



Income Risk and Aggregate Demand over the Business Cycle

Citation

Mericle, David. 2012. Income Risk and Aggregate Demand over the Business Cycle. Doctoral dissertation, Harvard University.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:9282602>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

©2012 — *David Alexander Mericle*

All rights reserved.

Income Risk and Aggregate Demand over the Business Cycle

ABSTRACT

This dissertation consists of three essays on income risk and aggregate demand over the business cycle, each addressing an aspect of the Great Recession.

The first chapter reframes the standard liquidity trap model to illustrate the costly feedback loop between idiosyncratic risk and aggregate demand. I first show that a liquidity trap can result from excess demand for precautionary savings in times of high uncertainty. Second, I show that the output and welfare costs of the ensuing recession depend crucially on how the drop in demand for output is translated into a reduction in demand for labor. Increased unemployment risk compounds the original rise in idiosyncratic productivity risk and reinforces precautionary motives, deepening the recession. Third, I show that increasing social insurance can raise output and welfare at the zero bound. I decompose these effects to distinguish the component unique to the liquidity trap environment and show that social insurance is most effective at the zero bound when it targets the type of idiosyncratic risk households face, which in turn depends on the labor market adjustment mechanism.

The second paper offers a novel model of the connection between the consumer credit and home mortgage markets through an individual's credit history. This paper introduces a novel justification for the home mortgage interest deduction. In an economy with both housing assets and consumer credit, the mortgage interest deduction is modeled as a subsidy for the accumulation of collateralizable assets by households who have maintained good credit. As such, the subsidy loosens participation constraints and facilitates risk-sharing. Empirical evidence and a calibration exercise reveal that the subsidy has a sizable impact on the availability of credit.

The third paper assesses the role of policy uncertainty in the Great Recession. The Great Recession features substantial geographic variation in employment losses, a fact that is often presented as a challenge to uncertainty-based models of the downturn. In this paper we show that there is a substantial correlation between the distribution of employment losses and the increases in local measures of both economic and policy uncertainty. This relationship is robust across a wide range of measures.

Contents

Abstract	iii
Acknowledgments	vi
1. Uncertainty, the Liquidity Trap and Social Insurance	1
1.1 Introduction	1
1.2 Model	3
1.2.1 The basic model with time-varying risk	4
1.2.2 The zero lower bound	9
1.3 Calibration	13
1.4 Results	15
1.4.1 Flexible case	15
1.4.2 The zero lower bound	18
1.4.3 Policy experiments	20
1.4.4 Discussion	25
1.5 Numerical solution method	27
1.5.1 Stationary states	27
1.5.2 Transition dynamics	30
1.6 Conclusion	32
2. The Credit Market Effects of the Home Mortgage Interest Deduction	34
2.1 Introduction	34
2.2 Simple model	36
2.2.1 Basic model	36
2.2.2 Redistribution model	37
2.3 Assumptions check	40
2.3.1 The HMID is large in size	40
2.3.2 Credit access is restricted in bankruptcy	40
2.4 Model	43
2.4.1 Model 1: Endogenous borrowing constraint	46
2.4.2 Model 2: Variable interest rate	46
2.4.3 Equilibrium	46
2.5 Numerical solution	47
2.5.1 Calibration	48
2.6 Results	49
2.6.1 Baseline results	49
2.6.2 Tax incidence by age and income	51
2.6.3 Welfare	53
2.7 Suggestive empirics	55
2.7.1 Credit limits and home prices	55
2.7.2 Bankruptcy and taxes: time and cross section	57
2.8 Conclusion	59

3. Uncertainty and the Geography of the Great Recession	61
3.1 Introduction	61
3.2 Uncertainty	62
3.2.1 Media	62
3.2.2 Internet searches	63
3.2.3 State budgets	64
3.2.4 Misleading indicators	68
3.2.5 State credit ratings	69
3.2.6 Uncertainty summarized	70
3.2.7 Horse race?	73
3.3 Conclusions	74
References	75

ACKNOWLEDGMENTS

I am grateful to my thesis committee, Professors N. Gregory Mankiw, Emmanuel Farhi and David Laibson, and to all of the participants in the Harvard macroeconomics seminar.

I thank Harvard University and the National Science Foundation for financial support.

1 Uncertainty, the Liquidity Trap and Social Insurance

1.1 Introduction

The global recession and the fall of interest rates in major economies to nearly zero poses a number of questions about the liquidity trap. This paper develops a theory of why an economy falls into a liquidity trap, shows that the costly feedback loop between idiosyncratic risk and aggregate demand at the zero lower bound can deepen the ensuing recession, and quantitatively assesses the effectiveness of monetary and fiscal policy interventions, in particular increases in social insurance.

Beginning with Krugman (1998), the liquidity trap literature is usually motivated with an ad-hoc demand shock. This paper joins recent work by Eggertsson and Krugman (2011) and Guerrieri and Lorenzoni (2011) in offering explicit microfoundations for the liquidity trap. While those papers propose a sudden tightening of borrowing constraints caused by a financial sector shock to push the economy to the zero lower bound, I instead emphasize the role of cyclical shocks to the level of idiosyncratic labor income risk. In particular, I microfound the liquidity trap as resulting from a rise in demand for precautionary savings in times of increased uncertainty about future income.

I develop a heterogeneous agent model in the tradition of Bewley (1977)-Huggett (1993)-Aiyagari (1994) with three additional elements. First, following Bloom, Floetotto, and Jaimovich (2010), the model features regime-switching between aggregate states of low and high uncertainty, modeled as the size of the conditional variance of the idiosyncratic labor income process. Second, the model introduces a government with access to two instruments of social insurance: unemployment insurance and progressive taxation. Third, I introduce price rigidities and a zero lower bound. I calibrate the model to roughly match the United States economy just prior to the crisis and use Storesletten, Telmer, and Yaron (2004)'s estimates of cyclical changes in idiosyncratic risk. In the high-uncertainty state, heightened demand for precautionary savings causes the market-clearing interest rate to fall below the zero bound, a result that survives several modifications of the baseline calibration.³ This finding that a rise in idiosyncratic labor income risk provides a robust and quantitatively plausible microfoundation for the liquidity trap is the first major contribution of the paper.

The model also shows how an economy endogenously exits from a liquidity trap. As in Guerrieri and Lorenzoni (2011), the interest rate overshoots because the wealth distribution is initially at a low-risk level of dispersion in the aftermath of the shock. As wealth-poor households build up their bond holdings in response to the increased level of risk, the market-clearing interest rate rises. For some calibrations, the interest rate eventually rises above zero so that exit from the liquidity trap is endogenous even within

³Gourinchas and Parker (2001) show that such increases in individual income uncertainty indeed lead to higher savings rates.

the high-uncertainty aggregate state. Capturing this mechanism by which the household distribution evolves to endogenously end the liquidity trap is a virtue of a framework featuring heterogeneity over representative agent liquidity trap models, which usually assume that exit is exogenous.

How does output adjust during the liquidity trap, and how deep is the recession in a demand-constrained economy? I compare two mechanisms by which the reduction in demand for output is translated into a reduction in demand for labor to illustrate the importance of the interaction between idiosyncratic risk and aggregate demand. The standard practice in the literature introduces an endogenous labor wedge at the zero lower bound. Firms bid down wages, and households work less in response. The limitation of this approach is evident in a heterogeneous agent model that emphasizes idiosyncratic risk: when the labor market clears at the zero bound through a labor wedge, the costs of the recession are shared proportionally by all employed households. The implication that output falls without creating any additional idiosyncratic risk neglects an important aspect of recessions.

I therefore consider a simple alternative inspired by Hall (2011). If wages are rigid, firms adapt to the reduction in demand at the zero bound by reducing employment. I model this possibility by finding the unemployment entry and exit probabilities (and implied unemployment levels), respectively, that achieve market clearing. Adjusting the entry probability corresponds to an increased likelihood of unemployment spells of normal duration, while reducing the exit probability corresponds to an increased expected duration of infrequent unemployment spells. In contrast to the labor wedge approach, adjustment through unemployment magnifies idiosyncratic income risk. This increase in risk reinforces the original sharpening of precautionary motives, creating a feedback loop between aggregate demand and idiosyncratic risk at the zero bound that is unique to the model.

The quantitative results of the model suggest that both the output and welfare costs of the liquidity trap are considerably larger when the labor market adjusts in a way that amplifies the rise in productivity risk, namely, by increasing employment risk. Adjustment through increased entry into unemployment proves worse than adjustment through reduced wages, and an increase in unemployment achieved through a reduction in the unemployment exit rate is even worse. The intuition for this result is that a decrease in labor supply achieved through longer unemployment spells is the most concentrated form of idiosyncratic risk and induces the strongest precautionary response and therefore the largest drop in output in a demand-constrained economy. This mirrors a result in Challe and Ragot (2010), who show that the precautionary savings response to the business cycle is especially strong when shocks lead to persistent unemployment. This finding – that the liquidity trap is more costly when it involves an amplification of idiosyncratic risk, especially concentrated risk - is the paper's second major contribution.

Finally, I consider a set of policy experiments to quantitatively assess the effectiveness of monetary

interventions and expansion of social insurance. The policy results suggest several general lessons that form the paper's third major contribution. First, policies that greatly expand the aggregate supply of bonds or loosen borrowing constraints are very effective in mitigating the output drop at the zero bound because they increase households' ability to smooth consumption. Second, while social insurance can also reduce the fall in output, social insurance policies are most effective when they successfully target the type of idiosyncratic risk generated by the liquidity trap. For example, extending the duration of unemployment benefits is very effective when the reduction in labor demand is manifested in lower unemployment exit rates, and quite ineffective otherwise. Third, a decomposition of the welfare effects of each policy into a component that exists in the fully flexible case – the standard distortionary effect – and a component unique to the liquidity trap shows that some policies that would be undesirable in normal times become output- or welfare-enhancing at the zero bound. More generally, these numerical experiments demonstrate the point, also made by Mertens and Ravn (2011), that the nature of the shock that generates the liquidity trap has implications for what policies are effective in softening the ensuing recession.

The broad goal of the paper is to ask how reframing the liquidity trap with an emphasis on the relationship between idiosyncratic risk and aggregate demand in times of high uncertainty affects both estimates of its output costs and the effectiveness of possible policy responses. The first section describes the model. The second section discusses the calibration. The third section shows the results of a high-uncertainty shock first for a fully flexible economy and then for a zero lower bound economy, and then discusses the quantitative effects on output and welfare of several policy interventions. The last section concludes, and an appendix discusses the numerical methods used to solve the model.

1.2 Model

I first present a simple model of idiosyncratic risk that features regime-switching in the conditional variance of labor productivity. A government sector is then introduced that provides social insurance through unemployment insurance and redistributive taxation. The emphasis on incomplete markets and precautionary savings builds on the class of models in the tradition of Bewley (1977), Huggett (1993), and Aiyagari (1994). However, as in Bloom, Floetotto, and Jaimovich (2010), risk is time-varying. The model provides microfoundations for the liquidity trap, which is reached in times of high uncertainty about future idiosyncratic labor productivity because of a rise in households' desire for precautionary savings. The possibility that the market-clearing gross interest rate is less than one is not extraordinary in this class of models.

1.2.1 The basic model with time-varying risk

This section describes a simple Bewley (1977) model with two added features, regime-switching in the conditional variance of the income process ("uncertainty shocks") and a government with two tax instruments. For now, prices are flexible and the interest rate is unconstrained.

Uncertainty The economy fluctuates stochastically between a state of low uncertainty, $v_t = 0$, and high uncertainty state, $v_t = 1$, according to the transition probabilities

$$\begin{aligned} P(v_{t+1} = 1|v_t = 1) &= v^{HH} \\ P(v_{t+1} = 0|v_t = 0) &= v^{LL}. \end{aligned} \tag{1}$$

Uncertainty here refers to the variance of the innovations to the idiosyncratic household productivity process. Household i 's productivity $\theta_{i,t}$, measured as efficiency units of labor, has a persistent component $z_{i,t}$ and transitory component $\varepsilon_{i,t}$. The log of productivity evolves as

$$\log \theta_{i,t} = z_{i,t} + \varepsilon_{i,t} \tag{2}$$

$$z_{i,t} = \rho z_{i,t-1} + \eta_{i,t}, \tag{3}$$

where the innovations exhibit a regime-switching conditional variance dependent on the uncertainty state,

$$\varepsilon_{i,t} \sim N\left(0, (1 - v_t) \sigma_{\varepsilon_{low}}^2 + v_t \sigma_{\varepsilon_{high}}^2\right) \tag{4}$$

$$\eta_{i,t} \sim N\left(0, (1 - v_t) \sigma_{\eta_{low}}^2 + v_t \sigma_{\eta_{high}}^2\right) \tag{5}$$

This decomposition of the earnings process into a transitory and permanent component is standard in the literature and in this particular formulation follows Storesletten, Telmer, and Yaron (2004), whose parameter estimates of idiosyncratic risk over the business cycle are used for calibration.

Aggregate productivity risk is ignored here to reduce computational demands. This can be justified by the minimal quantitative role of the aggregate relative to the idiosyncratic component in total household income risk.⁴

⁴In related work, Basu and Bundick (2011) also develop a model in which strengthened precautionary motives induce a liquidity trap. Their emphasis, however, is on uncertainty about aggregate productivity, calibrated to match stock market volatility, instead of idiosyncratic labor income risk.

Households The economy consists of a unit continuum of households i . Households have period utility

$$u(c_{i,t}, n_{i,t}) = \left(\frac{c_{i,t}^{1-\gamma^c}}{1-\gamma^c} + N \frac{(1-n_{i,t})^{1-\gamma^n}}{1-\gamma^n} \right), \quad (6)$$

where $c_{i,t}$ is consumption and $n_{i,t}$ is labor supply. Households have 1 unit of time each period, which they divide between labor and leisure.

Each household has idiosyncratic productivity $\theta_{i,t}$, defined as efficiency units of labor, which evolves as described in (2). Households also fall into unemployment, $\varepsilon_{i,t} = -\infty$, with likelihood,

$$\begin{aligned} P(\varepsilon_{i,t+1} = -\infty | \varepsilon_{i,t} \neq -\infty) &= \varphi^{in} \\ P(\varepsilon_{i,t+1} \neq -\infty | \varepsilon_{i,t} = -\infty) &= \varphi^{out}. \end{aligned} \quad (7)$$

These stochastic unemployment shocks are modeled as part of the transitory income process to limit the state space, and transitory risk for employed households is determined by (4). Unemployed households have leisure time equal to $1 - n^U$, where n^U parameterizes the amount of time spent searching for work while unemployed.

Employed households choose labor supply $n_{i,t}$. Output is a linear function of efficiency units of labor, so that pre-tax earnings are equal to $\theta_{i,t}n_{i,t}$. As in Heathcoate, Storesletten, Violante (2010), employed households pay a distortionary labor income tax,

$$T_{i,t} = \theta_{i,t}n_{i,t} - \lambda (\theta_{i,t}n_{i,t})^{1-\tau}, \quad (8)$$

where τ parameterizes the progressivity of taxation and λ the level. Total after-tax income is then $\lambda (\theta_{i,t}n_{i,t})^{1-\tau}$.

Unemployed households receive unemployment insurance equal to a replacement rate u on their income, subject to a maximum \bar{u} . Because it is infeasible to retain the history of each household's labor supply choices, I set the income level used to calculate unemployment benefits as the labor income that would be earned by a household in the same uncertainty state v_t with identical values of persistent income $z_{i,t}$ and savings $b_{i,t}$, but with transitory shock $\varepsilon_{i,t} = 0$. This "counterfactual" labor supply serves as an approximation of last period's labor supply. Unemployment insurance is then equal to

$$\kappa_{i,t}(v_t, z_{i,t}, \varepsilon_{i,t}, b_{i,t}) = \min(u e^{z_{i,t}} n(v_t, z_{i,t}, \varepsilon_{i,t} = 0, b_{i,t}), \bar{u}). \quad (9)$$

Unemployment insurance is of limited duration. Each period, an unemployed household has a stochastic probability ζ of losing its unemployment insurance. Unemployed households no longer receiving unemployment insurance receive a minimal transfer payment that is sufficient to prevent default, but cannot receive UI again during the current unemployment spell.

Combining earnings and unemployment insurance, total household income is

$$y_{i,t} = \lambda (\theta_{i,t} n_{i,t})^{1-\tau} + \kappa_{i,t}. \quad (10)$$

Asset markets are incomplete and there is no capital in the model, but households can smooth consumption using one-period bonds $b_{i,t}$ available in fixed net supply \bar{B} . The household budget constraint is then

$$c_{i,t} + \frac{b_{i,t+1}}{r_t} = \lambda (\theta_{i,t} n_{i,t})^{1-\tau} + \kappa_{i,t} + b_{i,t}, \quad (11)$$

where $r_{i,t}$ is the gross interest rate from period t to $t + 1$.

For simplicity, households are assumed to face a fixed exogenous borrowing constraint,

$$b_{i,t} \geq \underline{b} \forall i, t, \quad (12)$$

that applies to both the low and the high uncertainty states. Clearly, a change in the level of idiosyncratic income risk has strong implications for credit availability. I address the effects of time-varying income risk in an Eaton and Gersovitz (1981) framework with defaultable debt in a related paper (Mericle 2012), but keep the borrowing constraint fixed here to isolate the effects of income risk from those of credit shocks, discussed by Eggertsson and Krugman (2011) and Guerrieri and Lorenzoni (2011).

The household value function can now be written

$$\mathbf{V}(v_t, z_{i,t}, \varepsilon_{i,t}, b_{i,t}) = \max_{c_{i,t}, n_{i,t}, b_{i,t+1}} \frac{c_{i,t}^{1-\gamma^c}}{1-\gamma^c} + N \frac{(1-n_{i,t})^{1-\gamma^n}}{1-\gamma^n} + \beta E_t \mathbf{V}(v_{t+1}, z_{i,t+1}, \varepsilon_{i,t+1}, b_{i,t+1}), \quad (13)$$

subject to (11) and (12). The Euler equation is then

$$c_{i,t}^{-\gamma^c} \geq \beta r_t E_t (c_{i,t+1}^{-\gamma^c}), \quad (14)$$

with inequality only when (12) is binding. In periods when they are employed, households make the optimal labor-leisure trade-off

$$(1-\tau) \lambda \theta_{i,t}^{1-\tau} n_{i,t}^{-\tau} c_{i,t}^{-\gamma^c} = N (1-n)^{-\gamma^n}. \quad (15)$$

The household labor supply function $n(v_t, z_{i,t}, \varepsilon_{i,t}, b_{i,t})$ and consumption function $c(v_t, z_{i,t}, \varepsilon_{i,t}, b_{i,t})$ then yield a policy function for tomorrow's bond holdings $b_{i,t+1}$, $b(v_t, z_{i,t}, \varepsilon_{i,t}, b_{i,t})$.

Government The government has three fiscal instruments, the replacement rate on unemployment insurance u , the expected duration of benefits $1/\zeta$, and the progressivity of taxation τ . These two instruments can be thought of as principally targeting the level of transitory and persistent income risk, respectively.

Government revenues consist solely of labor income taxes. Government spending includes three components: exogenous government consumption G , unemployment insurance payments, and interest payments on aggregate outstanding debt \bar{B} . Government consumption does not enter into the utility function, but is included to induce realism in the marginal effect of distortionary taxation by roughly matching the level of the tax burden. The government's budget is balanced in every period,

$$\int y_{i,t} - \lambda y_{i,t}^{1-\tau} di = G + \int \kappa_{i,t} di + \left(1 - \frac{1}{r_t}\right) \bar{B}, \quad (16)$$

by setting the parameter λ to equate revenues and spending. I consider policy experiments later in which the parameters τ and u are allowed to vary with v_t , but the budget must be balanced in each uncertainty state.

The restriction of policy choices to balanced-budget outcomes is deliberate. In reality, increased unemployment insurance and transfer payments are often deficit financed in recessions. The primary reason for this modeling choice is again to isolate the effect of the social insurance policy by distinguishing it from the effect of government spending at the zero bound, on which there is a large literature.⁵ In addition, the literature on monetary policy in a liquidity trap centers on the commitment problem familiar since Krugman (1998). The issue of commitment looms equally large in considering fiscal responses, such as credibly-temporary stimulus spending. In our case, the emphasis on social insurance highlights the fiscal role in reducing future idiosyncratic risk, and a policy rule specifying indefinite debt-financed social insurance might create concern about sovereign debt risk, which is beyond the scope of this model.⁶

Equilibrium with regime-switching I now define what is meant here by an equilibrium in the context of a regime-switching model. The model never reaches a stationary state in the usual sense because of the regime-switching feature. That is, because the uncertainty state is an aggregate variable that applies to all households in any given period, it is not meaningful to think of a stationary distribution of households across uncertainty states. Instead, I define the low- and high-uncertainty stationary states corresponding to

⁵See, for example, Eggertsson (2006, 2009), Christiano, Eichenbaum, and Rebelo (2009), Cogan, Cwik, Taylor, and Wieland (2010), or the White House's official estimates in Romer and Bernstein (2009).

⁶Related to this, Baker, Bloom, and Davis (2011) studies the effect of recent policy uncertainty on aggregate outcomes.

$v = 0$ and $v = 1$ as the household joint distributions $\mathcal{H}_0(z, \varepsilon, b)$ and $\mathcal{H}_1(z, \varepsilon, b)$ reached when the economy has been in the low- (high-) uncertainty state for a sufficiently long time that the distribution is no longer changing.⁷ An equilibrium, then, is a pair of low- and high-uncertainty stationary states. This implies that while the household policy rules reflect the possibility of transitioning between uncertainty states, they are only approximations in the sense that they ignore the transition dynamics between stationary states. In other words, households act as if a switch in regime leads immediately to the other stationary state. A more thorough discussion of the solution method is provided in the appendix.

More precisely, an equilibrium consists of a pair of interest rates (r_0, r_1) , household policy functions $n(v, z_i, \varepsilon_i, b_i)$, $c(v, z_i, \varepsilon_i, b_i)$, and $b(v, z_i, \varepsilon_i, b_i)$, a set of government policy parameters $(\lambda_0, \lambda_1, \tau_0, \tau_1, u_0, u_1, \zeta_0, \zeta_1)$, and joint distributions of household characteristics $\mathcal{H}_0(z, \varepsilon, b)$ and $\mathcal{H}_1(z, \varepsilon, b)$ such that the following conditions hold.

1. *Household optimization.* Household decision rules $n(v, z_i, \varepsilon_i, b_i)$, $c(v, z_i, \varepsilon_i, b_i)$, and $b(v, z_i, \varepsilon_i, b_i)$ maximize (13) subject to (11) and (12) taking prices and government policies as given.
2. *Bond market clearing.* Total demand for bonds is equal to the fixed aggregate supply, $\int b(v, z_i, \varepsilon_i, b_i) di = \bar{B}$, for both $v_t = 0$ and $v_t = 1$.
3. *Balanced government budget.* The government policy set $(\lambda_0, \lambda_1, \tau_0, \tau_1, u_0, u_1, \zeta_0, \zeta_1)$ satisfies the balanced budget constraint (16) for both $\mathcal{H}_0(z, \varepsilon, b)$ and $\mathcal{H}_1(z, \varepsilon, b)$.
4. *Conditional stationarity of the joint distributions of household characteristics.* Define Γ as the law of motion of the joint distribution of household characteristics $\mathcal{H}(z, \varepsilon, b)$ that results from the exogenous processes (1), (3), and (4), and the household policy rules $n(v, z_i, \varepsilon_i, b_i)$, $c(v, z_i, \varepsilon_i, b_i)$, and $b(v, z_i, \varepsilon_i, b_i)$. Then the stationary state household distributions must satisfy

$$\begin{aligned}\mathcal{H}_0(z, \varepsilon, b) &= \Gamma(v = 0, \mathcal{H}_0(z, \varepsilon, b)) \\ \mathcal{H}_1(z, \varepsilon, b) &= \Gamma(v = 1, \mathcal{H}_1(z, \varepsilon, b)).\end{aligned}\tag{17}$$

That is, \mathcal{H}_0 is the low-uncertainty state stationary joint distribution if a population starting at \mathcal{H}_0 at time t remains unchanged after experiencing shock $v_{t+1} = 0$, and likewise for \mathcal{H}_1 .

This equilibrium concept of two mutually consistent stationary states defined by the limiting distributions has the advantage that the possibility of the shock is known. The policy rules in the two states differ to the extent that households discount the future and that the probability of transitioning to the other

⁷In a similar framework, Mehrling (1998) also uses the limiting distribution to define the distribution corresponding to a particular aggregate state.

state is small. Admittedly, considering the stationary distributions might seem odd in that the economy is unlikely to reach either limiting distribution, and abstracting from the transition between states in the equilibrium rules is a further simplification. However, this method is perhaps more plausible than the usual alternative framework in which an equilibrium economy is hit by an entirely unforeseen shock, and a new permanent state of the world follows. An alternative solution method could follow Krusell and Smith (1998), but the complexity of the model suggests that the forecasting rules would require more terms than usual, including interactions with the zero bound, and the need to impose market clearing in every period of that algorithm's simulation step using Young (2010)'s method further suggests that this would be computationally infeasible.

1.2.2 The zero lower bound

I now consider the model with a zero lower bound. To motivate the liquidity trap, I must also add nominal rigidities to the model to prevent the price level from immediately adjusting to deliver the market-clearing real interest rate. I follow Krugman (1998) and Guerrieri and Lorenzoni (2011) in assuming a fixed price level, $P_t = P_{t-1}$, so that the real and nominal rates are the same in what follows. These simplifications are necessary to make the model computationally tractable; without fixing the price level in a model with heterogeneous agents, we would have to search for multiple prices to satisfy multiple market-clearing conditions simultaneously.⁸

Firms Firms are added to the model to motivate price level rigidity. As in the standard new Keynesian setup, production consists of perfectly competitive final goods firms and a unit continuum of monopolistically competitive intermediate goods firms indexed by j whose profits are rebated lump sum to households. Final goods firms combine intermediates $Y_t(j)$ into a final consumption good Y_t using a Dixit-Stiglitz technology, $Y_t \equiv \left(\int_0^1 Y_t(j)^{\frac{\phi-1}{\phi}} dj \right)^{\frac{\phi}{\phi-1}}$, $\phi > 1$. Profit-maximization by final goods producers results in demand for intermediate j :

$$Y_t(j) = \left(\frac{p_t(j)}{P_t} \right)^{-\phi} Y_t, \quad (18)$$

where $p_t(j)$ is the price of intermediate j and P_t is the price of the final good, $P_t \equiv \left(\int_0^1 p_t(j)^{1-\phi} dj \right)^{\frac{1}{1-\phi}}$. Intermediate goods firms use the same linear production function as before, $Y_t(j) = \theta_t(j) n_t(j)$, where $\theta_t(j)$ is interpreted as a weighted average of the individual productivities of firm j 's workers, and pay

⁸Some justification for imposing the assumption of a flat price level in the context of a zero bound model comes from Meier (2010), who finds that disinflation during a diverse group of "persistent large output gaps" tends to fall off to low positive rates of inflation. The Bureau of Labor Statistics core inflation series (Consumer Price Index for All Urban Consumers: All Items Less Food & Energy) shows a similarly low but non-negative rate for the US during the current recession.

households a wage w_t per unit of productivity-augmented labor. Profit maximization,

$$\pi_t(j) = \max_{p_t(j)} p_t(j) Y_t(j) - w_t Y_t(j), \quad (19)$$

leads firms to set price $p_t(j) = \frac{\phi}{\phi-1} w_t$, so that under the fixed price assumption, aggregate profits as a share of total output equals $\frac{1}{\phi} Y_t$. Following the standard assumption, firm profits are rebated lump sum to households. However, evenly distributing firm profits to all households is unrealistic, and this rebate – effectively a linear income tax redistributed as a lump-sum payment – interferes quantitatively with the calibration of household income risk. I therefore assume that the profit share of output is negligible. This can be justified by assuming either that intermediates are nearly perfect substitutes so that price $p_t(j)$ is negligibly greater than w_t , that the government taxes nearly 100% of firm profits and rebates them as a linear earned income tax credit prior to applying the labor income tax described earlier, or that firms are workers’ cooperatives that likewise distribute profits in linear proportion to wage income.

Adjustment at the zero lower bound I now consider the mechanism by which output adjusts at the zero lower bound and the labor market implications for households. I consider both the standard and an alternative approach, and consider the implications of each for the magnitude of the output drop. A richer discussion of the labor market consequences of a drop in output at the zero bound is provided by Hall (2011).

In the simple Bewley (1977) model without nominal rigidities, market clearing is achieved by searching for the interest rate r that clears the bond market and therefore the goods market as well. However, as will be seen below, the high-uncertainty shock causes the market-clearing interest rate to fall well below the zero bound (the gross rate falls below 1). This cannot happen when r is constrained by the zero lower bound and prices are sticky. At $r = 1$, household demand for bonds exceeds the fixed aggregate supply. This disequilibrium in the bond market implies that households wish to produce more output than they wish to consume. For markets to clear, output must adjust downward instead. I consider two methods of modeling the drop in output. The first method follows the representative agent zero lower bound literature and Guerrieri and Lorenzoni (2011) in a heterogeneous agent environment in introducing an endogenous labor wedge that translates the drop in demand into a reduction in the wage, and then searches for the market-clearing path of wages $\{w_t\}$ over the transition. This method follows in spirit Galí, Gertler, and López-Salido (2007), who show that the labor wedge can be used to measure the costs of business cycle fluctuations. The second method assumes that wages are also rigid. Firms facing reduced demand must instead adjust their output by reducing employment.

Following a high uncertainty shock, each firm's price $p_{t+1}(j) = p_t(j)$ remains fixed by assumption but demand for the final good falls from Y_t to $Y_{t+1} = \int c(v_{i,t+1} = 1, z_{i,t+1}, \varepsilon_{i,t+1}, b_{i,t+1}) di$. Because prices are fixed, firms now produce $Y_{t+1}(j) = \left(\frac{p_t(j)}{P_t}\right)^{-\phi} Y_{t+1}$. How they adjust depends on whether wages are assumed to be flexible or nominally rigid.

Adjustment through a fall in wages In the first case, wages are assumed to be flexible. Firm j now faces profit-maximizing condition

$$\pi_t(j) = \max_{w_t(j)} p_{t+1}(j) Y_{t+1}(j) - w_{t+1} \theta_{t+1} n_{t+1} \quad (20)$$

and seeks to produce the level of output implied by the fixed price level as cheaply as possible. The disequilibrium in the goods market caused by the zero bound and price rigidity implies that households wish to supply more labor than is necessary to satisfy demand. As a result, firms are able to bid down wages to w^* , the wage at which labor supply equals labor demand. Aggregate firm profits are now necessarily non-trivial and a lump sum payment of $(P_{t+1} - w^*) Y_{t+1}$ is rebated to each household. The endogenous labor wedge reduces households' desire to work and equilibrates the goods market at a lower level of output. The wage takes the place of the interest rate in achieving market-clearing. Consequently, a solution to the model is now found by first searching for the path of interest rates $\{r_t\}$ that clears the market, and then searching for an additional path of wages $\{w_t\}$ where $w_t < \frac{\phi-1}{\phi}$ in any period t in which excess bond demand remains at $r_t = 1$. That is, we search for the market-clearing wage only if the zero lower bound prevents the interest rate from clearing the bond market.

The introduction of an endogenous labor wedge is standard in the literature, but it has strong and potentially misleading implications that are highlighted by this model's emphasis on idiosyncratic risk. A reduction in labor supply achieved through a shared reduction in wages implies that the output costs of the recession are shared evenly across all (employed) households. Each unit of effective labor now earns a wage less than 1, but there is no increase in idiosyncratic risk implied by the drop in output. This misses a potentially quantitatively important feedback mechanism between idiosyncratic risk and aggregate demand: a drop in output at the zero bound might increase idiosyncratic risk through higher unemployment, which in turn reinforces precautionary motives and further reduces output in a demand-constrained zero bound economy. In fact, not only does the use of an endogenous labor wedge mean that the output drop does not amplify idiosyncratic risk, but the standard assumption that profits are rebated lump-sum to households implies a reduction in idiosyncratic risk. As noted earlier, in a simple bond economy, profits serve as income redistribution financed by a linear labor income tax. While adjustment is usually achieved

with a modest level of profit, the profit rebate is quite large relative to the incomes of poor households. As a result, adjustment to reduced demand through an endogenous labor wedge implies that the recession is associated with a *reduction* in idiosyncratic risk. This is an odd way to think about the labor market during recessions and potentially distorts quantitative estimates of both the size of the drop in output following a demand shock and of the efficacy of social insurance policies that target idiosyncratic risk.

Adjustment through a rise in unemployment In the second case, wages are assumed to be rigid, $w_{t+1} = w_t$. Because both unit prices and costs are fixed and firms earn epsilon profit per unit produced, firms choose the labor input $n_{t+1}(j)$ necessary to satisfy the new level of demand $Y_{t+1}(j)$. The gap between the aggregate labor households would like to supply at the old wage and firms' aggregate labor demand now appears as an increase in unemployment; that is, the goods market now equilibrates by preventing some households from working. To solve the model when adjustment at the zero bound occurs through unemployment, I search for either the path of unemployment entry rates $\{u_t^{in}\}$ or exit rates $\{u_t^{out}\}$ consistent with market clearing. As in the labor wedge case, I first allow the interest rate to clear the market and only search for the market-clearing unemployment transition probabilities when a surplus demand for bonds remains when the zero bound is binding. In either case, I adjust only the entry or the exit rate and leave the other fixed at the low-uncertainty state value. The aggregate demand for bonds falls to clear the market because unemployed households have greatly reduced savings demand.

Adjustment through an increase in unemployment entry rates implies an increase in the incidence of unemployment spells of unchanged expected duration. Adjustment through unemployment exit rates implies that households have the same likelihood as before of entering unemployment, but that unemployment spells are now expected to last longer. How unemployment rises during recessions remains under debate. Hall (2005) and Shimer (2007) find that unemployment rises almost entirely through a reduction in the unemployment exit probability, while Fujita and Ramey (2008) find that the separation rate (entry into unemployment) accounts for up to half of the cyclical fluctuation. To illustrate clearly the implications of each, I consider adjustment mechanisms in which the output drop is translated entirely into increases in the entry rate or decreases in the exit rate, respectively.

I therefore consider three models in addition to the flexible case: adjustment through wages (an endogenous labor wedge), adjustment through the unemployment entry rate, and adjustment through the unemployment exit rate. The difference between adjustment through wages and adjustment through unemployment emphasizes the effect of the fall in output on idiosyncratic risk. The first approach is associated with a uniform drop in the level of income that affects all employed households, but no further increase in idiosyncratic risk. The second approach imposes the entire cost of the output drop on unem-

ployed households, and imposes no reduction in wages on the employed. The third approach further concentrates idiosyncratic risk: when increased unemployment is achieved by reducing the exit rate, its costs fall on a limited group of long-term unemployed. As will be seen below in the discussion of the results of the calibrated model, the output costs of the initial demand shock are considerably greater when market clearing is achieved in a manner that amplifies the initial rise in idiosyncratic risk.

1.3 Calibration

The main innovation in the model is the inclusion of time-varying risk in idiosyncratic labor productivity. The parameters of this process are taken from Storesletten, Telmer, and Yaron (2004), who estimate a similar regime-changing income process using data from the Panel Study of Income Dynamics. I convert to quarterly values their annual estimates of $\rho = 0.95$, $\sigma_{low}^{\eta} = 0.12$, $\sigma_{high}^{\eta} = 0.21$ and $\sigma^{\epsilon} = 0.25$. Because their criterion for dividing periods – whether GNP growth was below or above average – includes some mild years in the recession category, I use this set of estimates as the "conservative" parameterization of the income process. I use those parameters to extrapolate a second set of more "extreme" high-state parameters intended to represent the current recession. It is of course too early to estimate the parameters of the income process for the last couple years. As a crude approximation, I first compute the ratio of two ratios: the ratio of the average civilian unemployment rate in their recessions to that in their expansions, relative to the ratio of parameter estimates $\sigma_{high}^{\eta} / \sigma_{low}^{\eta}$. I use this ratio combined with the ratio of the average unemployment rate in the current recession to that in their average recession to settle on a scaling up of 1/3.

The probabilities of job separation and job finding are set as quarterly adjustments of the monthly values of 0.034 and 0.45 reported in Shimer (2005). The implied unemployment level is meant to represent that arising from labor market frictions, so that any cyclical unemployment imposed by adjustment at the zero bound is additional. I set the unemployment insurance replacement rate $u = 0.3$, with a maximum unemployment benefit \bar{u} of 60% of the median realized income. Both parameters vary considerably across US states, and these are slightly on the higher end of state averages. I set non-leisure hours for unemployed workers $n^U = 0.043$ to match Krueger and Mueller (2010)'s estimate of the 41 minutes the average unemployed worker spends looking for new employment, converted to a share of waking hours. As explained earlier, unemployment is modeled as a transitory shock so that a household's employment status has no effect on the evolution of the permanent component of income. While also a subject of debate, this assumption is in accordance with BLS data on workers' earnings following unemployment spells relative to their pre-unemployment wage, which show that slightly more than half of workers in the pre-recession

period found new employment at a higher wage.⁹ The probability of losing unemployment insurance while unemployed ζ is set to 0.5 in the baseline version, matching the pre-recession environment in which unemployment benefits extended for two quarters. In the policy experiments below, ζ is changed so that the expected duration of benefits is in line with recent policy changes.

The other parameters of the model are calibrated according to standard estimates or to match the pre-recession environment of the US. I set $\gamma^c = 2$ and $\gamma^n = 1.50$, which yields a realized Frisch elasticity of labor supply for a median-income household with the median wealth level roughly in keeping with estimates in Chetty, Guren, Manoli, and Weber (2011). The final parameter in the utility function, N , is set to 1 so that households work about half of their waking hours. The borrowing limit $\underline{b} = -0.75$ is set following the value in Chang and Kim (2006) and is the same for all households. Following Guerrieri and Lorenzoni (2011), I set aggregate debt \bar{B} in the economy to 170% of the model's realized GDP based on their estimate using Federal Reserve Flow of Funds data, and I set $\beta = 0.99$ (quarterly) to achieve a low-uncertainty interest rate in the baseline model with an annualized value of about 3%.




Figure 1.1

The government spending and baseline taxation parameters are also set to match recent levels for the United States. Government tax receipts G are set to 26% of GDP to capture the size of the tax burden net of unemployment benefits. The baseline tax progressivity parameter τ is set to 0.26 following the estimate in Heathcoate, Storesletten, and Violante (2010). Figure 1.1 illustrates the marginal and average tax rates implied by these values for the baseline flexible model under the conservative income parameterization over the range of realized household incomes.

The grid for the transitory shock ε has five points corresponding to bad, neutral, and good realizations,

⁹Worker Displacement News Release. Bureau of Labor Statistics. http://bls.gov/news.release/archives/disp_08202008.htm.

as well as unemployment and unemployment without unemployment benefits. The bad and good realizations are placed one standard deviation from the center. The grid for the persistent component z has nine points spaced according to the Tauchen procedure, but scaled down slightly at the top so that the expected value of income is identical in the low and high states. The ratio of the lowest income level to the median income level is set to mirror the ratio of the income of an employed worker earning the federal minimum wage to the median household income in the United States, about 29%. The bond grid has 250 points, concentrated where the saving rule exhibits sharper concavity, and extends high enough that fewer than one in a million households would otherwise exceed that wealth level. The uncertainty process is of course binary, and I rely on Bloom, Floetotto, and Jaimovich (2010) for estimates of the transition probabilities between uncertainty states, setting $v_{HH} = 0.885$ and $v_{LL} = 1 - 0.047$.

1.4 Results

1.4.1 Flexible case

I first consider the model without nominal rigidities and a zero lower bound. Understanding the fully flexible case helps to clarify household behavior in the low- and high-uncertainty states and the response of the economy to a high-uncertainty shock. This section uses the conservative income risk parameterization and the baseline policies described in the calibration section.




Figure 1.2




Figure 1.3

The labor supply policies of low-income, average-income, and high-income households in the low-uncertainty and high-uncertainty stationary states are shown in Figure 1.2. In this and subsequent figures, income level refers to the persistent component of productivity, z ; the chosen low and high levels are just

above and below the extremes of the grid. Data is taken from the corresponding policy rule with $\varepsilon = 0$. Middle-income households, whose incomes are equally likely to rise or fall, supply slightly more labor in the high uncertainty state. This parallels what Basu and Bundick (2011) calls "precautionary labor" in the face of increased uncertainty, an analogue of precautionary saving. The richest households increase labor supply even more, while the poorest households reduce labor supply when uncertainty rises. This is an unfortunate artifact of discretization of the income process – households already at the extremes of the income distribution can only move in one direction. I discuss the drawbacks of these implications at the end of the paper.

Figure 1.3 plots bond accumulation or decumulation, $b(v, z, \varepsilon, b) - b$, because the policy function itself is too linear for visual inspection. Households with low to medium wealth choose to accumulate at a higher rate in the high- than in the low-uncertainty state as a result of heightened precautionary motives. Because the asset is available in fixed net supply, their increased savings must be matched by reduced holdings of bonds by high-wealth households, who decumulate because of the reduced interest rate. The curves intersect at the wealth level where the effect of the lower interest rate overtakes that of the sharpened precautionary motive.




Figure 1.4




Figure 1.5

The reduction in the dispersion of bond holdings is shown in Figure 1.4, a Huggett (1993)-type plot of the share of the aggregate bond supply in the hands of households at or below a given level of bond holdings. The red curve representing the high uncertainty state is above the blue curve, indicating that the distribution of household bond holdings is more equal in the high uncertainty state, or equivalently that wealth is more dispersed in the low uncertainty state. Indeed, median bond holdings more than double from the low- to the high-uncertainty state. In addition to the precautionary channel, the distribution of

bond holdings is also affected by the income dispersion. An increase in the variance of labor productivity causes both an increase in labor income risk and a widening of the distribution of individual productivities. Households are more likely to reach extreme incomes, but less likely to stay there. The importance of the precautionary motive can be confirmed by checking that the aggregate bond demand of a population characterized by either state's stationary productivity distribution is greater under high- than under low-uncertainty policy rules derived from a common interest rate. Figure 1.5 shows that aggregate bond demand at a given interest rate is greater in the high-uncertainty state. (In Figure 1.5 and below, the interest rate is shown in annualized terms.) Because bonds are in fixed supply across states in the baseline calibration, the market-clearing interest rate in the low-uncertainty stationary state is higher than the high-uncertainty rate. The graph indicates that as the aggregate supply of bonds – which can be thought of as a vertical line on the graph – increases, the gap between the market-clearing rates in the two states narrows. A later section considers a policy experiment motivated by that observation.




Figure 1.6. Response to a high uncertainty shock.

Figure 1.6 shows the impact of a high-uncertainty shock on an economy originally in the low-uncertainty stationary state. The interest rate drops immediately from the low-uncertainty stationary level of 2.83% as a result of the sudden increase in income risk. Initially, the distribution of bond holdings is too dispersed and households with low savings seek to accumulate bonds to adapt to the higher-risk environment. As in Guerrieri and Lorenzoni (2011), this demand for bonds causes the interest rate to overshoot the new lower level. The red line shows the dispersion of bond wealth as it falls from the initial low-uncertainty level. As the dispersion falls, the downward pressure on the interest rate subsides and the interest rate levels off at the high-uncertainty stationary level. This dip and gradual rise clarifies the intuition behind the limited duration of the liquidity trap and endogenous exit from the trap, for certain calibrations. In

the immediate aftermath of the shock, however, the interest rate continues to fall for some time. This results from the rapid rise in the variance of household productivities, which initially dominates the effect of the falling bond dispersion. The population variance of households productivities, the black line in Figure 1.6, is initially low. Because bond accumulation is a convex function of labor productivity z , the quick dispersion of the productivity distribution tends to push the interest rate down as a larger share of households reach high incomes. When this productivity dispersion process ends, the wealth concentration effect drives interest rate dynamics.

Output expands in the first phase of the transition and eventually settles at a slightly higher level. This is achieved with a slightly lower aggregate labor supply because increased labor income volatility enables households to direct more of their labor effort to periods when they are highly productive. However, this effect is quantitatively modest and balanced by the increased desire of low-wealth households (who are more likely to also be low-income households) to work more. Output is temporarily high because in aggregate, the effect of the increase in labor supply is stronger than the dampening effect of increased precautionary motives on consumption. Because the interest rate is fully flexible here, the intertemporal price of consumption today falls and high-wealth households in particular increase spending. In contrast, in the next section the zero lower bound and price stickiness will make consumption today "too expensive", preventing the boom.

In contrast to models of forced deleveraging, the transition occurs more slowly here because the reallocation of bond-holdings occurs freely. The figures plot the long-term aftermath of the shock to illustrate the full transition, but it is probably unrealistic to think of the high-uncertainty state lasting this long. To the extent that the economy bounces more frequently between the low- and high-uncertainty states, the reallocation of bond holdings will be incomplete and the effect of switching uncertainty regimes will therefore be less severe.

To test the robustness of the result that the market-clearing interest rate falls below zero during the transition period, I also compute the stationary states and transition dynamics for $\gamma^c = 1.25$ and \underline{b} equal to 6 times median income, cases in which households have reduced risk aversion and looser borrowing limits. In both cases, the high-uncertainty stationary state is characterized by an interest rate above zero, but the zero bound does bind for part of the transition period, falling to a minimum of roughly 0.997 in each case.

1.4.2 The zero lower bound

I now consider the introduction of nominal rigidities and the zero lower bound. To capture a shock closer in magnitude to the current recession, I use the extreme income risk parameterization from here forward.

I first compare adjustment through an endogenous labor wedge with adjustment through increased likelihood of entry into unemployment by calculating the stationary state and then finding the transition dynamics following a high uncertainty shock. The comparisons in this section use the baseline parameters from the calibration section.




Figure 1.7

Figure 1.7 shows the transition paths following a high uncertainty shock. The frictionless case under the new extreme income risk calibration is an exaggerated version of the result in Figure 6, but output falls in both the labor wedge and unemployment cases. The green line shows that adjustment is more painful in output terms when it occurs through unemployment. This results from the further increase in precautionary saving demand induced by the adjustment mechanism, which worsens the paradox of thrift. Figures 1.8 and 1.9 explore the transition dynamics at the zero bound in greater depth. When adjustment occurs through a labor wedge, the wage level falls by up to 10%, mirroring the path of the interest rate in Figure 1.6. Output drops from the low-uncertainty stationary state immediately on impact, but the drop in the wage is considerably larger than the fall in output. When adjustment occurs through unemployment, shown in Figure 1.9, the initial spike in excess demand for savings causes a jump in unemployment that partially dissipates. For as long as prices and wages remain fixed, the economy remains in a state of heightened unemployment. The output drop is considerably worse than in the case of the labor wedge because households now face increased idiosyncratic risk from two sources: the increased volatility of the labor productivity process (the high-uncertainty shock) and the heightened risk of unemployment. Figure 1.9 also shows that excess unemployment, the amount by which the unemployed share of the population exceeds the baseline frictional level in the low-uncertainty state, is roughly 1.5-2 times the size of the output drop. The response of output is more muted because employed households choose to work more than

under low uncertainty.




Figure 1.8




Figure 1.9

The transition to the high-uncertainty stationary state occurs more gradually in the unemployment model. The intuition is straightforward: the transition is complete when households have arrived at their desired high-uncertainty levels of precautionary savings, but households that wish to increase their bond holdings have very limited ability to do so in periods in which they cannot earn labor income.

1.4.3 Policy experiments

I next compare the effects of various policy interventions and parameter experiments under each adjustment mechanism. I provide a comparison with the effects of these policies in the fully flexible case so that the effect of the policy in normal circumstances, the purely distortionary effect, can be disentangled from its impact on mitigating or worsening the output costs unique to the zero bound.

The results of the policy experiments are shown in Table 1.1 and Table 1.2. The first row of Table 1.1 uses the baseline calibration under extreme income risk. Row 2 considers a parameter experiment in which the unemployment entry and exit probabilities are halved, so that the expected duration of unemployment doubles, holding constant the unemployment rate. Row 3 shows the results of a large increase in borrowing limits. Rows 4 and 5 show the effect of increasing the aggregate supply of bonds in the high uncertainty state. Row 6 considers a policy experiment with a state-varying unemployment replacement rate that is higher in the high-uncertainty state, and rows 7 and 8 lower and raise the replacement rate across states, respectively.¹⁰ Rows 10-12 perform the same experiment for τ , the progressivity of the income tax. Row 9

¹⁰Engen and Gruber (2001) provide evidence of the effect of unemployment insurance on savings through differences in the programs of US states. They find that the effect of more generous unemployment insurance on household asset accumulation is negative and highly significant. The effect is modest in dollar terms but increases sharply when unemployment risk is higher.

considers an expected duration of unemployment benefits of 8 quarters to approximately match the recent US federal extension of benefits to 99 weeks. In each case, λ adjusts so that the government budget is always balanced.

The effects of policies should first be compared within columns, each of which corresponds to a particular model (labor market adjustment mechanism). All values are reported as percentage deviations from the baseline parameterization of the fully flexible model. The first row shows the output drop in the baseline case for each model relative to the level in the fully flexible case. For example, when adjustment occurs through the unemployment entry rate, output is slightly more than 4% lower than in the flexible case. We can now look down the column to compare the effects of different policies and parameter experiments with the baseline for the given model.

Table 1.1. Output in high uncertainty stationary state, measured as percentage deviation from flexible case no-policy baseline.				
	<i>Flexible</i>	<i>Labor wedge</i>	<i>Unemp. entry</i>	<i>Unemp. exit</i>
1. Baseline, extreme income risk parameterization	0	-2.38	-4.14	-22.03
2. Twice expected duration of unemployment	-0.04	-3.29	N/A	N/A
3. $\underline{b} = 6$ times median household earnings	0.32	-1.13	-1.82	-2.63
4. 20% increase in \bar{B} in high state	0.06	-1.91	-3.32	
5. 50% increase in \bar{B} in high state	0.04	-1.40	-2.16	-4.22
6. $u = 0.20$ in low state, $u = 0.40$ in high state	-0.12	-2.24	-3.59	-21.28
7. $u = 0.20$	0.15	-2.28	-4.32	-22.44
8. $u = 0.40$	-0.06	-2.47	-4.02	-21.75
9. $\zeta = \frac{1}{8}$ (expected UI duration of 8 quarters)	-0.04	-2.32	-3.92	-8.30
10. $\tau = 0.26 * 0.925$ in low state, $\tau = 0.26 * 1.075$ in high state	-0.89	-1.87	-2.62	
11. $\tau = 0.26 * 0.85$	1.34	-1.34	-3.42	
12. $\tau = 0.26 * 1.15$	-1.40	-3.50	-5.08	

Several results are noteworthy. First, a comparison across columns shows that while adjustment through increased entry into unemployment leads to a sharper drop in output than adjustment through reduced wages, the results are dramatically worse when adjustment comes through the unemployment exit rates. Intuitively, households with moderate savings are less concerned about a modest risk of a brief period of unemployment, during which they can decumulate bond holdings, than about a smaller risk of prolonged

unemployment, for which their savings are less adequate. This is a simple consequence of concavity. As a result, the extreme concentration of risk provokes a much sharper precautionary reduction in demand and therefore in output. The same result can be seen within a model by comparing rows 1 and 2 of the labor wedge case, which contrasts the baseline with a doubling of the expected duration of unemployment, holding constant the total number of unemployed. Figure 1.10 shows the transition path of output for this comparison.




Figure 1.10

A second lesson of Table 1.1 is that policies that greatly expand household borrowing limits or the aggregate supply of bonds are very effective in mitigating the drop in output. These policies reduce risk by increasing households' ability to smooth consumption. Note, however, that this result is achieved in the context of a class of model that does not allow for default.




Figure 1.11




Figure 1.12

A third lesson is that the effectiveness of particular social insurance policies depends on the manner in which output adjusts at the zero lower bound. Figures 1.11 and 1.12 plot the transition dynamics of output in the wedge and unemployment entry case for the baseline policy and a policy experiment that introduces a higher replacement rate in the high- than in the low-uncertainty state. The policy has a larger effect in the model in which output adjusts through unemployment because that model features greater unemployment risk. For a social insurance policy to have a substantial quantitative effect on output, it must reduce precautionary concerns by targeting the type of idiosyncratic risk created by the output drop. An even starker illustration of this point is provided by row 9. The extension of unemployment benefits has a very modest effect on output in the labor wedge and unemployment entry models because it is uncommon for households to remain unemployed for longer than they receive benefits, but it has a very large effect when unemployment rises through a reduction in exit probabilities. In the current recession, the average duration of unemployment spells rose to exceed 3 quarters in mid-2011. This is roughly double the longest average duration ever previously recorded in the US and is a salient characteristic of the current recession.

Similarly, a fourth lesson is that certain social insurance policies become more appealing at the zero lower bound when they have the additional benefit of restoring output by reducing precautionary savings demand. Consider, for example, row 10, which shows the effect of a time-varying progressivity of the labor income tax. The first column shows that this policy is unappealing in a fully flexible model, producing a reduction in output of 0.89% relative to the baseline. In the labor wedge case, however, the policy increases output relative to the model's baseline by about half a percentage point, from a drop of 2.38% to a drop with the policy of 1.87%. The results are even stronger when the labor market adjusts through entry into unemployment. This is one of the core implications of the model: there is an added benefit to social

insurance in a paradox of thrift.

Table 1.2. Welfare in high uncertainty stationary state, measured as percentage deviation of consumption equivalent from flexible case no-policy baseline.				
	<i>Flexible</i>	<i>Labor wedge</i>	<i>Unemp. entry</i>	<i>Unemp. exit</i>
1. Baseline	0	0.20	-0.82	-8.17
2. Twice expected duration of unemployment	-0.19	0.18	N/A	N/A
3. $\underline{b} = 6$ times median quarterly earnings	0.49	0.60	0.08	-0.13
4. 20% increase in \bar{B} in high state	0.22	-0.03	-0.47	
5. 50% increase in \bar{B} in high state	0.50	0.54	0.03	-0.67
6. $u = 0.20$ in low state, $u = 0.40$ in high state	0.02	0.19	-0.63	-7.74
7. $u = 0.20$	0.01	0.22	-0.90	-8.50
8. $u = 0.40$	0.04	0.23	-0.70	-7.91
9. $\zeta = \frac{1}{8}$ (expected UI duration of 8 quarters)	0.01	0.20	-0.74	-2.07
10. $\tau = 0.26 * 0.925$ in low state, $\tau = 0.26 * 1.075$ in high state	-0.14	-0.05	-0.48	
11. $\tau = 0.26 * 0.85$	-0.45	-0.17	-1.36	
12. $\tau = 0.26 * 1.15$	0.39	0.53	-0.34	

Table 1.2 presents the welfare effects of the various policy experiments in the high uncertainty stationary state measured in consumption equivalents. The consumption equivalence ratio between policy experiments A and B is

$$CE \text{ ratio} = \left(\frac{\iiint V^A(1, z, \varepsilon, b) * \mathcal{H}_1^A(z, \varepsilon, b) dzd\varepsilon db}{\iiint V^B(1, z, \varepsilon, b) * \mathcal{H}_1^B(z, \varepsilon, b) dzd\varepsilon db} \right)^{\frac{1}{1-\gamma^c}}$$

Once again, comparing with the first column allows us to decompose the welfare effects of policy interventions into two components, the distortionary effect of the policy and the effect unique to the zero lower bound. The second column of Table 1.2 highlights a disadvantage of the labor wedge approach. Despite the drop in output, welfare is higher in this model. As noted earlier, this is an unrealistic artifact of the rebating of profits, which effectively turns profits into a form of tax-financed lump-sum social insurance that considerably improves the welfare of the poor. In the baseline calibration, profits are just over 6% of output, a substantial increase in income for low-productivity households and the unemployed. Also of particular note are the large positive welfare effects of loosening borrowing constraints and expanding the

supply of bonds and the large gains from lengthening the duration of unemployment benefits in the fourth model.

Table 1.3. Percentage of households at the borrowing constraint or within 10% of it (in parentheses) in the high uncertainty state.					
		<i>Flexible</i>	<i>Labor wedge</i>	<i>Unemp. entry</i>	<i>Unemp. exit</i>
Baseline	<i>All households</i>	0.02 (0.33)	0.02 (0.34)	0.04 (0.78)	2.34 (4.22)
	<i>Unemployed</i>	0.32 (3.86)	0.35 (3.69)	0.30 (5.40)	7.90 (14.15)
$\zeta = \frac{1}{8}$	<i>All households</i>	0.05 (0.65)	0.05 (0.68)	0.12 (1.56)	0.23 (0.92)
	<i>Unemployed</i>	0.87 (7.23)	0.84 (6.29)	1.03 (9.38)	1.38 (5.16)

Table 1.3 compares the number of households at or near the borrowing constraint for the baseline case and for the policy experiment with extended unemployment benefits. When unemployment is of long duration (column 4), the share of constrained or nearly-constrained households is significantly higher. The introduction of the policy sharply reduces the share for whom the constraint binds. In the other models, the extension of unemployment benefits encourages unemployed households to decumulate savings more quickly because the probability of being unemployed and running out of benefits is very small with the policy.

1.4.4 Discussion

The output and welfare consequences of adjustment through unemployment, and especially through the unemployment exit rate, might appear implausibly strong. Much of the explanation is the assumption that each household faces equal transition probabilities into and out of unemployment, so that strong precautionary motives are widely shared. In reality, Sum and Khatiwada (2010) shows that unemployment is dramatically higher among the poor, and unemployment rates for those with high past income and high levels of education have been considerably more modest in the current recession. That the cyclical component of unemployment tends to be concentrated among the young, poor, and less educated also helps explain a seemingly counterfactual implication noted above, namely, that households already at or near the bottom of the productivity distribution benefit from increased income risk. In actual recessions, these households are usually at the greatest risk of transitioning into unemployment. Empirically, in the US economy unemployment is heavily concentrated both the poor and certain demographic groups, including the young and some racial minorities. To the extent that these factors are absent from the model, the re-

sults might overstate the strength of precautionary motives in large demographic groups that are relatively insulated from unemployment shocks.

In contrast, the costs of the zero bound when adjustment occurs through an endogenous labor wedge seem too modest. Clearly the implication that wages are flexible while prices are not is counterfactual, and more importantly here, it is clearly not true that the costs of the reduction in output are shared equally by all employed workers.

Nevertheless, the more modest costs at the zero bound associated with adjustment through wages might lend support to another class of policy akin to the German *Kurzarbeit* program of reduced working hours or work-sharing. Such policies reduce the additional idiosyncratic risk created by the drop in output and consequently alleviate the drop in demand arising from precautionary motives. In countries with such policies, the labor wedge model perhaps provides a better approximation of the costs of the zero bound. It is also worth noting that such "working-sharing" policies reduce risk without acting through wage reductions, which in richer models such as Eggertsson (2010) and Eggertsson (2009) are likely to be translated into deflation that worsens the zero bound problem by raising real rates.

Finally, while social insurance policies such as extended unemployment benefits can help to alleviate the paradox of thrift, we might worry that they are deeply distortionary. How troubling is this concern? As Landais, Michailat, and Saez (2011) suggest, the answer is likely not very. They construct a search-and-matching model in which aggregate shocks and wage rigidities lead to job rationing in recessions. They show that because the return to job-search effort is modest when jobs are limited, the job-search distorting costs of unemployment insurance are correspondingly modest in recessions. A similar result would hold in a richer search-and-matching model of a demand-constrained economy with rigid wages. A higher rate of exit from unemployment due to increased job-search effort must be balanced by increased entry as long as excess savings demand remains at the zero bound. There are two additional reasons why the distortionary costs of social insurance are modest in the model. First, at the modest replacement rate of about 30% of income, concavity implies that the drop in income to a share u of a household's optimal labor income is extreme in utility terms relative to the additional leisure above the optimal choice. Second, when unemployment duration is high, it is especially unlikely that households will choose to risk unemployment. Voluntarily spending down accumulated savings for a short period is only a sensible strategy when unemployment is not expected to last long. I verify that the distortionary effects are small by checking whether households would ever choose to feign unemployment. A household at state $(v_i, z_{i,t}, \varepsilon_{i,t}, b_{i,t})$ would voluntarily enter a state of unemployment if it were true that

$$V(v_i, z_{i,t}, -\infty, b_{i,t}) > V(v_i, z_{i,t}, \varepsilon_{i,t}, b_{i,t}). \quad (21)$$

This does not hold for any $(v_i, z_{i,t}, \varepsilon_{i,t}, b_{i,t})$ in any of the parameterizations, including the policy experiment in which u is increased or benefit duration is extended. Intuitively, the reason households never voluntarily enter the unemployment state is that the drop in income to a share u of a household's optimal labor income is extreme relative to the modest gains in utility from leisure above the optimal choice.

1.5 Numerical solution method

1.5.1 Stationary states

As described in the definition of the regime-switching equilibrium concept above, finding an equilibrium consists of finding two mutually consistent stationary states, one corresponding to the low uncertainty state, $v = 0$, and one to the high uncertainty state, $v = 1$. The equilibrium therefore requires solving for four elements: (1) a pair of interest rates (r_0, r_1) , (2) household policy functions $n(v, z, \varepsilon, b)$, $c(v, z, \varepsilon, b)$, and $b(v, z, \varepsilon, b)$ with decision rules for both states that recognize the possibility of transition between states, (3) a pair of tax levels (λ_0, λ_1) that balance the government budget given $(\tau_0, \tau_1, u_0, u_1)$, and (4) a pair of joint distributions of household characteristics $\mathcal{H}_0(z, \varepsilon, b)$ and $\mathcal{H}_1(z, \varepsilon, b)$.

I solve the model in the following steps:

1. Calculate the transition matrix \mathcal{P} over the exogenous state variables (v, z, ε) . Use the Tauchen method to space the productivity gridpoints using the variances of the innovations in the low state, and then calculate transition probabilities between the same gridpoints using the variances corresponding to the high state. This creates a single transition matrix that recognizes the possibility of moving across uncertainty states. Finally, because productivity is an exponential function of the productivity gridpoints, scale down the highest z gridpoint so that the expected value of productivity is the same in the low and high states.
2. Guess a pair of interest rates, (r_0, r_1) .
3. Guess a pair of balanced-budget tax levels, (λ_0, λ_1) .
4. Solve for the policy functions using the generalized endogenous gridpoints method developed by Barillas and Fernández-Villaverde (2007), who build on Carroll (2006). The method alternates between an endogenous gridpoint method (EGM) step and a value function iteration (VFI) step, using the much faster EGM step to provide a good approximation of the continuation value for the value function iteration step, which is also necessary because labor supply is endogenous. Because of the speed of the EGM algorithm, we can define a much denser grid for b in the EGM step so that the

continuation value used in the VFI step is defined on a very fine grid. Both bond grids concentrate points in the region where the policy functions are most skewed. Starting with an initial guess of $n(v, z, \varepsilon, b) = 0.5$ for all (v, z, ε, b) and an initial continuation value that is increasing in bonds, apply an initial round of EGM to find a guess of the continuation value \mathbf{V}_{t+1} , handling the nonlinearity created by the borrowing constraint in the manner suggested by Carroll (2011).

5. Solve for the policy functions $n(v, z, \varepsilon, b)$, $c(v, z, \varepsilon, b)$, and $b(v, z, \varepsilon, b)$ using VFI, and update \mathbf{V}_{t+1} .
6. Calculate the two stationary states, $\mathcal{H}_0(z, \varepsilon, b)$ and $\mathcal{H}_1(z, \varepsilon, b)$, in the manner suggested by Lundquist and Sargent (2000). Recall the meaning of the stationary states: the low-uncertainty stationary state is the household distribution reached when $v = 0$ for sufficiently long. That is, it is the invariant distribution corresponding to the subset of low-state policy rules, which recognize the possibility of transition to the high state. Calculating the stationary states on every EGM/VFI iteration of the policy function loop has modest computational costs and allows us to update the guess of (λ_0, λ_1) on each iteration. This is a much faster alternative to imposing an additional, outermost loop that searches for the correct parameter values around the interest rate loop beginning in step 2. We also update the calculation of unemployment insurance $UI_t(1, z, \varepsilon, b)$ using $n(v, z, \varepsilon, b)$ as the "counterfactual" labor supply.
7. Continue to alternate between steps of EGM and VFI until the policy functions have converged, interpolating between the defined grid and the endogenous grid as described by Barillas and Fernández-Villaverde (2007). Repeat step 6 on each iteration. The government budget will be balanced to within roughly the same convergence criterion, sometimes exceeding it mildly as a result of rounding error in the computation of the stationary states.
8. When the policy functions have converged, compare total bond demand in each uncertainty state to the fixed supply of bonds. Update the low-uncertainty state interest rate as

$$r'_0 = r_0 - \varepsilon_r \left(\int b(0, z, \varepsilon, b) di - \bar{B} \right),$$

where ε_r is an updating parameter, and likewise for r_1 . Note that the calculation of the deviation from market-clearing in one state is only correct to the extent that the interest rate in the other state is already correct; consequently, more efficient updating algorithms can be problematic.

9. Now return to step (5). We proceed directly to the VFI step on subsequent iterations of the interest

rate loop so that we can maintain our guess of the continuation value. We avoid starting with the EGM step because the EGM algorithm can be fragile in this case: it fails if the continuation value is ever non-monotonic in bonds, which might result from changes in the interest rates if the bond grid is sufficiently dense.

10. The equilibrium has been found when deviations from market clearing in both states satisfy a convergence criterion,

$$\int b(0, z, \varepsilon, b) di - \bar{B} < \varepsilon_b, \quad \int b(1, z, \varepsilon, b) di - \bar{B} < \varepsilon_b.$$

This equilibrium solution method involves a necessary approximation: household decision rules assume that the transition between uncertainty states is instantaneous. When the economy moves from the low uncertainty state to the high uncertainty state, it is immediately in the high-state equilibrium. This implies that equilibrium decision rules ignore the transitional dynamics between states discussed in the next section.

Flexible price and interest rate When prices are flexible and the interest rate is unconstrained, I apply the above steps with no limits on (r_0, r_1) and allow the interest rates to adjust to clear the bond market.

Nominal rigidities case, adjustment at the zero bound through a labor wedge We apply the same algorithm as above, but modify step 8 so that r_0 and r_1 are bounded at 1. In practice, the constraint only binds in the high-uncertainty state. When r_1 is already 1 and $\int b(0, z, \varepsilon, b) di > \bar{B}$, we instead adjust the wage w_1 (equal to one minus the wedge),

$$w_1' = w_1 - \varepsilon_\omega \left(\int b(1, z, \varepsilon, b) di - \bar{B} \right).$$

We continue to iterate until w_1 achieves market-clearing, holding r_1 to 1 on each iteration.

Nominal rigidities case, adjustment at the zero bound through unemployment When the zero bound prevents further interest rate adjustment and wages are fixed, market clearing is achieved by adjusting u^{in} or u^{out} , respectively. In the first case, u^{out} is held constant and the probability of entering the unemployment state is updated as

$$u^{in'} = u^{in} + \varepsilon_{u^{in}} \left(\int b(1, z, \varepsilon, b) di - \bar{B} \right).$$

In the second case, u^{in} is held constant and the probability of exiting the unemployment state is updated as

$$u^{out'} = u^{out} - \varepsilon_{u^{out}} \left(\int b(1, z, \varepsilon, b) di - \bar{B} \right).$$

We iterate until convergence and report the resulting stationary unemployment level implied by the transition probabilities.

1.5.2 Transition dynamics

I calculate the transition dynamics that follow a high-uncertainty shock to the low-uncertainty stationary state. To do this, I follow the standard method of calculating the policy rules backward and then calculating the population distribution forward, searching for the path of interest rates that achieves bond market clearing at every point in the transition.

The approximation used in the solution method for the transitional dynamics parallels that used in calculating the equilibrium. The policy rules that define the equilibrium ignore the transitional dynamics between uncertainty states, and the policy rules calculated over the transition path likewise assume that reverting to the low-uncertainty state would be instantaneous. This is necessary because both transition paths – from the low to the high state and the high to the low state – cannot be solved simultaneously. This form of approximation could potentially be avoided using the Krusell and Smith (1998) algorithm, which is intended for models with aggregate risk but could be used here as well. However, that algorithm is considerably slower here and the approximation involved in using a small number of moments of the household distribution is less transparent.

The transitional dynamics are calculated in the following steps:

1. Choose a number of periods T for the transition. We choose a value of T such that the income distribution in the low uncertainty state $\mathcal{H}_0(z, \varepsilon)$ converges to within ε_{TD} of $\mathcal{H}_1(z, \varepsilon)$ when the high-uncertainty transition rules for z and ε are applied repeatedly for T periods.
2. Guess a path of interest rates, $\{r_t\}$.
3. Solve for a path of household policy functions, $\{n_t(1, z, \varepsilon, b)\}_{t=1, \dots, T}$, $\{c_t(1, z, \varepsilon, b)\}_{t=1, \dots, T}$, and $b_t(1, z, \varepsilon, b)_{t=1, \dots, T}$. Note that we are only solving for the high-uncertainty policy rules along the transition path, and assume that the low uncertainty state policy rules would hold immediately if the transition were disrupted by a shock $v_t = 0$ at any point t along the path. Start by assuming $n_T = n_t(1, z, \varepsilon, b)$; that is, that by the end of the transition, the household distribution and interest rate path have converged to the high-uncertainty state levels so that the associated policy rules are the true

values for period T . Then solve backward from $T - 1$ to 1. Begin with a single iteration of EGM to find $c_{T-1}(1, z, \varepsilon, b)$; again, only one iteration is necessary because $c_T(v, z, \varepsilon, b)$ is assumed to be true. Then use the VFI step as described above to find $n_{T-1}(1, z, \varepsilon, b)$ and $b_{T-1}(1, z, \varepsilon, b)$, with the same correction as above at gridpoints where households are borrowing constrained. Then interpolate from the endogenous grid to tomorrow's grid and proceed to $t - 1$, continuing to $t = 1$.

4. Use the complete set of policy rules along the transition path to find the household distribution at each point along the transition path, $\{\mathcal{H}_t(z, \varepsilon, b)\}_{t=1, \dots, T}$. Start by setting $\mathcal{H}_{t=1}(z, \varepsilon, b) = \mathcal{H}_0(z, \varepsilon, b)$, the household distribution in the low uncertainty stationary state. To update the household distribution from $\mathcal{H}_t(z, \varepsilon, b)$ to $\mathcal{H}_{t+1}(z, \varepsilon, b)$, use the exogenous transition rules for z and ε and the policy rule $b_{t-1}(1, z, \varepsilon, b)$.
5. Use the full path of household distributions $\{\mathcal{H}_t(z, \varepsilon, b)\}_{t=1, \dots, T}$ to calculate the economy's aggregate variables. For each period, update the balanced-budget tax level, yielding a path $\{\lambda_t\}_{t=1, \dots, T}$, and use the labor supply policies $\{n_t(1, z, \varepsilon, b)\}_{t=1, \dots, T}$ to update the value of unemployment insurance $\{UI_t(1, z, \varepsilon, b)\}_{t=1, \dots, T}$ based on "counterfactual" labor supply.
6. At each point on the transition path, calculate total bond demand in the high uncertainty state. Update the interest rate in the same manner as before,

$$r'_t = r_t - \varepsilon_r \left(\int b_t(1, z, \varepsilon, b) di - \bar{B} \right),$$

at each point t along the transition path.

7. Check if the deviations from market clearing satisfy the convergence criterion at every point along the path,

$$\int b_t(1, z, \varepsilon, b) di - \bar{B} < \varepsilon_b \quad \forall t \in [1, T].$$

If convergence has not been achieved, return to step 3 and continue iterating.

8. Once convergence has been achieved, compute $|\mathcal{H}_T(z, \varepsilon, b) - \mathcal{H}_1(z, \varepsilon, b)|$ to check that T was large enough that the final household distribution is sufficiently close to the high-uncertainty stationary state. We can also compare r_T to r_1 and final output to output in the high uncertainty state.

Transition dynamics with nominal rigidities The solution method can be modified for the case with a zero bound, whether adjustment occurs through unemployment or a labor wedge, using a method parallel to that used for calculating the stationary states. We first search for the paths of market-clearing interest

rates, and adjust the labor wedge or unemployment transition probabilities at transition period t only if an excess of bond demand remains when $r_t = 1$.

Two changes are necessary in solving for unemployment rates over the transition path. First, aggregate bond demand in period t is now used to update the likelihood of entering unemployment from period $t - 1$ to t . Once the entry rate has been updated, new transition matrices must be computed for each period along the transition path. Second, unlike in the flexible interest rate or wedge cases, changing the unemployment transition probabilities between periods does not allow us to respond to total bond demand in the initial period. As a result, the initial population $\mathcal{H}_{t=1}(z, \varepsilon, b)$ must be changed directly. On each iteration, in addition to updating the full path of transition probabilities $\{u^{in}\}$ and $\{u^{out}\}$, I update the initial population according to

$$\mathcal{H}_{t=1}(z, \varepsilon, b) = \mathcal{H}_{t=1}(z, \varepsilon, b) * (1 - \varepsilon_{\mathcal{H}}) \left(\int b_1(1, z, \varepsilon, b) di - \bar{B} \right) \text{ for all } \varepsilon \neq -\infty.$$

This reduces the employed population at each point in proportion to the size of the excess demand for bonds. For each (z, ε, b) , $\varepsilon \neq -\infty$, I then transfer the corresponding population density to the point on the grid representing a newly-unemployed household with the same values of z and b ,

$$\mathcal{H}_{t=1}(z, \varepsilon \neq -\infty, b) = \mathcal{H}_{t=1}(z, \varepsilon \neq -\infty, b) + \mathcal{H}_{t=1}(z, \varepsilon, b) * \varepsilon_{\mathcal{H}} \left(\int b_1(1, z, \varepsilon, b) di - \bar{B} \right).$$

This ensures that the adjustment to the shock is immediate: the share of the population that is unemployed changes immediately, and then the transition path is found for the remaining periods. Once convergence is achieved, I use $\{\mathcal{H}_t(z, \varepsilon, b)\}_{t=1, \dots, T}$ to calculate the unemployment level over the transition path.

1.6 Conclusion

The model presented above makes three points. First, a liquidity trap can result from a sudden increase in precautionary savings demand prompted by higher idiosyncratic labor income uncertainty. A rise in idiosyncratic risk provides a robust microfoundation for the liquidity trap and explains how the initial large drop in the market-clearing interest rate dissipates over time. Second, the output and welfare effects of the liquidity trap can be large when the zero lower bound binds. However, the magnitude of the recession depends crucially on the adjustment mechanism by which the labor market clears when price stickiness and the zero bound prevent the real interest rate from falling. The output and welfare costs of the liquidity trap are higher when adjustment occurs through a mechanism that reinforces the original increase in idiosyncratic risk. In particular, a rise in unemployment, and especially in unemployment of long dura-

tion, concentrates idiosyncratic risk and provokes increased precautionary behavior and therefore a sharper drop in aggregate demand and thus output. Third, social insurance policies can be effective in reducing the output drop at the zero lower bound. In fact, some social insurance policies that are otherwise output- and welfare-reducing instead become output- and welfare-enhancing at the zero bound because of their effect on reducing precautionary motives and alleviating the paradox of thrift. However, social insurance policies are most effective at the zero bound when they succeed in targeting the type of idiosyncratic risk households face as a result of the labor market consequences of the output drop. Most strikingly, the extension of unemployment benefits has a very large positive effect in recessions characterized by heightened risk of long-term unemployment, a defining feature of the current recession.

2 The Credit Market Effects of the Home Mortgage Interest Deduction¹¹

2.1 Introduction

The US home mortgage interest deduction (HMID) is a large tax expenditure used to subsidize debt-financed home ownership. It is also one of the most heavily criticized tax-expenditures in the United States. Critics argue that it is expensive (estimated to cost almost \$100 billion in 2011), highly regressive, distortionary, ineffective at increasing home ownership, and bad for the environment. Proponents of the subsidy have thus far touted weak benefits, noting that home owners are more likely to maintain their homes, garden, vote in local elections, and deter crime. The vast majority of cost-benefit analyses, such as Glaeser and Shapiro (2003) and Gahvari (1985), conclude overwhelmingly that these externalities are swamped by the large costs of the exemption.

In this paper we explain a heretofore unidentified benefit to the subsidy, namely that it rewards good credit and continued access to the credit market. By doing so, it reduces agents' incentive to default and eases credit constraints. We show that this benefit is sizeable in a calibrated model of the US economy and meaningfully alters any cost-benefit analysis.

To explore this effect, we use a life-cycle model in which individuals have limited commitment (they can default) and where the penalty for defaulting on one's debt is the loss of assets and exclusion from the credit market. In this model, the incentive for individuals without seizable assets to repay their obligations is limited to the value of continued access to the credit market. Acknowledging the lack of commitment, lenders prevent borrowers from obtaining a level of debt that might lead to default. This means that a number of consumers will face binding credit constraints relative to the full-commitment outcome and will suffer a corresponding welfare loss. This constraint will bind most severely on assetless agents who place low value on future credit access, or assetless agents with high and growing incomes.

In this model (and in reality) the HMID is a deduction for *high income*, young (read *low asset*) people to buy a *collateralizable asset* with debt. Individuals are only eligible for this subsidy if they can obtain a mortgage, or if they have not been excluded from the credit markets by a prior default. After purchasing a home, individuals do not default due to a desire not to lose their houses. Thus the home mortgage interest is a conditional subsidy that provides strong incentives to maintain one's credit and is targeted towards agents most affected by the problem of limited commitment. The apparent regressivity of the deduction – that it targets the high income – is a feature, not a bug. We calibrate this model and explore the welfare

¹¹Joint with Danny Shoag.

and credit market effects of eliminating or reducing the subsidy. We find that the HMID enhances welfare along two dimensions. The realized income elasticity of housing demand at the upper end of the income distribution is less than one in our calibration. This means that although the HMID appears regressive, it actually makes progressive net transfers. In addition to this redistributive effect, the HMID loosens borrowing constraints by allowing agents to overcome commitment problems. The welfare benefits of both channels are concentrated on the poor, especially low-income homeowners.

The credit channel proposed here also offers a partial solution to the long standing puzzle of why more people do not file for bankruptcy. Though bankruptcy rates have risen considerably in the past two decades, the relatively generous bankruptcy laws of most states mean that filing is in the narrow economic interests of many more people than actually file. Estimates of the share of households that would be better off filing for bankruptcy range from 15% (White (1998)) to 50% (Cohen-Cole and Duygan-Bump (2009)). Researchers have reconciled these figures with observed bankruptcy rates by arguing that there is a stigma against filing (Fay, Hurst, and White (1998)), by referring to the option value of filing later (White (1998)) and by citing ignorance of bankruptcy law (Cohen-Cole and Duygan-Bump (2008)). The channel outlined in this paper can explain some of this 'under-filing' in a rational optimizing framework.

After exploring the credit market effects of the HMID in a calibrated model, we present suggestive empirical evidence documenting the link between the HMID and credit market access. We find that cities where house prices are higher relative to income (or where the HMID is more valuable) have higher average credit limits. Similarly, we find that cities where rental options are less attractive also have higher credit. This fits nicely with the channel outlined in the model, in that agents have more incentive to maintain good credit when the conditional subsidy is larger.

We also demonstrate that changes in state and federal tax laws affecting the HMID are correlated with bankruptcy rates. Using the NBER TAXSIM data developed by Feenberg and Coutts (1993), we calculate the benefit of the HMID for a fixed representative family across each state year. Since the representative family does not vary across states and years, changes in the benefits being measured are driven solely by changes to tax law. We then document how changes in the value of this subsidy are related to state level bankruptcy rates in both the cross section and time series. Consistent with the model, we find that larger subsidies meaningfully reduce default.

The paper proceeds as follows. In section 2.2, we present a simple model that outlines the basic intuition that the HMID can help overcome distortions caused by limited commitment. We also present an elaboration of this simple model, which demonstrates the robustness of the result even when the HMID cannot be funded by universal taxation. Section 2.3 presents evidence to support two crucial assumptions in our paper, namely that the HMID is large enough to affect incentives and that benefitting from the HMID is

partially contingent on one's default history. In section 2.4, we outline the full model and explain the calibration. Section 2.5 describes the numerical solution algorithm. Section 2.6 presents the results from the calibration and shows that the HMID leads to modest welfare gains, concentrated at the bottom of the income distribution. In section 2.7, we present suggestive empirical evidence that the channel outlined in this paper plays a role in regional credit markets. Section eight concludes.

2.2 Simple model

2.2.1 Basic model

To convey the intuition behind the general model, we first begin by considering a simple OLG endowment economy where each agent lives three period. Households have a concave period utility function, and for simplicity, agents do not discount utility between periods. Household endowment streams are given deterministically by $\{e_1, e_2, e_3\} = \{v, V, v\}$, where $v < V < 4v$. Households have access to a one-for-one storage technology across periods.

Agents are allowed to borrow from a bank that costlessly intermediates between households. Agents cannot commit to repayment, however, and cannot collateralize their endowment streams. There is no asymmetric information, and the banks refuse to lend to agents who will decide to default ex-post.

Given this set-up, the utility achieved by the households will be:

$$U^{No\ Commitment} = u(v) + 2u\left(\frac{V+v}{2}\right) < U^{Full\ Commitment} = 3u\left(\frac{V+2v}{3}\right).$$

Now suppose that households may purchase a collateralizable asset (a house), which returns a stream of consumption goods equal to $\frac{V+2v}{3}$ in both the period it is built and in the following period. Construction costs for the asset are $\frac{2}{3}(V+2v)$. We show that without a subsidy, the existence of such an asset is irrelevant.

Lemma 1 *Without a subsidy, the existence of the asset is irrelevant to the determination of consumption and welfare is unchanged.*

Proof. *Agents in the second period of life are indifferent, as they could achieve the same stream of consumption without the asset. If agents in the third period of life have funds less than $2(V+v)/3$, they cannot commit to repaying the amount they would need to borrow. If they do have these funds, they are indifferent to buying and reselling the asset. Lastly, agents in the first period of life cannot commit to repaying $(2V-v)/3$, the amount they need to borrow to finance construction, since the house will only be worth $2(V+v)/3$ in period 2. In other words, the house would be "under-water" the next period. ■*

Now suppose that there exists a subsidy $x > \frac{V+v}{3}$, $x < V$ which is given to households if they purchase a home using debt. The subsidy is financed by a lump sum tax evenly divided on all agents in period 2. We show that this subsidy induces the agent in only the second period of life to purchase the asset.

Lemma 2 *With the subsidy, agents in the second period of life purchase the asset and household welfare achieves the full commitment outcome.*

Proof. *The analysis for agents in the last period of life remains the same as before. Agents in the second period of life must choose between repaying their debts and purchasing a subsidized house if it is possible or not repaying and forgoing the subsidy. Note that no agent will repay and then either choose not to purchase a house or find it impossible to borrow a sufficient sum to do so. In those cases, repayment is strictly dominated. The benefit to repayment when purchasing will be an option, then, is $2u\left(\frac{V+2v}{3}\right)$. The benefit to default is $2u\left(\frac{V-x+v}{2}\right)$. This means that when $x > \frac{V+v}{3}$, the agent finds it optimal not to default. Given this, the agent can borrow in the first period up to the point where he will still be able to purchase a home in the second period. This means that his total net worth (absent the equivalent tax and subsidy) in period two must equal the price of the house, $V + v - b = \frac{2}{3}(V + 2v)$. This means that the maximum sustainable debt in period 1 is $\frac{V+2v}{3} - v$, or exactly enough to ensure perfect consumption smoothing as in the full commitment case. ■*

This simple model conveys the basic intuition. By subsidizing access to the credit market, the HMID reduces the frictions imposed by limited commitment and can increase welfare.

2.2.2 Redistribution model

One feature of the previous model is that agents are required to pay the tax used to finance the subsidy, whether or not they choose autarky. Though this is indeed one of the ways in which the HMID actually operates, it is important to show that this aspect of the model is not intrinsic to the underlying channel. Once again we consider a model in which agents live for three periods, however now we allow two types of agents. Each agent is endowed with an income stream which is either high (type H) or low (type L) with equal probability. The high income stream is given by $y_h = \{1, V^h, v\}$ and the low income stream is given by $y_l = \{1, V^l, v\}$, where $V^H > V^l > v \gg 1$. Types are revealed in period two, and unlike before, we let the period three endowment be collateralizable¹².

In addition to this process, agents may also receive an 'expenditure' shock at date 2 with probability μ . This shock requires them to forfeit X of the good at date 2, provides no utility, and is independent of the agent's type. Conceptually, this shock could represent a large medical expenditure, divorce, car accident or accidental pregnancy. I assume that $\frac{V^H-v}{2} > X > \frac{V^l-v}{2}$. The within-period timing at date 2 operates as

¹²This is required for the simple three period model, but can be relaxed in an infinite horizon version.

follows: first agents must decide whether or not to repay their loans from date 1, then the expenditure shock is realized, and then agents can decide how much to borrow (if they are able to do so) for consumption at date 2.

We once again consider a world with limited commitment and costless savings technology prior to the introduction of a collateralizable asset. Once again there is a bank that can intermediate borrowing and saving across households, but the bank is unable to lend to an agent that has defaulted. Further, the bank refuses to lend to an agent that may choose to default in period two.¹³

Given this setup, the rich agent has no value to having access to the credit markets in period two. The size of the shock is less than the optimal level of saving, and hence the high income agent can still perfectly smooth consumption between periods one and two. This is not true for the low income agent, who benefits from being able to borrow against period-three income in the event of a shock.

Consumption if there were no credit markets:

$$\{c_1, c_2, c_3\} = \begin{cases} \{1, \frac{V^H+v}{2}, \frac{V^H+v}{2}\} & \text{when high type, no shock} \\ \{1, \frac{V^H+v-X}{2}, \frac{V^H+v-X}{2}\} & \text{when high type, shock} \\ \{1, \frac{V^L+v}{2}, \frac{V^L+v}{2}\} & \text{when low type, no shock} \\ \{1, V^L - X, v\} & \text{when low type, shock} \end{cases}$$

Consumption with anonymous, no-default, limited commitment credit markets:

$$\{c_1, c_2, c_3\} = \begin{cases} \{1, \frac{V^H+v}{2}, \frac{V^H+v}{2}\} & \text{when high type, no shock} \\ \{1, \frac{V^H+v-X}{2}, \frac{V^H+v-X}{2}\} & \text{when high type, shock} \\ \{1, \frac{V^L+v}{2}, \frac{V^L+v}{2}\} & \text{when low type, no shock} \\ \{1, \frac{V^L+v-X}{2}, \frac{v+V-x}{2}\} & \text{when low type, shock} \end{cases}$$

In this economy, all agents under-consume relative to the optimum in the first period of life. This is because the high income agent has no value to maintaining access to the credit markets. Because agent types are unknown, the assumptions on bank behavior mean that no loans will be provided to agents in the first period.

Once again, we introduce a durable and collateralizable house that returns an actuarially fair stream of consumption. To keep this simple and analogous to the earlier example, we let the house product consumption good equal to $\frac{V^H+v-X}{2}$ for two periods at a cost $V^H + v - X$ to produce. We assume that $V^L + X < V^H$, so that only high types can purchase the house. As in the previous case, the asset will have no impact on

¹³The results derived below could be derived under different assumptions. We explore this issue in an appendix.

the equilibrium without a subsidy, as the only agent who can afford it (period 2, high, no shock) can already perfectly smooth between periods 2 and 3.

Now suppose there is a conditional subsidy x from the low type to the high type for the purchase of that asset when financed by debt. Given our assumptions on bank behavior, the high-type can only borrow the funds needed to purchase the asset provided he has not defaulted. This means that the high type is indifferent when

$$\begin{aligned}
& 2\mu u\left(\frac{V^H + v - X - b + x}{2}\right) + 2(1 - \mu) u\left(\frac{V^H + v - b + x}{2}\right) \\
& \qquad \qquad \qquad \text{Repayment payoff} \\
= & 2\mu u\left(\frac{V^H + v - X}{2}\right) + 2(1 - \mu) u\left(\frac{V^H + v}{2}\right). \\
& \qquad \qquad \qquad \text{Default payoff}
\end{aligned}$$

As is clear from the equation, the high type will prefer repayment for any $x > b$.

Given our assumption about the inability to tax autarky, the low type could choose to avoid paying the subsidy. However, unlike the high type, the low type values the credit market in the second period to help smooth consumption. This means that, conditional on having borrowed b , the low type will choose to repay his debts and the subsidy as long as the insurance value of the credit market exceeds the income loss from repayment.

This framework means that the agents' borrowing in period 1 will be constrained by their inability to commit to repayment at date 2. As demonstrated above, the poor agent does value his access to the credit market. Therefore, if agents' types were known at date 1, he could borrow (b) up to the point where

$$\begin{aligned}
& 2(1 - \mu)u\left(\frac{V^l + v - 2b}{2}\right) + 2\mu u\left(\frac{V^l + v - 2b - X}{2}\right) \\
& \qquad \qquad \qquad \text{Repayment Payoff} \\
= & 2(1 - \mu)u\left(\frac{V^l + v}{2}\right) + \mu u(V^l - X) + \mu u\left(\frac{V^l + v}{2}\right). \\
& \qquad \qquad \qquad \text{Default payoff}
\end{aligned}$$

This implies that there exists $b^* > 0$ for which this holds when X is sufficiently close to V^l . In other words, there will be a subsidy that the government can take from the non-defaulting low-types to subsidize the high type's credit use that will not force the low type into autarky.

With the subsidy, the high type has an incentive not to default on debt carried into period two, lest he lose his opportunity to buy the house. That means that there is a positive level of debt which neither agent will choose to default on in period two. Since neither agent will default on some level of debt b^* , the agents can borrow up to b^* in period one. If we assume that the endowment is sufficiently low in

period one, the benefits from making credit available where it was otherwise constrained may increase the utility of the low as well as the high type. In other words, though the subsidy is regressive and the benefits disproportionately favor the high type, some positive subsidy may nonetheless be Pareto optimal. Under some circumstances, the regressivity of the HMID may be a necessary feature of a welfare-improving subsidy.

2.3 Assumptions check

The channel described above relies on two important assumptions, namely that agents care about the HMID and that they would be less likely to benefit from it after a default. Before calibrating a full model to estimate the broader effects of the subsidy, it is worthwhile to demonstrate that these two assumptions are empirically reasonable.

2.3.1 The HMID is large in size

A number of papers have attempted to calculate household tax savings from the HMID. Poterba and Sinai (2008) calculate average savings using data from the 2004 SCF. Their estimates are presented below in Table 2.1.

Table 2.1. Average HMID deduction by age and income.

Age	<\$40K	\$40-\$75K	\$75-\$125	\$125-\$250K	\$250K+
25-35	212	571	1,801	3,468	\$7,711
35-50	244	747	1,525	3,534	\$6,575
50-65	161	491	1,034	2,095	\$5,741
65+	21	178	329	981	\$1,322

As is evident from the table, the annual savings for wealthier and younger households are considerable (between 1-2% of income). To get a sense of the value of the HMID over the course of a thirty-year mortgage, we calculate the savings on a \$200,000 mortgage with a 6.8% interest rate and a deductible tax rate of 30%. In this example, the first-year savings are \$13,535, and the total savings over the course of the mortgage are \$80,816. These calculations clearly indicate that the HMID is large enough to plausibly influence agents' incentives in the credit market.

2.3.2 Credit access is restricted in bankruptcy

A second central assumption in this paper is that the HMID is a conditional subsidy in that it is more difficult to receive after a default. The empirical relevance of this channel therefore hinges on whether or

not it is indeed more difficult to obtain credit following a default. This issue has been studied in a number of papers. Cohen-Cole, Duygan-Bump, and Montoriol-Garriga (2009) find that most individuals are able to obtain credit cards within two years of filing bankruptcy, though this does not seem true for larger loans. Fisher, Filer, and Lyons (2004), using data from the Survey of Consumer Finances, find that after controlling for a large number of demographic and financial variables, individuals who have declared bankruptcy in the past five years are 30-40% more likely to be turned down for a mortgage and 97% more likely to be turned down for a non-mortgage bank loan. These effects persist in the SCF data, and there remains a 50% increase in the probability of having been turned down for a loan 11-15 years later.

We test this result in the latest (2007) version of the SCF and find similar results. Table 2.2 demonstrates that, even after non-parametric controls for income, the probability of obtaining a mortgage, bank loan or owning your home drops dramatically after bankruptcy.

Table 2.2. Credit and prior bankruptcy.

	(1A)	(2A)	(3A)	(1B)	(2B)	(3B)
	Turned Down	Did Not Apply	Own Home	Turned Down	Did Not Apply	Own Home
Prior Bankruptcy Filing	0.83*** (.07)	0.65*** (.07)	-0.48*** (.07)	0.67*** (.08)	0.49*** (.07)	-0.23*** (.07)
Income quintile controls	No	No	No	Yes	Yes	Yes

The decision to file for bankruptcy is a highly endogenous one, however, and these results cannot be interpreted as causal. It is possible that the decision to file for bankruptcy is correlated with some trait observable by lenders but not by an econometrician. Therefore it is impossible to rule out that these traits, and not the bankruptcy itself, are driving the results.

Musto (2004) recovers a potentially exogenous source of variation by looking at the amount of credit available to households around the 10 year post-bankruptcy mark, or the period after which a bankruptcy is legally removed from one's credit report. Musto finds a significant increase in the amount of credit used by households immediately after the ten year mark, suggesting that the bankruptcy itself was responsible for constraining credit.

We take a different approach to this question. We use a unique data set of loan-level data from an online peer-to-peer lending company named The Lending Club. The Lending Club is a website that matches borrowers and lenders. Since opening in 2007, the website has made more than \$250 million in loans. Many

of the site's applicants have poor credit histories and struggle to get credit through traditional channels. Lenders screen loan applicants using basic data verified by the website (credit score, court records, etc.) and from online discussion with the borrower. The website posts detailed data on all partially or fully funded loan applications.

One nice feature of this data is that the econometrician is privy to all the verified data used by lenders to screen loans. Because the identity of the borrower is confidential to the lenders, there is no worry that unobserved traits correlated with bankruptcy might be confounding the estimated effect of a prior bankruptcy on credit availability. We use only data with available delinquency and public record data, and we code any agent with a delinquency or public record as a default. Restricting the sample in this way yields data on 16,887 loans of which, consistent with the target audience of the site, 92% had delinquency or bankruptcy on record.

Table 2.3 shows the effect of previous defaults have on credit access in this market. In column 1, the point estimate is small, indicating that those with prior defaults see the percentage of their request funded fall by 4 percentage points. This estimate is biased, though, in that the site targets risky borrowers. Those with no prior defaults appear to be more risky borrowers along other recorded dimensions (e.g. a debt to income ratio that is 1.2 percentage points higher). When we control for the range of borrower characteristics listed, either linearly where possible or by using 25 quantile dummies, we find that a prior bankruptcy still significantly reduces the funds obtained by borrowers. The baseline estimates indicate a reduction in funding for a given request by 17-18 percent. This is further evidence that a prior bankruptcy does, in fact, reduce the ability to borrow.

Table 2.3. Default and loan approval.

	(1)	(2)	(3)	(4)
	Percent Funded	Percent Funded	Percent Funded	Dollars Funded
Prior Default	-.04*** (.01)	-0.17*** (.01)	-0.18*** (.01)	-\$1,147*** (162.5)
Controls	None	Linear	25 Quintiles	25 Quintiles
R ²	0	.13	.15	.24
Obs	16,887	16,868	16,868	16,868

The variables being controlled for are Debt to Income, Monthly Income, Amt Requested, Home Ownership, College Education, Loan Length, FICO Range, Employment Tenure, Year, Loan Purpose, and State of residence. Robust SE in parenthesis.

2.4 Model

We consider a simple life-cycle model with uninsurable income risk and an option to declare bankruptcy. A unit continuum of households i are divided over T overlapping generations, $t = \{1, 2, \dots, T\}$. The probability of surviving from generation t to $t + 1$ is denoted ϕ_t , with ϕ_T corresponding to the probability of remaining in cohort H . The population shares of each generation remain constant and there is no population growth.

Households receive stochastic endowment income $y_{i,t}$,

$$\ln y_{i,t} = \ln z_{i,t} + \ln \varepsilon_{i,t} \quad (22)$$

$$\ln z_{i,t} = \ln \kappa_t + \ln z_{i,t-1} + \ln \eta_{i,t} \quad (23)$$

where $z_{i,t}$ is the permanent component of the income process with drift κ_t capturing the life-cycle trend of income and permanent shock $\ln \eta_t \sim N(0, \sigma_\eta^2)$, and $\varepsilon_{i,t}$ is a transitory shock with age-specific variance, $\ln \varepsilon_t \sim N(0, \sigma_\varepsilon^2)$. This decomposition of the earnings process into a transitory and permanent component is standard in the literature and in this particular formulation follows Storesletten, Telmer, and Yaron (2004) and Gourinchas and Parker *Econometrica* 2002.

Period T represents retirement. Households remain in period T (alive) with probability ϕ_T . During each period of retirement, households received a fixed income of y_T plus the annuitized value of their

current savings, an actuarially fair value derived from ϕ_T .

During each generation of the lifecycle, households can spend their endowment income on a consumption good $c_{i,t}$, a housing asset $h_{i,t}$, and one-period nominal bonds $a_{i,t}$ available in zero net supply that pay 1 in period $t + 1$. Period-utility is identical across households and cohorts and takes the form

$$U(c_{i,t}, h_{i,t}) = \frac{c_{i,t}^{1-\gamma^c}}{1-\gamma^c} + \lambda \frac{h_{i,t}^{1-\gamma^h}}{1-\gamma^h}, \quad (24)$$

where $\lambda > 0$ parameterizes the flow utility of a household's housing assets. Buying or selling a unit of housing incurs a linear adjustment cost of ζ , capturing the relative illiquidity of housing compared to the bond. Households earn mortgage interest tax credit $\theta r \max(0, \min(h_{i,t}, -a_{i,t}, H))$. That is, a household with positive housing $h_{i,t} > 0$ and debt $a_{i,t} < 0$ can deduct a fraction θ of the interest on its outstanding debt up to the value of its housing units, limited by H , the maximum home value to which the mortgage interest deduction applies. If a household has positive assets a_t , it cannot make a deduction.¹⁴ The tax credit is financed by a progressive income tax function $\tau(\cdot)$ that exactly funds the cost of the credit.

We assume that bonds are traded by perfectly competitive financial intermediaries that face transaction costs proportional to the size of the loan. This introduces a wedge between the deposit rate r^D and the borrower rate r^B , where

$$\begin{aligned} r &= r^D \text{ when } a_{i,t} > 0 \\ r &= r^B \text{ when } a_{i,t} < 0. \end{aligned}$$

As in Eaton and Gersovitz (1981), participants in lending markets are able to strategically default, repudiating outstanding debt at the cost of future exclusion from financial markets. At the beginning of each phase of the life cycle, agents learn the new value of their endowment and can choose to default on their outstanding debt $a_{i,t-1}$. Default resets $a_{i,t}$ to 0, but excludes the household from borrowing or lending for the next E generations, with the remaining periods of exclusion denoted $e_{i,t}$. If the agents owns housing, default causes it to be liquidated. In this sense, housing is a collateralizable asset.

We use the notation of Aguiar and Gopinath (2006) in defining the value functions. The value function of households with a good credit history is denoted V^G , while V^B denotes the value function of households that remain excluded from borrowing and lending. We can write the value of having a good credit history

¹⁴In reality, many home-owners have positive net savings as well as mortgage debt. For simplicity and computational feasibility we consider only the two-asset case.

as

$$V(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t-1}, h_{i,t-1}) = \max \left(V^G(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t-1}, h_{i,t-1}), V^B(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, E) \right). \quad (25)$$

If a household chooses to default, its asset and housing levels are set to zero and it is excluded from credit markets. In each subsequent period, it has a probability r of redemption. The value of having a bad credit history is

$$V^B(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t) = \max_{c_{i,t}} \left\{ U(c_{i,t}) + \beta \phi_t E_t V^B(\alpha_i, z_{i,t+1}, \varepsilon_{i,t+1}, t+1) \right\}. \quad (26)$$

When excluded from financial markets, the household simply consumes its entire endowment, $c_{i,t} = y_{i,t}$. When $e_{i,t} = 0$, the household's credit market standing is redeemed and $V^B = V(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, 0, 0)$. Redemption here corresponds to the "fresh start" system of bankruptcy used in the United States discussed in Livshits, MacGee, and Tertilt (2007).

Households that choose to remain in good standing choose $c_{i,t}$, $a_{i,t}$, and $h_{i,t}$ to maximize the value function

$$V^G(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t-1}, h_{i,t-1}) = \max_{c_{i,t}, a_{i,t}, h_{i,t}} \left\{ U(c_{i,t}, h_{i,t}) + \beta \phi_t E_t V(\alpha_i, z_{i,t+1}, \varepsilon_{i,t+1}, t+1, a_{i,t}, h_{i,t}) \right\} \quad (27)$$

subject to the budget constraint

$$\begin{aligned} c_{i,t} + \frac{1}{1+r} a_{i,t} + h_{i,t} + \zeta |h_{i,t} - h_{i,t-1}| + \tau \max(0, (y_{i,t} - \theta r \max(0, \min(h_{i,t}, -a_{i,t}, H)))) \\ = y_{i,t} + a_{i,t-1} + (1 + \sigma) h_{i,t-1} \end{aligned} \quad (28)$$

and the borrowing constraint

$$a_{i,t} > A, \quad (29)$$

where ζ is the housing adjustment cost and σ is the growth rate in house prices.

We now present two versions of the core model outlined thus far. The first version, hereafter Model 1, follows Kehoe and Levine (1993) in endogenizing the borrowing constraint by imposing a strict never-default rule. Model 2 follows Chatterjee, Corbae, Nakajima, and Ríos-Rull (2007) in allowing households to borrow at an interest rate proportional to their likelihood of default, a function of the household's current state and the size of the loan. The two variants of the model differ in the values of r^B and A .

2.4.1 Model 1: Endogenous borrowing constraint

We first consider an endogenous borrowing constraint as in Kehoe and Levine (1993). The borrower rate $r^B = R^B$ is assumed fixed for all households. Households are only allowed to borrow up to debt levels for which they would not default for *any* draw of the stochastic endowment process. In particular, $A(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, h_{i,t})$ is defined implicitly as follows:

$$\begin{aligned} A(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, h_{i,t}) &= \min a_{i,t} \text{ such that} \\ V^G(\alpha_i, z_{i,t+1}, \varepsilon_{i,t+1}, t+1, a_{i,t}, h_{i,t}) &> V^B(\alpha_i, z_{i,t+1}, \varepsilon_{i,t+1}, t+1) \quad \forall (z_{i,t+1}, \varepsilon_{i,t+1}) \end{aligned} \quad (30)$$

Note that the borrowing constraint depends on the current choice of housing $h_{i,t}$ meaning that housing is collateralizable, accounting for adjustment costs.

2.4.2 Model 2: Variable interest rate

In the second version, we adopt the variable interest rate framework of Chatterjee, Corbae, Nakajima, and Ríos-Rull (2007). Now the borrower interest rate $r^B(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t}, h_{i,t})$ depends on the probability of default,

$$1 + r^B(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t}, h_{i,t}) = \frac{1 + R^B}{1 - d(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t}, h_{i,t})}, \quad (31)$$

where $d(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t}, h_{i,t})$ is the likelihood of default implied by the stochastic endowment process and the value functions V^G and V^B . When there is zero likelihood of default, $r^B = R^B$ as in Model 1.

Note that in this version, A exists only because the income process is calibrated to have a finite upper bound in the numerical approximation.

2.4.3 Equilibrium

Similar to Livshits, MacGee, and Tertilt (2007), an equilibrium of the model is defined as a set of value functions V , V^G , and V^B , policy rules $a(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t-1}, h_{i,t-1})$ and $h(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t-1}, h_{i,t-1})$, and default probability function $d(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t}, h_{i,t})$ (in the case of Model 2 only) such that:

1. The value functions satisfy (25), (26), and (27) subject to (28) and (29).
2. The policy rules a and h correspond to those value functions.
3. (Model 1) The borrowing constraint $A(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, h_{i,t})$ is determined by (30).

4. (Model 2) The interest rate function $r^B(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t}, h_{i,t})$ satisfies (31) and the default function $d(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t}, h_{i,t})$ assigns the correct likelihood of repudiating debt corresponding to the value functions.

2.5 Numerical solution

We solve both models numerically using simple algorithms. Model 1 is solved in the following steps:

1. Guess a tax rate τ that will raise enough revenue to cover the cost of the HMID and the fixed tax rebate.
2. Guess a default matrix D_n defined on $(\alpha_i, z_{i,t+1}, \varepsilon_{i,t+1}, t, a_{i,t}, h_{i,t})$ that indicates where agents choose to default.
3. Given D_n , solve V^G . This does not require knowledge of V^B because the strict never-default lending rule prevents households in good credit standing from ever defaulting.
4. Given V^G , solve V^B .
5. Update D_{n+1} to 1 where $V^B(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t) > V^G(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t-1}, h_{i,t-1})$ and 0 elsewhere.
6. Repeat (2) – (4) until convergence to D_N .
7. "Infants" are born with $t = 1$, $a = 0$, $h = 0$, and α , z , and ε distributed according to their steady-state values in a number that exactly replace the share of the population that dies. Then use the policy rules $a(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t-1}, h_{i,t-1})$ and $h(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t-1}, h_{i,t-1})$ extracted from $V^G(D_N)$, the stochastic endowment process, and the mortality process to solve for the steady-state population distribution.
8. Use the population distribution to calculate the total cost of the HMID. If government revenues and government spending differ by more than 0.1% of GDP, update τ to exactly finance spending and return to step 1. Repeat until convergence.

Model 2 is solved through a similar set of steps:

1. Guess a tax rate τ that will raise enough revenue to cover the cost of the HMID and the fixed tax rebate.
2. Guess a default probability matrix $d_n(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t}, h_{i,t})$ and use it to calculate the interest rate function $r^B(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t}, h_{i,t})$.

3. Guess V^G and V^B and use the guesses to calculate $V = \max(V^G, V^B)$. Update the guesses and iterate until the value functions converge.
4. Use the stochastic income process and the policy rules $a(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t-1}, h_{i,t-1})$ and $h(\alpha_i, z_{i,t}, \varepsilon_{i,t}, t, a_{i,t-1}, h_{i,t-1})$ corresponding to the value functions to calculate a new default matrix. Return to step 2 and repeat (2) – (4) until convergence to d_N .
5. Follow steps (7) and (8) above.

2.5.1 Calibration

We calibrate the model as follows. We use mortality data from the Center for Disease Control’s Life Tables. Age is defined from 21 to 65, with $T = 65$ representing retirement. Each "generation" in the life-cycle therefore covers 10 years. Because we use 10-year periods, all annualized parameter estimates given below must be correspondingly adjusted.

The income process is parameterized using estimates from the literature on idiosyncratic income risk. In the current solution we leave the fixed effect α_i at 0 for all agents. We calibrate the age-trend of income κ_t using data from the Consumer Expenditure Survey. Specifically, we define income as the sum of wages, business income, farm income, pensions, pension, social security, unemployment insurance, worker’s compensation, welfare, and food stamps. We take the average income by age over ages 21-64 for households that were full income reporters, and define the retirement income y_T as the average income for ages 65 and older. We follow Gourinchas and Parker in setting $\sigma_\eta^2 = 0.0212$ and $\sigma_\varepsilon^2 = 0.0440$. These figures are consistent with several empirical studies on idiosyncratic income risk. We approximate the stochastic endowment process with a transition probability matrix following the Tauchen method.

In the household utility function, we use the estimates in Livshits, MacGee, and Tertilt (2007) of an annualized discount rate of 0.94 and $\gamma^c = 2$. We also use their estimates of the an annual deposit interest rate $r^D = 0.04$ and a borrower rate $r^B = 0.08$ for zero probability of default.

The housing parameters are taken from Lustig and Van Nieuwerburgh (2004). We set λ to match the aggregate ratio of spending on housing and consumption goods in the US economy. We set the deductible share of mortgage interest $\theta = 1$ and set $H = \$1,100,000$, the maximum home value including home improvements on which households can deduct interest under US law. The linear adjustment cost ζ is set to 0.065, reflecting an average 5% real estate brokerage fee plus 1 – 2% in taxes.

We set the probability of redemption to $r = 0.1$, reflecting an average 10-year period of exclusion from credit markets.

Finally, we calibrate the income tax function τ using US average tax rates for each node on the income grid. Because the cost of the HMID is modest, we add an additional 20% to the tax rate and rebate the total lump-sum in order to raise the marginal tax rate to a realistic level. Combined with the HMID, this roughly approximates the share of taxes in GDP in the US economy.

We define the permanent component of the income grid over the range 0.2-1, corresponding to a top income of \$150,000. The transitory component is defined as in Livshits, MacGee, and Tertilt (2007), with 80% of the population receiving no shock and 10% receiving positive and negative shocks, respectively. We also use their definition of expenditure shocks, assigning a small shock of \$32,918 corresponding to minor medical expenses or divorce or an unplanned pregnancy with probability 0.07104 and a large shock corresponding to major medical expenses of \$102,462 with probability 0.0046. We define the asset grid over the range -1 to 1, with a larger share of the grid points concentrated near 0. The lower bound for assets is low enough that it always exceeds the borrowing constraint in absolute value. The upper bound is high enough that almost no agents reach it. The housing grid has a lower bound of 0, an upper bound of H , and a midpoint of \$225,000, roughly the median home price in the US, around which nodes are disproportionately concentrated.

2.6 Results

2.6.1 Baseline results

Table 2.4 reports the baseline results with and without the HMID. The total tax cost of the HMID as a share of GDP, 0.659%, corresponds closely to the actual figure for the United States of about 0.7%. With the HMID, asset ownership shifts moderately away from non-housing wealth toward housing wealth and total wealth falls as the borrowing constraint is relaxed. Aggregate welfare is slightly higher with the policy (0.03%). The overall welfare gain reflects two changes caused by the policy: first, a redistributive effect caused by the unequal distribution of the tax costs and tax benefits of the policy over the income and age distribution, and second, the loosening of the borrowing constraint caused by the policy. Finally, the variance of period-utility across the population falls modestly. This effect can be similarly decomposed.

Table 2.4. Summary statistics.

	$\theta = 0$	$\theta = 1$
GDP	0.473	0.473
Tax rate	20%	20.66%
Total wealth	0.170	0.157
Aggregate non-housing wealth	-0.331	-0.365
Housing wealth	0.500	0.522
Tax cost of HMID / GDP	0%	0.66%
Total welfare	-7.874	-7.872
Population variance of period utility	1.441	1.440

Figure 2.1 plots the average borrowing constraint by age and shows that it loosens with the HMID. This loosening of the borrowing constraint is caused by the increased value of access to the credit market with the HMID. The gap widens in middle-age when more households have purchased homes, which are confiscated under default. Even in the first period of life, when households are born with no assets of either type, the borrowing constraint is looser under the policy because of the increased value of maintaining a good credit history. As households reach old age and sell their houses, the worlds with and without the policy become more similar.




Figure 2.1. Average borrowing constraint by age.

Figures 2.2, 2.3 and 2.4 show asset ownership over the life cycle. As expected, the HMID leads households to choose more housing and take on more non-housing debt to finance it. Net asset ownership is lower without the policy because of the loosening of the borrowing constraint: with greater ability to bor-

row in response to shocks, households have a weaker precautionary motive to save and fewer people are credit constrained. The figures also show the expected life-cycle properties of asset accumulation: households initially borrow to finance consumption and housing when they are young, but begin accumulating net positive assets as their income rises in middle age to prepare for retirement. At the end of their lives, households liquidate their housing to finance consumption both because they are credit constrained and because they do not want to die with positive assets.




Figure 2.2. Average total bond and housing assets by age.




Figure 2.3. Average bond ownership by age.




Figure 2.4. Average housing by age.

2.6.2 Tax incidence by age and income

Table 2.5 shows the tax incidence of the HMID by age and income. These tables can be compared with the estimates in Poterba and Sinai (2008), who perform the same exercise using US data. Like them, we find that the subsidy grows with income and drops off in retirement. The subsidy appears regressive

and has been widely derided for this property because the tax benefit grows considerably with income. However, the tax cost of funding the HMID also grows considerably with income. Even with a flat income tax, the increased weight of the HMID's tax burden on the rich mostly outweighs their higher share of its benefits. The table shows that the net tax benefit is actually progressive, even when abstracting from the credit market effects.

Table 2.5. Tax incidence of the HMID.

DEDUCTION							
	0.2	0.3823	0.523	0.6	0.677	0.8177	1
15-24	0	0	0	0	0	0	0
25-34	0.0007	0.0015	0.0022	0.0027	0.0034	0.0049	0.007
35-44	0.0014	0.0024	0.0033	0.004	0.0048	0.0061	0.0074
45-54	0.0027	0.0041	0.0052	0.0059	0.0064	0.0072	0.0072
55-64	0.0044	0.0058	0.0065	0.0067	0.0068	0.0067	0.0059
65-74	0.0053	0.0057	0.0054	0.0052	0.0049	0.0045	0.004
75-84	0.0053	0.0042	0.0036	0.0032	0.0029	0.0025	0.0021
85+	0	0	0	0	0	0	0
TAXES							
	0.0013	0.0025	0.0034	0.004	0.0045	0.0054	0.0066
NET TRANSFER							
	0.2	0.3823	0.523	0.6	0.677	0.8177	1
15-24	-0.0013	-0.0025	-0.0034	-0.004	-0.0045	-0.0054	-0.0066
25-34	-0.0006	-0.0011	-0.0013	-0.0012	-0.001	-0.0005	0.0005
35-44	0.0001	-0.0001	-0.0001	0	0.0003	0.0007	0.0008
45-54	0.0014	0.0016	0.0018	0.0019	0.002	0.0018	0.0006
55-64	0.0031	0.0033	0.003	0.0027	0.0023	0.0013	-0.0007
65-74	0.004	0.0032	0.002	0.0012	0.0004	-0.0009	-0.0026
75-84	0.004	0.0017	0.0001	-0.0007	-0.0015	-0.0029	-0.0045
85+	-0.0013	-0.0025	-0.0034	-0.004	-0.0045	-0.0054	-0.0066

Figure 6 shows the average net tax benefit of the HMID by age. The very young and very old face the highest net costs of the HMID because they do not own homes. Among homeowners, the net tax

benefit falls with income because the richest age categories pay a higher share of the tax cost. The greatest beneficiaries are low-income older groups that pay a small share of the HMID's tax cost but have not yet sold their homes.




Figure 2.5. Net of HMID tax benefit minus tax cost of HMID by age.

2.6.3 Welfare

In this section we first examine the welfare effects of the HMID over the life-cycle and across the income distribution and then decompose these effects into a redistributive component and a credit loosening component. We perform this decomposition by applying the borrowing constraints from the model without the HMID to the world with the HMID. Finding the equilibrium of the model where households get a tax credit for homeownership but do not experience any loosening of their borrowing constraints isolates the redistributive impact of the policy.

We first show aggregate welfare under the three regimes. Because the policy loosens borrowing constraints and is essentially progressive, it is welfare improving. The total welfare gains for the population are split roughly evenly between the redistributive component, reflected in the difference in the first two columns, and the loosening of the borrowing constraint, reflected in the difference between the second and third column.

Table 2.6. Welfare decomposition.

No HMID	HMID, no credit effects	HMID
-7.8737	-7.8726	-7.8716

Figure 2.6 shows the expected lifetime welfare of the first generation by initial income level. The graph shows that welfare is higher for most of the income distribution under the policy, and is only lower at the very top income brackets. Again, this reflects the higher tax burden faced by the rich needed to finance the HMID and the reduced benefits the rich receive from a looser borrowing constraint.




Figure 2.6. Value function of "infants" by initial income level.

Table 2.7 decomposes the changes in lifetime welfare into its two components. The benefit from income redistribution falls mostly on the lower half of the population and is quite large in magnitude, over 2% for the very poor. The benefit from the broadening of credit access is also large, especially for the poor, reaching 1% of welfare for the poorest category. This loosening of the borrowing constraint is actually costly for the richest income groups because the gains from increased access are swamped by the rise in the tax burden associated with wider homeownership.

Table 2.7. Percentage change in value function

	Benefit from Redistribution	Benefit from Credit Market
0.2957	2.12	1.02
0.3823	1.66	0.90
0.4586	1.34	0.74
0.523	1.08	0.59
0.5728	0.88	0.47
0.6	0.78	0.41
0.6272	0.67	0.35
0.677	0.42	0.24
0.7414	0.19	0.05
0.8177	-0.05	-0.04
0.9043	-0.36	-0.12
1	-0.64	-0.15

2.7 Suggestive Empirics

2.7.1 Credit Limits and Home Prices

The value of the HMID increases when house prices are high and where renting is unattractive. When prices are high relative to income, the amount that is saved through the deduction is greater, all else equal. Similarly when rental prices are high or when rental units are scarce, individuals will have a greater incentive to capture a subsidy for homeownership. Our model detailed the process through which the HMID expanded credit access. Therefore credit access should be greater in areas where the subsidy associated with the HMID is larger, or where house prices are high relative to income and rental units are expensive and scarce. In this subsection, we present evidence of this relationship. Although there are admittedly many potential confounding factors, the relationship provides suggestive evidence that the mechanism outlined above is empirically meaningful.

To test this proposition, we collected data on average credit limits for 258 metropolitan levels from Experian, a large credit reporting bureau. We then regress the ratio of average credit limit to median income on the ratio of median house prices to median income. The results of this baseline regression are presented in Column 1 of Table 2.8. The coefficient, which is statistically significant at the 5% level, indicates that if house prices rose by 100% of income then the average credit limit would likewise expand by 4.5% of median income.

Lefgren and McIntyre (2008) have shown that credit availability is heavily influenced by state bankruptcy and wage garnishment laws. To control for this, in column 2, we include state dummies. As demonstrated in the table, these controls have little effect on the estimated relationship. Lastly in column 3, we add 25 dummies corresponding to the percentile of median income within our sample. That the coefficient remains similar and highly significant after these flexible controls for state and income suggests that confounding factors along these dimensions are not driving the results.

In the second panel of Table 2.8, we explore the same relationships using geographic features to instrument for the ratio of housing prices to income. We use the instrument constructed by Saiz (2010) at the metropolitan level, which uses cross-city variation in amount of land geographically suited for development. These instruments are highly predictive of housing prices and offer a more plausibly exogenous source of identification. In Column 3, we repeat the same regression as in Column 1. The estimate effect is larger and highly significant, again confirming the prediction of the model. This effect remains after controlling for state and income dummies, as demonstrated in Column 4.

Table 2.8. Credit limits and relative home prices.

	(1)	(2)	(3)	(4)
	<i>Credit Limit</i> <i>Income</i>	<i>Credit Limit</i> <i>Income</i>	<i>Credit Limit</i> <i>Income</i>	<i>Credit Limit</i> <i>Income</i>
<i>Median Home Price</i> <i>Income</i>	.045** (.02)	.064*** (.01)	0.13*** (.04)	0.075*** (.02)
First Stage F			31	23.6
Controls	-	State Income Quintile	IV	IV State, Inc Qnt
Obs	258	258	258	258

Robust standard errors in parentheses

The second prediction discussed in this section is that agents in places where renting is less attractive will have greater access to credit because the the HMID is more important. We test this hypothesis using the same data as in Table 2.9. The first column demonstrates that when rental units are expensive relative to housing prices, credit limits are higher. The coefficient implies that were median rental payments to increase by 1% of median house values, credit limits would expand by 0.4% of income. This result holds even after controlling for state and income dummies, as demonstrated in Column 2. Column 2 also includes a variable designed to measure the availability of rental units, namely the share of rental units in the housing stock. When rental units are easier to find, the HMID will be less important for agents, and therefore we expect

credit limits to be reduced. This prediction is borne out in the data. Column 2 estimates that an increase in the share of rental units by 1% reduces credit limits by -0.5% of income.

Table 2.9. Credit limits and the rental market.

	(1)	(2)
	<i>Credit Limit</i> <i>Income</i>	<i>Credit Limit</i> <i>Income</i>
$\frac{\text{Median Monthly Rent}}{\text{Median Home Price}}$	41.0***	36.0***
	(2.7)	(3.7)
$\frac{\text{Rental Units}}{\text{Total Housing Units}}$		-.005*
		(.003)
		State
		Income Quintile
	258	258

2.7.2 Bankruptcy and Taxes - Time and Cross Section

The previous section suggested that areas where the subsidy associated with the HMID is larger tend to have higher credit limits with respect to income. In this section we explore a second prediction of the model, namely that areas with higher HMID subsidies should be associated with lower bankruptcy rates.

As noted above, the size of the HMID subsidy for any individual depends on a host of factors which independently affect credit market variables. These factors include things like age, income, asset holdings, etc. Identifying the impact of the HMID on bankruptcy rates therefore requires some finesse.

In addition to the clearly confounding factors mentioned above, another determinant of the size of the HMID subsidy is state-level tax rates. Income tax rates, brackets and exemptions vary across states and time and interact with the HMID in complicated ways. For example, thirty-four states allow some deduction for mortgage interest in state income taxes, while the remaining 16 states either do not have personal income taxes or do not permit an exemption. Changes in tax laws can be used to identify changes in the size of the HMID subsidy. While changes in tax law are not completely exogenous to the factors that would otherwise determine state bankruptcy rates, they do provide a relatively cleaner and more valid source of variation.

To measure the impact of tax law on the HMID, we create a representative household and calculate the households tax bill given two levels of mortgage interest (\$10,000 and \$7,000 in 2008 dollars).¹⁵ The

¹⁵The representative household earns \$85,000 - \$50,000 from the primary earner, \$30,000 from the secondary and \$5,000 from interest income. The household has one dependent child, contributes \$3,000 a year to a 401K and pays \$2,500 in property tax. The mortgage

difference between the two (the tax savings) is calculated for each state and year using the NBER TAXSIM V9 program. In the following graph, we plot the real dollar value of this tax savings against bankruptcy rates. Each observation is a state-year and the time period runs from 1980 to 2008.




Figure 2.7. Bankruptcy rates and tax savings from the HMID.

The clear negative relationship demonstrated in the graph indicates that bankruptcy rates were lower when the subsidy was larger and vice versa. A great deal of the variation stems from national changes in bankruptcy rates and tax law. Bankruptcy rates have been rising over this period, while the subsidy has declined due to factors like the 1986 tax reform. While the time-series information is suggestive, it is possible to partial out potentially confounding national shocks by using the panel dimension. The regression presented in Table 2.10 below does just that, by regressing the five-year percentage change in state bankruptcy rates on year dummies and the five-year percentage change in the tax savings. The coefficients displayed give an elasticity. The first column presents the results using changes from 1980-1985, 1985-1990 and so on up to 2005. The coefficient is negative, as expected, but not statistically significant. One potential problem with this procedure is that the final observation is the change from 2000 to 2005. Bankruptcy rates exhibited a dramatic jump in 2005 as households rushed to file in advance of the 2006 bankruptcy reform. When this problematic year is removed from the data, as in column 2, the estimated elasticity jumps to 1.6 and is significant at the 5% level. This indicates that bankruptcy rates are highly sensitive to the size of the HMID. To check that this result is not sensitive to years other than 2005, we rerun the regression using interest varies between \$10,000 and \$7,000 (or roughly the interest on a \$150,000 or \$100,000 mortgage). All figures are in 2008 dollars.

every possible starting point and five year division (i.e. 1981-1986, 1986-1991, etc., then 1982-1987, etc.). We then use the Fama-MacBeth procedure to combine these point estimates and calculate a standard error. The results, presented in column 3, are very similar to the results in column 2. This indicates that the baseline results, which show that the HMID has a surprisingly large effect on bankruptcy rates, are not being driven by the choice of starting point.

Table 2.10. Changes in bankruptcy rates and HMID savings.

	(1)	(2)	(3)
	$\frac{\Delta B_{t,t-5}}{B_{t-5}}$	$\frac{\Delta B_{t,t-5}}{B_{t-5}}$	$\frac{\Delta B_{t,t-5}}{B_{t-5}}$
$\frac{\Delta T_{t,t-5}}{T_{t-5}}$	-0.47	-1.60**	-1.77***
	(.40)	(.76)	(.44)
Sample	80-85,85-90,..2000-2005	80-85,85-90,..1995-2000	Fama-McBeth
Controls	Year	Year	24 5-year Cross Sections
Obs	255	204	24

The evidence presented in this section is by no means conclusive. While we feel that these points suggest that the channel outlined in this paper is empirically relevant, we acknowledge that more work can and should be done to measure this effect. We are in the process of collecting panel data on credit limits that will enable us to test the model's implications for credit constraints more directly. We are also in the process of using the TAXSIM identification procedure applied to zip code level bankruptcy filings. By creating representative households designed to match the distribution of households in each zip code, we can capture how changes in tax law determine the size of the HMID at the zip code level. We can then use these measures to explore changes in bankruptcy filings at the ZIP code level. By increasing the number of observation, this approach should allow us measure the desired elasticity more precisely.

2.8 Conclusion

The HMID has been widely derided as an ineffective and regressive tax expenditure. These criticisms, however, only consider the partial equilibrium effects of the subsidy. In this paper, we argue that the HMID has two important general equilibrium effects. The policy broadens credit access and makes the tax code more progressive. Both of these effects increase welfare modestly in a calibrated model of the US economy.

The mechanism for expanding credit access is intuitive. Households can only take advantage of the HMID if they can borrow in the credit markets. If default restricts access to credit markets in the future,

the HMID helps households overcome their commitment problems by creating an extra incentive to repay outstanding debt. As such, it meaningfully expands credit constraints. We present new evidence that default does indeed restrict future credit access using a unique data set of online loans. This data allows us to observe all the information available to the actual lenders, greatly reducing the problem of omitted variable bias.

In our baseline calibrations, the HMID increases the progressivity of the tax code (even with a linear tax rate) because the realized income elasticity of housing demand at the upper end of the income distribution is below one. This is consistent with the literature, which generally finds elasticity estimates between .5 and 1 (Malpezzi and Mayo (1987)). With a progressive income tax, this effect would be even larger. The apparent regressive nature of the deduction belies the fact that the net transfer is actually progressive under the baseline calibration.

These general equilibrium effects have a significant effect on the welfare of young and poor agents, increasing lifetime utility by 2-3% for those at the bottom of the distribution. These gains are balanced by losses at the top of the distribution, meaning that aggregate utility increases by only .03% in the baseline calibration.

In addition to our numerical results, we also present some suggestive empirical evidence that the HMID expands credit access. Using data at the MSA level, we show that cities with higher house prices relative to income (and hence higher HMID values) have higher credit limits. Similarly, cities where the HMID is valuable because renting is less attractive also have higher limit.

Lastly, we show a relationship between changes in state level tax laws pertaining to the HMID and state level bankruptcy rates. Using the NBER TAXSIM program, we calculate the tax savings associated with the policy for a sample household with fixed properties across state years. We find that increases in this tax savings are associated with decreases in bankruptcy rates. Ongoing work is under way to establish a similar relationship between zip code level bankruptcy rates and changes for a fixed zip-code specific representative household's tax saving.

To conclude, we find the apparent partial equilibrium effects of the HMID differ greatly from its effects in a calibrated model with limited credit market commitments. Though more research is certainly needed, these results suggest that cost-benefit analyses of the HMID must consider the effects on the tax code and credit markets before advocating for its repeal.

In future extensions of this work, we plan to evaluate a policy frequently advocated by critics of the HMID, namely that the HMID be replaced with a flat one-time homeownership credit better targeted at boosting homeownership rates. As our model makes clear, such a proposal would significantly change the credit-market effects of the HMID.

3 Uncertainty and the Geography of the Great Recession¹⁶

3.1 Introduction

Macroeconomists have advanced a number of hypotheses to explain the severity of the 2007-2009 recession and the persistent decline in employment. These stories, which are certainly not mutually exclusive, include insufficient aggregate demand due to household deleveraging, gradual structural adjustment to sector-specific shocks (in particular, construction), lack of credit access caused by the disruption in the financial sector, and increased policy and general economic uncertainty. Distinguishing the quantitative impact of each of these channels is not straightforward.

One aspect of the Great Recession that might shed light on the causes of its severity is the substantial geographic variation in employment losses. The five states most deeply affected by the recession experienced increases in their unemployment rates of 6 percentage points or more (with the largest increase, in Nevada, exceeding 7.5 percentage points). Conversely, the five states least affected by the downturn saw their unemployment rates increase by less than 2.1 percentage points. Given the importance of this geographic variation, it is desirable that theories of the recession are consistent with this cross-sectional pattern.

The differential effect of the recession across geographic areas was not random. In line with a structural explanation, states with larger housing price run-ups prior to the recession and larger price declines during the recession suffered the largest employment losses. In particular, overall employment losses across states and counties are highly correlated with losses in the construction sector. Recent papers have demonstrated that geographic variation in household deleveraging is also correlated with employment losses. Mian and Sufi (2011) and Philippon and Midrigan (2011) both show that employment losses are most severe in areas with initially high and subsequently falling household debt to income ratios. Both the sector reallocation and deleveraging channels are consistent with the observed geographic dispersion in the effects of the recession.

Because financial markets are less locally segmented and much important economic policy is set at the federal level, the credit and uncertainty channels are viewed as less capable of explaining the wide geographic dispersion of unemployment levels in the US during the Great Recession. These facts have led some to argue that the credit constraints and policy uncertainty channels are not consistent with a central feature of the recession. This paper looks deeper into the potential role of uncertainty, which does in fact exhibit substantial geographic variation across the US. We find that increases in policy uncertainty over the 2006-2009 period are correlated with the effects of the recession. The bivariate correlation between policy uncertainty at the state level and employment losses is quite strong, and highly robust to alternate measures

¹⁶Joint with Danny Shoag and Stan Veuger.

of uncertainty.

The implication of this finding is that, like the structural and demand-driven channels, the policy uncertainty story is consistent with the geographic pattern of the recession. Moreover, the policy channel performs remarkably well in direct horse races against the other channels. This suggests that it is premature to dismiss the importance of a policy uncertainty channel in explaining the severity and persistence of the Great Recession. The remainder of this paper describes the various data series we use to measure policy uncertainty and their correlation with job losses.

3.2 Uncertainty

3.2.1 Media

Following Baker, Bloom, and Davis (2011), hereafter BBD, we first measure uncertainty using occurrences of the word "uncertain" in searches of media source databases. As BBD show, there is a pronounced secular trend in coverage by Google News, and we were not able to satisfactorily control for differences in this trend across states. As a result, we do not use Google News, probably the largest compilation of news sources, because its rapidly growing coverage over this period might bias geographic trends.

To capture information specific to local circumstances, we use data from an internet archive of local newspapers (NewsLibrary.com). The web site covers more than 4000 news sources and contains state identifiers. We collect all articles from 1/1/2006 through 12/31/2010 containing the word "uncertainty." We then filter these results to remove papers with spotty online coverage (i.e. restricting ourselves to papers with at least 25 articles a year from 2006 through 2010 with the word "uncertainty"). We also remove television transcripts, non-English language papers, and national papers such as USA Today. In the end, we are left with 116,120 articles from 42 states from 226 newspapers. We then aggregate the articles, counting by state and year, and construct an index of the level of uncertainty during the recession.

Since the downturn began at different times in different states, we use the 2006 measures as the baseline for each index. The results are not significantly changed by altering the starting date. The outcome variable is the average value of the index – current year articles relative to 2006 articles – over the recession period. The mean value is 1.08, indicating that there was an 8% increase in the number of articles containing the word uncertainty in this period when state-year observations are weighted equally. The measure ranges from a high of 1.8 for Rhode Island and 1.6 for Nevada and Arizona to lows of 0.62 for Louisiana and 0.73 for Maine. This means that there was a 60% increase in the number of articles containing the word uncertainty in Nevada and a 28% decrease in Maine in this period. These measures are noisy, but this range of variation appears reasonable relative to the within-state year-to-year variation during this period.




Figure 3.1. Uncertainty in the news and unemployment.

Figure 3.1 plots this relationship against the change in unemployment rates across states from 2006 to 2009. As is evident in the graph, the relationship is quite strong, with an R-squared of 0.22. This correlation is driven partially by the extreme case of Nevada, but the correlation remains at 0.4 even if Nevada is excluded. The relationship holds when looking at article word counts (R-squared = 0.24), when including 2010 (R-squared = 0.18), and when using the change in unemployment rates from 2007 to 2009 (R-squared = 0.22).

3.2.2 Internet Searches

A similar approach can be carried out using data on Google searches. In particular, we have in mind searches by households for information about local economic policy decisions, such as tax and spending levels. To construct a proxy for local uncertainty, we look at the increase in Google searches for the terms "state budget" and "state tax." We use Google Insights for Search database, looking at the relative search frequency for these terms in all states in 2009 relative to the level in 2006.




Figure 3.2. Google searches and unemployment.

Figure 3.2 shows the association of this search-based, or news demand-based, measure of uncertainty. As the figure shows, the correlation between the two is strong (R -squared = 0.14). Taken together, the news source results reveal a cross-sectional relationship between local measures of uncertainty and employment outcomes.

3.2.3 State Budgets

We next turn to more direct and straightforward measures of local policy uncertainty. In this section, we test the relationship between uncertainty about state spending and taxation and economic outcomes and similarly find a strong and robust relationship.

When exploring the relationship between policy uncertainty and economic outcomes, it is important not to rely upon policy measures that could be deterministically related to the outcomes in question. For example, income taxes receipts will of course be strongly correlated with employment rates in the absence of any changes to tax rates, but it would be a mistake to attribute this relationship to any notion of policy uncertainty. Therefore, instead of relying on Census data for realized receipts and spending, we use measures coded from National Association of State Budget Officers (NASBO) reports that record *changes* in legislated tax and spending levels. The NASBO estimates are also useful in that they are recorded in dollars, making it easy to compare disparate policy changes.

State Spending The first measure of uncertainty we use is mid-year budget cuts. State budgets are generally passed annually or biennially, but can be amended in mid-cycle to satisfy state balanced budget requirements or for other reasons. NASBO records these mid-year budget cuts in its *Fall Fiscal Survey of the States*. These cuts are a good measure of the policy uncertainty associated with state spending, as these measures could not be incorporated in the standard budget cycle.

In Figure 3.3, we aggregate the spending cuts in fiscal year 2008 and fiscal year 2009 (most states' fiscal years end in June, meaning fiscal year 2007 did not overlap with the recession), and express them in per-capita terms. We then plot these mid-year or "surprise" cuts against the change in unemployment rate across states from 2006 through 2009.




Figure 3.3. Budget cuts and unemployment.

Figure 3.3 shows a strong positive relationship between surprise budget cuts and employment losses in the recession. The R-squared for the relationship shown is 0.24, with each percentage point of unemployment associated with an extra \$30 of surprise cuts. This relationship is virtually unchanged when using the change in unemployment from 2007 to 2009 (R-squared = 0.25) or when scaling by total state government spending (R-squared = 0.37), as shown in Figure 3.4.




Figure 3.4. Scaled budget cuts and unemployment.

State Taxes NASBO also records legislated tax increases and decreases. Again, this measure is at least less directly endogenous than tax receipts and could be seen as introducing uncertainty about tax levels. Figure 3.5 shows the same relationship with changes in the unemployment rate as the previous graphs.




Figure 3.5. Tax policy and unemployment.

The R-squared in the relationship is 0.17. Excluding North Dakota, the R-squared is still 0.10 between the two variables. The relationship holds when scaling the tax changes by total state government spending

in the period (R-squared = 0.12) and when using the change in unemployment from 2007 through 2009 (R-squared = 0.14). These results document that states that suffered disproportionately in the recession saw larger legislated tax increases.

State Budgets Households might worry not only about uncertainty in future local tax rates, but also about uncertainty in future state spending. We therefore introduce as a measure of spending policy uncertainty the (absolute) gap between realized state government revenues and the projections used at the start of the budget cycle. We use the absolute value to measure the difficulty in predicting aggregate outcomes for the state in advance, rather than simply using the gap, which is more obviously endogenous. Figure 3.6 shows that the absolute value of these deviations is also highly correlated with the severity of the recession. The R-squared is 0.17 with North Dakota, which has been an outlier in its relative immunity to the Great Recession in large part because of an oil boom, and 0.38 when excluding it.




Figure 3.6. Revenue deviations and unemployment.

Figure 3.7 shows that the correlation between these deviations from projected values and the severity of the recession is largest for FY 2009, well after the recession had begun. The correlation for just that year is 0.72.




Figure 3.7. Revenue deviations and unemployment, 2009.

3.2.4 Misleading Indicators

We next consider another measure of uncertainty, also calculated on the basis of deviations from forecasts, using the Federal Reserve Bank of Philadelphia’s (FRBP) leading indicators of trends in labor markets in each of the 50 states.¹⁷ These leading indicators are meant to predict, six months in advance, how the coincident employment indices that the FRBP publishes on a monthly basis will change. Figure 3.8 shows that the leadings indicators’ absolute deviation from realized changes shows a strong association with changes in unemployment over the crisis period. The R-squared for this relationship is 0.17.

¹⁷“The coincident indexes combine four state-level indicators to summarize current economic conditions in a single statistic. The four state-level variables in each coincident index are nonfarm payroll employment, average hours worked in manufacturing, the unemployment rate, and wage and salary disbursements deflated by the consumer price index (U.S. city average). The trend for each state’s index is set to the trend of its gross domestic product (GDP), so long-term growth in the state’s index matches long-term growth in its GDP,” Federal Reserve Bank of Philadelphia (2012).




Figure 3.8. Coincident index deviations.

3.2.5 State Credit Ratings

In addition to worrying about incremental tax changes and sudden alterations to state government spending, households might also be concerned about uncertainty surrounding the long-term solvency of their state government. Households within the state are likely to be disproportionately concerned about solvency uncertainty if they hold an above-average share of state debt or if a solvency crisis could lead to dramatic and unpredictable changes in state government spending or tax levels. We use changes in the credit ratings provided by Standard & Poor's as our measure of uncertainty here¹⁸. Our meaning of uncertainty is admittedly different here than in the previous cases where we measured dispersion; now we are concerned with uncertainty about states being able or unable to live up to their financial obligations. As Figure 3.9 shows, a one-notch change in credit rating is associated with almost a full percentage point change in the unemployment rate (R-squared: 0.18).

¹⁸Similar results hold for the credit ratings given by Moody's and Fitch, though both of these credit agencies rate the debt issued by fewer states than S&P does.




Figure 3.9. State credit ratings and unemployment.

3.2.6 Uncertainty Summarized

The results discussed in the above sections are summarized with additional controls and robustness checks in Tables 3.1, 3.2 and 3.3. Again, these results are not meant to imply causality nor to suggest that the potential importance of the policy uncertainty channel precludes the importance of other channels, such as deleveraging and structural transformation. Rather, by demonstrating the link between the severity of the recession and various forms of (policy) uncertainty, we highlight the roles of all of the hypotheses in generating the regional dispersion of the recession's severity.

Table 3.1. Taxes, government revenues and unemployment

	(1)	(2)	(3)	(4)	(5)
	Legislated	Legislated	Legislated	Deviation	Deviation
	tax increases	tax increases	tax increases	from revenue	from revenue
	FY2008+2009	FY2008+2009	FY2008+2009	projections	projections
	(per capita)	(per capita)	(% of total	FY2008+2009	FY2008+2009
			expenditure)	(%)	(%)
ΔU , 2009-2006	15.6***	15.93***	0.12***	0.85***	1.27***
	(5.1)	(5.1)	(.04)	(0.31)	(0.30)
Lagged Control	N	Y	Y	N	Y
DV Mean	-\$5.74	-\$5.74	0.05%	4.30%	4.30%
Observations	50	50	50	50	50
R-squared	0.17	0.18	0.17	0.17	0.37

Standard errors in parentheses. A control for the lagged value of the dependent variable is included, but not reported, where indicated. *** 1% significance, ** 5% significance, * 10% significance.

Table 3.2. News articles, budget cuts and unemployment

	(1)	(2)	(3)	(4)	(5)
	"Uncertainty" articles relative to 2006 (2007-2009)	"Uncertainty" articles relative to 2006 (2008-2009)	"Uncertainty" articles relative to 2006 (2007-2009)	Surprise budget cuts FY2008-2009 (% of total expenditure)	Surprise budget cuts FY2008-2009 (% of total expenditure)
ΔU , 2009-2006	0.08***	.09***	0.09***	0.30***	0.30***
	-0.02	-0.03	-0.02	-0.06	-0.06
Lagged Control	N	N	Y	N	Y
DV Mean	1.11	1.15	1.11	0.84%	0.84%
Observations	42	42	42	50	50
R-squared	0.22	0.17	0.31	0.37	0.37

Standard errors in parentheses. A control for the lagged value of the dependent variable is included, but not reported, where indicated. *** 1% significance, ** 5% significance, * 10% significance.

Table 3.3. Google searches, leading indicators and unemployment

	(1)	(2)	(3)	(4)
	State policy searches relative to 2006 (2007-2009)	State policy searches relative to 2006 (2007-2009)	Leading indicator deviations 2006-2009 (2007-2009)	Leading indicator deviations 2006-2009 (2007-2009)
ΔU , 2009-2006	0.03*** 0.01	.02 0.01	10.10*** 3.29	7.55*** 1.73
Lagged Control	N	Y	N	Y
DV Mean	0.83	0.83	81.94	81.94%
Observations	46	46	50	50
R-squared	0.14	0.31	0.16	0.78

Standard errors in parentheses. A control for the lagged value of the dependent variable is included, but not reported, where indicated.

*** 1% significance, ** 5% significance, * 10% significance.

3.2.7 Horse Race?

One obvious way to assess which of the different explanations for the wide variety in unemployment changes across the country fits the data the best is to run a direct horse race. When we do this, for example, for our first uncertainty proxy, the newspaper article count, and the change in an index of housing prices, we find that both are significantly associated with the variation in unemployment rates. A doubling of the number of newspaper articles mentioning "uncertainty" in 2007-2009 compared to 2006 is associated with an increase in the unemployment rate of 1.5 percentage points after controlling for the impact of the decline in housing prices. As shown earlier, the measured uncertainty indices contain at least that much variation from the lowest to the highest uncertainty-level states, suggesting that the effect of the uncertainty channel is not only statistically significant but also quantitatively important even after controlling for the other most frequently cited channels.

3.3 Conclusions

This paper provides suggestive evidence that policy uncertainty is a contributing factor to the depth of the Great Recession. We show that a variety of measures of general uncertainty and policy uncertainty in particular are strongly correlated with increases in unemployment at the state level in the US during this period. Furthermore, even after controlling for other commonly cited central causes of the downturn, namely the collapse of the housing bubble, uncertainty remains quantitatively significant. Untangling the causal significance of uncertainty – or, for that matter, of deleveraging or structural reallocation – remains a hard question because of the difficulty of identification in this context. Nevertheless, the data suggest that the role of uncertainty should be taken seriously and in particular should not be downplayed on the grounds that it cannot possibly explain the observed geographic variation in unemployment across regions of the US.

In future extensions of this research, we plan to further the task of decomposing the causes of the recession by asking what types of policy uncertainty are likely to have the largest effect. In particular, we are developing a quantitative general equilibrium model to estimate the effect on aggregate demand and output of uncertainty about future tax rates, transfer payments, government spending levels, and potential default on state government debt. The purpose is twofold. First, we would like to know what types of policy uncertainty are likely to have the largest quantitative impact on inducing households to reduce consumption. Frequently, policy uncertainty is used as a catch-all term, and the mechanism of its impact is left vague. Second, we would then like to know whether realistic calibrations of cross-state variation in each type of policy uncertainty can explain a substantial component of the cross-state variation in output and employment outcomes. The aim is again to further quantitatively decompose the effects of commonly cited causes for the severity of the Great Recession.

References

- AGUIAR, M., AND G. GOPINATH (2006): "Defaultable Debt, Interest Rates and the Current Account," *Journal of International Economics*, 69(1), 64–83.
- AIYAGARI, S. R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 109(3), 659–684.
- BAKER, S. R., N. BLOOM, AND S. J. DAVIS (2011): "Measuring Economic Policy Uncertainty," available online.
- BARILLAS, F., AND J. FERNÁNDEZ-VILLAYERDE (2007): "A Generalization of the Endogenous Grid Method," *Journal of Economic Dynamics and Control*, 31, 2698–2712.
- BASU, S., AND B. BUNDICK (2011): "Uncertainty Shocks in a Model of Effective Demand," available online.
- BEWLEY, T. F. (1977): "The Permanent Income Hypothesis: A Theoretical Formulation," *Journal of Economic Theory*, 16(2), 252–292.
- BLOOM, N., M. FLOETOTTO, AND N. JAIMOVICH (2010): "Really Uncertain Business Cycles," available online.
- CARROLL, C. D. (2006): "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, 91(3), 312–320.
- (2011): "Solution Methods for Microeconomic Dynamic Stochastic Optimization Problems," available online.
- CHALLE, E., AND X. RAGOT (2010): "Precautionary Saving in the Business Cycle," available online.
- CHANG, Y., AND S. KIM (2006): "From Individual to Aggregate Labor Supply: A Quantitative Analysis Based on a Heterogeneous Agent Macroeconomy," *International Economic Review*, 47(1), 1–27.
- CHATTERJEE, S., D. CORBAE, M. NAKAJIMA, AND J.-V. RÍOS-RULL (2007): "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default," *Econometrica*, 75(6), 1525–1589.
- CHETTY, R., A. GUREN, D. MANOLI, AND A. WEBER (2011): "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins," *American Economic Review: Papers and Proceedings*, (101:3), 471–475.
- CHRISTIANO, L., M. EICHENBAUM, AND S. REBELO (2009): "When Is the Government Spending Multiplier Large?," available online.
- COGAN, J. F., T. CWIK, J. B. TAYLOR, AND V. WIELAND (2010): "New Keynesian versus old Keynesian government spending multipliers," *Journal of Economic Dynamics and Control*, (34), 281–295.
- COHEN-COLE, E., AND B. DUYGAN-BUMP (2008): "The Household Bankruptcy Decision: The Role of Social Stigma vs. Information Sharing," mimeo.
- (2009): "Social Influence and Bankruptcy: Why do so Many Leave so Much on the Table," mimeo.
- COHEN-COLE, E., B. DUYGAN-BUMP, AND J. MONTORIOL-GARRIGA (2009): "Forgive and Forget: Who Gets Credit after Bankruptcy and Why?," FRB Boston QAU Series paper no. QAU09-2.
- EATON, J., AND M. GERSOVITZ (1981): "Debt with Potential Repudiation: Theoretical and Empirical Analysis," *Review of Economic Studies*, 48(2), 289–309.
- EGGERTSSON, G. B. (2006): "Fiscal Multipliers and Policy Coordination," Federal Reserve Bank of New York Staff Report No. 241.
- (2009): "What Fiscal Policy is Effective at Zero Interest Rates?," Federal Reserve Bank of New York Staff Report No. 402.

- (2010): “The Paradox of Toil,” available online.
- EGGERTSSON, G. B., AND P. KRUGMAN (2011): “Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo approach,” available online.
- ENGEN, E. M., AND J. GRUBER (2001): “Unemployment Insurance and Precautionary Saving,” *Journal of Monetary Economics*, 47, 545–579.
- FAY, S. A., E. HURST, AND M. J. WHITE (1998): “The Bankruptcy Decision: Does Stigma Matter?,” mimeo.
- FEDERAL RESERVE BANK OF PHILADELPHIA (2012): “State Coincident Indexes,” <http://www.philadelphiafed.org/research-and-data/regional-economy/indexes/coincent/>.
- FEENBERG, D., AND E. COUTTS (1993): “An Introduction to the TAXSIM Model,” *Journal of Policy Analysis and Management*, 12(1).
- FISHER, J., L. FILER, AND A. LYONS (2004): “Is the Bankruptcy Flag Binding? Access to Credit Markets for Post-Bankruptcy Households,” American Law and Economics Association Annual Meetings 1041.
- FUJITA, S., AND G. RAMEY (2008): “The Cyclicalities of Separation and Job Finding Rates,” available online.
- GAHVARI, F. (1985): “Taxation of Housing, Capital Accumulation and Welfare: a Study in Dynamic Tax Reform,” *Public Finance Review*, 13(2), 132–160.
- GALÍ, J., M. GERTLER, AND J. D. LÓPEZ-SALIDO (2007): “Markups, Gaps, and the Welfare Costs of Business Fluctuations,” *The Review of Economics and Statistics*, 89(1), 44–59.
- GLAESER, E., AND J. SHAPIRO (2003): “The Benefits of the Home Mortgage Interest Deduction,” *Tax Policy and the Economy*.
- GOURINCHAS, P.-O., AND J. A. PARKER (2001): “The Empirical Importance of Precautionary Saving,” *The American Economic Review*, 91(2), 406–412.
- GUERRIERI, V., AND G. LORENZONI (2011): “Credit Crises, Precautionary Savings and the Liquidity Trap,” available online.
- HALL, R. E. (2005): “Job Loss, Job Finding, and Unemployment in the U.S. Economy over the Past Fifty Years,” *NBER Macroeconomics Annual*, 20, 101–137.
- (2011): “Clashing Theories of Unemployment,” available online.
- HEATHCOATE, J., K. STORESLETTEN, AND G. VIOLANTE (2010): “Redistributive Taxation in a Partial-Insurance Economy,” available online.
- HUGGETT, M. (1993): “The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies,” *Journal of Economic Dynamics and Control*, 17, 953–969.
- KEHOE, T. J., AND D. K. LEVINE (1993): “Debt-Constrained Asset Markets,” *Review of Economic Studies*, 60(4), 865–888.
- KRUEGER, A. B., AND A. MUELLER (2010): “Job Search and Unemployment Insurance: New Evidence from Time Use Data,” *Journal of Public Economics*, 94(3-4), 298–307.
- KRUGMAN, P. (1998): “It’s Baaack! Japan’s Slump and the Return of the Liquidity Trap,” *Brookings Papers on Economic Activity*, (2), 137–205.
- KRUSELL, P., AND A. A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106(5), 867–896.
- LANDAIS, C., P. MICHAILLAT, AND E. SAEZ (2011): “Optimal Unemployment Insurance over the Business Cycle,” available online.

- LEFGREN, L., AND F. MCINTYRE (2008): "Explaining the Puzzle of Cross-State Differences in Bankruptcy Rates," mimeo.
- LIVSHITS, I., J. MACGEE, AND M. TERTILT (2007): "Consumer Bankruptcy: A Fresh Start," *Journal of Monetary Economics*, 97(1), 402–418.
- LUNDQUIST, L., AND T. SARGENT (2000): *Recursive Macroeconomic Theory*. MIT Press.
- LUSTIG, H., AND S. VAN NIEUWERBURGH (2004): "Housing Collateral and Consumption Insurance Across US Regions," mimeo.
- MALPEZZI, S., AND S. MAYO (1987): "The Demand for Housing in Developing Countries: Empirical Estimates from Household Data," *Economic Development and Cultural Change*, 35(4), 687–721.
- MEHRLING, P. (1998): "Idiosyncratic Risk, Borrowing Constraints and Asset Prices," *Metroeconomica*, 49(3), 261–283.
- MEIER, A. (2010): "Still Minding the Gap—Inflation: Dynamics during Episodes of Persistent Large Output Gaps," IMF Working Paper.
- MERTENS, K., AND M. O. RAVN (2011): "Fiscal Policy in an Expectations Driven Liquidity Trap," Available online.
- MIAN, A., AND A. SUFI (2011): "What Explains High Unemployment? The Deleveraging-Aggregate Demand Hypothesis," mimeo.
- MUSTO, D. (2004): "What Happens when Information Leaves a Market? Evidence from Post-Bankruptcy Consumers," *Journal of Business*, 77(4), 725–748.
- PHILIPPON, T., AND V. MIDRIGAN (2011): "Household Leverage and the Recession," NBER Working Paper No. 16965.
- POTERBA, J., AND T. SINAI (2008): "Tax Expenditures for Owner-Occupied Housing: Deductions for Property Taxes and Mortgage Interest and the Exclusion of Imputed Rental Income," *American Economic Review Papers and Proceedings*, 96(2), 84–89.
- ROMER, C., AND J. BERNSTEIN (2009): "The Job Impact of the American Recovery and Reinvestment Plan," available online.
- SAIZ, A. (2010): "The Geographic Determinants of Housing Supply," *Quarterly Journal of Economics*, 125, 1253–89.
- SHIMER, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95(1), 25–49.
- (2007): "Reassessing the Ins and Outs of Unemployment," available online.
- STORELLETTEN, K., C. I. TELMER, AND A. YARON (2004): "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, 112(3), 695–717.
- SUM, A., AND I. KHATIWADA (2010): "Labor Underutilization Problems of U.S. Workers Across Household Income Groups at the End of the Great Recession: A Truly Great Depression Among the Nation's Low Income Workers Amidst Full Employment Among the Most Affluent," Center for Labor Market Studies Publications.
- WHITE, M. J. (1998): "Why Don't More Households File for Bankruptcy," *Journal of Law, Economics and Organization*, 14(2).
- YOUNG, E. R. (2010): "Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell-Smith Algorithm and Non-Stochastic Simulations," *Journal of Economic Dynamics and Control*, 34, 36–41.