OF OPTIMALITY. An optimal policy has the property that whatever the initial state and initial decisions are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decisions.

The practical implication of this principle is that we can rule out many decision paths when solving multi-period problems. This reduces the number of functional evaluations we must perform in order to determine the optimal choice of controls at a given state. For example, consider a decision maker who is in state $s_0 \in \mathbf{s}$ in a deterministic choice problem; and is considering choosing controls $c_1 \in \mathbf{c}$, which will move her to state s_1 . In order for her to know the discounted lifetime utility derived from that choice, she must not only be able to compute the instantaneous utility associated with the choice, but she must also know the choices she will make in all future states, contingent upon arriving at s_1 . The optimality principle makes this choice problem simpler by ruling out all future decision paths that are suboptimal. That is, she can now compare the discounted lifetime utilities associated with different control choices in the present period, under the assumption that she will follow an optimal path after her choice of controls move her to the next state. Furthermore, she can compute the

value of being in the current state, s_0 , under the assumption that she will follow the optimal path starting in the current period.

To understand how this simplifies a typical multi-period choice problem, consider how deterministic, finite horizon choice problems are solved using dynamic programming. We start with the final period, where the path of optimal future decisions is clear. That is, the agent dies and there are no decisions to make. Thus, utility is zero in all future periods. This suggests the optimal decision in this period is to maximize instantaneous utility, given the resources available (state) at the start of the period. Furthermore, it allows us to move to the second-to-last period of time—and solve the same problem again, under the assumption that we will maximize instantaneous utility in the period following. We can now follow this backwards iteration process until the initial period.

As an example of the principle of optimality, consider a three-period, finite horizon problem, where the only state variable is the amount of the consumption good stored in the previous period, s . For simplicity, assume that a household can choose to either store one unit ($s = 1$) or none ($s = 0$). Furthermore, assume that the household starts with no stored units of the consumption good ($s = 0$), receives a deterministic endowment of the consumption good, y , in the first two periods, and must choose how much to consume, c , in each period. If we solve this problem without the Principle of Optimality, we do not rule out any possible decision paths, which means that we must enumerate and evaluate four decision paths in order to choose controls in the initial period: $\{s_1, s_2\} = \{1, 0\}, \{1, 1\}, \{0, 1\}, \{0, 0\}$. However, if we apply the Principle of Optimality, then we can rule out paths where the household stores consumption in the third period—and thus gives up third period consumption with no future gain—leaving us with the following possible decision paths in the initial period: $\{s_1, s_2\} = \{1, 0\}, \{0, 0\}$. Notice that applying the Principle of Optimality has

reduced the dimensionality of the household’s second period decision problem by half.

While the Principle of Optimality allows us to reduce the dimensionality of DP problems, it does not eliminate the problem entirely. In fact, numerical dynamic programming is said to suffer from “the curse of dimensionality” when there are many state variables. That is, the number of function evaluations needed to identify the optimal controls at a given state will grow rapidly in the number of state variables. For example, if we have N state variables that can each take on two values, then the dimensionality of the decision problem¹ at a given state is 2^N . That is, each additional state variable will double the dimensionality of the problem. Furthermore, if we add a continuous state variable, such as the capital stock, the problem will become infinite-dimensional, which will require us to use approximation methods. The most common way to deal with this problem is to discretize the continuous state variables. However, this approach quickly becomes computationally infeasible when the model has many continuous state variables. This problem—the curse of dimensionality—places a limit on the amount of detail that can be incorporated into DSGE models.

One common strategy for reducing dimensionality is to use a coarse grid to approximate continuous state variables. If, for instance, we want to target a program run time of 12 hours, then we might successively decrease the number of discrete states that represent continuous state variable until the program can reliably converge within our targeted time limit. The tradeoff, of course, is that the accuracy of the solution will be lower than what could be achieved if the continuous state were discretized more finely.² Alternatively, we might consider some minimum tolerance

¹To be more precise, the dimensionality of the problem should also incorporate the size of the action space—that is, the dimensionality of the controls. We can write the dimensionality of the problem more precisely as $2^N x |c|$, where $|c|$ is the dimensionality of the controls.

²For example, the first order conditions might imply that the household’s bond holdings should be .55, but if continuous bond holdings are represented by a five-node discrete grid over the interval $[0,1]$ that does not include .5 (e.g. $\{0,.25,.5,.75,1\}$), then this won’t be possible. As a result, bond holdings will deviate slightly from optimality.

level or criterion for accuracy; and then set the coarseness of the grid in order to achieve it. We can then compare the time needed to achieve this criterion under different approaches.

Another strategy for dimensionality-reduction involves using a highly coarse state space grid, but interpolating between nodes in the state space. The previous approach—sometimes referred to as a “look-up table” approach—restricts the set of maxima candidates to a discrete set of states. In contrast, this approach allows the maximum to fall between two nodes. In many cases, this approach will outperform the “look-up table” approach by allowing us to satisfy the same convergence criterion or tolerance in a shorter period of time. However, it will become increasingly costly in stochastic problems as more exogenous states are added.

Other approaches to dimensionality-reduction involve constructing policy function approximations (i.e. functions that map states to decisions) that minimize the ex-post residuals of the Euler equations. Often, such approximations are done using a family of orthogonal polynomials, such as Chebychev polynomials, to transform the state variables. Some other options involve using collocation methods, splines, or finite element methods to approximate the policy or value functions. One benefit to this approach is that it gives us a continuous approximation of the policy or value function, rather than a “look-up table.” A major drawback, however, is that it is often difficult to choose a good initial parameterization.

Finally, when dynamic programming problems are stochastic, we face another major challenge: the time t state and choice of controls will not be sufficient to determine the state at time $t+1$, since shocks will arrive after choices are made in period t or at the start of period $t+1$. Rather, the state and choice of controls—combined with knowledge of the data generating processes for the exogenous variables—will determine the *distribution* of future states—and, therefore, distribution of future values.

In order to solve such problems, we will have to take an expectation over the value of the next period state. Typically, when the exogenous processes in the model are continuous, we will deal with the expectation by using quadrature methods or discretizing the process. The discretization approach, which is more computationally tractable (and, relatedly, more popular), works well when there is a simple exogenous process with known properties; however, it becomes increasingly ineffective and computationally costly when there are multiple exogenous processes with complicated dependence structures. This places substantial limitations on what can be done with DSGE models moving forward.

The purpose of this paper is to explore a recently-developed technique from the Operations Research (OR) literature (Powell 2007) and the Engineering literature (Bertsekas 2011) as a method for solving DSGE models; and as means of overcoming the aforementioned problems associated with more traditional approaches to numerical dynamic programming. In particular, this paper will describe the technique, compare its performance with common solution methods in computational macroeconomics, suggest where it can be most usefully employed, and examine several algorithmic refinements. In the OR literature, this method is referred to as approximate dynamic programming with post-decision states (hereafter, ADP); and entails writing the Bellman Equation around post-decision state variables, rather than pre-decision state variables, as is typically done in macroeconomic applications.

2 Literature Review

In this paper, I will introduce and evaluate several new algorithms for solving dynamic programming problems nested within DSGE models. In particular, these algorithms will be built around a post-decision state version of the Bellman equation, rather

than the pre-decision state version. This approach will enable practitioners to solve two classes of models that are typically computationally intractable: 1) models with many continuous state variables; and 2) models with several correlated exogenous processes. However, the formal tests will focus on the former. Furthermore, this method is substantially easier to implement than other solution methods that do not suffer from the curse of dimensionality, such the parameterized expectations approach and projection methods.

The baseline method used in this paper was introduced by Van Roy et al. (1997) in a neuro-dynamic programming application. It was first used to solve a retail inventory management problem with 33 state variables. Standard dynamic programming methods, such as value function iteration (infinite horizon) and backwards recursion (finite horizon) could not feasibly solve problems with such large state spaces. Even if all of the variables were discretized into the coarsest possible grid, it would still contain 2^{33} (8.59 billion) nodes. In comparison, the neuro-dynamic programming algorithm they constructed around post-decision states delivered decision rules that yielded substantial efficiency improvements over more heuristic approaches that were used in the absence of a formal decision rule.

This technique first entered the economics literature through Judd (1998), which explained how to construct the post-decision state Bellman equation. However, Judd (1998) did not exploit this method to construct any solution algorithms. Similarly, Sutton and Barto (1998) introduced the technique to the reinforcement learning literature, but did so originally without any substantial applications. Rather, it was used primarily as a strategy for absorbing the action space into the state space.³

³In many dynamic programming applications, the endogenous component of the state space is determined by the realizations of the shock at the start of the period and the controls selected thereafter. In some of these applications, the dimensionality of the problem can be reduced simply by eliminating the choice of controls; and instead considering only the result of those choices (i.e. the end-of-period endogenous states).

Powell (2007) developed the method introduced by Van Roy et al. (1997) into a set of solution algorithms for high dimensionality optimization problems. However, all of the algorithms in Powell (2007) were designed for a partial equilibrium context in OR; and most focused on a particular type of problem that is not common in general equilibrium macroeconomic models: inventory management. In problems of this variety, an agent faces random demands for a product and must choose how much inventory to hold. The inventory constraint in such models looks as follows:

$$I_{t+1} = \max\{I_t + R_t - \epsilon_{t+1}, 0\}$$

Here, I_t is the current period inventory level, R_t is the amount of product the agent adds to inventory in the current period, and ϵ_{t+1} is the amount of product consumers want to purchase at the start of the next period. Powell (2007) suggests that this variety of problem lends itself well to a post-decision representation of the state space, since both the control and shock affect the state in a simple, additive way. Using this approach allows us to write the post-decision state as $\hat{R}_t = I_t + R_t$. It is important to note that, in general, macroeconomic dynamic programming problems do not take this form. With the possible exception of models that include multiple assets, this particular benefit of writing problems in terms of post-decision states will not apply to most work in macroeconomics.

Within OR, Powell (2012) provides the most recent survey of work that has been done using post-decision state dynamic programming algorithms. Papers such as Maxwell (2011) and Simao et al. (2009) use the post-decision state formulation to solve problems that involve ambulance deployment and large-scale fleet management respectively. Papers such as Powell and Ryzhov (2010) extend the initial algorithm by demonstrating how Bayesian methods can be used to update the value function

at all states, rather than just the states visited on the simulation path.

Within the computational macroeconomics literature, my paper is most closely associated with work that attempts to overcome the curse of dimensionality. In particular, techniques such as the Parameterized Expectations Approach (PEA), introduced by Sargent (1987), Marcet (1988), and Den Haan and Marcet (1990), and Projection Methods, introduced by Judd (1992) and McGrattan (1999), provide general strategies for solving optimization problems with large state spaces. More recent work, such as Maliar and Maliar (2005) and Judd et al. (2011), demonstrates how to improve the stability and speed of such approaches. Papers such as Krusell and Smith (1998), Den Haan and Rendahl (2010), Maliar, Maliar, and Valli (2010), and Algan, Allais, and Den Haan (2008) provide algorithmic strategies for solving models that have large state spaces as a consequence of having many heterogeneous agents.

Heer and Maussner (2008) evaluates existing solution algorithms for standard business cycle models with flexible labor supply. In particular, they compare value function iteration, the extended deterministic path method, log-linearization around the steady state, the parameterized expectations approach, and projection methods. They conclude that log-linearization around the steady state is sufficiently accurate if the practitioner is primarily interested in generating business cycle moments. The other approaches mentioned presented substantial computational challenges; and yielded very little gain in terms of improvements in simulated business cycle moment accuracy.

I extend the aforementioned literature by 1) introducing a global solution method for DSGE models from the OR literature, 2) providing several refinements to the solution method, and 3) evaluating it relative to common global solution methods. In addition to this, I demonstrate when the algorithm can most usefully be employed; and also discuss when it is unlikely to outperform more traditional approaches.

3 Mathematical Preliminaries: The Pre-Decision State Bellman Equation

I will start by describing the formulation of the Bellman Equation with pre-decision state variables; and will then move on to the approach for post-decision state variables. This construction of the Bellman Equation will focus on an infinitely-lived agent who faces a deterministic choice problem, which requires her to choose a set of controls, c_t . A set of states, s_t , summarizes all information needed to make a choice at time t . Combined with c_t , s_t pins down the future state according to a transition function: $s_{t+1} = T(c_t, s_t)$. Finally, a set of constraints, $c_t \in \Gamma(s_{t-1})$, limits the choice of controls. After applying the Principle of Optimality as described in the introduction and originally in Bellman (1954), the value of being in the initial state can be written as follows:

$$\begin{aligned}
 V(s_0) &= \max_{\{c_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c_t) & (1) \\
 \text{s.t. } & s_{t+1} = T(c_t, s_t), c_t \in \Gamma(s_{t-1}), \forall t = 0, 1, \dots
 \end{aligned}$$

That is, the value of being in the initial, pre-decision state (i.e. before controls are selected) is the maximum, discounted lifetime utility that can be achieved through an infinite sequence of control choices. Note that computing $V(s_0)$ is a difficult task, since it depends on the optimal, infinite sequence of controls. However, Bellman's important insight was that we could rewrite $V(s_0)$ as follows:

$$\begin{aligned}
 V(s_0) &= \max_{c_0} \{u(c_0) + \beta[\max_{\{c_t\}_{t=1}^{\infty}} \sum_{t=1}^{\infty} \beta^t u(c_t)]\} & (2) \\
 \text{s.t. } & s_{t+1} = T(c_t, s_t), c_t \in \Gamma(s_{t-1}), \forall t = 0, 1, \dots
 \end{aligned}$$

Now, notice that we can exploit the recursive property of this equation to rewrite it as follows:

$$V(s_0) = \max_{c_0} \left\{ u(c_0) + \beta \underbrace{\left[\max_{\{c_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c_{t+1}) \right]}_{V(T(s_0, c_0))} \right\} \quad (3)$$

s.t. $s_{t+1} = T(c_t, s_t), c_t \in \Gamma(s_{t-1}), \forall t = 0, 1, \dots$

$$\rightarrow V(s_0) = \max_{c_0} u(c_0) + \beta V(s_1) \quad \text{s.t. } s_1 = T(c_0, s_0), c_0 \in \Gamma(s_0) \quad (4)$$

This equation—the Bellman equation—is much simpler in the sense that it permits us to break a multi-period decision problem down into a series of smaller, tractable steps, where only one set of controls must be selected at a time. However, it is more complicated in the sense that it now requires knowledge of an unknown function, V . Analytical dynamic programming solves this problem by recovering the algebraic expression for the value function, V , but is limited in application. In contrast, numerical dynamic programming—the focus of this paper—constructs an approximation of V and is applied to a wide variety of multi-period choice problems.

Next, I will augment household's problem to incorporate exogenous processes, which are denoted by $x_{t+1} = f(x_t, \epsilon_{t+1})$. The subscript $t+1$ indicates that the exogenous shocks arrive at the start of period $t+1$ and before time $t+1$ controls are chosen. Note that the joint transition function for states now captures the impact of exogenous shocks: $s_t = T(c_t, s_{t-1}, x_{t-1}, \epsilon_t)$.⁴ Furthermore, I assume that the feasible set of control choices also depends on the exogenous process: $c_t \in \Gamma(s_{t-1}, x_t)$. The

⁴This transition function is written as a joint process; however, it is important to note that s_{t+1} depends on x_t and s_t , but x_{t+1} is independent of s_t and c_t .

value of being in the initial state may now be written as follows:

$$V(s_0, x_1) = \max_{\{c_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t E[u(c_t)|s_0, x_1] \quad (5)$$

s.t. $s_t = T(c_t, s_{t-1}, x_t)$, $c_t \in \Gamma(s_{t-1}, x_t)$, $\forall t = 0, 1, \dots$

Again, I will exploit the recursive nature of $V(s_0, x_1)$ to rewrite it as follows:

$$V(s_0, x_1) = \max_{c_0} \left\{ u(c_0) + \beta \underbrace{\left[\max_{\{c_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^{t-1} E[u(c_{t+1})|s_0, x_1] \right]}_{E[V(T(c_1, s_0, x_1))]} \right\} \quad (6)$$

s.t. $s_t = T(c_t, s_{t-1}, x_t)$, $c_t \in \Gamma(s_{t-1}, x_t)$, $\forall t = 0, 1, \dots$

This yields the stochastic version of the Bellman equation:

$$V(s_0, x_1) = \max_{c_0} u(c_0) + \beta E[V(s_1, x_2)] \quad (7)$$

s.t. $s_1 = T(c_1, s_0, x_1)$, $c_1 \in \Gamma(s_0, x_1)$, $x_2 \in f(x_1, \epsilon_2)$

More generally, the Bellman equation can be written as follows:

$$V(s_{t-1}, x_t) = \max_{c_t} u(c_t) + \beta E[V(s_t, x_{t+1})] \quad (8)$$

s.t. $s_t = T(c_t, s_{t-1}, x_t)$, $c_t \in \Gamma(s_{t-1}, x_t)$, $x_{t+1} \in f(x_t, \epsilon_{t+1})$

The derivations above yield both the deterministic and stochastic versions of the standard Bellman Equation in terms of pre-decision state variables. In the next section, I will explain how to rewrite the Bellman equation in terms of post-decision states. This mathematically simple modification will afford computational modelers a tremendous amount of flexibility; and will render otherwise intractable DSGE models solvable.

4 The Post-Decision State Bellman Equation

In this section, I will construct the Bellman equation in terms of post-decision state variables, roughly following (but also expanding on) the treatment of this subject in Powell (2007) and Bertsekas (2011). Additionally, I will focus on building an adaptation that is well-suited to the timing conventions and model properties of standard DSGE models, rather than OR or Engineering applications. This exposition will begin with a comparison of pre-decision and post-decision states to make this distinction clear. A pre-decision state consists of last period's endogenous state variables, s_{t-1} , and this period's exogenous processes, x_t . This state is "pre-decision" because it is determined entirely before the time t controls, c_t , are selected. In contrast, the post-decision state consists of this period's endogenous state variables, s_t , and exogenous processes, x_t . Recall from the construction of the Bellman equation in the previous section that $s_t = T(c_t, s_{t-1}, x_t)$. That is, near the start of time t , s_{t-1} and x_t are pinned down. Once c_t is chosen, s_t will be pinned down, too. This yields the post-decision state—that is, the state immediately after we have selected controls. Figure 1 in the appendix shows this relationship through the arrival of new information in discrete time model.

Note that the *value* of a post decision state is the maximum, expected, discounted utility that the agent can achieve immediately after controls have been selected. Using this definition, the post-decision state value function, $V^x(s_{t-1}, x_{t-1})$, can be written as follows:

$$V^x(s_{t-1}, x_{t-1}) = E\{\max_{\{c_{t+s}\}_{s=0}^{\infty}} \sum_{s=0}^{\infty} \beta^s E[u(c_{t+s}) | s_{t-1}, x_t] | s_{t-1}, x_{t-1}\} \quad (9)$$

$$\text{s.t. } s_{t+s} = T(c_{t+s}, s_{t+s-1}, x_{t+s}), \quad c_{t+s} \in \Gamma(s_{t+s-1}, x_{t+s}), \quad x_{t+s} \in f(x_{t+s-1}, \epsilon_{t+s}), \quad \forall s = 0, 1, \dots$$

Notice that the agent will not choose c_t until x_t has been realized. This is why $V^x(s_{t-1}, x_{t-1})$ can be written as the expectation, conditional on the information set

s_{t-1}, x_{t-1} , of the maximum, expected, discounted utility that the agent will receive after x_t arrives.

With this definition in place, it is now possible to simplify the problem by reorganizing the terms inside of the outermost expectation, just as we did for the standard pre-decision Bellman equation.

$$V^x(s_{t-1}, x_{t-1}) = E\{max_{c_t}\{u(c_t) + max_{\{c_{t+s+1}\}_{s=0}^{\infty}} \beta E[\sum_{s=0}^{\infty} \beta^s u(c_{t+s+1})|s_{t-1}, x_t]\}|s_{t-1}, x_{t-1}\} \quad (10)$$

s.t. $s_{t+s} = T(c_{t+s}, s_{t+s-1}, x_{t+s}), c_t \in \Gamma(s_{t-1}, x_t), x_t \in f(x_{t-1}, \epsilon_t), \forall t = 0, 1, \dots$

Next, using the Law of Iterated Expectations (LIE) yields the following:

$$max_{\{c_{t+s+1}\}_{s=0}^{\infty}} E[\sum_{s=0}^{\infty} \beta^s u(c_{t+s+1})|s_{t-1}, x_t] = \quad (11)$$

$$max_{\{c_{t+s+1}\}_{s=0}^{\infty}} E\{E[\sum_{s=0}^{\infty} \beta^s u(c_{t+s+1})|s_t, x_{t+1}]|s_{t-1}, x_t\}$$

This expression may be rewritten:

$$max_{\{c_{t+s+1}\}_{s=0}^{\infty}} E\{E[\sum_{s=0}^{\infty} \beta^s u(c_{t+s+1})|s_t, x_{t+1}]|s_{t-1}, x_t\} = \quad (12)$$

$$E\{\max_{\{c_{t+s+1}\}_{s=0}^{\infty}} E[\sum_{s=0}^{\infty} \beta^s u(c_{t+s+1})|s_t, x_{t+1}]|s_{t-1}, x_t\}$$

Next, notice that the last expression can be rewritten in terms of a value function:

$$E\{\max_{\{c_{t+s+1}\}_{s=0}^{\infty}} E[\sum_{s=0}^{\infty} \beta^s u(c_{t+s+1})|s_t, x_{t+1}]|s_{t-1}, x_t\} = \quad (13)$$

$$E\{V(s_t, x_{t+1})|s_{t-1}, x_t\}$$

This permits a rewriting of equation (10) as follows:

$$V^x(s_{t-1}, x_{t-1}) = E\{\max_{c_t}\{u(c_t) + \beta E[V(s_t, x_{t+1})|s_{t-1}, x_t]\}|s_{t-1}, x_{t-1}\} \quad (14)$$

$$\text{s.t. } s_t = T(c_t, s_{t-1}, x_t), c_t \in \Gamma(s_{t-1}, x_t), x_t \in f(x_{t-1}, \epsilon_t),$$

Finally, (14) can be written in terms of $V(s_{t-1}, x_t)$ and simplified by the the Law of Iterated Expectations:

$$V^x(s_{t-1}, x_{t-1}) = E\{V(s_{t-1}, x_t)|s_{t-1}, x_{t-1}\} \quad (15)$$

$$\text{s.t. } s_t = T(c_t, s_{t-1}, x_t), c_t \in \Gamma(s_{t-1}, x_t), x_t \in f(x_{t-1}, \epsilon_t),$$

Stepping forward in time yields the following:

$$V^x(s_t, x_t) = E\{V(s_t, x_{t+1})|s_t, x_t\} \quad (16)$$

$$\text{s.t. } s_t = T(c_t, s_{t-1}, x_t), c_t \in \Gamma(s_{t-1}, x_t), x_{t+1} \in f(x_t, \epsilon_t),$$

Finally, to demonstrate the connection between the pre-decision and post-decision Bellman equations, I prove the following claim:

$$V(s_{t-1}, x_t) = \max_{c_t} u(c_t) + \beta V^x(s_t, x_t) \quad (17)$$

Start by recalling the standard Bellman equation for a stochastic dynamic programming problem:

$$V(s_{t-1}, x_t) = \max_{c_t} u(c_t) + \beta E[V(s_t, x_{t+1})|s_{t-1}, x_t] \quad (18)$$

$$\text{s.t. } s_t = T(c_t, s_{t-1}, x_t), c_t \in \Gamma(s_{t-1}, x_t), x_{t+1} \in f(x_t, \epsilon_{t+1})$$

Next, note that (s_{t-1}, x_t) —the pre-decision state—pins down c_t , which then pins down s_t :

$$s_t = T(\Gamma(s_{t-1}, x_t), s_{t-1}, s_t) \quad (19)$$

This suggests that the information set can be written as (s_t, x_t) :

$$\begin{aligned} V(s_{t-1}, x_t) &= \max_{c_t} u(c_t) + \beta E[V(s_t, x_{t+1}) | s_t, x_t] \\ \text{s.t. } s_t &= T(c_t, s_{t-1}, x_t), c_t \in \Gamma(s_{t-1}, x_t), x_{t+1} \in f(x_t, \epsilon_{t+1}) \end{aligned} \quad (20)$$

Recalling (16) yields the following:

$$\begin{aligned} V(s_{t-1}, x_t) &= \max_{c_t} u(c_t) + \beta V^x(s_t, x_t) \\ \text{s.t. } S_t &= T(c_t, s_{t-1}, x_t), c_t \in \Gamma(s_{t-1}, x_t), x_{t+1} \in f(x_t, \epsilon_{t+1}) \end{aligned} \quad (21)$$

These three derived properties can be used to make statements about the relationship between the pre-decision and post-decision Bellman equation. If (16) is substituted into (21), it yields the standard, pre-decision Bellman equation. On the other hand, if (21) is substituted into (15), it yields the following post-decision state Bellman equation:

$$\begin{aligned} V^x(s_{t-1}, x_{t-1}) &= E\{\max_{c_t} u(c_t) + \beta V^x(s_t, x_t) | s_{t-1}, x_{t-1}\} \\ \text{s.t. } s_t &= T(c_t, s_{t-1}, x_t), c_t \in \Gamma(s_{t-1}, x_t), x_t \in f(x_{t-1}, \epsilon_t), \forall t = 0, 1, \dots \end{aligned} \quad (22)$$

There are two important things to note. First, the expectations operator is outside of the maximum operator. This differs from the standard, pre-decision state Bellman equation. And second, the problem within the expectation is deterministic.

Van de Roy (1997), Judd (1998), and Barto and Sutton (1998) all separately

identified this alternative form of the Bellman equation, but did not use it to develop solution algorithms. Powell (2007) and Bertsekas (2011) provided the first general descriptions of how it could be used to construct solution algorithms for OR and Engineering applications respectively. In particular, they demonstrated that it was most valuable in high-dimensional dynamic programming problems, as well as problems where the stochastic processes were complicated or unknown. Modern, computational textbook approaches to solving DSGE models, such as those in Heer and Maussner (2009) and DeJong and Dave (2007), do not discuss either the post-decision state Bellman equation or the associated algorithms.

This paper adds four things to the literature. First, it provides formal derivations for the three previously mentioned useful relationships between pre-decision and post-decision Bellman equations.⁵ Second, it introduces the post-decision state Bellman equation and associated algorithms as a general approach to solving DSGE models. Third, it provides a formal comparison of this approach to more common approaches in the literature for global DSGE solution methods. And fourth, it proposes new algorithms for solving DSGE models that use this method. In the following section, I will describe the baseline algorithm that will be used to apply this method.

5 The Post-Decision State Dynamic Programming Solution Algorithm

In the previous section, I demonstrated how to construct the post-decision state Bellman equation. Next, I will discuss two algorithms that take advantage of this alternative approach to dynamic programming. Both of these solution methods fall under

⁵While the derivations are relatively simple, they are not given explicitly in Judd (1998), Barto and Sutton (1998), Powell (2007) or Bertsekas (2011).

the family of solution methods referred to as ADP with post-decision state variables. The first approach is forward dynamic programming with post-decision states.⁶ The second approach is a refinement that incorporates a “backwards propagation” step. The algorithms described in this section will follow those outlined in Powell (2007) and Bertsekas (2011).

Before we begin the explicit construction of the algorithm, it is useful to recall the form of the post-decision Bellman equation from the previous section:

$$\begin{aligned}
 V^x(s_{t-1}, x_{t-1}) &= E\{\max_{c_t} u(c_t) + \beta V^x(s_t, x_t) | x_{t-1}, x_{t-1}\} \\
 \text{s.t. } s_t &= T(c_t, s_{t-1}, x_t), c_t \in \Gamma(s_{t-1}, x_t), x_t \in f(x_{t-1}, \epsilon_t),
 \end{aligned}
 \tag{23}$$

Notice again that the maximization step is inside of the expectation. This property of the post-decision state Bellman equation initially appears to make things more complicated, but actually makes it substantially easier to solve certain classes of models. In particular, it will make it possible to construct algorithms that first perform the maximization step; and then compute the expectation afterwards. This first step is then combined with draws from the exogenous processes, which transition the agent to the next step; and then a smoothing step, which computes the expectation outside of the maximization step. The baseline algorithm is outlined below and follows Powell (2007).

Algorithm 1.1: Forward Dynamic Programming with Post-Decision State Variables (Infinite Horizon, Representative Agent)

Step 0: i. Initialize the “look-up table” approximation of the value function in all

⁶This is sometimes referred to as neuro-dynamic programming or reinforcement learning with post-decision states.

states. That is, for each state, ($s \in \mathbf{s}, x \in \mathbf{x}$), assign a starting value to $\bar{V}^0(s, x)$, where the superscript 0 indicates that this is an approximation taken at the 0th iteration.

ii. Initialize the state, ($s \in \mathbf{s}, x \in \mathbf{x}$). For representative agent, infinite horizon models, it may be natural to assume the the agent starts in the steady state.

iii. Increment the iteration counter to $n=n+1$.

Step 1: Choose a simulation period length, T . Generate a sample of the exogenous processes, X_t , for the simulation period: $1, \dots, T$.

Step 2: For all periods, $t=1, \dots, T$, perform the following three steps:

i. Choose controls to maximize the expression inside of the expectation of $V(s_{t-1}, x_{t-1})$ for a particular realization of x_t :

$$\tilde{V}^n(s_{t-1}, x_{t-1}) = \max_{c_t} u(c_t) + \beta \bar{V}^{n-1}(s_t, x_t)$$

ii. Compute the expectation by updating the value function approximation, V^{n-1} :

$$\bar{V}^n(s_{t-1}, x_{t-1}) = (1 - \alpha_{n-1}) \bar{V}^{n-1}(s_{t-1}, x_{t-1}) + \alpha_{n-1} \tilde{V}^n(s_{t-1}, x_{t-1})$$

Note that α denotes the step size. We index it with n to indicate that it may change with the iteration. It may also be stochastic.

iii. Use the result from the maximization step, s_t , and the realization of the exogenous processes, x_{t+1} , to compute the next period's pre-decision state.

Step 3: Check convergence criterion. If satisfied, go to Step 4. If not satisfied, increment n and go to Step 1.

Step 4: Return the value function approximation, $V(s, x)^N$, in the form of a look-up table.

Now that the basic algorithm has been described, it will be useful to examine Steps 1 and 2 in greater detail. Start with Step 1, where a sample of the exogenous processes is drawn. In a standard DSGE model with a representative agent, these processes might capture things like productivity shocks or monetary shocks. Usually, it is assumed that these exogenous variables follow a continuous process, which can be approximated through discretization. Often, Tauchen's method (1986) is employed to construct a Markov chain and transition matrix that approximate the continuous analog. Alternatively, quadrature methods can be used as a separate strategy for approximation.

One of the benefits of using Algorithm 1.1 is that it is possible to avoid computing expectations using either of these approaches. That is, the household's dynamic programming problem can be solved without ever employing quadrature methods or transforming the assumed continuous, exogenous processes into discrete Markov processes. To understand why this is important, consider the two examples given below.

Example #1: Correlated Exogenous Processes

In a model with several sectors or several assets, it is common to have two exogenous processes that are positively correlated empirically. Solving computational models will often require the assumption that processes are orthogonal for the sake of simplicity. Assuming this makes it possible to take expectations without explicitly accounting for the dependence between the processes. That is, if z and q are the

exogenous variables in the household's problem, then z_{t-1} will yield the pre-shock distribution of z_t , and q_{t-1} will yield the pre-shock distribution q_t . However, if we assume that these processes are dependent, then z_t depends not only on z_{t-1} , but also on q_t and, therefore, q_{t-1} . This not only requires us to use more technically sophisticated quadrature methods or discretization techniques (i.e. conditional Markov processes), but also requires us to use a higher dimensional approximation in order to capture the dependence relationship accurately.

In contrast, if Algorithm 1.1 is employed, then it is not necessary to devise a strategy for computing the expectation. Being able to simulate the exogenous process will be sufficient. This could be done, for instance, by using a copula simulator to generate series for the exogenous processes before the algorithm is even initiated. Step 2(ii) will then compute an approximation of the expectation that will improve as the simulation length increases. This suggests that constructing and solving a model with many exogenous processes that depend on each other will be no more difficult than constructing and solving a model with the same number of independent processes.

Example #2: Exogenous Processes with Unknown Functional Form

Consider a model in which a household may hold many classes of assets. Assume further that the returns to those assets depend on exogenous processes; and that those processes exhibit dependence. Working with a pre-decision state value function, it will be necessary to use technically sophisticated methods (i.e. copula estimation or multiple conditional Markov processes); and then make functional form assumptions about the joint distribution from which shocks are drawn. In contrast, using post-decision state variables, a practitioner can simply perform simulations with the actual asset data without ever performing the intermediate steps.

In short, examples 1 and 2 demonstrate one of the major benefits of using the

post-decision state value function: it makes it possible to bypass difficult expectations without ever explicitly computing them. Furthermore, in some situations, it makes it possible to avoid making any assumptions at all about the exogenous processes in the model, which can both limit model error and use of highly sophisticated techniques.

I will now return to Step 2 in greater detail, which is the other major departure from standard dynamic programming algorithms used in macroeconomics. In particular, in substep (i), controls were selected to maximize a *deterministic* function, which is a substantial departure from what is normally done in stochastic problems. This has the primary benefit of reducing the dimensionality of the problem. Substep (ii) may also look unfamiliar, even though it is similar to smoothing step often used in pre-decision dynamic programming algorithms. Here, it not only plays the role of smoothing changes in the value function approximation, but it also implicitly computes the expected value of being in state (s_{t-1}, x_{t-1}) . To understand why this is the case, notice that the maximum value of $\tilde{V}(s_{t-1}, x_{t-1})$ will depend on the realization of x_t . Thus, arriving at (s_{t-1}, x_{t-1}) for a second, third, and fourth time, but with a different realization of s_t will provide more information about the value of being in state (s_{t-1}, x_{t-1}) . Furthermore, the frequency with which agents encounter certain realizations of x_t will provide information about the probabilities of realizations.

We will now consider a refinement of Algorithm 1.1 that will improve the approximation of the value function, enabling faster convergence for a particular level of accuracy. This approach is best suited to OLG models, but I will also suggest an extension for the infinite horizon case. Furthermore, I will discuss a second refinement that does not appear in Powell (2007) or Bertsekas (2011). Note that tests will only be performed for the infinite horizon case; and that algorithms for both the OLG case are provided for completeness.

**Algorithm 1.2: Double-Pass Forward Dynamic Programming with
Post-Decision State Variables (OLG)**

Step 0: i. Initialize the “look-up table” approximation of the value function in all states. That is, for each state, $(s \in \mathbf{s}, x \in \mathbf{x}, a \in \mathbf{a}, S \in \mathbf{S}, X \in \mathbf{X})$, assign a starting value to $\bar{V}_a^0(S, X, s, x)$, where the superscript 0 indicates that this is an approximation taken at the 0th iteration and where a is the household’s age.⁷ Note that capital letters represent aggregate variables.

ii. Initialize the state, $(s_0 \in \mathbf{s}, x_0 \in \mathbf{x}, a_0 \in \mathbf{a}, S_0 \in \mathbf{S}, X_0 \in \mathbf{X})$. For overlapping generations models, there may be natural initial states for idiosyncratic shocks and endogenous, individual-specific state variables (i.e. a household may start with no assets), but aggregate-level state variables will depend on period when the household enters the model—and, thus, must be initialized separately for each agent.⁸

iii. Initialize the aggregation method states using the Krusell and Smith algorithm (1998) or Den Haan and Rendahl’s explicit aggregation method (2010).

iv. Increment the iteration counter to $n=n+1$.

Step 1: Choose a simulation period length, T . Generate a sample of the exogenous processes, $(X_t, x_{a,t})$, for the simulation period $(1, \dots, T)$ and for all ages $(1, \dots, \bar{a})$, where $x_{a,t}$ denotes the set of individual-specific state variables for the age a household

⁷Recall that finite horizon dynamic programming problems require us to use different approximations of the value function at each age, since the structure of the choice problem fundamentally changes over the lifecycle. Additionally, note that we could also include $a \in \mathbf{a}$ as a deterministic state variable.

⁸For a more complete treatment of initializing state variables in OLG macro models, see Heer and Maussner (2009).

at time t .

Step 2: For all periods, $t=1,\dots,T$, and all ages, $1,\dots,\bar{a}$, perform the following three steps:

i. Perform the “forward pass” portion of the algorithm. Choose controls to maximize the expression inside of the expectation of

$V_a(S_{t-1}, X_{t-1}, s_{a,t-1}, x_{a,t-1})$ for a particular realization of $(X_t, x_{a,t})$:

$$\tilde{V}_a^n(S_{t-1}, X_{t-1}, s_{a,t-1}, x_{a,t-1}) = \max_{C_t} u(C_t) + \beta \bar{V}_{a+1}^{n-1}(S_t, X_t, s_{a,t}, x_{a,t})$$

ii. Perform the backwards propagation through time (BTT) for each agent. This involves updating V_a^n using information gained about the values of states from future periods, but on the same state trajectory. This can be accomplished by recursively computing \tilde{V}_a in each period each agent is alive, starting in their respective terminal periods:

$$\tilde{V}_a^n(S_{t-1}, X_{t-1}, s_{a,t-1}, x_{a,t-1}) = u(\hat{C}_t) + \beta \tilde{V}_{a+1}^{n-1}(S_t, X_t, \hat{s}_{a,t}, x_{a,t})$$

Note that carets are used to denote endogenous state variables that were selected on the forward-pass step of the algorithm.

iii. Compute the expectation by updating the value function approximation, V_a^{n-1} :

$$\bar{V}_a^n(S_{t-1}, X_{t-1}, s_{a,t-1}, x_{a,t-1}) = (1-\alpha_{n-1})\bar{V}_a^{n-1}(S_{t-1}, X_{t-1}) + \alpha_{n-1}\tilde{V}_a^n(S_{t-1}, X_{t-1}, s_{a,t-1}, x_{a,t-1})$$

iv. Use the result from the maximization step, $s_{a,t}$, the realization of the exogenous processes, $X_{t+1}, x_{a,t+1}$, and the aggregate state, S_t , to compute the next period’s pre-decision state.

Step 3: Check convergence criterion. If satisfied, go to Step 4. If not satisfied, increment n and go to Step 1.

Step 4: Update the paths and laws of motion for the aggregate states. Check the aggregate convergence criterion. If satisfied, go to Step 5. If not satisfied, go to Step 1.

Step 5: Return the value function approximation, $V_a(S, X)^N$, in the form of a look-up table.

The most important departure from Algorithm 1.1 can be found in Step 2(ii). Here, we perform the BTT step, which permits us to construct an unbiased approximation of the value function. When the “forward pass” portion of the algorithm is performed, it generates a value function approximation that depends heavily on the value function initialization. However, when \tilde{V}_{a+1}^n replaces \bar{V}_{a+1}^{n-1} as an update to \tilde{V}_a^n , the value function approximation incorporates information about the utility derived in future states on this particular simulation trajectory.

Finally, we will consider two algorithm refinements that do not appear in Powell (2007) or Bertsekas (2011). The first extends Algorithm 1.2 to the infinite horizon case; and the second adds information sharing as a strategy for solving heterogeneous agents models using either Algorithm. For the sake of parsimony, truncated versions of each algorithm will be presented; each omits steps that included in either Algorithm 1.1 or 1.2.

**Algorithm 1.3: Double-Pass Forward Dynamic Programming with
Post-Decision State Variables (Infinite Horizon)**

Step 0: See Algorithm 1.2.

Step 1: Choose a simulation period length, T , and a cutoff length, \bar{T} , where $T > \bar{T}$. Generate a sample of the exogenous processes, X_t , for the simulation period: $1, \dots, T$.

Step 2: For periods, $t=1, \dots, T$, perform the following three steps:

i. Choose controls to maximize the expression inside of the expectation of $V(S_{t-1}, X_{t-1})$ for a particular realization of X_t :

$$\tilde{V}^n(S_{t-1}, X_{t-1}, s_{it-1}, x_{it-1}) = \max_{c_t} u(c_t) + \beta \bar{V}^{n-1}(S_t, X_t, s_{it}, x_{it})$$

ii. Perform the backwards propagation through time (BTT) step. This involves updating V^n using information gained about the values of states from future periods, but on the same state trajectory. This can be accomplished by recursively computing \tilde{V}^n in each period, starting in period T :

$$\tilde{V}^n(S_{t-1}, X_{t-1}, s_{it-1}, x_{it-1}) = u(\hat{c}_t) + \beta \tilde{V}^{n-1}(S_t, X_t, \hat{s}_{it}, x_{it})$$

Note that carets are used to denote household-level endogenous state variables that were selected on the forward-pass step of the algorithm.

iii. For $1, \dots, \bar{T}$, set $\tilde{V}^n(S_{t-1}, X_{t-1}, s_{it-1}, x_{it-1}) = \tilde{V}^n(S_{t-1}, X_{t-1}, s_{it-1}, x_{it-1})$.

iv. Compute the expectation by updating the value function approximation, V^{n-1} :

$$\bar{V}^n(S_{t-1}, X_{t-1}, s_{it-1}, x_{it-1}) = (1 - \alpha_{n-1}) \bar{V}^{n-1}(S_{t-1}, X_{t-1}, s_{it-1}, x_{it-1}) + \alpha_{n-1} \tilde{V}^n(S_{t-1}, X_{t-1}, s_{it-1}, x_{it-1})$$

v. Use the result from the maximization step, s_{it} , the realization of the exogenous

processes, X_{t+1} , x_{it+1} , and the aggregate state, S_t , to compute the next period's pre-decision state.

Step 3: Check convergence criterion. If satisfied, go to Step 4. If not satisfied, increment n and go to Step 1.

Step 4: Return the value function approximation, $V(s, x)^N$, in the form of a look-up table.

Algorithm 1.3 differs from Algorithm 1.2 in two important ways. First, it covers the infinite horizon case. And second, it only performs the BTT step on a subset of the data: $1, \dots, \bar{T}$, where $T > \bar{T}$. The reason why only a subset of the data is used is because \tilde{V}^n is recursively constructed, which means that it will be downward-biased in finite sample approximations, since it will miss the value attributable to periods $T + 1, \dots, \infty$. Now, if an agent is in period 10 and T is large, then the recursively computed value function approximation in that period will not be substantially downward-biased, since $T-10$ is large, which means that any value from $T + 1, \dots, \infty$ will be heavily discounted and therefore negligible. On the other hand, if the agent is at $T-1$, then the future streams of utility from $T + 1, \dots, \infty$ will be important, which will yield a substantial downward-bias. For this reason, the algorithm only performs the BTT step on $1, \dots, \bar{T}$. Furthermore, note that the larger $T - \bar{T}$ is, the less biased the value function approximations will be, but the longer it will take to achieve convergence, since it will not be possible to fully exploit BTT.

Finally, I will discuss an algorithm refinement that is useful for both infinite horizon and OLG macro models. The focus of this algorithm will be to integrate information from all agents' choices into the value function approximation. Since

computational approaches to heterogenous agents models may require households to solve problems in parallel, it is an open question how and when each agent should assimilate information from other agents' choices. I will proceed with the infinite horizon case, but it is important to note that this approach can easily be generalized to OLG models.

Algorithm 1.4: Double-Pass Forward Dynamic Programming with Post-Decision State Variables and Information-Sharing (Infinite Horizon, Heterogenous Agents)

Step 0-Step2(iii): See Algorithm 1.3.

Step 2(iv): Compute the agent-specific expectation by updating the value function approximations, V_i^{n-1} , separately for each household: $\bar{V}_i^n(S_{t-1}, X_{t-1}, s_{it-1}, x_{it-1}) = (1 - \alpha_{n-1})\bar{V}_i^{n-1}(S_{t-1}, X_{t-1}, s_{it-1}, x_{it-1}) + \alpha_{n-1}\tilde{V}_i^n(S_{t-1}, X_{t-1}, s_{it-1}, x_{it-1})$

Note that α denotes the step size. We index it with n to indicate that it may change with the iteration. Note that i indexes the households.

iv. Average the value function approximations across households to share all available information:

$$\bar{V}^n(S_{t-1}, X_{t-1}, s_{it-1}, x_{it-1}) = \sum_{i \in I} \bar{V}_i^n(S_{t-1}, X_{t-1}, s_{it-1}, x_{it-1}) f(i)$$

Note that I is the set of all households and $f(i)$ is the mass of household i . For example, if there are K agents of equal mass in the model, then $f(i) = \frac{1}{K}$, for all i . Alternatively, the distribution might be written as function of state variables.

Step 2(v) - Step 5: See Algorithm 1.3.

It is important now to gain a full understanding of the merits of Algorithm 1.4. A good place to start is by comparing it to an extension of Algorithm 1.3 with heterogeneous agents, but without information sharing. Here, households solve their decision problems simultaneously, but do so without the benefit of what other agents have learned. If information-sharing is imposed in a model with 10,000 heterogeneous agents, then each agent gains information about the value of 10,000 state-path trajectories after each iteration—a substantial improvement over one.

Another possible approach to this problem would involve sharing information across agents within iterations. In the limit, this approach might entail using the same value function approximation for all agents. This approach has two benefits: first, it provides agents with more information about the values of states within iteration, leading to faster convergence if all other things are equal; and second, in the limiting case (one value function approximation only), it substantially reduces the amount of data that must be held in memory (down from 10,000 arrays to 1). However, there is a substantial drawback to using this approach, rather than the one outlined in Algorithm 1.4: it does not lend itself to parallelization. That is, it is not possible to solve each household’s problem in isolation on a separate CPU or GPU core if household i ’s value function approximation depends on household j ’s approximation. For this reason, the algorithm defers information-sharing and performs it as an end-of-iteration task. This will not prevent further parallelization, since heterogeneous agent models require aggregation and law of motion estimation at the end of each iteration, regardless of how the household problem is solved.

I have now outlined the four primary algorithms that are the focus of this paper. In the next section, I will briefly review global solution methods that are common

in the DSGE literature, which can be used as performance benchmarks. In particular, I will focus on value function iteration, backwards recursion, interpolation, the parameterized expectations approach, and projection methods.

6 Review of Global Solution Methods for DSGE Models

In this section, I will briefly review common global solution methods for DSGE models. I will then evaluate the performance of some of these methods in the following sections; and will compare the results to those of algorithms that make use of post-decision state variables. I will begin the overview with a review of value function iteration.

6.1 Value Function Iteration

Value function iteration (VFI) is used to solve infinite horizon dynamic programming problems. It involves discretizing the state space, iterating over all states, and then updating the value function approximation until some measure of convergence is achieved. It can be used as a standalone solution method or it can serve as a first step for more advanced solution methods. The basic algorithm for VFI is presented below. Since it can easily be extended to the heterogenous agents case, I do not provide a separate solution algorithm.

Algorithm 2.1: Value Function Iteration

Step 0: Discretize the state space. For a representative agent model, choose a grid for endogenous state variables, $\mathbf{s} = \{s_1, \dots, s_N\}$, where $s_h < s_j$ if $h < j$. Use Tauchen's method (1986) to discretize the exogenous processes, \mathbf{x} , into a K-state Markov chain

with an associated transition matrix, P , where element $p_{k,l}$ is the probability of transitioning from state k to l . For heterogeneous agents models, do this for both aggregate and individual-level state variables.

Step 1: Initialize the value function, $V^0(s_i, x_j)$, at all grid points, $\mathbf{s} \times \mathbf{x}$.

Step 2: For each grid point, (i,j) , compute the value of choosing all possible post-decision values of \mathbf{s} , s_m : $w_m(s_i, x_j) = u(s_i, x_j, s_m) + \beta \sum_{l=1}^K p_{j,l} V^0(s_m, x_l)$.

Step 3: For each grid point, (i,j) , identify the index associated with the choice of endogenous state variables that maximizes $w_m(s_i, x_j)$, m^* . Construct V^1 as follows for all (i,j) : $V^1(s_i, x_j) = w_{m^*}$.

Step 4: Check for convergence using the following metric:

$$\max_{i=1, \dots, N, j=1, \dots, K} |V^1(s_i, x_j) - V^0(s_i, x_j)| \leq \epsilon$$

If the convergence criterion is satisfied, stop the algorithm and compute the policy function. If not, set $V_0 = V_1$ and return to Step 1.

Next, I will discuss the basic algorithm for finite horizon choice problems for OLG models. Additionally, I will discuss refinements to both VFI and backwards recursion in later sections.

6.2 Backwards Recursion

Backwards recursion is the finite horizon analog of VFI. That is, it is a solution method that relies on the discretization of the endogenous state variables; and the

conversion of the continuous exogenous processes into discrete Markov chains with associated transition probability matrices. In comparison to VFI, it has one computational advantage and one disadvantage. The advantage is that the agent lives a finite number of periods, which means that there will be a set of terminal states whose value we can compute exactly by maximizing the instantaneous utility function in the final period. Furthermore, this set of terminal states will serve as a starting point for the backwards recursion—that is, the construction of all previous period value functions. The disadvantage of using an OLG model that requires backwards recursion as a solution method is that it requires us to solve for separate value function approximations for each age. Alternatively, this can be accomplished by introducing a deterministic state for age. This is a computational trade-off that should be considered carefully when choosing an OLG model, rather than an heterogeneous agents model with infinitely-lived agents. The basic algorithm for backwards recursion is outlined below.

Algorithm 2.2: Backwards Recursion

Step 0: Discretize the state space. Choose a grid for endogenous, aggregate state variables, $S = \{S_1, \dots, S_N\}$, where $S_h < S_j$ if $h < j$. Discretize the exogenous processes, X , into a K -state Markov chain with an associated transition matrix, P , where element $p_{k,l}$ is the probability of transitioning from state k to l . For heterogeneous agents models, do this for both aggregate and individual-level state variables. Choose a maximum age, \bar{a} , where $A = 1, \dots, \bar{a}$.

Step 1: At \bar{a} , solve $V_{\bar{a}}(s_i, x_j) = \max_{s_m} u(s_i, x_j, s_m)$, at all grid points, $\mathbf{s} \times \mathbf{x}$.

Step 2: Using the value of each state at age \bar{a} determined in Step 1, compute the

value of choosing all possible post-decision state, s_m , at age $\bar{a} - 1$:

$$w_m^{\bar{a}-1}(s_i, x_j) = u(s_i, x_j, s_m) + \beta \sum_{l=1}^K p_{j,l} V^{\bar{a}}(s_m, x_l).$$

Step 3: For each grid point, (i,j) , identify the index associated with the choice of endogenous state variables that maximizes $w_m(s_i, s_j)$, m^* . Construct $V^{\bar{a}-1}$ as follows for all (i,j) : $V^{a_{max}-1}(S_i, X_j) = w_{m^*}$.

Step 4: Repeat Steps 2-3 for $\bar{a} - 2, \dots, 1$. Use the resulting age-specific value functions to determine the age-specific policy functions. For heterogenous agents models, nest Steps 1-4 within an aggregation algorithm, as described earlier.

Since VFI and backwards recursion both revolve around discretizing the state space, they also both suffer from the curse of dimensionality when the household's choice problem involves many continuous state variables. Similarly, improvements to both approaches rely on the same family of techniques. In particular, linear and cubic interpolation are two common strategies for overcoming the curse of dimensionality when such solution methods are applied. For the sake of parsimony, I will avoid discussing a separate algorithm for these approaches, but will instead discuss how interpolation can be added to 2.4 and 2.5.

6.3 Interpolation

The primary purpose behind interpolation is to reduce the number of nodes in the state space without incurring a corresponding reduction in accuracy. This is accomplished by dividing the maximization routine of Algorithms 2.1 and 2.2 into two steps. In the first step, maximization is performed over a sparse grid. Once the node that maximizes the value function is identified, interpolation is performed. The purpose

behind this is to provide a continuous approximation to the value function in the area between the candidate maximum and its neighboring nodes. A maximization routine for continuous functions can then be employed to identify the true maximum.

One important thing to note about this method is that it will become increasingly costly as the dimensionality of the exogenous component of the state space increases. In particular, if a problem has K possible exogenous states the agent can transition into from her current state, then K interpolations must be performed. Moving forward, I will not perform tests for interpolation or other improvements for discretization methods, since they can also be applied to the new algorithms proposed in the paper; and, thus, do not constitute an advantage of VFI or backwards recursion.

6.4 Parameterized Expectation Approach (PEA)

The Parameterized Expectations Approach (PEA) differs from the solution methods we have considered so far in that it does not involve the discretization of the state space. Instead, it requires us to construct continuous approximations of the expectational components of the first order conditions. These approximations are created through the repetition of three steps. First, a set of approximating functions are chosen and parameterized. Second, decisions are simulated based on the chosen parameterization. And third, some convergence criterion is checked to determine if the simulated decisions were sufficiently close to optimal.

One important thing to note about PEA is that it is not vulnerable to the curse of dimensionality in the same sense that VFI and backwards recursion are. In solution methods that use discretization, adding more state variables—and, in particular, continuous state variables—substantially increases the number of function evaluations that must be performed to identify the maximum. In contrast, with PEA, adding another state variable adds no such requirement. It does, however, make it more

difficult to choose an initial parameterization for the expectational terms.

The algorithm below constructs a general description of PEA, following Heer and Maussner (2009) and Den Haan and Marcet (1990). While the algorithm is for the infinite horizon case, it can easily be extended to capture finite horizon applications.

Algorithm 2.3: Parameterized Expectations Approach (PEA)

Step 0: Identify all model equations, including first order conditions, constraints, and exogenous processes.

Step 1: Choose an initial functional form and parameterization for the expectational components of the first order conditions. Here, we will represent this by $G_h(S_{t-1}, X_t, s_{it-1}, x_{it}; \theta_h^0)$, where G_h is a function that approximates the h^{th} expectational term and using parameters θ_h^0 . Typically, high order polynomials are used to construct G . Assume that the realization of the expectational component at time $t+1$ is ψ_{it+1}^h , which means that $\psi_{it+1}^h = G_h(S_{t-1}, X_t, s_{it-1}, x_{it}) + u_{it}$ and $u_{it} \sim iid$.

Step 2: Using the expectational term approximations, the initial values, and T periods worth of simulated exogenous processes, compute the values of the endogenous, individual-level state variables in each period and for each agent. Increase the value of n , the parameter value iteration counter.

Step 3: Discard observations for periods \tilde{T} and lower to eliminate bias introduced by initial values. Use the data for all periods after \tilde{T} , all agents, and all expectational components to estimate the parameter values using the following objective function:

$$\hat{\theta}_n = \underset{\theta^n}{\operatorname{argmin}} \sum_{i \in I} \sum_{t=\tilde{T}+1}^{T-1} \sum_{h \in H} (\psi_{it+1}^h - G_h(S_{t-1}, X_t, s_{it-1}, x_{it}; \theta^n))^2$$

Step 4: Set $\theta^n = (1 - \alpha_n)\theta^{n-1} + \alpha_n\hat{\theta}^n$. Check the following convergence criterion: $\max(\hat{\theta}^n - \theta^{n-1}) < \epsilon$, where $\max()$ is performed element-wise. If the condition is satisfied, then stop. Otherwise, return to Step 2.

In the next section, we will consider a class of solution methods called Projection Methods (PM). This approach nests PEA; and provides computational modelers with a highly flexible tool that can be used to solve dynamic choice problems.

6.5 Projection Methods (PM)

Projection Methods (PM) are difficult to characterize because they cover a wide variety of approaches to solving multi-period choice problems. Broadly, PM consists of many tools that can be used to approximate common functions in choice problems (i.e. value functions, decision rules, expectational equations, etc.). The PM approach is typically performed by choosing the functional form for an approximation, computing the residual of that approximation, and then using some method to minimize the residual. Below, we will consider a basic algorithm for PM, which follows Heer and Maussner (2009) and Heer and Maussner (2008).

Algorithm 2.4: Projection Methods (PM)

Step 0: Identify a set of optimality conditions or functional equations, $\hat{g}(S, X, s, x; \theta)$, that have an associated residual. Choose a family of basis functions, $\xi(S, X, s, x)$, and a degree of approximation, ρ , where $\hat{g}(S, X, s, x; \theta) = \sum_{k=0}^{\rho} \phi_k \xi_k(S, X, s, x)$. Note that Chebychev Polynomials are a common choice in the literature; and have the benefit of being an orthogonal class of polynomials, which reduces multicollinearity in Step 3.

Step 1: Define the residual function $R(S, X, s, x; \theta) = G(\hat{g}(S, X, s, x; \theta))$.

Step 2: Choose a projection function, χ_i , and a weighting function, w . Compute the inner product for all of the H optimality conditions or functional equations, where $i=1, \dots, H$: $\chi_i = \int_S \int_X \int_s \int_x w(S, X, s, x) R(S, X, s, x; \theta) \chi_i dS dX ds dx$

Step 3: Choose the value of θ that yields $\chi_i = 0$ for $i=1, \dots, H$. Alternatively, find θ that minimizes $\int_S \int_X \int_s \int_x R(S, X, s, x; \theta)^2 dS dX ds dx$

Step 4: Check the accuracy of the solution, $\hat{\theta}$. If the measure of error is too large, then increase ρ and return to Step 1. For heterogenous agents models, nest Steps 1-4 within an aggregation algorithm.

For Step 0, consider the Bellman equation for an infinite horizon choice problem. We know that $V(s_i, x_j) = \max_{s_m} u(s_i, x_j, s_m) + \beta \sum_{l=1}^K p_{j,l} V(s_m, x_l)$, which implies that $R(s_i, x_j; \theta) = \max_{s_m} [V(s_i, x_j) - u(s_i, x_j, s_m) - \beta \sum_{l=1}^K p_{j,l} V^0(s_m, x_l)]$. Here, an ideal choice of $g()$ might be $V()$. However, we might think about choosing policy functions or Euler equations for $g()$ instead.

In the next section, I will discuss the metrics that can be used to measure the accuracy of all of the aforementioned solution methods; and make it possible to make meaningful comparisons across alternatives.

7 Accuracy Measures

Four sufficient statistics are often used to measure the accuracy of solution methods: 1) second moments; 2) Euler equation residuals; 3) the DM statistic; and 4) law of motion fit. This section will present each of these methods.

7.1 Second Moments

Heer and Maussner (2008) use second moments as a means of evaluating the accuracy of solution methods for standard business cycle models with flexible labor supply. The approach they use is implemented as follows: first, generate a complete set of exogenous shocks for the simulation period length. Second, solve the model under all of the different solution methods you wish to test. Third, simulate the model for each solution method, using the same set of shocks that was generated in the first step. Fourth, compare the second moments of the aggregate variable time series.

The rationale behind this approach is as follows: since the exogenous processes are identical in all simulations, the differences in second moments should arise entirely from differences in the solution methods. This is especially useful in cases where one solution method is faster, but the other is more accurate. If there are no statistically significant differences in second moments, then it may make sense to use the faster approach—especially if the research agenda only requires the computation of simulated moments.

As an example of such a comparison, Heer and Maussner (2008) take the same model and solve and simulate it using two broad classes of methods. The first class is log-linearization around the steady state, which is fast but inaccurate. The second class is the broad family of global solution methods, which are slower, but more accurate for a given grid size. They find that the volatility of investment is different under log-linearized solution, but all other differences are statistically insignificant. They suggest that log-linearization is probably the better option for solving such models, since there is a tremendous speed gain; and no substantial trade-off in terms of accuracy.

While using second moments might be an appropriate strategy for comparing global and non-global (local) solution methods, they become less useful when both

methods are of the same type. For instance, if we have two global solution methods—and it is not clear which is more accurate for a given grid size—then the most we can say from second moment differences is that the solution methods differ. Additionally, second moments often do not capture important differences between solution methods. For this reason, we will mostly focus on other accuracy measures, as outlined in the following three sections.

7.2 Euler Equation Residuals

Heer and Maussner (2009) find that there are substantial differences in accuracy between global and local solution methods when Euler equation residuals are used as the metric for accuracy. The general strategy behind Euler equation residuals is to test the accuracy of the decision rules by determining whether or not they are broadly consistent with the model’s Euler equations.

The Euler equation residual method is implemented in five steps. First, identify the model’s Euler equations and write them in terms of residual equations that are equal to zero if the agent has made an optimal decision. Second, solve the model and recover the decision rules. Third, plug the decision rules into the Euler equations. And fourth, evaluate the residual equations over a discrete state space grid. Finally, compute descriptive statistics for the Euler equation residuals to use as a basis for comparison.

It is important to note that there have been several notable refinements of the original Euler equation residual approach, including Christiano and Fisher (2000). However, for the purposes of this paper, we will focus only on the approach described in this section.

7.3 The DM Statistic

Den Haan and Marcet (1994) offer an alternative approach for evaluating the accuracy of the solution method that also uses Euler equation residuals. This approach proceeds in the following steps: first, solve and simulate the model. Second, construct an ex-post forecast error series, using the household's simulated decisions. Third, regress the error series on the previous period state variables. And fourth, perform an F-test to check whether the coefficients are statistically significantly different from zero.

The general strategy behind the DM Statistic is to test whether or not the ex-post forecast error is consistent with an optimal decision. That is, we want to know if deviations are due to shocks or due to optimization errors. If deviations are due only to shocks, then the decision rules are accurate.

7.4 Law of Motion Fit

Solution methods for models with many heterogeneous agents must incorporate a step that approximates the behavior of the aggregate, endogenous state variables. Such a step will enable agents to pin down contemporaneous factor prices and make predictions about the future values of aggregate, endogenous states. For the purposes of this paper, I will use the Krusell and Smith algorithm (1998) for this part of the solution method. This entails estimating a separate law of motion for each productivity state. In the case where productivity follows a two-state Markov process and where aggregate capital is the variable of interest, the following pair of regressions are run, where K is the mass-weighted sum of the individual capital stocks:

$$\ln(K_{t+1}) = \begin{cases} B_{1,H} + B_{2,H}\ln(K_t), & \text{if productivity is high} \\ B_{1,L} + B_{2,L}\ln(K_t), & \text{if productivity is low} \end{cases} \quad (24)$$

According to Krusell and Smith (1998), there are several ways in which the accuracy of the laws of motion can be tested. One way is to check the value of R^2 for each of the above regressions, which I denote as $R^2(H)$ (high productivity state) and $R^2(L)$ (low productivity state) in the test results. A high value of R^2 indicates that the law of motion accurately predicts the next period value of aggregate endogenous state variables.

A poor fit for the law of motion may indicate two problems with the solution method. First, it might represent instability in the convergence of individual-level decision rules. Large shifts in decision rules will lead to corresponding shifts in aggregated individual capital stocks, which has the potential to slow convergence and weaken fit. And second, a poor fit might indicate that decision rules perform poorly in a way that is not measured by Euler equation residuals or other individual-level measures of decision rule quality. In particular, crude decision rule parameterizations may perform poorly far away from the steady state, which has the potential to introduce influential observations into the Krusell and Smith algorithm (1998) regressions and subsequently lower R^2 .

It is important to note, however, that R^2 may be a weak measure of the accuracy of the law of motion. As is pointed out by Den Haan (2010), R^2 only captures the accuracy of a one period ahead forecast. This is because the next period value of aggregate capital (i.e. the dependent variable in the law of motion regressions) is updated by aggregating individual capital stocks at the end of each period. This prevents long term forecast deviations that would arise from repeated deviations from accuracy from showing up in R^2 . Den Haan (2010) suggests several alternatives to R^2 , including the mean squared error (MSE) of a multi-period forecast. In future work, I will expand this paper by considering alternative measures of law of motion accuracy; however, for the purposes of this paper, I will use R^2 , since it is the most

popularly cited measure in the literature.

Overall, measures of fit for the law of motion provide another means of testing the accuracy of different solution methods for the household's optimization problem. A poor fit indicates that the decision rules that emerged from a particular solution method exhibit bad properties in the aggregation phase of the algorithm.

8 Results

In this section, we will test the algorithms that were described earlier. This will proceed as follows. First, the baseline model will be presented. Next, the model will be solved using the aforementioned algorithms. Third, the results will be compared using the measures of accuracy described above. Moving forward, I will start by outlining a standard business cycle model with no labor supply, which will be used in the first set of tests. In the second set of tests, I will focus models with incomplete markets, aggregate uncertainty, and many heterogeneous agents.

8.1 Neoclassical Model (Infinite Horizon)

Households in the baseline model are assumed to maximize expected, discounted utility from consumption, c_t :

$$\max_{c_0} E_0 \sum_{t=0}^{\infty} \beta^t [\log(c_t)] \quad (25)$$

It is assumed that $\beta \in (0, 1)$. Additionally, the household faces the following budget constraint, where k_t is the household's capital stock, and r_t is the return to capital:

$$k_{t+1} = (1 + r_t)k_t - c_t \quad (26)$$

The firm produces output using aggregate capital, K_t with the following production function:

$$Y_t = Z_t K_t^\alpha \quad (27)$$

Capital is assumed to depreciate at rate δ . Productivity, Z_t is assumed to follow an AR(1) process in logs:

$$\ln Z_t = \rho \ln Z_{t-1} + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma) \quad (28)$$

The firm maximizes profit, yielding the following factor price for capital:

$$r_t = \alpha Z_t K_t^{\alpha-1} - \delta \quad (29)$$

The economy is subject to an aggregate resource constraint:

$$Y_t = K_{t+1} - (1 - \delta)K_t - C_t \quad (30)$$

Additionally, we assume that all individual variables are equal to aggregate variables in equilibrium: $c_t = C_t$ and $k_t = K_t$. Combining the first order conditions for the household, the equilibrium conditions, and the first order conditions for the firm, we get the following equations that characterize the model:

$$\frac{1}{c_t} = \lambda_t \quad (31)$$

$$\lambda_t = \beta E_t \lambda_{t+1} [\alpha Z_{t+1} K_{t+1}^{\alpha-1} + (1 - \delta)] \quad (32)$$

$$K_{t+1} + C_t = Z_t K_t^\alpha + (1 - \delta)K_t \quad (33)$$

Using the description above and the methods outlined earlier, it will now be

possible to simulate, and test the accuracy of the model under each of the different candidate solution methods. In particular, results will be presented for value function iteration, parameterized expectations, projection methods, and ADP with post-decision state variables (both forward-pass and double-pass).

8.2 Full Depreciation Case

In the case where $\delta = 1$, the model has a known, closed-form solution. In particular, the decision rule for capital is given as follows:⁹

$$k_{t+1} = \alpha\beta z_t k_t^\alpha \tag{34}$$

This information could be used to perform two different tests. First, it could be used to determine the optimal value of next period capital, k' , for each node in the state space. This could then be compared to the values of next period capital implied by a particular solution method at each node. This would yield a measure of accuracy that differs from Euler equation residuals, the DM statistic, and second moments of the time series; and provides an exact deviation from optimality. And second, it could be used to perform a simulation-based test that captures multi-period deviations from optimality that occur as a result of suboptimal choices in previous periods. This could be done by setting k_0 , simulating a path for Z_t , and then comparing the decision paths for the exact solution and the solution methods under consideration.

In this paper, I do not perform tests using the full depreciation case. However, in future work, I will exploit this method to provide an expanded evaluation of the accuracy of the ADP algorithm with post-decision state variables.

⁹See the appendix for a proof.

8.3 Heterogeneous Agents Model with Incomplete Markets and Aggregate Uncertainty (Infinite Horizon)

The second set of tests will expand on the representative agent model by introducing incomplete markets, aggregate uncertainty, and many heterogeneous agents. This will have three effects on the tests. First, it will substantially increase the dimensionality of the state space. Second, it will introduce a source of instability (the aggregation step) that may cause less stable algorithms to periodically diverge, resulting in higher run times or reduced accuracy. And finally, it will make it possible to incorporate the information-sharing algorithm described earlier to generate performance improvements in the ADP tests. The model used for the second set of tests is given below.

In the model, all household's assumed to maximize expected, discounted utility from consumption, c_t :

$$\max_{c_{i0}} E_0 \sum_{t=0}^{\infty} \beta^t [\log(c_{it})] \quad (35)$$

It is assumed that $\beta \in (0, 1)$. Additionally, each household faces the following budget constraint, where k_{it} is household i 's capital stock, r_t is the factor price for capital, w_{it} is the wage if employed, and \bar{w} is the home production wage for the unemployed:

$$k_{it+1} = \begin{cases} (1 + r_t)k_t + w_{it} - c_{it}, & \text{if } e_{it} = 1 \\ (1 + r_t)k_t + \bar{w} - c_{it}, & \text{if } e_{it} = 0 \end{cases} \quad (36)$$

As shown by the constraint above, each household supplies labor perfectly inelastically, but is subject to an unemployment shock ($e = 0$). For simplicity, this shock will depend only on the current level of productivity and not the previous employment

state. As in the representative agent model, it will be assumed that productivity, Z_t follows an AR(1) process in logs:

$$\ln Z_t = \rho \ln Z_{t-1} + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma) \quad (37)$$

The firm produces output using aggregate capital, K_t with the following production function:

$$Y_t = Z_t K_t^\alpha N_t^{1-\alpha} \quad (38)$$

The economy is subject to an aggregate resource constraint:

$$Y_t = K_{t+1} - (1 - \delta)K_t - C_t \quad (39)$$

In contrast to the representative agent model, individual variables do not equal aggregate variables. Rather, aggregate consistency conditions must be imposed:

$$K_t = \sum_{i=1}^N k_{it} \mu_i \quad (40)$$

$$C_t = \sum_{i=1}^N c_{it} \mu_i \quad (41)$$

$$N_t = \sum_{i=1}^N n_{it} e_{it} \mu_i \quad (42)$$

Finally, the model is closed with a law of motion to approximate aggregate capital stock movements. In the simplest case where the productivity shock has two states (low and high), the Krusell and Smith (1998) algorithm can be used to estimate and update the following law of motion:

$$\ln(K_{t+1}) = \begin{cases} B_{1,H} + B_{2,H}\ln(K_t), & \text{if productivity shock is high} \\ B_{1,L} + B_{2,L}\ln(K_t), & \text{if productivity shock is low} \end{cases} \quad (43)$$

See Figure 2 for an example of this law of motion's fit in a model solved using PM. Note that the solid, blue line indicates the aggregate path of capital, simulated according to the estimated law of motion above. The dotted, red line indicates the actual path of capital, aggregated from individual household's choices.

8.4 Model Calibration

Both the representative agent model and the heterogenous agents model were calibrated using the parameter values given in Table 6. Note the model is calibrated for an annual time period, which offers two benefits. First, it allows me to use an autoregressive parameter for the productivity process that is further from unity.¹⁰ And second, it makes it possible to use the ADP-BTT algorithm without a longer simulation period. A model calibrated to a quarterly time period, for instance, would require $T - \bar{T}$ to be roughly four times as long in order to ensure that streams of utility that arrive after $T+1$ are sufficiently heavily discounted.

8.5 Algorithm Refinements and Details

In earlier sections, I presented general algorithms for all of the solution methods that will be tested in the following section. However, for the purposes of this paper, I made several algorithmic refinements to ensure convergence in all of the tests. In particular, the algorithm for PEA was adjusted to incorporate dynamic bounds for

¹⁰Many discrete approximations to continuous processes perform poorly when the series is highly persistent. In this case, I have used the Rouwenhorst method (1995), which tends to perform better than other discretization methods when series are persistent, but is still subject to the same problem.

the household-level capital grid. To understand how this works, consider a grid for capital, $k \in [1, \dots, k_{kn}]$, which is centered around its steady state value, $k_{kn/2+1}$. When the algorithm is initiated, use the smaller grid, $k \in [k_{n/2-j}, \dots, k_{kn/2+j}]$, which is more tightly centered around the steady state value of capital. After each iteration, expand the grid by adding more nodes if any endpoint nodes were chosen in the previous step. Without this refinement, the performance of PEA was substantially worse; and was prone to divergence in heterogenous agents models.

In addition to the PEA refinement, I also incorporated my own refinement for the ADP algorithms. Here, I exploited the value function's monotonicity and concavity in individual-level capital to update more than just the state visited. In particular, the algorithm identifies states that have not been visited, but fall between states that have. It then revises the values of those states to reflect the monotonicity of the value function in individual capital. In particular, it scales the value of the state, so that it falls below the value of the node above it, but above the node below it. It exploits concavity by placing greater weight on the node above the state being updated. This procedure is similar to linear interpolation, but is intended to update unvisited states, rather than providing a value function estimate off of the grid. This procedure substantially improved convergence speed for a given level of accuracy.

Finally, it is important to note that the PM approach used in this paper approximates the value function, rather than the Euler equations. The purpose behind this decision was to make the PM algorithm sufficiently distinct from the PEA algorithm, rather than performing the same set of tests twice. The algorithm works by first performing VFI on a small grid. The resulting look-up table approximation is then combined with the fixed point condition implied by the functional equations to perform the projection. In the heterogenous agents model, both the PM and PEA algorithms use third order Chebychev polynomials of the first type to transform state

variables in the approximation equations. Additionally, all continuous state variable transformations are interacted with the agent's employment state. In the representative agent model, a simple, exponential polynomial of state variables is used to parameterize PEA.

8.6 Simulation Results

We will now examine the results by comparing the solution methods using run time, accuracy, and stability as criteria. Tables 1-2 in the appendix show the results for the standard, representative agent business cycle model without labor. Tables 3-5 show the results for the heterogeneous agents model with incomplete markets and aggregate uncertainty. All simulations were performed using a quad-core 3.40 gigahertz processor.

Note that "MAX," "MEAN," and "STD," refer to the maximum, mean, and standard deviation of the Euler equation residuals, which were computed on a grid around the steady state. RT stands for run time and has the following format: hours:minutes:seconds:hundredths of a second. In heterogeneous agent model results, $R^2(H)$ and $R^2(L)$ are the measures of fit for the laws of motion for capital in high and low productivity states respectively. Finally, DMS denotes the Den Haan and Marcet (1994) statistic. In heterogeneous agent model tests, the test statistic is expensive to calculate and less informative, so it is dropped. In representative agent model tests, the fraction of DM test statistics above the 97.5% critical value in 1000 simulations are provided for the first simulation.¹¹

Table 1 provides the results for the standard business cycle model without labor.

¹¹The distribution of the DM statistics is used to judge if a solution method is sufficiently accurate, but is not a useful determinant of the degree of accuracy. For this reason, the DM statistic is only used to establish that the least accurate set of solutions is sufficiently accurate. It is then dropped in later simulations.

All simulations were performed on a 500-node grid: 100 individual capital nodes, and 5 productivity nodes. As expected, VFI is the most accurate solution method for this grid size (as measured by max and mean Euler equation residual size), but is substantially slower than the alternatives. For instance, it takes VFI 5.05 times longer to converge than PM; and 13.23 times longer to converge than ADP. However, it ultimately yields a solution that is an order of magnitude more accurate than PEA, ADP, and ADP-BTT. Finally, note that the fraction of DM statistics that exceeded the .975 critical value were under 2.5% for all methods other than VFI and ADP-BTT. Moving forward, we will drop the DM statistic in favor of the more informative (and less costly) Euler equation residual summary statistics.

Table 2 provides the results for the standard business cycle, but now with a 2000-node grid: 100 individual capital nodes, and 20 productivity nodes. The purpose behind this simulation was to examine algorithmic performance when computing expectations became more computationally expensive. As the table shows, all solution methods other than PEA improved in accuracy by roughly one half an order of magnitude. The limited parameterization of PEA (a simple, exponential polynomial) fostered stability and rapid convergence, but prevented substantial improvements in accuracy, even after the number of nodes in the state space was increased by a factor of four. Conversely, both ADP and ADP-BTT have experienced substantial gains in improvement; and at a low cost. Each algorithm took roughly five times longer to converge. In contrast, PM took eight times longer to converge; and VFI took 26 times longer to converge. This suggests that ADP and ADP-BTT will perform favorably—in terms of the accuracy-run time tradeoff—as expectations become increasingly costly. That is, if the stochastic component of the state space is very large, then ADP or ADP-BTT should be able to achieve convergence at the desired level of accuracy in a shorter run time than VFI or PM. Furthermore, without an improved parameteri-

zation for PEA, it will be unable to generate the degree of accuracy obtainable with PM, VFI, ADP, or ADP-BTT.

Next, we'll consider the results for the model with incomplete markets and aggregate uncertainty, which are given in Table 3. All simulations listed in the table were performed on a 200-node grid: 10 individual capital nodes, 5 aggregate capital nodes, 2 productivity nodes, and 2 employment nodes. Additionally, the simulations were performed for 3000 periods with a relatively small number of agents (200). From the table, we can see that VFI yields the most accurate solutions with respect to Euler equation residuals. Both the mean and max Euler equation residuals are nearly an order of magnitude smaller than ADP with post-decision states (hereafter, ADP) and ADP with post-decision states and backwards propagation through time (hereafter, ADP-BTT). Next to VFI, PEA yields the most accurate individual-level solutions; and is followed by ADP, ADP-BTT, and then PM. Additionally, with such a small state space, VFI is also the fastest solution method, achieving convergence within 40 seconds. It is followed by ADP-BTT, PEA, ADP, and then PM. Additionally, among all the solution methods, ADP, ADP-BTT, and VFI achieve the best law of motion fit. PEA, while fast, is less stable at the aggregate level and achieves a substantially worse fit.

The next set of simulations, shown in Table 4, expands the grid to 1000 nodes: 25 individual capital nodes, 10 aggregate capital nodes, 2 productivity nodes, and 2 employment nodes. Relative to the 200-node results, each solution method achieves an accuracy increase of about one order of magnitude, measured by mean and max Euler equation residual reductions. VFI, again, is the most accurate at the individual-level and most stable at the aggregate level. However, VFI now takes 11.15 times longer to converge, even though the number of nodes in the state space has only increased by a factor of five. In contrast, ADP-BTT takes only 1.45 times longer to convergence, but

also retains stability at the aggregate level while improving individual-level accuracy by an order of magnitude. Furthermore, note that all solution methods now converge faster than VFI; and that VFI now takes roughly 7 times as long to converge as ADP-BTT.

Finally, Table 5 expands the grid to 4000 nodes: 50 individual capital nodes, 20 aggregate capital nodes, 2 productivity nodes, and 2 employment nodes. Again, all solution methods experience an improvement in accuracy. This time, it is less than an order of magnitude. VFI remains the most accurate and stable. PEA is the second most accurate and also the fastest, but is substantially less stable at the aggregate level. ADP-BTT is both accurate and stable and converges faster than anything other than PEA by a substantial margin. Importantly, VFI takes 34.33 times as long to converge as ADP-BTT.

These findings suggest that ADP-BTT may provide substantial improvements in convergence time over common global solution methods when the state space is large. In particular, ADP-BTT converges at roughly the same rate as PEA and is slightly less accurate at the individual level, but is substantially more stable at the aggregate level. Additionally, it is roughly as accurate as PM, but is substantially faster and slightly more stable. Finally, in comparison to VFI, it is roughly as stable at the aggregate level, but is approximately one order of magnitude less accurate. Importantly, however, it becomes substantially faster than VFI as the state space grows. In particular, the difference in run times between the algorithms grows faster than the difference in accuracy. This suggests that ADP-BTT will tend to dominate VFI in problems with many continuous state variables, since VFI will become infeasible if it isn't performed on a highly coarse grid.

9 Conclusion

The ADP-BTT algorithm provides a promising alternative to currently-used global solution methods for DSGE models. For the state space discretizations tested, it was roughly one order of magnitude less accurate than VFI, which was the most accurate solution method tested. However, it was substantially faster than VFI; and the difference in convergence speeds grew rapidly in the number of nodes in the state space, while the difference in accuracy remained largely unchanged. This suggests that ADP-BTT will tend to outperform VFI when a model contains many state variables (e.g. models with heterogenous agents) and would require a coarse discretization for VFI to converge within a reasonable amount of time.

The only solution method that outperformed ADP-BTT on individual-level measures of accuracy and also maintained a low run time was PEA. However, PEA was substantially less stable than ADP-BTT at the aggregate level in models with many heterogenous agents. Additionally, PEA is both sensitive to parameterization and sensitive to the choice of parameterized equations.¹² Furthermore, it can become increasingly difficult to parameterize PEA in large models; and accuracy gains beyond a certain magnitude will require the addition of higher order polynomial transformations of the state variables. Finally, PEA is designed explicitly to minimize Euler equation residuals—unlike all other solution methods tested—which may cause the measures of accuracy considered (Euler equation residuals and the DM statistic) to overstate its accuracy.

Furthermore, ADP-BTT tends to outperform VFI and PM when a model contains many stochastic processes; or requires a fine discretization of one stochastic process. This is because ADP-BTT does not require the explicit computation of expectations.

¹²Maliar and Maliar (2005) show that the choice of equation can lead to dramatic differences in convergence speed

Furthermore, the algorithm is arguably easier to implement than VFI; and is substantially easier to implement than PM and PEA for new practitioners of computational methods.

Finally, it is important to note that this paper does not fully showcase the set of tools that ADP algorithms with post-decision states offer. In several places in the paper, I have mentioned that ADP and ADP-BTT are particularly well-suited to dealing with problems that have multiple, correlated exogenous processes or exogenous processes for which empirical data is available. In the paper, I provide only algorithms for implementing these ideas—and not test results to demonstrate their advantages. In cases where exogenous processes are complicated and exhibit dependence, ADP’s value goes well beyond reductions in run time. It makes solving computational models more accessible to non-computational economists; and permits computational economists to solve increasingly sophisticated models by greatly simplifying the programming task. For this reason, formal test results for accuracy and run time may be less useful for these applications.

10 References

Algan, Y., Allais, O., & Den Haan, W.J. 2008. “Solving Heterogeneous-Agent Models with Parameterized Cross-Sectional Distributions.” *Journal of Economic Dynamics and Control* Vol. 32(3), 875-908.

Bellman, R. 1954. “The Theory of Dynamic Programming.” *Bulletin of the American Mathematical Society* Vol 60, 503-515.

Bertsekas, D. 2011. “Chapter 6: Approximate Dynamic Programming.” In *Dynamic Programming and Optimal Control*. Vol. II, 3rd Edition. MIT. Draft.

Christiano, L. & Fisher, J. 2000. “Algorithms for Solving Dynamic Models with Occasionally Binding Constraints.” *Journal of Economic Dynamics and Control* Vol. 24, 1179-1232.

DeJong, D., & Dave, C. 2007. *Structural Macroeconometrics*. Princeton, New Jersey: Princeton University Press.

Den Haan, W.J. 2010. “Assessing the Aggregate Law of Motion in Models with Heterogeneous Agents.” *Journal of Economic Dynamics and Control* Vol. 34, 79-99.

Den Haan, W.J. & Marcet, A. 1994. “Accuracy in Simulations.” *Review of Economic Statistics* Vol. 8, 31-34.

Den Haan, W.J. & Marcet, A. 1990. “Solving the Stochastic Growth Model by

Parameterizing Expectations.” *Journal of Business and Economic Statistics*, Vol. 8, 31-34.

Den Haan, W.J. & Rendahl, P. 2010. ”Solving the Incomplete Markets Model with Aggregate Uncertainty Using Explicit Aggregation.” *Journal of Economic Dynamics and Control* Vol. 34(1), 69-78.

Duffy, J. & McNelis, P. 2001. “Approximating and Simulating the Stochastic Growth Model: Parameterized Expectations, Neural Networks, and the Genetic Algorithm.” *Journal of Economic Dynamics and Control* Vol. 25, 1273-1303.

Heer, B., & Maussner, A. 2008. “Computation of Business Cycle Models: A Comparison of Numerical Methods.” *Macroeconomic Dynamics* Vol. 12, 641-663.

Heer, B., & Maussner, A. 2009. *Dynamic General Equilibrium Modeling: Computational Methods and Applications*. 2nd Edition. Berlin: Springer.

Judd, K. 1998. *Numerical Methods in Economics*. Cambridge, MA, London: MIT Press.

Judd, K. 1992. “Projection Methods for Solving Aggregate Growth Models.” *Journal of Economic Theory* Vol. 58(2), 410-452.

Judd, K., Maliar, L., & Maliar, S. 2011. “Numerically Stable and Accurate Stochastic Simulation Approaches for Solving Dynamic Economic Models.” *Quantitative Economics* Vol. 2(2), 173-210.

Krusell, P. & Smith, A. 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy* Vol. 106, 867-96.

Marcet, A. 1988. "Solving Nonlinear Models by Parameterizing Expectations." Working Paper, Carnegie-Mellon University.

McGrattan, E. 1999. "Application of Weighted Residual Methods to Dynamic Economic Models." In Ramon Marimon and Andrew Scott, *Computational Methods for the Study of Dynamic Economies*. Oxford and New York: Oxford University Press, 114-142.

Maliar, L. & Maliar, S. 2005. "Parameterized Expectations Algorithm: How to Solve for Labor Easily." *Computational Economics, Society for Computational Economics* Vol. 25(3), 269-274.

Maliar, L., Maliar, S., & Valli, F. 2010. "Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell-Smith Algorithm." *Journal of Economic Dynamics and Control* Vol. 34(1), 42-49.

Maxwell, M. 2011. "Approximate Dynamic Programming Policies and Performance Bounds for Ambulance Redeployment." Dissertation, Graduate School of Cornell University.

Powell, W. 2007. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley Series in Probability and Statistics. John Wiley Sons, Inc.,

Hoboken, New Jersey.

Powell, W. 2012. "Perspectives of Approximate Dynamic Programming," *Annals of Operations Research*, Springer, DOI 10.1007/s10479-012-1077-6.

Powell, W., & Ryzhov, I. 2010. "Approximate Dynamic Programming With Correlated Bayesian Beliefs," *Proceedings of the 48th Allerton Conference on Communication, Control, and Computing*.

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

Sargent, T.J. 1987. *Dynamic Macroeconomic Theory*. Cambridge, MA: Harvard University Press.

Simao, H., Day, J., Abraham, G., Gifford, T., Nienow, T., and Powell, W. (2009). An Approximate Dynamic Programming Algorithm for Large-Scale Fleet Management: A Case Application. *Transportation Science* Vol. 43(2), 178-197.

Sutton, R., & Barto, A. 1998. *Reinforcement Learning*. Cambridge, Massachusetts: The MIT Press.

Tauchen, G. 1986. "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions." *Economic letters* Vol. 20, 177-181.

Van Roy, B., Bertsekas, D., Lee, Y., & Tsitsiklis, J.N. 1997. A Neurodynamic Programming Approach to Retailer Inventory Management. In *Proceedings of the 36th IEEE Conference on Decision and Control* Vol. 4, 4052-4057.

11 Appendix

11.1 Figures and Tables

Figure 1: Pre-Decision State vs. Post-Decision State

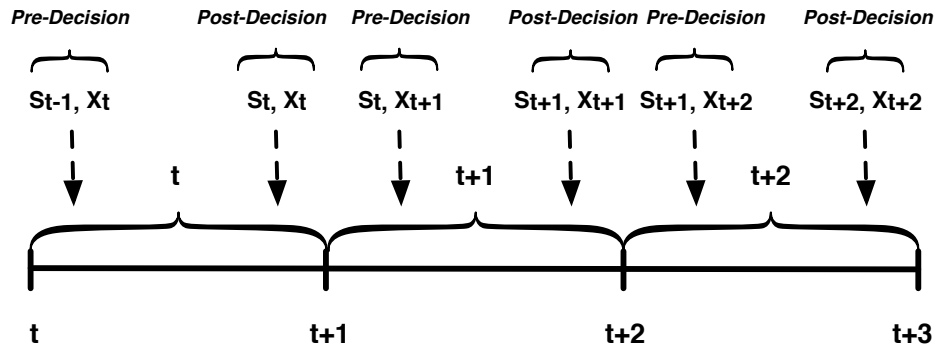


Figure 2: Law of Motion for Capital




Table 1: Standard Business Cycle Model (100 Individual Capital Nodes, 5 Productivity Nodes, 3000-Period Simulation Length)

METHOD	MAX	MEAN	STD	RT	DMS (.975)
VFI	1.700E-3	3.500E-3	1.700E-3	00:02:18:24	0.034
PM	7.500E-3	1.360E-2	7.500E-3	00:00:15:48	0.000
PEA	3.800E-2	3.640E-2	3.800E-2	00:00:00:65	0.003
ADP	2.620E-2	5.200E-2	2.620E-2	00:00:05:91	0.000
ADP (BTT)	2.820E-2	6.970E-2	2.100E-2	00:00:06:52	0.060

Table 2: Standard Business Cycle Model (100 Individual Capital Nodes, 20 Productivity Nodes, 3000-Period Simulation Length)

METHOD	MAX	MEAN	STD	RT
VFI	8.6594E-4	4.528E-4	2.463E-4	01:00:60:44
PM	5.800E-3	3.700E-3	1.700E-3	00:02:26:17
PEA	3.490E-2	3.410E-2	5.878E-4	00:00:00:80
ADP	2.060E-2	1.01EE-2	6.000E-3	00:00:35:47
ADP (BTT)	1.810E-2	8.500E-3	5.200E-3	00:00:31:73

Table 3: Heterogenous Agents Model (10 Individual Capital Nodes, 5 Aggregate Capital Nodes, 3000-Period Simulation Length, 200 Agents)

METHOD	MAX	MEAN	STD	RT	R^2 (H)	R^2 (L)
VFI	5.800E-2	2.480E-2	1.810E-2	00:00:40:00	0.9876	0.9641
PM	4.043E-1	1.586E-1	1.240E-1	00:01:21:50	0.9414	0.9397
PEA	1.033E-1	5.730E-2	2.940E-2	00:00:54:00	0.9079	0.9045
ADP	2.340E-1	1.114E-1	7.370E-2	00:00:88:20	0.9974	0.9996
ADP (BTT)	3.336E-1	1.764E-1	1.043E-1	00:00:44:02	0.9956	0.9939

Table 4: Heterogenous Agents (25 Individual Capital Nodes, 10 Aggregate Capital Nodes, 3000-Period Simulation Length, 200 Agents)

METHOD	MAX	MEAN	STD	RT	R^2 (H)	R^2 (L)
VFI	4.800E-3	2.500E-3	1.400E-3	00:07:25:82	0.9947	0.9825
PM	1.660E-2	7.600E-3	4.400E-3	00:06:12:52	0.9534	0.9625
PEA	2.700E-2	1.420E-2	7.600E-3	00:02:55:91	0.8515	0.8391
ADP	3.680E-2	1.810E-2	1.110E-2	00:05:05:40	0.9735	0.9512
ADP (BTT)	4.640E-2	2.490E-2	1.340E-2	00:01:04:98	0.9907	0.9822

Table 5: Heterogenous Agents (50 Individual Capital Nodes, 20 Aggregate Capital Nodes, 3000-Period Simulation Length, 200 Agents)

METHOD	MAX	MEAN	STD	RT	R^2 (H)	R^2 (L)
VFI	7.300E-4	3.833E-4	2.172E-4	01:43:25:55	0.9991	0.9987
PM	1.000E-2	5.200E-3	3.100E-3	00:09:19:36	0.9660	0.9652
PEA	7.500E-3	3.700E-3	2.300E-3	00:01:10:05	0.8942	0.8914
ADP	1.430E-2	7.300E-3	4.500E-3	00:10:01:49	0.9905	0.9905
ADP (BTT)	1.360E-2	7.900E-3	3.600E-3	00:03:12:13	0.9941	0.9919

Table 6: Baseline Parameter Values

Parameter	Description	Value
β	Discount Factor	.95
α	Capital's Share in Production	.33
σ	Standard Deviation of Technology Shock	0.035
ρ	Technology Level AR(1) Coefficient	0.8
δ	Physical Capital Depreciation Rate	.1

11.2 Proof: Closed-Form Solution for Full Depreciation Case

Assume that z is a Markov process. Additionally, let next period variables be denoted with primes. That is, $x_{t+1} = x'$. Finally, assume that $\delta = 1$ and recall the budget constraint and Euler equation for capital from the representative agent model described earlier:

$$c = k' - zk^\alpha \quad (44)$$

$$\frac{1}{c} = \beta E \left\{ \frac{\alpha z'}{c'} \middle| z \right\} \quad (45)$$

Now, assume that the decision rule for capital takes the following form:

$$\Xi(k, z) = g_0 + g_1 zk^\alpha \quad (46)$$

Next, plug the budget constraint into the Euler equation for capital:

$$\frac{1}{k' - zk^\alpha} = \beta E \left\{ \frac{\alpha z' k'^{\alpha-1}}{k'' - z' k'^{\alpha}} \middle| z \right\} \quad (47)$$

Now, add the assumed decision rule:

$$\frac{1}{g_0 + (g_1 - 1)zk^\alpha} = \beta E \left\{ \frac{\alpha z' (g_0 + g_1 zk^\alpha)^{\alpha-1}}{g_0 + g_1 z' (g_0 + g_1 zk^\alpha)^\alpha - z' (g_0 + g_1 zk^\alpha)^\alpha} \middle| z \right\} \quad (48)$$

To simplify things, we will look for a solution where $g_0 = 0$.¹³ This allows us to simplify the above equation as follows:

$$\frac{1}{(g_1 - 1)zk^\alpha} = \beta E \left\{ \frac{\alpha z' (g_1 zk^\alpha)^{\alpha-1}}{g_1 z' (g_1 zk^\alpha)^\alpha - z' (g_1 zk^\alpha)^\alpha} \middle| z \right\} \quad (49)$$

¹³It can be shown that there will be no solution for the case when $g_0 \neq 0$.

$$\rightarrow \frac{1}{(g_1 - 1)zk^\alpha} = \beta E \left\{ \frac{\alpha}{(g_1 - 1)B_1zk^\alpha} \middle| z \right\} \quad (50)$$

$$\rightarrow B_1 = \alpha\beta \quad (51)$$

We can then substitute this into our initial guess for the decision rule, yielding the following:

$$k' = \alpha\beta zk^\alpha \quad (52)$$

Note that we have demonstrated that the decision rule given above is compatible with the model's first order conditions.

CHAPTER 3

A PREDICTIVE MODEL OF FINANCIAL EPIDEMICS

Isaiah Hull

ABSTRACT

The SIR model is one of the most thoroughly developed frameworks for predicting the dynamics of infectious disease. This paper demonstrates how the SIR model can be microfounded in an economic context to make predictions about financial epidemics, such as the spread of asset-backed securities (ABS) and exchange-traded funds (ETFs), the proliferation of zombie financial institutions, and the expansion of bubbles and mean-reverting fads. The paper will proceed by developing the 1-host SIR model for economic and financial contexts; and will then move on to demonstrate how to work with the multi-host version of the model. In addition to showing how the SIR framework can be used to model economic interactions, it will also: 1) show how it can be simulated; 2) use it to develop and estimate a sufficient statistic for the spread of a financial epidemic; and 3) show how policymakers can impose the financial analog of herd immunity—that is, prevent the spread of a financial epidemic without completely banning the asset or behavior associated with the epidemic. Importantly, the paper will focus on developing a neutral framework to describe financial epidemics that can be either “bad” or “good.” That is, the general framework can be applied to epidemics that constitute a mean-reverting fad or an informational bubble, but ultimately yield little value and shrink in importance; or epidemics that are long-lasting and yield a new financial instrument that generates permanent efficiency gains or previously unrealized hedging opportunities.

JEL Classification: C510, C530, C580, G140

Keywords: Econometric Modeling, Econometric Forecasting, Financial Econometrics, Information Bubbles, Bubbles, Fads

1 Introduction

In a recent op-ed for Project Syndicate, Robert Shiller (2012) described a speculative bubble as “...a social epidemic whose contagion is mediated by price movements.” He then explained how initially high prices result in positive investor narratives, which spread to other investors, and result in the expansion of the bubble. He suggests that these phenomena are largely social, are divorced from fundamentals, and have the potential to result in substantial asset pricing anomalies.

Shiller’s work on speculative bubbles confirms this casual description (Shiller, 1990). In a series of questionnaires, Shiller surveyed wealthy individuals, individual investors, and institutional investors about three different events: 1) the stock market crash of 1987; 2) the real estate booms and busts of the late 1980s; and 3) the underpricing of initial public offerings (IPOs). One common thread in Shiller’s findings is that the decision making of investors often revolves around simple heuristics; and is influenced not only through social dialogue, but also through the perception of what others investors or institutions might do. For instance, he found that roughly two-thirds of investors surveyed in the U.S. and three-fourths of investors surveyed in Japan viewed psychological factors as more important than fundamentals when explaining the 1987 stock market crash. Additionally, when it came to real estate booms, he found that homeowners were substantially more likely to discuss the housing market—and, therefore, propagate positive narratives—when house prices were rising than when they were flat. Finally, with respect to IPOs, he found that wealthy individual investors were often more concerned with which types of companies other investors would find compelling than they were with the fundamentals of the company.

This suggests that the investors (individual and institutional) may use internal

models that rely on their perceptions of what others will do, rather than on fundamentals alone. Thus, when modeling the adoption of a financial instrument or increased demand for existing instruments, it is important to incorporate channels for investor-to-investor, investor-to-institution, and institution-to-institution interactions. In fact, in the last decade, these channels for investor interaction (or, more generally, for any agent interaction) have become increasingly developed in both the macroeconomics and finance literatures.

In the finance literature, Hong et al. (2010) embed an epidemiological model of opinion transmission within a standard asset pricing and trading framework. They assume that the agents can be divided into “informed” and “uninformed” groups. The informed group has recently received news or a common opinion about the value of an asset being traded; and the assimilation of this information changes their demand for the asset. The uninformed group has not yet received the news; and, thus, maintains the same level of demand, conditional on price. Over time, the news about the asset passes through the investor population, moving individuals from the uninformed to informed group, and leading to changes in aggregate investor demand that affect price. Hong et al. (2010) refer to this effect as “opinion diffusion”; and use it to explain non-linearities in price drift and trading volume.

Shive (2010) tests whether epidemiological models have predictive power for stock choice. The author uses municipality-level data in Finland to identify the number of holders of the 20 most active stocks in Finland over a 9 year horizon. She then attempts to predict the total number of buy orders placed per investor, using a several controls, along with a common measure of disease transmission: the number of investors who hold the stock, multiplied by the number of investors who do not. In the SIR model, this is equivalent to the number of susceptible individuals, multiplied by the number of infected individuals. The estimated coefficient gives us a measure of

how social interaction affects stock purchase decisions. The author finds that socially motivated stock trades not only affect the number of buy orders placed, but also the returns to the stock.

In addition to Shive (2010), other recent finance papers have demonstrated the importance of word-of-mouth effects and social interaction on buying decisions. Frieder and Subrahmanyam (2005), Feng and Seasholes (2004), and Hong et al. (2005) all document micro-level interactions between investors that are socially motivated. The units of observation range from mutual fund managers to individual investors—all of which were found to respond to other investors and other institutions in a way that was divorced from fundamentals.

In the macroeconomics literature, Carroll (2003) suggests that expectations of inflation are unlikely to be exactly rational. Rather, it is more likely that they fall somewhere between adaptive and rational, given the empirical evidence. He then uses a simple model from epidemiology—the single source model—to build expectations that are forward-looking, but not perfectly rational. In his model, only a fraction of households, λ , read the latest inflation prediction from a rational expectations forecaster in each period. The rest of the households retain the forecast they used in the previous period. He then shows that adjustment yields substantial performance improvements generating model with contractionary disinflations, lagged effects from monetary policy, inflation from rising growth rates, and increases in the natural rate of unemployment from productivity slowdowns. He also demonstrates that this idea is plausible by showing that the Philadelphia Fed’s Survey of Professional Forecasters can empirically be seen as the rational expectations forecast; whereas, population measures of expectations, such as the University of Michigan’s Inflation Expectations Survey, only contain a subset of that information.

Since Carroll (2003), several new additions to the macroeconomics literature have

incorporated epidemiological modeling as a means of capturing expectations formation. Pfajfar and Santoro (2012) and Sommer and Carroll (2004), for instance, use epidemiological modeling to capture the impact of news on inflation expectations and the impact of expectations on consumption, respectively. Additionally, Capistran and Timmerman (2006) and Brissimis and Magginas (2006) build off of Carroll (2003) to consider whether or not there is substantial heterogeneity in inflation expectations; and, if so, what impact it has on monetary policy.

My paper differs from the existing literature across several important dimensions: first, it nests a complete SIR model from epidemiology within a finance model; second, it generalizes the model to the case where there is more than just one class of agents; third, it derives R_0 , the sufficient condition for the take off of a bubble, within a financial context; fourth, it shows how R_0 can be estimated in financial models and used to predict the spread of new financial innovations; and fifth, it develops the concept of “herd immunity” as a strategy for bubble prevention.

More generally, this paper is not only designed to develop a sufficient statistic for bubble prediction, but also to provide a toolset for modeling contagion within an economic and financial context. Those who have an interest in developing models of contagion with multiple classes of agents may find this paper especially helpful.

Finally, it is important to note that the framework in this paper is designed to be sufficiently general to cover bubbles, mean-reverting fads, informational bubbles, and the permanent adoption of new financial instruments. That is, the purpose of using R_0 will be to detect a take-off, rather than to determine how such a take-off will end. It may be possible to develop strategies to separately identify each of these categories of event; however, that will not be the focus of the remaining sections.

This paper will proceed as follows. In Sections 2.1 and 2.2, I will provide a brief overview of the general methods used to solve epidemiological models for both the

1-host and 2-host case. I will show how the model parameters can be used to compute R_0 and to check the herd immunity condition. In section 2.3, I will build a financial model that explicitly nests the SIR model; and yields an estimable value for R_0 . In section 2.4, I will show how comparative statics can be used to determine whether herd immunity holds. I will then discuss and prove 12 theorems, which can be used to simplify the estimation of R_0 or to check the herd immunity condition in a 2-host model. Finally, in section 2.5, I estimate R_0 and check herd immunity in several applications of my model.

2 An Overview of the SIR Model

In this section, I will discuss the basic mathematical tools that form the SIR model in epidemiology. Beyond this paper, the toolkit described in this section may lend itself to useful modeling and econometric applications elsewhere in economics and finance; and, indeed, has made an appearance in a limited fashion in some of the work referenced in the introduction. The purpose of this section will be to provide a formal introduction to the toolkit that will be used later in this paper.

2.1 The Single-Host SIR Model

The single-host SIR model is one of the core analytical frameworks of mathematical epidemiology. It places structure on the flow of individuals into one of three groups based on pathogenic status: “susceptible,” “infected,” or “recovered” (sometimes called “removed”). Typically, an epidemic is simulated in the SIR framework by initially placing all individuals into the susceptible group. Next, an infected (or some mass of infecteds) is introduced into the population. In SIR models with “perfectly mixed” populations, it is assumed that each infected individual has the same

probability of transmission, c , per contact with a susceptible individual.¹ It is further assumed that the frequency of contact between individuals in a given period of time is q . Thus, the disease transmission rate is given by $\beta = qc$.

For simplicity, the mass of individuals is typically normalized to 1 in SIR models. This provides us with a convenient way to express the flow of individuals from the susceptible group to the infected group: βIS , where S is the mass of susceptibles and I is the mass of infecteds. Additionally, in some models, the population is assumed to grow at a fixed rate, μ_g , and individuals in the susceptible group are assumed to perish at rate μ_p . This leads to the following equation, which is one of the three that will describe the single-host SIR model:

$$\frac{\partial S}{\partial t} = \mu_g - \beta IS - \mu_p S \quad (1)$$

Notice that μ_g is not multiplied by S , since we assume that the birth rate is determined by the mass of the population—and not the mass of the susceptibles. On the other hand, we must multiply μ_p by S to get the flow of individuals who exit the susceptibles group, since the flow depends on the mass of S .

Next, we may write down the equation for the flows into and out of the infected group:

$$\frac{\partial I}{\partial t} = \beta IS - (\mu_p + \gamma + m)I \quad (2)$$

Here, μ_p is the mortality rate for individuals who perish for reasons other than exposure to the pathogen, γ is rate at which individuals recover from the disease and gain immunity, and m is the disease-induced mortality rate. If we want the population to

¹For those interested in further applications of the SIR model, the modern epidemiology literature contains many strategies for dealing with heterogenous populations, one of which will be applied in this paper.

retain a mass of 1 at all times, we can simply set $\mu_p = \mu_g$.

Finally, we may close the model by specifying an equation for the flows into and out of the recovered group:

$$\frac{\partial R}{\partial t} = \gamma I - \mu_p R \quad (3)$$

This equation simply states that the net flow of individuals into the recovered group (which is now immune to the pathogen²) consists of two components: 1) the mass of individuals who left the infected group by recovering γI ; and 2) the mass of individuals in the recovered group who died for reasons other than exposure to the disease, $\mu_p R$.

Finally, note that the continuous time SIR model is typically numerically simulated using a variation of the Runge-Katta integration scheme (Keeling & Rohani 2008). In order to further illustrate the purpose and output from this exercise, two graphs are provided in the appendix (Figure 1 and Figure 2) that illustrate the dynamics of a single-host SIR model. These simulations were performed in Matlab using the Runge-Katta routine. For simplicity, we assumed that $\mu_g = 0$ and $\mu_p = 0$ (i.e. the natural birth and death rates are zero). We also assumed that the disease was non-lethal (i.e. $m=0$). Figure 1 shows the dynamics of an epidemic following the introduction of an infected individual with mass of .01 into a population of susceptibles when $\beta = 1$ and $\gamma = .1$. Figure 2 shows the results for the same simulation, but with $\beta = 1$ and $\gamma = 1.01$.

Notice that the disease “invades” in Figure 1. That is, by the end of the epidemic, every person in the population has been infected by the disease, and has recovered, gaining immunity.

To the contrary, in Figure 2, the disease never invades. That is, not every sus-

²In some applications of the SIR model, immunity is only temporary. While this device may have useful applications in economics, we will not discuss this further in this subsection.

ceptible in the population becomes infected. Instead, the disease dies off before it reaches most of the population. To see why this is the case, consider the differential equation for the flows into the infected category at $t=0$:

$$\frac{\partial I}{\partial t} = \beta IS - \gamma I \quad (4)$$

If $\beta IS < \gamma I$, then more individuals will flow out of the infected group than will flow into it. This means that the group of infected individuals will start to shrink as soon as the disease is introduced into the population. This is called the “threshold phenomenon” (Keeling & Rohani 2008), and is the basis for constructing the basic reproductive number, R_0 :

$$\mathbb{R}_0 = \frac{\beta}{\gamma} \quad (5)$$

\mathbb{R}_0 is an important ratio because it can be used to determine whether or not a population of susceptibles has “herd immunity.” In particular, from the equation for the infecteds, we can see that a group will have herd immunity (that is, a new disease will not spread to the entire group of susceptibles) if the following condition holds:

$$\frac{1}{\mathbb{R}_0} > S \quad (6)$$

Furthermore, we can extend the concept of herd immunity by determining the minimum mass of susceptibles that must be immunized in order to achieved herd immunity. We do this by rewriting $S(0)$ as $S(1-P)$, where P is the fraction of immunized susceptibles

$$\frac{1}{\mathbb{R}_0} > S(1 - P) \quad (7)$$

Thus, if the condition below holds, then the population has herd immunity, and the

infectious disease will not invade. Notice that P must increase as \mathbb{R}_0 increases in order for the condition to hold:

$$P > 1 - \frac{1}{\mathbb{R}_0 S} \quad (8)$$

The basic reproductive number and the associated concept of herd immunity are critically important ideas in epidemiology. In the following sections, we will apply these concepts to economics by constructing a sufficient statistic for the inflation of bubbles. We will also see how the concept of herd immunity can be used to construct policies that prevents the inflation of bubbles efficiently.

2.2 SIR Model Extension: Multi-Host

In this section, I will construct a general model of financial bubbles using a 2-host version of the SIR model. From this, I will derive the basic reproductive number for the 2-host version—and, relatedly, determine the conditions under which herd immunity holds. Finally, I will explain how to construct the basic reproductive number for the general multi-host case. All of this will be done through the framework in McCormack (2007). To this author’s knowledge, this has not yet been done in economics.

Figure 3 in the appendix shows the basic structure of a 2-host SIR model. Note there are now two different categories of hosts, which are generically indexed 1 and 2. Individuals in each group may move from susceptible to infected and from infected to removed, as indicated by the solid arrows. Dotted arrows indicate that the mass of the infected group impacts the flows of individuals between groups. In particular, as the mass of infecteds of category 1 grows, susceptibles of category 1 and category 2 become more likely to encounter infecteds of category one. As a result, the flow of individuals from the susceptible to the infected category increase for both categories.

Using the same notation from the previous section, we may write down the six equations for a generic, two-host SIR model:

$$\frac{\partial S_1}{\partial t} = \mu_1^g - \beta_{1,1}I_1S_1 - \beta_{2,1}I_2S_1 - \mu_1^p S_1 \quad (9)$$

Note that μ_1^g is the growth rate of individuals in category 1, $\beta_{1,1}$ is the transmission rate between category 1 infecteds and category 1 susceptibles, $\beta_{2,1}$ is the transmission rate between category 2 infecteds and category 1 susceptibles, and μ_1^p is the mortality rate for reasons other than exposure to the infectious disease.

The remaining five equations follow a similar format and adopt identical subscripting conventions, so I will comment only on the differences from the single-host SIR model:

$$\frac{\partial S_2}{\partial t} = \mu_2^g - \beta_{2,2}I_2S_2 - \beta_{1,2}I_1S_2 - \mu_2^p S_2 \quad (10)$$

$$\frac{\partial I_1}{\partial t} = \beta_{1,1}I_1S_1 - \beta_{2,1}I_2S_1 - (\gamma_1 + \mu_1^p + m_1)I_1 \quad (11)$$

$$\frac{\partial I_2}{\partial t} = \beta_{2,2}I_2S_2 - \beta_{1,2}I_1S_2 - (\gamma_2 + \mu_2^p + m_2)I_2 \quad (12)$$

$$\frac{\partial R_1}{\partial t} = \gamma_1 I_1 - \mu_1^p R_1 \quad (13)$$

$$\frac{\partial R_2}{\partial t} = \gamma_2 I_2 - \mu_2^p R_2 \quad (14)$$

One important difference from the single-host SIR model is that we now have subscripts on γ , μ , m . This is because we are allowing the removal (or recovery) rate, the population growth rate, the population death rate (from reasons other than exposure to the disease), and the infectious disease mortality rate to vary according to

category.

With the definitions now in place for the six equations that characterize a generic SIR model, I will now discuss how to compute the basic reproduction number (BRN) for a two-host SIR model using the techniques described in Diekmann and Heesterbeek (1990), van den Driessche and Watmough (2002), and McCormack (2007).

Using the framework in the aforementioned papers, \mathbb{R}_0 (the BRN) can be computed from the spectral radius of the nxn matrix, M_n , where $M_n = (\mathbb{R}_{j,k})_{j,k=1}^n$. Each $\mathbb{R}_{j,k}$ in the matrix is given as follows:

$$\mathbb{R}_{j,k} = \frac{K_j \beta_{j,k}}{K_k (\gamma_k + m_k + \mu_k^p)} \quad (15)$$

Here, K_j and K_k are the carrying capacities for population groups j and k. Typically, the population masses for category 1 and category 2 individuals are normalized to 1, causing the carrying capacity terms will drop out of the equation. However, even without this normalization, the terms will cancel out in the special two-host case when \mathbb{R}_0 is computed.

Note that P_q refers to the fraction of the population of category q that has been immunized against the infection. These individuals are effectively removed from the SIR model dynamics. From this, the special case for the two-host SIR model is given as follows:

$$M_2 = \begin{pmatrix} \mathbb{R}_{1,1} & \mathbb{R}_{1,2} \\ \mathbb{R}_{2,1} & \mathbb{R}_{2,2} \end{pmatrix}$$

Taking the spectral radius of M_2 , BRN is computed as follows:

$$\mathbb{R}_0 = \frac{\mathbb{R}_{1,1} + \mathbb{R}_{2,2} + \sqrt{(\mathbb{R}_{1,1} - \mathbb{R}_{2,2})^2 + 4\mathbb{R}_{1,2}\mathbb{R}_{2,1}}}{2} \quad (16)$$

Re-writing this expression in terms of the parameters and variables in the two-host model yields the following:

$$\begin{aligned} \mathbb{R}_0 = & \frac{1}{2} \sqrt{\left(\frac{\beta_{1,1}}{\gamma_1 + m_1 + \mu_1^p} - \frac{\beta_{2,2}}{\gamma_2 + m_2 + \mu_2^p} \right)^2 + \frac{4\beta_{1,2}\beta_{2,1}}{(\gamma_2 + m_2 + \mu_2^p)(\gamma_1 + m_1 + \mu_1^p)}} \\ & + \frac{1}{2} \left(\frac{\beta_{1,1}}{\gamma_1 + m_1 + \mu_1^p} + \frac{\beta_{2,2}}{\gamma_2 + m_2 + \mu_2^p} \right) \end{aligned} \quad (17)$$

Finally, I will construct a modified version of herd immunity that does not include S (as it did in the previous subsection). Instead, I will use the fact that \mathbb{R}_0 is the average number of secondary cases of the infection that arise from a primary case. This means that the disease will invade (i.e. the number of individuals in the infectious category will rise at $t=0$) if the following condition holds:

$$\mathbb{R}_0 > 1 \quad (18)$$

That is, if each primary case tends to generate more than one secondary case, then the infectious groups will tend to grow when the initial infected individual is introduced into the group. Otherwise, the epidemic will not invade.

3 A Two-Host SIR Model of Financial Epidemics

Using the basic structure given in the previous section, I will now construct a model of financial epidemics. In particular, this model will feature interactions between two types of economic agents (i.e. the two types of hosts)—financial intermediaries and investors—that can lead to the inflation of bubbles, the creation of a fad, or another

variety of financial epidemic.

First, assume that there exists a continuum of both financial intermediaries and investors; and that their respective masses are normalized to one. Additionally, assume that each financial intermediary receives deposits from investors, and may choose to invest those deposits in one of two assets: 1) a conventional asset, which generates a normalized return of 1 with certainty, or 2) an unconventional asset, which has a uniformly-distributed, time-varying return:

$$b(t, i) \sim U[b^L(t), b^H(t)] \quad (19)$$

For simplicity, assume that the upper and lower boundaries decline linearly over time:

$$b^H(t) = b^H \max\{(1 - \delta t), 0\} \quad (20)$$

$$b^L(t) = b^L \max\{(1 - \delta t), 0\} \quad (21)$$

In the above two equations, b^H , b^L , and δ are constants. Given the assumptions made so far, the probability that a financial intermediary receives a return lower than some arbitrary threshold, \bar{T} , as given as follows:

$$pr\{b(t, i) < \bar{T}\} = \frac{\bar{T} - b^L(t)}{b^H(t) - b^L(t)} \quad (22)$$

Next, two additional assumptions are needed: 1) in each period of time, an investor who is initially holding conventional assets will search a fixed number of financial intermediaries, q , to determine the return they offer on investment. If an investor searches an intermediary that is investing in unconventional assets, it will switch to this intermediary with a fixed probability, c . Taken together, this gives us the transmission rate from financial intermediaries (category 2) to investors (category 1):

$$\beta_{2,1} = qc.$$

Additionally, assume that investors watch the behavior of other investors, and choose whether or not to lend to financial intermediaries based on that.³ Similarly, the transmission rate will be defined as the product of the frequency with which investors observe each others' investment decisions and the probability of switching from conventional to unconventional assets after observing a investor who invests in an unconventional asset intermediary: $\beta_{1,1}$.

Next, assume that financial intermediaries monitor a fixed number of other financial intermediaries in each period; and have a constant probability of copying the ones who invest in unconventional assets if they monitor them in that period. This gives us the intermediary-to-intermediary transmission rate: $\beta_{2,2}$. Finally, intermediaries come into contact with investors who are searching rates at different intermediaries. Observing the outcome of this process spreads information about investor preference for the unconventional asset—and results in conventional asset intermediaries becoming unconventional asset intermediaries with rate $\beta_{1,2}$.

With a theoretical foundation for the transmission rates in place, the remaining dimensions of the model can now be calibrated:

1. μ_1^g, μ_2^g - Investor and financial intermediary entry rates.
2. μ_1^p, μ_2^p - Investor and financial intermediary exit rates.
3. m_1, m_2 - The rates at which investors and financial intermediaries go bankrupt as a result of the unconventional asset's poor performance.
4. γ_1, γ_2 - The rates at which investors and financial intermediaries stop investing in the unconventional asset and instead invest in the conventional asset.

³Papers such as Shive (2010) and Hong et al. (2010) suggest that there is empirical evidence for socially-motivated trades.

Note that the rates described do not have to be structural parameters; and, in fact, could emerge from optimization problems. In order to limit the scope of this paper, however, I will assume that entry and exit rates are set to zero or will be estimated using data.

Next, I microfound bankruptcy rates induced by poor unconventional asset performance. I do this by assuming that all investors start at $t=0$ with D units of currency to deposit; and that they place those funds in financial intermediaries that invest in conventional assets, since the unconventional asset has not yet been invented at $t=0$. Furthermore, I assume that all banks start with τD in capital. Applying the Law of Large Numbers (LLN) yields the following bankruptcy rate, m_2 , for financial intermediaries:

$$m_2 = pr\{(1 - b(t))I_1D > \tau I_2D\} = pr\{1 - b(t) > \tau I_2/I_1\} \quad (23)$$

$$= pr\{b(t) < (I_1 - \tau I_2)/I_1\} \quad (24)$$

$$\rightarrow m_2 = \frac{\frac{I_1 - \tau I_2}{I_1} - b^L(t)}{b^H(t) - b^L(t)} \quad (25)$$

Note that the conventions from the SIR model have been retained, even though we are now using an entirely financial model. That is, the bankruptcy rate, m_2 , is the mortality rate from the SIR model. Additionally, I_1 denotes the fraction of investors who are investing in financial intermediaries that purchase unconventional assets and I_2 denotes the fraction of financial intermediaries that purchase unconventional assets. Similarly, in the SIR model, this would denote the fraction of infecteds in each population.

With this in mind, the interpretation of equation 25 is simple: if an intermediary invests in an unconventional asset, and that asset generates a return lower than the principal invested, then the intermediary will have to use capital to repay investors.

If the capital is not sufficient to repay principal, then the financial intermediary will go bankrupt after dividing all remaining funds among investors.⁴

One important feature of this setup is that bankruptcy rates depend on financial intermediary leverage. As the mass of unconventional asset investors grows relative to capital holdings in unconventional asset intermediaries (i.e. as a rough measure of leverage increases), the bankruptcy rate for intermediaries accelerates.

Next, consider the bankruptcy rate for investors. Note that investors in this model do not face a capital constraint; and, thus, will not go bankrupt if they lose some of the principal invested. Instead, investors will only go bankrupt if they lose their initial endowment, which happens with zero probability:

$$m_1 = 0 \tag{26}$$

I will now place structure on the rate at which unconventional asset intermediaries switch back to purchasing conventional assets permanently. For the sake of simplicity, I will retain the term “recovery rate.” In this model, I assume that this happens when such an intermediary experiences a loss that reduces its capital holdings, but isn’t sufficiently large to trigger bankruptcy. As a result of the intermediary losing some of its initial capital holdings, they choose not to continue purchasing unconventional assets permanently⁵—and instead dump the remaining principal in conventional assets.

Next, the recovery rate for the financial intermediary is computed. Two conditions must be satisfied in order for an intermediary to move from the infected to recovered category: 1) the intermediary must experience a loss—that is, $b(t) < 1$; and 2) the loss

⁴Of course, this is a strong condition for bankruptcy. Firms may often choose to initiate a strategic bankruptcy, even if their capital holdings are not completely depleted; however, I will maintain this assumption in the model for tractability.

⁵Note that this “recovery” could be temporary—or it could allow them to issue conventional assets again with the possibility of moving back to unconventional ones. There are modifications of the SIR model that permit both options.

must be smaller than the amount of capital the intermediary holds: $I_1 D(1 - b(t)) < \tau D I_2$. Again, taking advantage of the LLN and the uniformly distributed returns assumption yields the following rate of return for financial intermediaries:

$$\gamma_2 = pr\{I_1(1 - b(t)) < \tau I_2 | b(t) < 1\} \quad (27)$$

$$\rightarrow \gamma_2 = \frac{\frac{\tau I_2}{I_1}}{b^H(t) - b^L(t)} \quad (28)$$

Finally, the rate at which investors switch from unconventional assets to conventional assets (i.e. the investor recovery rate) is computed. Recall that a financial intermediary will repay principal until it has exhausted all of its capital. This means that an investor will not incur a loss unless the financial intermediary becomes bankrupt. Thus, the fraction of investors who hold unconventional assets and then “recover,” switching back to conventional assets permanently, is equal to the fraction of firms that become bankrupt. Recall that this was computed earlier as the bankruptcy rate for financial intermediaries:

$$\gamma_1 = \frac{\frac{I_1 - \tau I_2}{I_1} - b^L(t)}{b^H(t) - b^L(t)} \quad (29)$$

Taken together, equations (9)-(14), (19)-(21), (25), (26), (28), and (29) specify a complete, two-host SIR model of financial bubbles. In the remaining sections, I will 1) use the fully-specified model to perform comparative statics on \mathbb{R}_0 and 2) estimate the model parameters and compute \mathbb{R}_0 for three applications.

3.1 Comparative Statics

Recall that herd immunity requires the following condition on \mathbb{R}_0 to hold:

$$\begin{aligned} \mathbb{R}_0 = & \frac{1}{2} \sqrt{\left(\frac{\beta_{1,1}}{\gamma_1 + m_1 + \mu_1^p} - \frac{\beta_{2,2}}{\gamma_2 + m_2 + \mu_2^p} \right)^2 + \frac{4\beta_{1,2}\beta_{2,1}}{(\gamma_2 + m_2 + \mu_2^p)(\gamma_1 + m_1 + \mu_1^p)}} \\ & + \frac{1}{2} \left(\frac{\beta_{1,1}}{\gamma_1 + m_1 + \mu_1^p} + \frac{\beta_{2,2}}{\gamma_2 + m_2 + \mu_2^p} \right) < 1 \end{aligned} \quad (30)$$

Furthermore, note that herd immunity in this context refers to a state in which an unconventional asset will not spread to the entire population of intermediaries and investors that could potentially hold it.

Now, assume that $\mu_1^p, \mu_2^p = 0$. That is, investors and intermediaries do not become bankrupt for reasons other than receiving poor returns to the unconventional asset.⁶ Additionally, recall that $m_1 = 0$ —that is, investors are assumed not to become bankrupt as a result of the unconventional asset. The condition on \mathbb{R}_0 now is now simplified to the following:

$$2 > \sqrt{\left(\frac{\beta_{1,1}}{\gamma_1} - \frac{\beta_{2,2}}{\gamma_2 + m_2} \right)^2 + \frac{4\beta_{1,2}\beta_{2,1}}{(\gamma_2 + m_2)\gamma_1}} + \left(\frac{\beta_{1,1}}{\gamma_1} + \frac{\beta_{2,2}}{\gamma_2 + m_2} \right) \quad (31)$$

Using assumptions and simple algebraic manipulations, it is now possible to determine the conditions under which markets will be susceptible to the spread of the unconventional asset. For expository purposes, I will relegate most of the algebra to the appendix and focus primarily on the results.

⁶This is certainly not an accurate assumption; however, in more rigorous applications, this rate can be calibrated to match real bankruptcy rates.

3.1.1 Cross-Transmission Only

If it is assumed that cross-transmission (i.e. intermediary-to-investor and investor-to-intermediary) is the only channel for persuading intermediaries and investors to switch to the unconventional asset, the following assumptions are imposed on the model: 1) $\beta_{1,1}=0$, and 2) $\beta_{2,2} = 0$. This reduces the herd immunity condition to the following:

$$2 > \sqrt{\frac{4\beta_{1,2}\beta_{2,1}}{(\gamma_2 + m_2)(\gamma_1)}} \quad (32)$$

$$\rightarrow (\gamma_2 + m_2)\gamma_1 > \beta_{1,2}\beta_{2,1} \quad (33)$$

The above equation provides proof for our first theorem:

Theorem 1. *If $\beta_{1,1} = 0$ and $\beta_{2,2} = 0$, then herd immunity will hold whenever $(\gamma_2 + m_2)\gamma_1 > \beta_{1,2}\beta_{2,1}$.*

This result suggests that higher recovery and bankruptcy rates will drain the pool of unconventional investors and intermediaries faster. If they do this sufficiently fast—relative to the cross-transmission rates—then the unconventional asset will never take off. That is, an epidemic invasion—as described in the SIR model—will not occur.

Next, I will consider placing additional restrictions on the parameter values. This may be helpful to do if certain parameters are inestimable; and we need to determine them by theory in order to check herd immunity.

First, we note that $\gamma_2 + m_2 = \max\{\frac{1-b^L(t)}{b^H(t)-b^L(t)}, 0\}$ and $\gamma_1 = \max\{\frac{\frac{I_1-\tau I_2}{I_1}-b^L(t)}{b^H(t)-b^L(t)}, 0\}$ because both are probabilities. This is used to construct the next theorem:

Theorem 2. *If $\beta_{1,2} > 0, \beta_{2,1} > 0, \beta_{2,2} = 0, \beta_{1,1} = 0$, and either $b^L(t) = 1$ or $b^L(t) \geq \frac{I_1-\tau I_2}{I_1}$ and $b^L(t) < 1$, then herd immunity is violated.*

This theorem states that the unconventional asset will spread to all susceptible conventional intermediaries and investors if 1) cross transmission is the only channel for switching to the unconventional asset, and 2) the lower bound for the return on the unconventional asset is high enough to ensure that no financial intermediaries investing in it become bankrupt.

Theorem 3. $\beta_{1,2} > 1, \beta_{2,1} > 1, \beta_{2,2} = 0$, and $\beta_{1,1} = 0$, then herd immunity is violated if $b^H(t) \geq 1$.

Since $\gamma_2 + m_2 \leq 1$ and $\gamma_1 \leq 1$ (since it is a probability), the inequality in equation 33 will not hold if $\beta_{1,2} > 1$ and $\beta_{2,1} > 1$. This suggests that if the strength of cross-transmission is sufficiently large, then the unconventional asset will continue to spread, regardless of bankruptcy and recovery rates.

3.1.2 No Intermediary-to-Intermediary Transmission

With no intermediary-to-intermediary transmission, $\beta_{2,2}$ is effectively set to zero. That is, we assume that financial intermediaries consider only the decisions of investors when deciding whether or not to switch to the unconventional asset; and do not pay attention to the choices of other intermediaries. This reduces the condition for herd immunity to the following:

$$2 > \sqrt{\left(\frac{\beta_{1,1}}{\gamma_1}\right)^2 + \frac{4\beta_{1,2}\beta_{2,1}}{(\gamma_2 + m_2)\gamma_1}} + \frac{\beta_{1,1}}{\gamma_1} \quad (34)$$

Some algebraic manipulation yields the following, simpler condition:

$$(\gamma_1 - \beta_{1,1})(\gamma_2 + m_2) > \beta_{1,2}\beta_{2,1} \quad (35)$$

This yields the first theorem for the case with no intermediary-to-intermediary transmission:

Theorem 4. *If $\beta_{2,2} = 0$ and $(\gamma_1 - \beta_{1,1})(\gamma_2 + m_2) > \beta_{1,2}\beta_{2,1}$, then herd immunity holds.*

Next, note that $1 > (\gamma_2 + m_2) \geq 0$, as was shown earlier. Since $\beta_{1,2} > 0$, $\beta_{2,1} > 0$, and $\beta_{1,1} > 0$, the left hand side will be smaller whenever $\beta_{1,1} > \gamma_1$. This gives yields the next three theorems.

Theorem 5. *If $\beta_{2,2} = 0$ and $\gamma_1 < \beta_{1,1}$, then herd immunity is violated if $\beta_{1,2}\beta_{2,1} > (1 - \gamma_1)\frac{\beta_{1,1}^2}{4\gamma_1}(\gamma_2 + m_1)$.*

Theorem 6. *If $\beta_{2,2} = 0$ and $\beta_{1,1} \geq 1$, then herd immunity is violated if $\beta_{1,2}\beta_{2,1} > (1 - \gamma_1)\frac{\beta_{1,1}^2}{4\gamma_1}(\gamma_2 + m_1)$*

Theorem 7. *If $\beta_{1,2} > 0, \beta_{2,1} > 0, \beta_{2,2} = 0$, and $b^L(t) \geq 1$, then herd immunity is violated.*

3.1.3 No Investor-to-Investor Transmission

If there is no investor-to-investor transmission, then $\beta_{1,1} = 0$. This reduces the condition for herd immunity to the following:

$$2 > \sqrt{\frac{\beta_{2,2}^2}{(\gamma_2 + m_2)^2} + \frac{4\beta_{2,1}\beta_{1,2}}{(\gamma_2 + m_2)\gamma_1}} + \frac{\beta_{2,2}}{\gamma_2 + m_2} \quad (36)$$

$$\rightarrow (\gamma_2 + m_2 - \beta_{2,2})\gamma_1 > \beta_{2,1}\beta_{1,2} \quad (37)$$

This expression gives us our first investor-to-investor theorem:

Theorem 8. *If $\beta_{1,1} = 0$ and $(\gamma_2 + m_2 - \beta_{2,2})\gamma_1 > \beta_{1,2}\beta_{2,1}$, then herd immunity holds.*

Next, consider the conditions under which the left hand side will be nonpositive (and, thus, herd immunity will be violated).

Theorem 9. *If $\beta_{1,1} = 0$ and $\beta_{2,2} > (\gamma_2 + m_2)$, then herd immunity is violated.*

Theorem 10. *If $\beta_{1,2} > 0, \beta_{2,1} > 0, \beta_{1,1} = 0$, and $b^L(t) \geq 1$, then herd immunity is violated.*

Theorem 11. *If $\beta_{1,2}\beta_{2,1} > 2, \beta_{1,1} = 0$, then herd immunity is violated.*

Theorem 12. *If $\beta_{2,2} > 2, \beta_{1,1} = 0$, then herd immunity is violated.*

We now have 12 theorems that can be used to place theoretical restrictions on parameters. In some cases, making a couple of weak theoretical restrictions will make it possible to check the herd immunity condition by estimating only one parameter.

3.2 Empirical Applications

In this section, I will discuss three applications of the SIR model that involve estimating R_0 . In the first, I will 1) estimate one equation from the SIR model, 2) use the results to compute R_0 , 3) use the delta method to compute the standard error of R_0 , and then 4) perform a hypothesis test to determine whether or not the new financial security will spread. In the second application, I will 1) estimate two separate 1-host models for a related phenomena, 2) impose theoretical restrictions, and then 3) determine whether or not herd immunity is violated in the 2-host version of the model. In the final application, I will show how estimating R_0 25-quarters after ETFs were released in the U.S. could have been used to predict their widespread popularity. In addition to this, I will demonstrate an alternative method of estimating R_0 directly, rather than constructing it using β and γ estimates.

3.2.1 Detecting the Spread of a Financial Innovation

In the first application, I consider the introduction of non-agency, non-GSE securitized debt. The purpose of this application will be to demonstrate how to make predictions about the spread of new financial securities or the inflation of bubbles using the 1-host SIR framework.

The dataset was constructed using three series: 1) the level of GSE-securitized debt; 2) the level of household debt; and 3) the level of non-GSE securitized debt. The series for S (i.e. the fraction of debt that is susceptible to private securitization) is constructed as follows:

$$S = \frac{\text{Total Debt} - \text{GSE-Securitized Debt} - \text{Privately-Securitized Debt}}{\text{Total Debt} - \text{GSE-Securitized Debt}} \quad (38)$$

For simplicity, I will assume that there is no recovery in the model. That is, financial firms that stop securitizing do not do so permanently. Instead, they simply stop securitizing for now, but remain open to the possibility in the future. This simplifies the model by allowing us to compute I (the fraction of securitizers) as follows:

$$I = 1 - S \quad (39)$$

This framework is called the SIS model, and has an R_0 that is identical to the one derived earlier for the 1-host SIR model. Without placing any additional structure on the problem, it is possible to estimate equation (4):

$$\frac{\partial I}{\partial t} = \beta IS - \gamma I \quad (40)$$

Notice that γ measures the rate at which financial firms exit the securitization business and return to issuing non-AB securities. Note also that β is the contact rate multiplied by the probability of becoming a securitized. That is, in each period, a financial firm scans some fraction of other financial firms to determine which products they offer. If the firm encounters a securitizer, then it also becomes a securitizer with some fixed probability. The product of these two numbers is β .

Equation (5) was estimated with IV on quarterly data from 1983-1998, using lags of I, SI, and S as instruments. The results are given in Table 1 in the appendix. The transmission rate ($\beta = 0.2525$) was found to exceed the mortality rate ($\gamma = .2309$), indicating that an epidemic invasion (i.e. the spread of ABS) was highly likely. Table 1 also provides R_0 and its standard error.

Note that R_0 was computed as $\frac{\beta}{\gamma}$, yielding 1.0936. R_0 's standard error was then computed as follows: let Σ represent the covariance matrix and let $D = [\frac{1}{\gamma}, -\frac{\beta}{\gamma^2}]$ be the row vector of partial derivatives of R_0 . Then, by the delta method, the variance of R_0 is $var(R_0) = D\Sigma D'$. In this case, $SE(R_0) = 0.1836$. Now, recall that herd immunity will hold (i.e. the innovation will not spread) if:

$$\frac{1}{R_0} > S(1 - P) \tag{41}$$

Since we have the standard error for R_0 , we can perform a hypothesis test using data on S and P to see if herd immunity is violated. If it is assumed that P=0 and the test is performed at the end of the period (1998), then we may reject the null hypothesis with 95% confidence. That is, in 1998, an estimate of R_0 would have provided support for the claim that private securitization was likely to spread rapidly, which is precisely what happened until the financial crisis occurred in 2007.

In practice, we could use new data to update S, P, and R_0 in each period; and

then check whether or not herd immunity holds. If herd immunity does not hold, then it may make sense to monitor the financial instrument carefully, as it could play an increasingly important role in determining the stability of the financial system. Alternatively, if herd immunity does hold, then this could suggest that the spread of the instrument will cease in the future, potentially bursting a bubble.

From a policy perspective, we might consider adjusting P . That is, policy-makers might consider increasing P if they worry that a risky and untested financial instrument is likely to spread rapidly (i.e. has a high R_0). They could achieve this by placing temporary legal limits on the issuance of the security until we have a longer period of time to evaluate its performance.

Finally, notice that we completely ignored the returns to privately-issued asset backed securities when we performed this analysis, even though this is likely to play an important role in determining entry and exit. In the next section, we will consider an example that makes a more serious attempt to use the theory developed earlier to guide estimation.

3.2.2 Checking the Susceptibility of FDIC-Insured Banks to a Zombie Invasion

During the S&L Crisis of the 1980s and 1990s, many savings and loan institutions became insolvent, but were able to continue operating because of implicit and explicit government guarantees, such as deposit insurance. Kane (1987) coined the phrase “zombie bank” to describe such institutions, which were effectively dead, but were able to continue to lend and to collect deposits.

Kane (1987) points out that zombie institutions were particularly worrisome because they were tempted to “gamble for resurrection.” That is, knowing that equity holders would get wiped out by an asset revaluation with certainty, they would find

ways to attract depositors (e.g. by offering higher interest rates); and then invest the new funds in increasingly risky projects. If the bank was lucky and the investments paid off, then the institution would not be insolvent at revaluation, saving the equity-holders. If, on the other hand, the investments did poorly, then the new depositors would effectively subsidize the equity holders by bearing some of the losses.

The behavior of zombies is problematic for two reasons. First, it encourages institutions to lend to increasingly risky borrowers. And second, it encourages nearby institutions to do the same. If they did not, then they could not compete with the rates zombie banks offer depositors (and creditors more generally), which results in a shift of deposits away from healthy institutions.

While the industry has undergone substantial reform since the S&L crisis, zombie banks still exist; and still have the bad incentives that come from implicit and explicit government guarantees. In a recent paper, Zwick (2012) describes a strategy for identifying zombie banks in FDIC data; and uses it to examine their impact on the Great Recession.⁷

I follow Zwick’s zombie-detection approach by constructing an unbalanced panel out of 12 years (2001-2012) worth of institution-level cross-sections from the FDIC’s call reports. I then use the FDIC certificate from the call reports to match the institutions to the ones listed in CalculatedRisk’s unofficial record of Prompt Corrective Actions (PCA) and Desist orders. The combined dataset allows me to identify banks by location (state and county), balance sheet data, and zombie status (i.e. if they were given an order). It also permits me to link banks across time.

From there, I construct the model variables, S and I. Initially, I assume that all banks are susceptible to zombification, and identify the institutions that are given a

⁷It is important to note that Zwick (2012) is more concerned with the “evergreening” phenomena that is described by Caballero (2008), rather than risky lending. In fact, his paper suggests that banks may tend to repeatedly rollover loans, rather than originating new, risky loans.

PCA or Desist order as infected institutions. The reason I label them as infecteds prior to the order is because this is when banks are most likely to engage in high risk behavior in an attempt to divert debt-financed funding from other institutions. Furthermore, I define the mass of infecteds, I , as the fraction of deposit volume held by these institutions. I then estimate the following equation with OLS to recover the transmission and mortality rates:

$$I_{i,t+1} - I_{i,t} = \beta I_{i,t} S_{i,t} - \gamma I_{i,t} + \phi_l + u_t + \epsilon_{i,t} \quad (42)$$

Note that ϕ_l is a county fixed effect and u_t is a time effect. The fixed effects will allow us to sweep out any variation that is attributable to county-specific idiosyncrasies and has little to do with the interaction between zombies and non-zombies. The time effects will sweep out the effect of national trends, such as expansions and recessions that might create a common time trend in the fraction of deposits held by zombies.

The results are given in Table 2 of the appendix. Notice that there are separate regressions in this table: one for the fraction of deposit volume held by zombies and another for the fraction of asset volume held by zombies within a county. The purpose behind this is to test whether zombie banks create a pure deposit-diversion epidemic or whether zombie behavior also has a substantial impact on assets.

The results suggest that there is evidence for strong deposit-side transmission at the county level. Additionally, the relative rates of transmission and mortality of infecteds suggests that deposit-side zombification (i.e. diversion of deposits away from healthy banks) is susceptible to an epidemic if some shock introduces a small number of zombies initially. In particular, recall that a county will have herd immunity if the

following condition holds:

$$\frac{1}{R_0} \geq S_t(1 - P) \quad (43)$$

R_0 is given in Table 2, and was constructed from β and γ , which were recovered from the regression. For deposits, $R_0 = 1.2087$ and has a standard error of 0.0825, which was calculated using the delta method. This implies that a county will lose herd immunity with 95% confidence if $S_t(1 - P) > .8183$. That is, if a county initially has no zombies $S = 1$, but then a financial shock hits and causes an institution to have negative net worth (i.e. become a zombie), then that institution will tend continue to divert deposits away from healthy institutions and convert healthy institutions into zombies. Additionally, this process can be affected by regulatory measures that reduce the stock of susceptible deposits by affecting P .

It is important to note that the “epidemic” described is with respect to volume of assets being held by zombies, rather than the proportion of institutions that are zombies. It is entirely possible, for instance, that measures targeted at regulating non-zombies more thoroughly (i.e. increasing P) could increase the deposit flow volume to zombies, which would not necessarily help to impose herd immunity. However, if herd immunity is achieved, then the interpretation behind the epidemic’s end is simple: zombies will be identified and dismantled by regulatory agencies faster than they are able to divert deposits, resulting in a net decrease in deposit inflow to zombie institutions.

Next, consider the results for the asset regression. Here, we can see that transmission is not particularly strong (i.e. statistically insignificantly different from zero). That is, an increase in contact between zombies and non-zombies does not lead to particularly large increases in zombie asset volume. Furthermore, the zombie mortality rate is large and statistically significant, which leads to a low R_0 of .4623. In

this case, a zombie-driven epidemic is simply not possible on the asset-side of bank balance sheets, even if $P=0$.

With these results in hand, we may now return to the theoretical model of “financial intermediaries” and “investors” to take advantage of the theorems derived earlier. In the zombie bank model, the depositors are equivalent to the investors and the banks are the financial intermediaries that choose between normal and risky assets. Furthermore, assume that we want to know if a joint epidemic (deposits and assets) is possible if we allow for cross-transmission—that is, allow the deposit epidemic to affect the asset epidemic.

From the asset regression in Table 2, we know that there is no statistically significant intermediary-to-intermediary transmission (i.e. $\beta_{2,2} = 0$). This means that we can use the theorems derived in section 2.4.2. In particular, Theorem 7 states the following: if $\beta_{1,2} > 0$, $\beta_{2,1} > 0$, and $\beta_{2,2} = 0$, and $b^L \geq 1$, then herd immunity does not hold. In words, this means that a county will be susceptible to an epidemic when a zombie is introduced if two conditions are met: 1) all zombie institutions in the county are receiving weakly positive excess returns; and 2) there is positive cross-transmission from financial intermediaries to investors and vice versa (even if it is very small).

3.2.3 Predicting the Displacement of Existing Funds by ETFs

As a final example, I considered the expansion of exchange-traded funds (ETFs), which were introduced in the United States in 1993, and then steadily grew in popularity thereafter. In particular, I use data from the 1990s to estimate R_0 , compare it to the fraction of other fund volume that was susceptible to being converted into an ETF (e.g. close ETF substitutes), and then check whether or not the ETF takeoff could be predicted.

For the purposes of this exercise, I used Flow of Funds (FoF) data from financial businesses, which was listed in the Federal Reserve Board of Governors' data download program. I computed the total volume of funds susceptible to ETF conversion by adding up the volumes for all assets in the same category as ETFs. This included several varieties of money market mutual funds, mutual funds, and closed-end funds—all predecessors to ETFs and also likely substitute investment vehicles.

I then used the total volume of funds to compute S and I . S was defined as the fraction of asset volume that was concentrated in non-ETFs. And I was defined as the fraction of total asset volume concentrated in ETFs. With these definitions in place, I used the first 25 quarters of data to recover an early estimate of the basic reproductive ratio and then used the delta method to compute its standard error. I found that $R_0 = 1.0267$ with a standard error of 0.1510. When R_0 was computed, the fraction of susceptible funds (0.9817) was greater than $\frac{1}{R_0}$, which suggests that herd immunity was already violated and an invasion was likely. Figure 4 in the appendix plots the growth of ETFs as a fraction of total fund volume over time. The dotted black line marks the data that was used to make the prediction.

Finally, note that ETFs were later introduced in Europe in 1999. This means that European regulators could have used the estimate of R_0 from 1998 to determine how to handle the introduction of ETFs. In particular, knowing that they were likely to become a popular asset class might have encouraged regulators to handle their introduction with greater scrutiny; and to dedicate more resources to their monitoring.

3.3 Estimating R_0 Directly

Notice that we can rearrange the flow equation for infecteds as follows:

$$I_{t+1} - I_t = \gamma(R_0 I_t S_t - I_t) \quad (44)$$

Now, assume that $I_t S_t$ and I_t are I(1) processes and $I_{t+1} - I_t$ is an I(0) process. Instead of estimating the above equation, we may estimate the implied cointegration vector, $(R_0, 1)$. In this case, the normalization choice is clear, since we wish to recover R_0 directly. We may do so by estimating the following equation via OLS:

$$I_t = R_0 I_t S_t + \epsilon_t \quad (45)$$

There are two important benefits of this approach. First, under the null hypothesis of cointegration, estimates of R_0 will be superconsistent, converging at rate T , rather than $T^{1/2}$. This eliminates the need for IV. The second benefit is that we do not need to compute the standard errors for R_0 separately with the delta method, since R_0 is estimated directly using OLS.

In order to provide a complete demonstration of this approach, I use it to repeat the estimation exercise for ETFs and to confirm the value of R_0 . In addition to this, I also check the assumptions of the approach by performing nonstationarity tests for I_t and $I_t S_t$, and by testing whether the two series are cointegrated under the vector $(\hat{R}_0, 1)$.

I start by performing the following two regressions:

$$I_t = \alpha + \rho I_{t-1} + \hat{u}_t \quad (46)$$

$$I_t S_t = \alpha + \rho I_{t-1} S_{t-1} + \hat{u}_t \quad (47)$$

I then use the residuals from each regression to perform two separate Phillips-Perron

tests, where the null hypothesis is that the series is nonstationary. As described in Hamilton (1994), the Phillips-Perron Z_ρ test statistic can be computed in the following set of steps:

$$s^2 = (T - 2)^{-1} \sum_{t=1}^T \hat{u}_t^2 \quad (48)$$

$$\hat{c}_j = (T)^{-1} \sum_{t=j+1}^T \hat{u}_t \hat{u}_{t-j} \quad (49)$$

$$\hat{\lambda}^2 = \hat{c}_0 + 2 \sum_{j=1}^4 [1 - (j/5)] \hat{c}_j \quad (50)$$

$$Z_\rho = T(\hat{\rho} - 1) - \frac{1}{2} \left(\frac{T\hat{\sigma}_\rho}{s} \right)^2 (\hat{\lambda}^2 - \hat{c}_0) \quad (51)$$

The test statistic, Z_ρ , is then compared to the relevant 5% critical value, which is -12.5 for $T=25$ and $K=2$. For equation (46), $Z_\rho = 3.5132 > -12.5$, which means we fail to reject the null of nonstationarity. For equation (47), $Z_\rho = 3.4670 > -12.5$, which again means we fail to reject the null hypothesis of nonstationarity.

Next, equation (45) can be estimated using OLS to recover \hat{R}_0 . Here, I find that $\hat{R}_0 = 1.0697$, which is close to the previous subsection's estimate of 1.0267. With \hat{R}_0 recovered, the cointegration of $I_t S_t$ and I_t can be tested next. To do this, I first compute the series, z , using \hat{R}_0 :

$$z_t = \hat{R}_0 I_t S_t - I_t \quad (52)$$

Next, I perform a Phillips-Perron test using the residual from the following regression:

$$z_t = \alpha + \rho z_{t-1} + \hat{u}_t \quad (53)$$

The null hypothesis is that the cointegrating relationship does not hold. For equation (53), $Z_\rho = -18.2163 < -12.5$, which means we may reject the null hypothesis of no cointegration. Combined with our earlier findings, this suggests that the estimate of R_0 is superconsistent.

It is important to note that there are alternative approaches to estimating R_0 that may have better statistical properties in certain circumstances. However, if I_t and $I_t S_t$ are cointegrated, then recovering R_0 becomes simple. All a practitioner must do is define the three groups (S, I, and possibly R), transform the data to match the group definitions, and then estimate equation (45). If there is any doubt about the underlying assumptions needed for superconsistency, then the set of tests performed in this subsection can be employed.

4 Conclusion

The SIR model is one of the most thoroughly developed frameworks in mathematical epidemiology. It has many tools for prediction that can readily be applied to problems in economics and finance after appropriate adjustments have been made for context. In this paper, I describe how the multi-host version of the SIR model can be used in economics to make predictions about financial innovations, bubbles, and fads.

In addition to importing the multi-host SIR model, I also develop its tools for use in economics. In particular, I construct a financial model that yields R_0 as a predictive measure of the take off of new financial innovations; and I explain how R_0 can be computed when the model has more than one host.

Beyond importing the basic tools, I prove 12 theorems that can be used to work with the finance model. These theorems permit us to place very weak restrictions on certain model parameters in the 2-host case, making it possible to determine whether

the herd immunity condition has been violated, even if we can only identify a subset of the model's parameters.

Finally, I provide empirical evidence for the approach by using R_0 to predict the take-off of asset-backed securities and exchange-traded funds in the 1990s. I also use the tools developed in this paper to demonstrate that the introduction of a zombie bank (i.e. a bank with negative net worth) into county with a sufficiently high fraction of susceptible banks will lead to a zombie bank epidemic, where an increasing fraction of deposits are diverted to zombie institutions. I also showed that, in insolation, the introduction of a zombie bank will not lead to an increasing concentration of assets within zombie institutions; however, if zombie banks are receiving weakly positive excess returns and cross-transmission is positive, then the epidemic will span both sides of the balance sheet, concentrating both deposits and assets in the hands of zombie institutions.

Future work could expand on the approach introduced in this paper by constructing a richer and more detailed financial model; and by relaxing the assumption about the data generating process for returns. In addition to this, future empirical work could identify R_0 for many classes of assets that were not considered in this paper, yielding useful information for regulators.

5 References

Brissimis, S. & Magginas, N. 2006. “Monetary Policy Rules under Heterogeneous Inflation Expectations.” Working Papers 35, Bank of Greece.

Camerer, C. 1989. “Bubbles and Fads in Asset Prices.” *Journal of Economic Surveys* Vol. 3(1).

Capistran, C. & Timmermann, A. 2006. “Disagreement and Biases in Inflation Expectations.” *Computing in Economics and Finance 2006 3*, Society for Computational Economics.

Carroll, C. 2003. “Macroeconomic Expectations Of Households And Professional Forecasters.” *The Quarterly Journal of Economics* Vol. 118(1), 269-298.

Diekmann, O. & Heesterbeek, J.A.P. 1990. “On the Definition and the Computation of the Basic Reproduction Ratio R_0 in Models for Infectious Diseases in Heterogenous Populations,” *Journal of Mathematical Biology* Vol. 28, 365-382.

van den Driessche, P. & Watmouth, J. 2002. “Reproduction Numbers and Sub-Threshold Endemic Equilibria for Compartmental Models of Disease Transmission.” *Mathematical Biosciences* Vol. 180, 29-48.

Feng, L. & Seasholes, M.S. 2004. “Correlated Trading and Location. *Journal of Finance* Vol. 59, 2117-2144.

Frieder, L. & Subrahmanyam, A. 2005. "Brand Perceptions and the Market for Common Stock." *Journal of Financial and Quantitative Analysis* Vol. 40(1), 57-85.

Hamilton, J. 1994. "Time Series Analysis." Princeton, NJ: Princeton University Press.

Hong, H., Kubik, J.D., & Stein, J.C. 2005. "Thy Neighbors Portfolio: Word-of-Mouth Effects in the Holdings and Trades of Money Managers." *Journal of Finance* Vol. 60, 2801-2824.

Hong, D., Hong, H., & Ungureanu, A. 2010. "An Epidemiological Approach to Opinion and Price-Volume Dynamics." *AFA 2012 Chicago Meetings Paper*, Available at SSRN: <http://ssrn.com/abstract=1569418> or <http://dx.doi.org/10.2139/ssrn.1569418>.

Ivkovic, Z. & Weisbenner, S. 2007. "Information Diffusion Effects in Individual Investors Common Stock Purchase: Cover the Neighbors Investment Choices." *Review of Financial Studies* Vol. 20, 1327-1357.

Kane, E. 1978. "Dangers of Capital Forebearance: The Case of the FSLIC and 'Zombie' S&Ls." *Contemporary Economic Policy* Vol. 5(1), 77-83.

Keeling, M. Rohani, P. 2008. *Modeling Infectious Diseases in Humans and Animals*, Princeton, NJ: Princeton University Press.

Lux, T. 1995. "Herd Behaviour, Bubbles and Crashes." *The Economic Journal* Vol. 105(431), 881-896.

McCormack, R. 2007. “Multi-Host and Multi-Patch Mathematical Epidemic Models for Disease Emergence with Applications to Hantavirus in Wild Rodents.” *Texas Tech University Dissertation in Mathematics*.

Pfajfar, D. & Santoro, E. 2012. “News on Inflation and the Epidemiology of Inflation Expectations.” Discussion Paper 2012-048, Tilburg University, Center for Economic Research.

Shiller, R. 2012. “Bubbles without Markets.” *Project Syndicate*, <http://www.project-syndicate.org/commentary/bubbles-without-markets>.

Shiller, R. 1990. “Speculative Prices and Popular Models.” *The Journal of Economic Perspectives* Vol. 4(2), 55-65.

Shive, S. 2010. “An Epidemic Model of Investor Behavior.” *Journal of Financial and Quantitative Analysis* Vol. 45(1), 169-198.

Sommer, M. & Carroll, C. 2004. “Epidemiological Expectations and Consumption Dynamics.” Money Macro and Finance (MMF) Research Group Conference 2003 92, Money Macro and Finance Research Group.

Zwick, E. 2012. “Regulators vs. Zombies: Loss Overhang and Lending in a Long Slump.” Harvard FAS Working Paper.

6 Appendix

6.1 Figures and Tables

Figure 1: Epidemic Invasion




Figure 2: No Invasion




Figure 3: Two-Host SIR Model




Figure 4: ETF Invasion Prediction

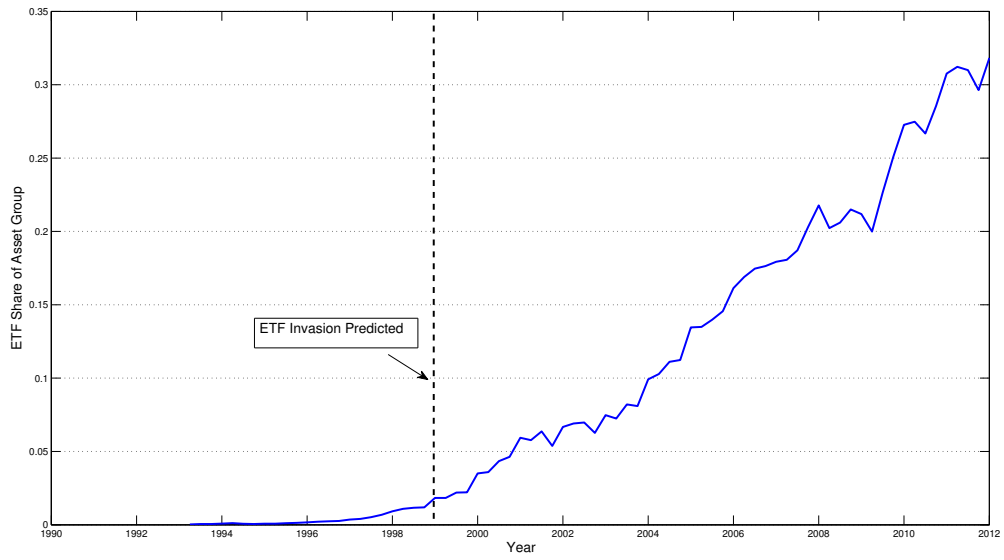


Table 1: Asset-Backed Securities Bubble Regression

	ABS
β	0.2525** (0.0278)
γ	0.2309** (0.0261)
R_0	1.0936** (0.1836)
Robust SE	X
N	60
R^2	0.25

* Note that $p < .01$ is denoted by ** and $p < .05$ is denoted by *.

Table 2: County-Level Regression Results for Zombie Banks

	Deposits	Assets
β	0.3036** (0.0684)	0.1022 (0.0709)
γ	0.2512** (0.0313)	0.2211** (0.0307)
R_0	1.2087** (0.0825)	0.4623** (0.1077)
County FE	X	X
Time FE	X	X
Robust SE	X	X
N	7131	7131
R^2	0.25	0.20

* Note that $p < .01$ is denoted by ** and $p < .05$ is denoted by *.

6.2 Proof: Theorem 1

We start with condition for herd immunity:

$$2 > \sqrt{\left(\frac{\beta_{1,1}}{\gamma_1} - \frac{\beta_{2,2}}{\gamma_2 + m_2}\right)^2} + \frac{4\beta_{1,2}\beta_{2,1}}{(\gamma_2 + m_2)\gamma_1} + \left(\frac{\beta_{1,1}}{\gamma_1} + \frac{\beta_{2,2}}{\gamma_2 + m_2}\right) \quad (54)$$

Next, we impose that $\beta_{1,1} = 0$ and $\beta_{2,2} = 0$ —that is, we have cross-transmission only.

This gives us the following, simplified condition for herd immunity:

$$2 > \sqrt{\frac{4\beta_{1,2}\beta_{2,1}}{(\gamma_2 + m_2)\gamma_1}} \quad (55)$$

$$\rightarrow 4 > \frac{4\beta_{1,2}\beta_{2,1}}{(\gamma_2 + m_2)\gamma_1} \quad (56)$$

$$\rightarrow (\gamma_2 + m_2)\gamma_1 > \beta_{1,2}\beta_{2,1} \quad (57)$$

6.3 Proof: Theorem 2

As in Theorem 1, we assume cross-transmission only. However we add to this three more assumptions: 1) $\beta_{1,2} > 0$, 2) $\beta_{2,1} > 0$, and 3) either $b^L(t) = 1$ or $b^L(t) \geq \frac{I_1 - \tau I_2}{I_1}$.

From here, we may skip to the final equation in the proof of Theorem 1:

$$\rightarrow (\gamma_2 + m_2)\gamma_1 > \beta_{1,2}\beta_{2,1} \quad (58)$$

Now, recall the definitions for γ_1 , γ_2 , and m_2 :

$$\gamma_1 = \max \left\{ \frac{\frac{I_1 - \tau I_2}{I_1} - b^L(t)}{b^H(t) - b^L(t)}, 0 \right\} \quad (59)$$

$$\gamma_2 = \max \left\{ \frac{\frac{\tau I_2}{I_1}}{b^H(t) - b^L(t)}, 0 \right\} \quad (60)$$

$$m_2 = \max \left\{ \frac{\frac{I_1 - \tau I_2}{I_1} - b^L(t)}{b^H(t) - b^L(t)}, 0 \right\} \quad (61)$$

After substituting in these definitions and performing some algebraic manipulations, we get the new condition for herd immunity:

$$\frac{(1 - b^L(t))\left(\frac{I_1 - \tau I_2}{I_1} - b^L(t)\right)}{(b^H(t) - b^L(t))^2} > \beta_{1,2}\beta_{2,1} \quad (62)$$

Since $\beta_{1,2} > 0$, $\beta_{2,1} > 0$, and $(b^H(t) - b^L(t))^2 \geq 0$, we know that herd immunity will be violated whenever:

$$(1 - b^L(t)) \left(\frac{I_1 - \tau I_2}{I_1} - b^L(t) \right) \leq 0 \quad (63)$$

Thus, if $\frac{I_1 - \tau I_2}{I_2} < b^L(t)$ and $b^L(t) < 1$, then herd immunity will not hold. Alternatively, if $b^L(t) = 1$, then $(1 - b^L(t)) = 0$, so herd immunity will not hold.

6.4 Proof: Theorem 3

Again, we assume that cross-transmission is the only channel that causes investors and intermediaries to switch to the unconventional asset. This allows us to start with the simplified condition for herd immunity:

$$(\gamma_2 + m_2)\gamma_1 > \beta_{1,2}\beta_{2,1} \quad (64)$$

Now, we assume that $\beta_{1,2} > 1$ and $\beta_{2,1} \geq 1$ or $\beta_{2,1} > 1$ and $\beta_{1,2} \geq 1$ —that is, the right hand side of the equation is greater than 1. Since γ_1 is a probability, it will be no greater than 1. Thus, if $\gamma_2 + m_2 \leq 1$, then herd immunity will not hold. Note that we may rewrite this as follows:

$$\frac{1 - b^L(t)}{b^H(t) - b^L(t)} \leq 1 \quad (65)$$

$$\rightarrow 1 \leq b^H(t) \quad (66)$$

6.5 Proof: Theorem 4

We start by setting $\beta_{2,2} = 0$. This gives us the following condition for herd immunity:

$$2 > \sqrt{\left(\frac{\beta_{1,1}}{\gamma_1}\right)^2 + \frac{4\beta_{1,2}\beta_{2,1}}{(\gamma_2 + m_2)\gamma_1}} + \frac{\beta_{1,1}}{\gamma_1} \quad (67)$$

$$\rightarrow \frac{(2\gamma_1 - \beta_{1,1})^2}{\gamma_1} > \beta_{1,1}^2 + \frac{4\beta_{1,2}\beta_{2,1}}{(\gamma_2 + m_2)} \quad (68)$$

$$(\gamma_1 - \beta_{1,1})(\gamma_2 + m_2) > (\gamma_1 - 1)\frac{\beta_{1,1}^2}{4\gamma_1}(\gamma_2 + m_2) + \beta_{1,2}\beta_{2,1} \quad (69)$$

Note that γ_1 , γ_2 , and m_2 are probabilities. Thus,

$$(\gamma_1 - 1)\frac{\beta_{1,1}^2}{4\gamma_1}(\gamma_2 + m_2) \leq 0 \quad (70)$$

This means that herd immunity will hold if the following condition is true:

$$(\gamma_1 - \beta_{1,1})(\gamma_2 + m_2) > \beta_{1,2}\beta_{2,1} \quad (71)$$

Note that the above condition is stronger than the one in equation 57 and does not allow us to avoid estimating any parameters; however, it will be useful for proving the remaining theorems.

6.6 Proof: Theorem 5

We start from equation (69):

$$(\gamma_1 - \beta_{1,1})(\gamma_2 + m_2) > (\gamma_1 - 1)\frac{\beta_{1,1}^2}{4\gamma_1}(\gamma_2 + m_2) + \beta_{1,2}\beta_{2,1} \quad (72)$$

Since $\beta_{1,1} > \gamma_1$, the left hand side will be negative. This suggests that herd immunity will be violated if the right hand side is positive:

$$(\gamma_1 - 1) \frac{\beta_{1,1}^2}{4\gamma_1} (\gamma_2 + m_2) + \beta_{1,2}\beta_{2,1} > 0 \quad (73)$$

Re-arranging the equation yields the proof:

$$(\gamma_1 - 1) \frac{\beta_{1,1}^2}{4\gamma_1} (\gamma_2 + m_2) < \beta_{1,2}\beta_{2,1} \quad (74)$$

6.7 Proof: Theorem 6

Since γ_1 is a probability, $1 \geq \gamma_1$. Using this, the proof for Theorem 6 is identical to the proof for Theorem 5.

6.8 Proof: Theorem 7

Recall the definitions for γ_1 , γ_2 , and m_2 from the proof for Theorem 2. If $b^L(t) \geq 1$, then $\gamma_1 = 0$, $\gamma_2 = 0$, and $m_2 = 0$. Thus, the condition for herd immunity will be violated trivially.

6.9 Proof: Theorem 8

Theorem 8 states that if $\beta_{1,1} = 0$ and $(\gamma_2 + m_2 - \beta_{2,2})\gamma_1 > \beta_{1,2}\beta_{2,1}$, then herd immunity holds. To prove this, we start with the condition for herd immunity, which we get from imposing $\beta_{1,1} = 0$:

$$2 > \sqrt{\frac{\beta_{2,2}^2}{(\gamma_2 + m_2)^2} + \frac{4\beta_{2,1}\beta_{1,2}}{(\gamma_2 + m_2)\gamma_1}} + \frac{\beta_{2,2}}{\gamma_2 + m_2} \quad (75)$$

$$\rightarrow \left(2 - \frac{\beta_{2,2}}{\gamma_2 + m_2}\right)^2 > \frac{\beta_{2,2}^2}{(\gamma_2 + m_2)^2} + \frac{4\beta_{2,1}\beta_{1,2}}{(\gamma_2 + m_2)\gamma_1} \quad (76)$$

$$\rightarrow 4 - \frac{4\beta_{2,2}}{\gamma_2 + m_2} + \frac{\beta_{2,2}^2}{(\gamma_2 + m_2)^2} > \frac{\beta_{2,2}^2}{(\gamma_2 + m_2)^2} + \frac{4\beta_{2,1}\beta_{1,2}}{(\gamma_2 + m_2)\gamma_1} \quad (77)$$

$$\rightarrow (\gamma_2 + m_2 - \beta_{2,2})\gamma_1 > \beta_{1,2}\beta_{2,1} \quad (78)$$

6.10 Proof: Theorem 9

Theorem 9 claims that herd immunity will be violated if $\beta_{1,2} > 0$, $\beta_{2,1} > 0$, $\beta_{1,1} = 0$ and $\beta_{2,2} > (\gamma_2 + m_2)$. From equation (64), we can see that the left hand side will be negative, violating the condition for herd immunity.

6.11 Proof: Theorem 10

Theorem 10 states that if $\beta_{1,2} > 0$, $\beta_{2,1} > 0$, $\beta_{1,1} = 0$, and $\beta^L(t) \geq 1$, then herd immunity is violated. To begin the proof, we start with the simplified condition for herd immunity:

$$\rightarrow (\gamma_2 + m_2 - \beta_{2,2})\gamma_1 > \beta_{1,2}\beta_{2,1} \quad (79)$$

Next, we use the definitions for γ_2 and m_2 to re-write the condition:

$$\rightarrow \left(\max \left\{ \frac{1 - b^L(t)}{b^H(t) - b^L(t)}, 0 \right\} - \beta_{2,2}\right) \gamma_1 > \beta_{1,2}\beta_{2,1} \quad (80)$$

Since $b^L(t) \geq 1$, we may re-write this again:

$$\rightarrow -\beta_{2,2}\gamma_1 > \beta_{1,2}\beta_{2,1} \quad (81)$$

Since $\beta_{2,2}$ and γ_1 are nonnegative and $\beta_{1,2}$ and $\beta_{2,1}$ are positive, it must be the case that:

$$\beta_{1,2}\beta_{2,1} > -\beta_{2,1}\gamma_1 \quad (82)$$

Thus, herd immunity is violated.

6.12 Proof: Theorem 11

Theorem 11 states that if $\beta_{1,2}\beta_{2,1} > 2$ and $\beta_{1,1} = 0$, then herd immunity is violated. Again, recall the herd immunity condition from equation (64):

$$\rightarrow (\gamma_2 + m_2 - \beta_{2,2})\gamma_1 > \beta_{1,2}\beta_{2,1} \quad (83)$$

Now, if $\beta_{1,2}\beta_{2,1} > 2$, then herd immunity will be violated if $(\gamma_2 + m_2 - \beta_{2,2})\gamma_1 \leq 2$. Additionally, note that γ_1 , γ_2 , and m_2 are probabilities, which means that the left hand side will be no greater than $(2 - \beta_{2,2})$. This is maximized when $\beta_{2,2} = 0$; and herd immunity is violated here. Thus, herd immunity will always be violated under the stated conditions.

6.13 Proof: Theorem 12

Theorem 12 states that if $\beta_{2,2} > 2$ and $\beta_{1,1} = 0$, then herd immunity is violated. Again, using equation (71), if $\beta_{2,2} > 2$, then $(\gamma_2 + m_2 - \beta_{2,2})\gamma_1 \leq 0$, which means that:

$$\rightarrow (\gamma_2 + m_2 - \beta_{2,2})\gamma_1 \leq \beta_{1,2}\beta_{2,1} \quad (84)$$

Thus, the condition for herd immunity is violated.