ABSTRACT

Title of dissertation: CREDIT AND LIQUIDITY IN THE MACROECONOMY

Jonathan Kreamer, Doctor of Philosophy, 2015

Dissertation directed by: Dr. Anton Korinek
Department of Economics
Johns Hopkins University

This dissertation studies the role of credit and liquidity in macroeconomic fluctuations.

Chapter 1 analyzes the effect of endogenous unemployment risk on the dynamics of recovery from a liquidity trap. In a liquidity trap, an adverse demand shock raises unemployment and produces a period of slow hiring. Slow hiring further reduces demand, both for standard precautionary reasons and because credit conditions endogenously worsen, reducing households' ability to borrow and consume. Multiple equilibrium paths exist, and which one the economy follows depends on household expectations and the policy rule adopted by the central bank after the economy exits the trap. Employment remains depressed for a substantial period after an adverse shock because high unemployment increases the dispersion of household debt holdings, slowing the recovery of demand. I find that the initial household debt distribution significantly affects the economy's sensitivity to a demand shock, and study the role of central bank policy in mitigating the initial fall in employment and promoting faster recovery.

Chapter 2 explores a novel channel through which financial shocks affect the real economy through the supply of liquidity. I consider a model in which firms require uncertain ongoing financing, and agency costs limit their ability to raise new funds. To secure future financing, firms hold assets to sell if needed, and purchase credit lines from financial intermediaries. I collectively refer to these instruments as liquidity. Financial intermediaries' ability to commit future funds depends on their capital. This creates a linkage between bank balance sheets and the aggregate supply of liquidity. Bank losses raise the liquidity premium and reduce investment. I analyze the optimal supply of public liquidity, and find that when private liquidity is scarce the government should issue bonds for their liquidity properties. I further find that the optimal supply of government debt is decreasing in bank capital. This suggests that in the wake of a financial crisis in which financial intermediaries suffer large losses, governments should increase debt issuance.

Chapter 3 considers the distributive implications of financial regulation.[1] It develops a model in which the financial sector benefits from financial risk-taking by earning greater expected returns. However, risk-taking also increases the incidence of large losses that lead to credit crunches and impose negative externalities on the real economy. A regulator has to trade off efficiency in the financial sector, which is aided by deregulation, against efficiency in the real economy, which is aided by tighter regulation and a more stable supply of credit.

[1]Chapter 3 was coauthored with Professor Anton Korinek, and was published as Korinek and Kreamer (2014)

CREDIT AND LIQUIDITY IN THE MACROECONOMY

by

Jonathan Kreamer

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2015

Advisory Committee:
Professor Anton Korinek, Co-chair/Advisor
Professor John Shea, Co-chair
Professor Borağan Aruoba
Professor Luminita Stevens
Professor Phillip Swagel
Professor Katharine Abraham

Acknowledgments

Completing a dissertation is difficult, and I couldn't have done it alone. I would like to express my gratitude too all the people who helped me through.

First, I owe a great debt of gratitude to my advisor Anton Korinek for his invaluable guidance and wise counsel throughout my graduate studies. He has been unfailingly generous to me with his time and energy, and has provided unwavering support and encouragement.

I would also like to thank the other members of my committee for their advice and support throughout my studies. I would like to thank John Shea for providing detailed and insightful comments on many drafts of my work. I would like to thank Borağan Aruoba for his feedback on my work, and for the invaluable education he provided me in the three courses I took from him. I would like to thank Luminita Stevens for excellent comments on my research. I would like to thank Phillip Swagel for agreeing to serve on my committee and provided helpful comments. Finally, I would like to thank Katharine Abraham for serving on my committee on short notice.

In addition to the academic support and guidance provided by my advisors, I have benefited immensely from the loving support of my family. Most of all, I would like to thank my wife Stefanie for her love, for her wisdom, and for helping me through the hard times. I have been truly blessed by her love and support. I would like to thank the rest of my family, particularly my mother for her constant prayers on my behalf. I have benefited greatly from my parents' love and support throughout my life, and for this I thank them.

Finally, I would like to thank my friends for providing me with advice, support, and breaks from the work of graduate school. I would like to thank Nick for his long-time friendship, wisdom, and generous spirit; Eddie for his faithfulness and wise experience; and my classmates and friends at Maryland, particularly Pablo and Naomi, for many delightful conversations over the course of our studies.

Table of Contents

List of Tables

List of Figures

Chapter 1: Household Debt, Unemployment, and Slow Recoveries

1.1 Introduction

In the wake of the 2007 – 2009 recession, the U.S. experienced a period of persistently high unemployment with short-term interest rates at the zero lower bound. Several researchers have linked the severity of the recession and the slow recovery of hiring to weak consumer demand, possibly due to high levels of household debt entering the recession.[1] In this paper I offer an alternative hypothesis for persistent weak demand: that slow hiring itself weakens demand, amplifying shocks and producing self-fulfilling slow recoveries.

Slow hiring weakens demand for two reasons. First, slow hiring implies longer spells of unemployment, which prompt greater saving by households for standard precautionary reasons. Second, high unemployment endogenously tightens the financial constraints of households because longer expected spells of unemployment

[1] Mian and Sufi (2010), and Mian, Rao, and Sufi (2013) argue that high household leverage contributed to the decline in employment and consumption in 2008. Dynan (2012) and Hall (2011) argue that high levels of household debt contributed to the slow recovery since 2008. Reinhart and Rogoff (2009), Hall (2010), and Jordà, Schularick, and Taylor (2013) argue that financial crises are often preceded by increases in leverage and followed by slow recoveries, but for a contrary finding see Romer and Romer (2015).

increase the incentives for households to default on their debt. When the economy is demand-constrained, as in a liquidity trap, the increase in desired saving lowers demand, causing employment to fall further.

To explore this mechanism, I develop a model in which households face idiosyncratic endogenous unemployment risk. Employed households face a constant exogenous rate of job loss, while the rate at which unemployed households find jobs is time-varying and endogenous. Employed households save and unemployed households dissave in order to smooth consumption over unemployment spells. Unemployed households face a borrowing constraint, set to prevent default. As in a standard New Keynesian model, prices are sticky and the levels of output and employment are determined by demand. Unlike in a standard model, the path of output in turn determines the rate of job-finding, which affects demand through the precautionary and credit channels described above.

I consider a persistent demand shock in this setting, which in my baseline experiment is an exogenous worsening of commitment problems in credit markets that reduces households' ability to borrow. This shock reduces aggregate spending for a given path of the interest rate and the job-finding rate, but does not affect the hiring incentives of firms. If prices were flexible, the interest rate would fall to raise desired consumption and maintain the rate of hiring that prevailed before the shock. If instead prices are sticky, the central bank will try to replicate the flexible-price equilibrium by lowering the interest rate. However, if the economy is in a liquidity trap, the central bank will not be able to lower the interest rate sufficiently to offset the demand shock, and employment falls instead.[2]

My first main result is that the presence of unemployment risk amplifies the fall of employment following this shock. Since the demand shock is persistent, employment remains low for a period of time. This implies that the hiring rate falls, which reduces demand via the precautionary and credit channels discussed above. This feedback mechanism amplifies the employment effect of the initial demand shock, producing a larger fall in employment than would otherwise occur. In my baseline experiment, a temporary shock to household credit access, I calculate that the initial fall in demand (and thus employment) is about twice as large as in a complete-markets model that experiences a demand shock of equivalent magnitude.[3]

My second main result is that unemployment risk can produce a slow recovery of employment following a demand shock. This occurs because high unemployment increases the dispersion of asset holdings, both by increasing saving by employed households, and by increasing the number and length of unemployment spells. Since unemployed households borrow, this produces more households with large debt. Greater asset dispersion depresses demand because poor households close to the borrowing constraint reduce their consumption by more than wealthy households increase theirs. Thus the increase in asset dispersion keeps demand low after the shock has dissipated, producing a slow recovery of employment.

The speed of recovery varies quite a bit depending on the persistence of the

[2] This result is similar to Eggertsson and Woodford (2003), Werning (2011), and Eggertsson and Krugman (2012).

[3] For details on this comparison, see section 1.5.3.

initial shock. The baseline credit shock does not produce a slow recovery because its duration is fairly short. The tighter borrowing constraint facing households causes highly indebted households to reduce their debt holdings, a period of deleveraging that reduces demand and causes employment to fall in the short run. However, the reduced debt burden allows demand to recover quickly when the shock begins to fade. Intuitively, a shock directly to household credit purges bad balance sheets, so that after the initial deleveraging households are able to rapidly increase spending as credit conditions recover. By contrast, a permanent credit shock induces a lengthy period of deleveraging as the economy transitions to a new asset distribution. This lengthy period of high unemployment causes households to accumulate debt, slowing the deleveraging process.

In a similar manner, the *initial* asset distribution significantly affects the economy's response to demand shocks. Greater dispersion in assets implies there are more households close to the constraint, which increases the economy's sensitivity to shocks. This both increases amplification, since the households close to the constraint are forced to deleverage when the constraint tightens following a shock, and slows the recovery of employment because the greater debt burden requires a longer period of deleveraging.

A further result is that there are many equilibrium paths following any particular shock. If households are pessimistic about the rate of recovery of employment following the demand shock, the initial fall in demand will be greater, and a slower recovery will follow. Such a multiplicity of equilibria is quite common in New Keynesian models, particularly those with a zero lower bound on the nominal interest rate.[4] Typically determinacy is achieved by assuming that the central bank follows a conditional interest rate policy that rules out undesirable equilibria. However, in a liquidity trap the central bank cannot lower interest rates, and so may not be able to rule out slumps generated by self-fulfilling pessimism about the rate of recovery. Thus any demand shock that pushes the economy into a liquidity trap could bring about a persistent slump.

Not only may pessimistic expectations alter the dynamics of the economy following a demand shock, they may constitute a shock by themselves. If households anticipate a period of reduced hiring, they will reduce their spending. A sufficiently large decrease in consumption would require a negative real interest rate to offset, causing the economy to fall into a liquidity trap and inducing a persistent slump.

The existence of multiple equilibria suggests a role for policy to offset demand shocks. Several authors, starting with Krugman (1998), have suggested that the central bank can reduce the fall in employment during a liquidity trap by keeping interest rates low after the recovery. These low interest rates generate a post-recovery boom in consumption, which raises demand during the slump as households smooth their consumption.[5] In my formulation, the analogous mechanism is a boom in *hiring*, which raises demand by reducing precautionary saving. However, as emphasized by several authors, engineering such a boom requires the central bank to commit to

[4] See, for example, Bernanke and Woodford (1997), Benhabib, Schmitt-Grohé, and Uribe (2001), Clarida, Galí, and Gertler (2000), and Cochrane (2011).

[5] Werning (2011) first clarified that this mechanism does not rely on inflation, but is purely a consumption boom.

following an ex post suboptimal policy once the economy exits the trap.

Literature. This paper is part of a rapidly growing literature on the dynamics of the economy in a liquidity trap. The early papers in this literature, Krugman (1998) and Eggertsson and Woodford (2003), were the first to consider the possibility of a binding zero lower bound on nominal interest rates in a modern context, and suggested the possibility of forward guidance in substituting for conventional monetary policy in these circumstances. Werning (2011) investigates optimal policy in this setting, and clarifies the dynamics of a liquidity trap in a simple and elegant model. Cochrane (2011) and Aruoba and Schorfheide (2013) discuss the problem of determinacy in the presence of a zero lower bound, and the possibility of deflationary traps, which is related to my finding of multiple equilibrium paths of recovery.

A number of papers have highlighted the role of household debt and deleveraging as the cause of low demand leading to a liquidity trap. Eggertsson and Krugman (2012) first explored this deleveraging channel in the context of a simple model with a zero lower bound. Hall (2011) develops a similar hypothesis, with particular reference to high levels of household debt producing a slow recovery of demand. Korinek and Simsek (2014) study the role of macroprudential regulation in reducing the effects of this deleveraging channel. Guerrieri and Lorenzoni (2011) investigate the effect of precautionary saving behavior on the evolution of the distribution of debt holdings during a deleveraging episode. This paper contributes to this literature by introducing endogenous time-varying unemployment risk into a model of a liquidity trap and studying the resulting feedback from high unemployment to weak demand through precautionary saving and endogenous borrowing constraints.

Several papers have empirically investigated the household debt / aggregate demand hypothesis of the 2007 – 2009 recession. Mian and Sufi (2010) and Mian, Rao, and Sufi (2013) use detailed county-level data to show that counties with high household leverage and large declines in house prices before the crisis had larger declines in employment and output during the crisis. Mian and Sufi (2014) shows that these differential employment declines are driven by the hiring decisions of nontradable good firms, suggesting that the mechanism operates through a demand channel. Dynan (2012) finds that households with high leverage saw larger declines in consumption in 2007 – 2009, despite a smaller decline in net worth, indicating the existince of a household credit channel rather than a wealth channel. Baker (2014) finds that households with higher levels of debt adjust their consumptiion more in response to changes in income, suggesting a higher marginal propensity to consume for highly indebted households.

Several authors have recently offered alternative models of aggregate demand channels, many involving matching in product or labor markets. Kocherlakota (2012) analyzes a so-called incomplete labor market model, in which the real interest rate is set by the central bank and the labor supply condition may not hold. This formulation is similar to a New Keynesian model with fixed prices, except that in a standard New Keynesian model the labor demand condition is dropped instead of labor supply. Chamley (2014) investigates the possibility of self-fulfilling precautionary demand for savings in a model with money and bonds, in which the economy may become stuck in a low demand saving trap, or may converge very slowly towards full employment. Michaillat (2012) argues that matching frictions are insufficient to explain unemployment and develops a model of job rationing during recessions. Michaillat and Saez (2013) develop a model of aggregate demand with matching frictions in both labor and goods markets, and show that tightness in one market affects tightness in the other, generating an aggregate demand channel for employment fluctuations. I view these approaches as complementary to the New Keynesian formulation of aggregate demand used in this paper.

Since this paper's primary mechanism operates through precautionary saving by households, it is related to the empirical on precautionary saving behavior. Carroll and Samwick (1997) and Carroll and Samwick (1998) find that households that face greater income uncertainty hold more wealth, and estimate that precautionary saving accounts for $39 - 46\%$ of household asset holdings. Gourinchas and Parker (2002) find that young households target a buffer of precautionary wealth, and estimate a coefficient of relative risk aversion of $0.5 - 1.4$. Parker and Preston (2005) also find a significant and strongly countercyclical precautionary saving motive, that is similar in magnitude to the interest-rate motive. Carroll, Slacalek, and Sommer (2012) find that a significant portion of the consumption decline since 2008 was attributable to precautionary effects.

A few papers have considered the interaction of precautionary savings and endogenous unemployment risk in the presence of incomplete markets. Challe et al. (2014) develop a model that combines these ingredients with sticky prices, and show that variations in precautionary saving over the business cycle amplify employment fluctuations. Ravn and Sterk (2013) likewise study the effect of unemployment risk on demand, although their focus is on an exogenous increase in labor market mismatch increasing long-term unemployment. Caggese and Perez (2013) study a model with unemployment risk and credit constraints facing both firms and households. They show that precautionary behavior by households and firms can interact to generate a negative demand externality that significantly increases the volatility of employment over the business cycle. Beaudry, Galizia, and Portier (2014) study endogenous unemployment risk in a model with decentralized markets and search frictions. However, they consider contemporaneous employment risk only, and do not focus on the dynamics of household asset holdings.

The papers cited above generate amplification due to precautionary saving behavior. This paper sheds additional light on this mechanism, showing that it operates through a depressed job-finding rate following a persistent demand shock. It also demonstrates the critical role of the asset distribution in determining these dynamics, which the papers cited above do not study. This paper also differs by focusing on the liquidity trap case, by studying the role of expectations in generating multiple equilibrium paths of recovery, and by studying the role of endogenous credit conditions facing consumers.

1.2 Households

The model is set in continuous time with a single non-storable consumption good. There are three types of agents: households, final-good firms, and intermediate-good firms. I first discuss the household problem, and then turn to the rest of the model.

There is a measure 1 of households, indexed by $i \in [0,1]$. Household i has expected lifetime utility

$$E_0\left[\int_0^\infty e^{-\rho t}u(c_i(t))dt\right]$$

where $u(\cdot)$ is a standard utility function, and ρ is the household subjective discount rate. I assume throughout that $u(c)$ exhibits constant relative risk aversion γ, i.e. $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ for $\gamma \neq 1$ and $u(c) = \log(c)$ for $\gamma = 1$.

1.2.1 Income Process

Households receive a flow of nonlabor income $e(t)$, which is identical across all households, and is assumed to be constant over time. This income allows households to carry debt while maintaining positive consumption. Household i receives a flow of labor income $w(t)h_i(t)$, where $h_i(t) \in \{0,1\}$ is household i's employment status at time t, and $w(t)$ is the wage of employed workers.[6]

Since households suffer no disutility from labor, all households would like to work, implying an inelastic labor supply at $n = 1$. This would imply full employment in a Walrasian labor market, but here there is a hiring cost that produces equilibrium unemployment. Employment and unemployment shocks arrive with hiring probability $p(t)$ and separation (job loss) probability $s(t)$ per unit of time. These are flow probabilities, i.e. the job-finding rate $p(t)$ means that the probability that a worker who is unemployed at time t finds a job during the interval $[t, t+dt]$

[6] It is common in models of the labor market that unemployed households receive some income, which is generally interpreted as unemployment benefits. One could interpret $e(t)$ in my formulation as unemployment benefits, and $w(t) + e(t)$ as the wage.

is approximately $p(t) \cdot dt$, which becomes exact as $dt \to 0$. Likewise, the total probability that an unemployed household will become employed at least once during the interval $[t_0, t_1]$ is

$$1 - e^{-\int_{t_0}^{t_1} p(t)dt}$$

and likewise for job separation probabilities, with $s(t)$ in place of $p(t)$. With a constant hiring rate p, the length of an unemployment spell has an exponential distribution, with expected value $1/p$.

Household i enters period t with net assets $a_i(t)$, and can save or borrow at net interest rate $r(t)$. Assuming no default, household assets evolve according to

$$\dot{a}_i(t) = r(t)a_i(t) + e(t) + w(t)h_i(t) - c_i(t) \tag{1.1}$$

where $\dot{a}_i(t)$ is household i's net saving.

Households face a time-varying borrowing constraint, which they take as exogenous. This borrowing limit takes the form of a lower bound on asset holdings $\underline{a}(t)$, such that asset holdings of household i must satisfy

$$a_i(t) \geq \underline{a}(t)$$

at every time t. The borrowing constraint satisfies $\underline{a}(t) < 0$, so that it is always possible for households to carry some debt.[7]

[7]Technically, employed households could face a looser borrowing constraint than unemployed households. However, assuming a single borrowing constraint rules out default following job loss. In any case, the borrowing constraint generally does not bind for employed households.

1.2.2 Household Problem Under Repayment

There is no aggregate uncertainty, so that the paths of all variables are known at time 0 except for the employment status and asset holdings of individual households, $h_i(t)$ and $a_i(t)$. Since households are identical except for asset holdings and employment status, we need only derive household decisions at each point (h, a, t).

Let $V(a, t)$ be the value function of an employed household in period t with assets a, and let $U(a, t)$ be the value function of an unemployed household. Employed households choose current consumption to maximize the Hamiltonian

$$\rho V(a,t) = \max_c \{u(c) + V_a(a,t) \cdot \dot{a}_e + V_t(a,t) + s(t)\left(U(a,t) - V(a,t)\right)\} \tag{1.2}$$

where $\dot{a}_e = r(t)a + w(t) + e(t) - c$, and where V_a and V_t denote the partial derivatives of the value function. I denote the optimal consumption decision rule by $c_e(a, t)$, which implies the saving decision rule $\dot{a}_e(a, t) = r(t)a(t) + w(t) + e(t) - c_e(a, t)$.

Unemployed households similarly choose current consumption to maximize the Hamiltonian

$$\rho U(a,t) = \max_c \{u(c) + U_a(a,t) \cdot \dot{a}_u + U_t(a,t) + p(t)\left(V(a,t) - U(a,t)\right)\} \tag{1.3}$$

where $\dot{a}_u = r(t)a + e(t) - c$. I again denote the optimal consumption decision rule by $c_u(a, t)$, which implies saving decision rule $\dot{a}_u(a, t) = r(t)a(t) + e(t) - c_u(a, t)$.

The optimal consumption choices $c_e(a, t)$ and $c_u(a, t)$ satisfy

$$u'(c_e(a,t)) = V_a(a,t) \tag{1.4}$$

$$u'(c_u(a,t)) = U_a(a,t) \tag{1.5}$$

for $a > \underline{a}$. Intuitively, the shadow value of wealth equals the marginal utility of consumption, since under optimal consumption unconstrained households are indifferent between consuming or saving the marginal unit of wealth.

When $a = \underline{a}$, the borrowing constraint may bind. In this case, c is chosen such that $\dot{a} = \dot{\underline{a}}$, where $\dot{\underline{a}} \neq 0$ may hold as the borrowing constraint tightens and loosens with changing macroeconomic conditions. Then the optimality conditions satisfy

$$u'(c_e) \geq V_a(\underline{a})$$

$$u'(c_u) \geq U_a(\underline{a})$$

Intuitively, when the constraint binds the household would like to borrow more, but is prevented from doing so by the constraint. This implies that the marginal utility from consumption is strictly greater than the marginal utility from saving, i.e. $u'(c) > V_a$.

The costate equations of the Hamiltonian at $a > \underline{a}$ are

$$(\rho + s - r)\,\lambda = \lambda_a \dot{a}_e + \lambda_t + s\kappa$$

$$(\rho + p - r)\,\kappa = \kappa_a \dot{a}_u + \kappa_t + p\lambda$$

where $\lambda = V_a$ and $\kappa = U_a$ are the costate variables.

Let $\dot{\lambda} = \lambda_a \dot{a}_e + \lambda_t$ and $\dot{\kappa} = \kappa_a \dot{a}_u + \kappa_t$, so that $\dot{\lambda}(a,t)$ and $\dot{\kappa}(a,t)$ are the instantaneous rates of change of the costate variables of employed and unemployed households, respectively, when these households do not change employment status.[8]

[8]This is an abuse of notation, since $\dot{c}$ refers to the rate of change of consumption by a particular household, rather than the evolution of the household decision rule defined at a particular point (a,t).

Using these terms, we can express the costate equations as

$$
\begin{aligned}
-\frac{\dot{\lambda}}{\lambda} &= r - \rho + s\left(\frac{\kappa - \lambda}{\lambda}\right) \\
-\frac{\dot{\kappa}}{\kappa} &= r - \rho + p\left(\frac{\lambda - \kappa}{\kappa}\right)
\end{aligned}
$$

1.2.3 Household Euler Equation

When $a > \underline{a}$, we can express the costate equations in terms of the consumption decision rules. First taking the expression in terms of $\dot{\lambda}$, and observing that $-\dot{\lambda}/\lambda = \gamma \dot{c}_e / c_e$, where $\dot{c}_e(a,t) = \frac{\partial}{\partial a} c_e(a,t) + \frac{\partial}{\partial t} c_e(a,t)$ is the instantaneous rate of change of consumption of an employed household that does not lose its job, we find that $\dot{c}_e$ satisfies

$$
\frac{\dot{c}_e}{c_e} = \gamma^{-1}(r - \rho) + \gamma^{-1} s \left(\frac{c_u^{-\gamma}}{c_e^{-\gamma}} - 1\right) \tag{1.6}
$$

This is a form of the continuous time consumption Euler equation. When $s = 0$, there is no unemployment risk, and the expression simplifies to the familiar consumption Euler equation under certainty:

$$
\frac{\dot{c}}{c} = \gamma^{-1}(r - \rho) \tag{1.7}
$$

This implies an increasing path of consumption when $r > \rho$, and a decreasing path when $r < \rho$. This captures the intertemporal substitution of consumption in response to the interest rate. A low interest rate prompts households to shift consumption towards the present, implying a low growth rate of consumption, whereas a high interest rate prompts households to reduce current consumption in order to save more, which implies a higher rate of consumption growth.

When $s > 0$, households face unemployment risk. The term $\gamma^{-1} s \left(c_u^{-\gamma}/c_e^{-\gamma} - 1\right)$ is positive because $c_u < c_e$, implying a faster rate of consumption growth corresponding to greater saving. This term does not perfectly capture the precautionary motive, however, because $\dot{c}_e$ only contains changes in consumption for households that remain employed, whereas we also should consider changes in consumption due to job loss. In order to isolate the precautionary saving effect, we separate this expression into a precautionary term related to the *volatility* of future consumption, and a term corresponding to the *expected level* of future consumption. To do this, we simply add the expected change in consumption due to job loss to each side of (1.6) to obtain:

$$\underbrace{\frac{\dot{c}_e + s\left(c_u - c_e\right)}{c_e}}_{E[\dot{c}/c]} = \gamma^{-1}\left(r - \rho\right) + \underbrace{s\left[\gamma^{-1}\left(\frac{c_u^{-\gamma}}{c_e^{-\gamma}} - 1\right) - \left(1 - \frac{c_u}{c_e}\right)\right]}_{\text{precautionary motive}} \tag{1.8}$$

We may likewise express the Euler equation for unemployed households as:

$$\frac{\dot{c}_u + p\left(c_e - c_u\right)}{c_u} = \gamma^{-1}\left(r - \rho\right) + p\left[\gamma^{-1}\left(\frac{c_e^{-\gamma}}{c_u^{-\gamma}} - 1\right) - \left(1 - \frac{c_e}{c_u}\right)\right] \tag{1.9}$$

Equations (1.8) and (1.9) correspond more precisely to our idea of an Euler equation. They give expected consumption growth as a function of intertemporal substitution and a precautionary saving term. Let

$$T(x) = \gamma^{-1}\left(x^{-\gamma} - 1\right) - (1 - x) \tag{1.10}$$

so that the the precautionary term for employed households is $sT(c_u/c_e)$, and for unemployed households is $pT(c_e/c_u)$. Thus the precautionary motive is a simple function of the percentage change in consumption due to a change of employment status, times the probability of this change. Since $T(x) = 0$ at $x = 1$, the precautionary motive approaches zero as $c_u \to c_e$ — intuitvely, households would not be concerned about unemployment spells if they could perfectly smooth their consumption. The derivative of T is $T'(x) = (x^{\gamma+1} - 1)/x^2$. This is strictly negative for $x \in (0, 1)$, and strictly positive for $x > 1$. Therefore $T(x)$ obtains a unique minimum at $x = 1$ on the interval $(0, \infty)$, so that the precautionary motive term is always positive.

To gain a little intuition for the precautionary motive term, we can take a second-order Taylor expansion around $x = 1$. Then the precautionary motive term is approximately

$$T(x) \approx \frac{1}{2}(1+\gamma) \times (1-x)^2$$

This expression has a simple intuition. Taking the problem of a currently employed household, let dc be the change in consumption of this household over a small time interval dt. Then for dt small, dc behaves like a binary random variable with probability distribution

$$dc = \begin{cases} \dot{c}_e \cdot dt & \text{with probability } 1 - sdt \\ c_u - c_e & \text{with probability } sdt \end{cases} \tag{1.11}$$

We used a similar concept above when we observed that $E[\dot{c}] = [\dot{c}_e + s(c_u - c_e)]$ in equation (1.8). Since $\dot{c} = dc/dt$, we can multiply by dt to obtain $E[dc] = [\dot{c}_e + s(c_u - c_e)]\,dt$, which gives the expectation of the random variable dc. We can likewise compute the variance, which, again neglecting higher order terms, is:

$$Var(dc) = s(c_u - c_e)^2\,dt$$

or in terms of the rate of change in consumption, $s(c_u - c_e)^2 = Var(dc/dt) \cdot dt$.[9]

By analogy to the Euler equation, we are interested in the variance of the growth *rate* of consumption, dc/c. Thus we divide dc be c, which at time t is a known constant. Then we see that the Taylor expansion of the precautionary motive term for employed households is just:

$$sT\left(\frac{c_u}{c_e}\right) \approx \frac{1+\gamma}{2} \times Var\left(\frac{dc/dt}{c}\right) dt \tag{1.12}$$

The term $1+\gamma = -\frac{cu'''(c)}{u''(c)}$, and so $1+\gamma$ is the relative prudence of the utility function, as defined by Kimball (1990). Thus the precautionary saving term is just (one-half) prudence times a term proportional to the variance of the consumption growth rate. This yields a very intuitive expression for the Euler equation:

$$E\left[\frac{dc/dt}{c}\right] \approx \frac{1}{\gamma}(r-\rho) + \frac{1}{2}(1+\gamma) \cdot Var\left(\frac{dc/dt}{c}\right) \cdot dt \tag{1.13}$$

1.2.4 Aggregate Euler Equation

We can aggregate the Euler equations of individual households to obtain an aggregate Euler equation. The household Euler equations (1.8) and (1.9) can be written as

$$E[\dot{c}] = \gamma^{-1}(r-\rho)\, c + qT(c_{-h}/c) \cdot c$$

where $T(x)$ is the precautionary motive term (1.10), q is the probability of an employment transition, and c_{-h} is consumption if the household switches employment status.

[9]Note that this implies that as $dt \to 0$, so that $dc/dt \to \dot{c}$, the variance of dc/dt approaches infinity. This reflects the possibility of discrete jumps in consumption following employment shocks.

There is also a mass of constrained households, who are unemployed at asset level $\underline{a}$. If they remain unemployed, their consumption grows at $\dot{c} = -\ddot{\underline{a}}$, but if they become employed, which happens with probability pdt in time interval dt, they will increase their consumption by amount $c_e(\underline{a}) - c_u(\underline{a})$. Therefore their expected consumption growth is:

$$E\left[\dot{c}_{\underline{a}}\right] = p\left(c_e(\underline{a}) - c_u(\underline{a})\right) - \ddot{\underline{a}}$$

The expected rate of change of aggregate consumption $\dot{C}$ is just the weighted average of the expected changes of consumption of individual households. Moreover, because there is no aggregate risk in the economy, the actual growth of aggregate consumption equals its expectation. Therefore, letting χ be the share of households that are constrained, aggregate consumption growth satisfies

$$\frac{\dot{C}}{C} = \underbrace{(1-\chi)\,\gamma^{-1}\,(r-\rho)}_{\text{interest-rate substitution}} + \underbrace{T(\sigma_C^2)}_{\text{precautionary saving}} + \underbrace{\chi\left(\frac{p\Delta c(\underline{a}) - \ddot{\underline{a}}}{C}\right)}_{\text{constrained households}} \tag{1.14}$$

where $\Delta c(\underline{a}) = c_e(\underline{a}) - c_u(\underline{a})$ is the increase in consumption when a constrained household finds a job, and

$$T(\sigma_C^2) = \int \left[m_e \frac{c_e}{C} \cdot sT\left(\frac{c_u}{c_e}\right) + m_u \frac{c_u}{C} \cdot pT\left(\frac{c_e}{c_u}\right)\right] \tag{1.15}$$

is the consumption-weighted average of the precautionary saving terms of individual households.[10]

[10] $m_e(a,t)$ and $m_u(a,t)$ give the mass of employed and unemployed households, respectively, with assets a at time t. The integral is taken over assets $a > \underline{a}$, meaning that it excludes the point mass of constrained unemployed households households with $a = \underline{a}$. Since employed households are not constrained, there is a negligible mass of employed households with $a = \underline{a}$.

We can derive a more intuitive expression for aggregate consumption growth using the second-order Taylor approximation of $T(x)$ given above. Let $\sigma_e^2(a,t)$ and $\sigma_u^2(a,t)$ be the consumption growth volatility facing a particular employed or unemployed household with assets a at time t. Then, as we showed above, $sT\left(c_u/c_e\right) \approx \left(\frac{1+\gamma}{2}\right)\sigma_e^2$ and $pT\left(c_e/c_u\right) \approx \left(\frac{1+\gamma}{2}\right)\sigma_u^2$. Then we can write the aggregate consumption Euler equation as:

$$\frac{\dot{C}}{C} \approx (1-\chi)\,\gamma^{-1}\,(r-\rho) + \left(\frac{1+\gamma}{2}\right)\sigma_C^2 + \chi\left(\frac{p\Delta c(\underline{a}) - \ddot{a}}{C}\right) \tag{1.16}$$

where $\sigma_C^2 = \int_a \left[m_e \sigma_e^2 \frac{c_e}{C} + m_u \sigma_u^2 \frac{c_u}{C}\right]$ is the consumption-weighted average variance of the growth rate of consumption facing unconstrained households.

Observe that the interest-rate term in (1.14) and (1.16) is multiplied by the share of unconstrained households $1-\chi$. This reflects that, whereas unconstrained households balance consumption volatility and intertemporal substitution, constrained households are restricted to simply consume their current income. The result is that future interest rates are discounted at the rate $1-\chi$, and for this reason I call (1.14) the *discounted* aggregate Euler equation. This discounting reflects that binding constraints on some households make the Euler equation less forward-looking, as emphasized by McKay, Nakamura, and Steinsson (2015). However, (1.14) makes clear that while incomplete markets reduce the sensitivity of aggregate consumption to future interest rates, they also introduce another forward-looking term through precautionary responses to future consumption volatility. Thus whether incomplete markets make the Euler equation more or less forward-looking overall depends on the relative strength of these effects. For instance, if a relatively low fraction of households are constrained at any point in time, but idiosyncratic consumption volatility varies considerably with the business cycle and with monetary policy, then an incomplete markets model may be *more* forward-looking than a complete markets model.

1.2.5 Importance of the aggregate Euler equation

The aggregate Euler equation is critical to the dynamics of a New Keynesian model of the liquidity trap. To understand the source of this model's dynamics, it is useful to contrast the Euler equation above to the standard Euler equation under certainty:

$$\frac{\dot{C}}{C} = \gamma^{-1}\left(r - r^*\right) \tag{1.17}$$

r^* is often called the "natural rate of interest", which under complete markets is equal to the rate of time preference. In the simplest model, the path of output is completely determined by the path of this gap. With no capital, the path of consumption determines the paths of output and employment. Under the maintained assumption that the economy returns to its steady state in the long run, the path of consumption growth rates will then also determine the path of employment.

In this setting, a stylized way of capturing a demand shock is as a fall in the natural rate of interest ρ. If the central bank succeeds in setting $r = \rho$, the only effect will be a fall in the interest rate — in this case the central bank succeeds in fully offsetting the demand shock. However, if r does not fall enough, e.g. because of the zero lower bound on nominal interest rates, then $r > \rho$ will result. This implies a positive growth rate of aggregate output, and under the assumption of convergence to the steady state this likewise implies a large current fall in output.

Rearranging (1.14) slightly, we can express our aggregate Euler equation in terms of a "natural rate of interest" as well:

$$\gamma \frac{\dot{C}}{C} = r - \underbrace{\left[\rho - \gamma T(\sigma_C^2) - \gamma \chi \mu\right]}_{\text{Natural rate of interest}} \tag{1.18}$$

where $\mu = \left(p \Delta c(\underline{a}) - \ddot{\underline{a}}\right)/C - \gamma^{-1}\left(r - \rho\right)$ is the wedge due to binding constraints.

This expression for the natural rate differs in two respects from the complete markets benchmark. First, it contains a precautionary term arising from idiosyncratic consumption volatility due to incomplete markets. If all households were unconstrained, so that $\chi = 0$, this would be the only difference. However, as long as there is a non-trivial mass of households at the constraint, these households will mechanically consume their liquid wealth every period. Thus there are two departures from the Euler equation under certainty, corresponding to the precautionary motive of households responding to consumption volatility, and a further reduction of spending by households who are currently constrained.

What are the consequences of these changes for the model's dynamics? The first consequence is to lower the steady state of interest to $r < \rho$. In (1.18), note that both precautionary terms are positive, implying a higher rate of consumption growth than under certainty. In steady state, we need $\dot{C} = 0$, and so we need $r < \rho$. Further, the magnitude of $\rho - r$ is growing in the degree of income risk in the steady state. By analogy to the Euler equation under certainty, we can speak of the natural rate of interest as the path of r that achieves $\dot{C} = 0$.[11]

While a lower natural rate of interest might seem a mere curiosity, one implication is that the zero lower bound is potentially a greater problem. For a given rate of time preference, the natural rate of interest is lower, and thus $r \geq 0$ is a tighter constraint. Moreover, an increase in steady state income risk or a tightening of borrowing constraints will lower the natural rate still further, pushing the economy closer to $r = 0$.

Perhaps more interesting are the consequences for the dynamic response of consumption to shocks. Any shock that temporarily raises income risk or tightens borrowing constraints will raise the growth rate of consumption and therefore lower current output. Since both events typically happen during a recession, this is a generic amplification channel as long as demand shocks cannot be fully offset by monetary policy. Moreover, there is the potential for additional feedback effects, since this fall in demand will typically raise income risk and tighten borrowing constraints further. This is the amplification channel developed in this paper.

1.2.6 Household saving rules

Above we discussed household decision rules in terms of consumption decision rules. We can equivalently represent household decision rules in terms of the rate of as-

[11] Since income risk and the share of constrained households are endogenous, there may not be a unique natural rate of interest at each point in time. But we can still speak of a natural path of interest rates, and for ease of terminology I will refer to the current rate on one such path as the "natural rate of interest".

set accumulation for employed and unemployed households, denoted by $\dot{a}_e$ and $\dot{a}_u$ respectively. As implied by the discussion above, the rate of asset accumulation is decreasing in asset holdings a, because as households gain wealth, their precautionary saving motive declines (in the limit as a becomes arbitrarily large, $c_u \to c_e$ and the smoothing motive goes to zero).

Assuming $r < \rho$, which will hold in equilibrium, there is a maximum target asset level $\bar{a}$ for employed households that satisfies

$$s\left[\left(\frac{c_u(\bar{a})}{c_e(\bar{a})}\right)^{-\gamma} - 1\right] = \rho - r$$

This is the asset level at which the precautionary motive to accumulate assets is offset by the impatience of consumers relative to the interest rate. At any asset level $a < \underline{a}$, employed households accumulate assets towards the target.

Figure 1.1 represents the saving decision rules of employed and unemployed households with different levels of asset holdings.[12] Unemployed households choose to dissave until they reach the borrowing constraint $\underline{a}$, at which point they choose $\dot{a} = 0$. Employed households choose to accumulate assets up to their target level of assets $\bar{a}$. For both employed and unemployed households, $\dot{a}$ is strictly decreasing in asset holdings a. These are general features of the household decision rules.

An important feature of the household decision rules is that the desired rate of asset accumulation is convex in asset holdings. This is particularly pronounced for unemployed households, since they rapidly decrease borrowing as their total debt approaches the borrowing constraint.This convexity implies that an increase in

[12]This figure depicts decision rules at the steady state defined in section 1.4, using the parameters discussed in section 1.3.4.

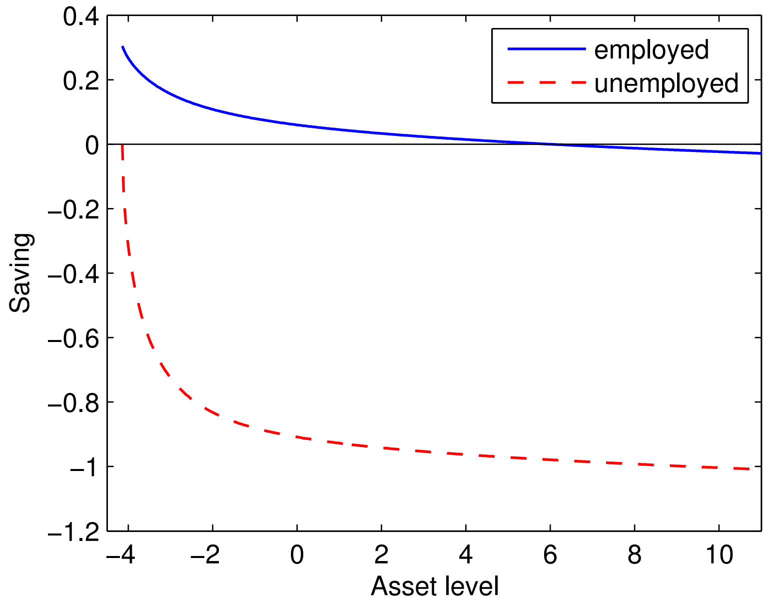

Figure 1.1: Household decision rules.

dispersion of asset holdings will raise aggregate desired savings, or equivalently will decrease desired aggregate consumption. This implies that the distribution of asset holdings is an important determinant of aggregate demand. Since the curvature of the saving rule is greater for unemployed households close to the borrowing limit, the share of households with very low wealth (close to being constrained) is of particular importance. A corollary is that the determination of the borrowing constraint is also of great importance.

The saving decision rule of households depicted in Figure 1.1 matches some features of the data. First, it shows that highly indebted unemployed households have a higher marginal propensity to consume out of wealth than less indebted households. This matches the finding of Mian, Rao, and Sufi (2013) that highly levered households have a higher marginal propensity to consume out of changes in housing wealth, and that this is governed by credit access. Also, the difference between the saving of unemployed and employed households decreases greatly as households become indebted, implying that high-debt households reduce their consumption much more when they become unemployed. This matches the findings of Baker (2014) that the consumption of highly indebted households is more sensitive to changes in their income.

Nevertheless, the saving behavior implied by the model does not perfectly match the stylized facts found in the data. In particular, the decision rule implies that saving rates are decreasing in wealth, whereas they are increasing in the data, as documented by Dynan, Skinner, and Zeldes (2004) and others. More precisely, Figure 1.1 shows that saving rates *conditional on employment status* are decreasing in wealth. However, because employed households are on average wealthier and have greater saving than unemployed househods, in the steady state of the baseline calibration the correlation between wealth and saving is effectively 0. Still, there is no positive relation as documented in the data.

1.2.7 Borrowing Constraint

I assume that indebted households make the decision to repay or default at every point in time t. If a household defaults, it suffers a fixed utility penalty D and is able to borrow again immediately. Thus the value function of an unemployed household that repays at time t is $U(a,t)$, whereas its value function under default is $U(0,t) - D$, and so an unemployed household will repay if and only if

$$U(a,t) \geq U(0,t) - D$$

Since the value function $U(a,t)$ is strictly increasing in a, there is a unique threshold level of assets below which unemployed households default. I assume that lenders set the borrowing constraint equal to this threshold level $\underline{a}(t)$, which is implicitly defined by

$$U(\underline{a}(t),t) = U(0,t) - D \tag{1.19}$$

There is a similarly-defined default threshold for employed households defined by $V(\underline{a}^e,t) = V(0,t) - D$. However, this default threshold will be lower than that of unemployed households, and employed households will choose to accumulate assets at $\underline{a}$. Thus we need only worry about the borrowing constraint of unemployed households, since no household will ever hold assets below this level. Moreover, since employed households will never choose to hold assets $a < \underline{a}$, and unemployed households cannot choose to do so, no default will occur in equilibrium.

Since the cost of default is fixed, endogenous variation in the borrowing constraint is driven by changes in the benefit from defaulting. This benefit is $U(\underline{a},t) - U(0,t)$, which we can write as $\int_{\underline{a}}^{0} U_a da$. Thus the borrowing constraint is defined by

$$\int_{\underline{a}(t)}^{0} \kappa(a,t) da = D$$

where $\kappa(a,t) = U_a(a,t)$ is the marginal value of wealth at asset level a. We can write this is more intuitively as $|\underline{a}| \cdot \bar{\kappa} = D$, where $\bar{\kappa}$ is the average marginal value of wealth over the interval $[\underline{a}, 0]$. Thus the benefit of defaulting is the increase in wealth from defaulting times the average value of wealth.

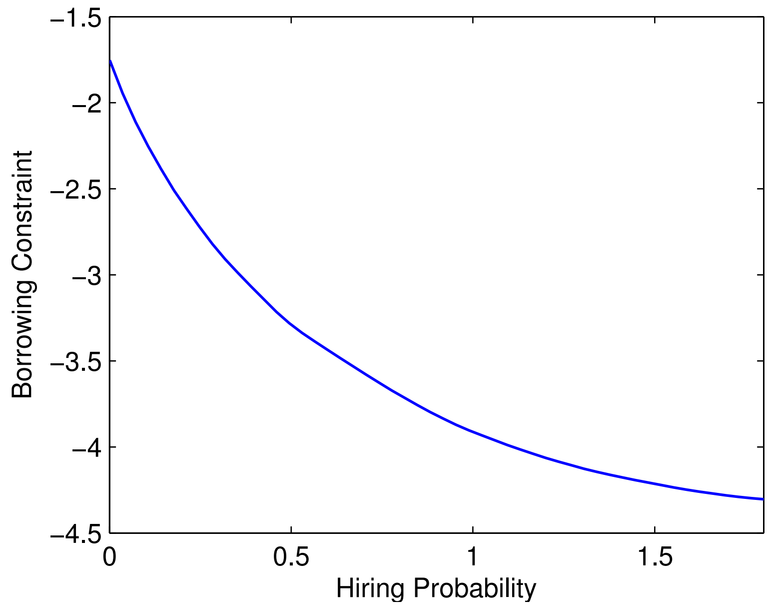

Figure 1.2: Borrowing Constraint for various hiring probabilities.

Holding $\underline{a}$ fixed, any increase in the marginal value of wealth will raise the benefit of defaulting. This will prompt an increase in $\underline{a}$, i.e. a tightening of the borrowing constraint. Intuitively, if wealth is more valuable, there is a greater benefit to defaulting, and so lenders must tighten the borrowing constraint to prevent defaults. One of the chief determinants of the marginal value of wealth for unemployed households is the hiring probability in the near future. When the hiring probability is high, households expect that they will not be unemployed for very long, and so choose a relatively high level of consumption in anticipation of their greater future wealth. This higher consumption implies a lower marginal value of wealth, since $\kappa = c_u^{-\gamma}$. This is the mechanism by which high unemployment causes tighter borrowing constraints, since higher unemployment implies a lower rate of job-finding.

Figure 1.2 depicts $\underline{a}(p)$ for various levels of the job-finding rate p.[13] The figure illustrates that borrowing constraints can vary substantially with the hiring probability. From the beginning of 2008 to the end of 2009, the quarterly job-finding rate in the U.S. fell from about 1.9 to 0.75, which according to Figure 1.2 would tighten the borrowing constraint by over 20%.

1.3 Production, Monetary Policy, and Equilibrium

Now that we have discussed the household problem, I describe the rest of the model and define the equilibrium.

1.3.1 Firms

Final goods are produced from intermediate goods using a Dixit-Stiglitz aggregation technology

$$Y = \left(\int_i y_i^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}$$

where $\epsilon > 1$ is the elasticity of substitution between inputs. Consider the problem of a representative final good firm that purchases intermediate goods from intermediate good producers at prices p_i. The aggregator firm chooses inputs y_i to maximize profits

$$\Pi = P \left(\int_i y_i^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}} - \int_i p_i y_i$$

which yields optimality condition

$$y_i = Y \left(p_i / P \right)^{-\epsilon} \tag{1.20}$$

[13] Parameters for this case are given in section 1.3.4.

Assuming zero profits, the aggregate price level is $P = \left(\int_i p_i^{1-\epsilon}\right)^{\frac{1}{1-\epsilon}}$. The level of output is not pinned down by the price distribution, but is instead determined by the aggregate demand for final goods. Thus the final good producer takes aggregate demand as given, and purchases intermediate goods according to (1.20) to meet this demand. Since in equilibrium all intermediate good prices will be identical, the purchase of intermediate goods will satisfy $y_i = Y$.[14]

Intermediate Good Firms. There is a measure 1 of identical intermediate good firms that face demand function (1.20). Firm i with n_i employees produces output

$$y_i(t) = z(t)n_i(t) - C_i(t)$$

where $C_i(t)$ is the hiring cost paid by the firm for new workers, which takes the form of a loss in output.

Firms perceive equilibrium wage $w^e(t)$ and per-worker hiring cost $\vartheta(t)$. Firms lose workers at an exogenous hazard rate of s and the firm must pay the hiring cost on its gross hires $\dot{n} + sn$. Thus the total hiring cost is:

$$C_i(t) = \vartheta(t) \cdot [\dot{n}_i(t) + s(t)n_i(t)]$$

I consider firm behavior under two pricing regimes: fixed prices and fully flexible prices. The fixed price regime will be the baseline case and may be interpreted as a demand-constrained environment with anchored inflation expectations. This

[14] All intermediate good prices will be identical because we will only consider the extremal cases of perfectly flexible or perfectly fixed prices. If prices were partially flexible, we would have price dispersion, with resulting aggregate productivity losses.

seems a good approximation to the circumstances that prevailed in the US and Europe in 2008 – 2014. The flexible price case is useful mainly as a benchmark for optimal behavior, and will be used to define the central bank policy target. I assume that the central bank would like to replicate the flexible price benchmark.[15]

Fixed Prices. Suppose that firms' prices are fixed at $p(t) = 1$, which implies that the aggregate price level is $P(t) = 1$. As is standard in New Keynesian models, firms are assumed to meet the demand that they face at this price. Thus firm i chooses a path of $n_i(t)$ that satisfies

$$d_i(t) = y_i(t) \tag{1.21}$$

where $d_i(t)$ is demand for the firm's output.

Thus at every point in time, I assume that firms set employment $n(t)$ equal to demand. This may imply an instantaneous adjustment of aggregate employment at time $t = 0$, as initial employment $n(0)$ jumps to the equilibrium path. However, I assume that the path of demand $d(t)$ and therefore of employment $n(t)$ is continuous and differentiable for $t > 0$, so that $\dot{n}(t)$ and $p(t)$ are defined for all $t \geq 0$.[16] I further assume that firms will be able to hire sufficient workers to meet the demand they face, which will be true in all the cases we consider.

I assume that wages are a constant markdown of labor productivity, $w(t) =$

[15] This corresponds to setting the interest rate equal to the "natural rate". In a standard New Keynesian model, such a policy is optimal since it produces zero inflation and a zero output gap. In the present model this policy is suboptimal, because there are frictions in the labor and capital markets. Nevertheless, this assumption provides a simple and intuitive policy target.

[16] $\dot{n}(0)$ is defined as the right-derivative of $n(t)$ at $t = 0$, i.e. as the limit of $\dot{n}(t)$ as $t \to 0$.

$\varphi z(t)$ for $\varphi \in (0,1)$. This implies that wages are unaffected by fluctuations in demand, although they respond to changes in productivity.[17] Under this assumption, firm i's flow profits are

$$\pi_i(t) = (1-\varphi)\, z(t) n_i(t) - \vartheta(t) \cdot [\dot{n}_i(t) + s n_i(t)] \tag{1.22}$$

where $\vartheta(t)$ is the hiring cost, and total hires are $\dot{n}(t) + sn(t)$ since firms must hire a flow of $s \cdot n(t)$ workers to replace workers lost to exogenous separation. I assume that households all own equal shares of firms, and that net profits are paid to households every period, so that nonlabor income is $e = \bar{e} + \pi$.[18]

Flexible Prices. Now suppose that firms may adjust their prices and choose any (p_i, y_i) consistent with (1.20). Further suppose that the government provides a production subsidy to firms, financed by a lump-sum tax, chosen to exactly offset the incentive for underproduction due to market power. If the government sets the production subsidy equal to $\tau = \frac{1}{\epsilon - 1}$, then the firm problem is isomorphic to maximizing discounted flow profits (1.22).[19]

[17] Many researchers consider sticky wages to be a promising mechanism to produce greater responses of employment to shocks (e.g., Hall (2005), Pissarides (2009), and Galí (2011)). However, in the present model sticky wages are not necessary to produce employment volatility because under fixed prices the path of employment is determined by demand regardless of the incentives for job creation. The main consequence of fixed wages is to simplify the central bank policy rule.

[18] In the baseline calibration, steady state profits amount to just 0.25% of GDP, so they make little difference to the analysis.

[19] Specifically, firm revenue net of taxes is $R = (1+\tau)\, PY^{\frac{1}{\epsilon}} (zn)^{\frac{\epsilon-1}{\epsilon}} - \tau PY$, and the marginal revenue product of labor is $R_n = \frac{\epsilon-1}{\epsilon} (1+\tau)\, PY^{\frac{1}{\epsilon}} (zn)^{\frac{\epsilon-1}{\epsilon}} \frac{1}{n}$. In a symmetric equilibrium ($P = 1$ and $Y = zn$), this is equivalent to $R = zn$ and firms perceive $R_n = z$.

Suppose that firms discount flow profits at the same rate as households ρ. Firms choose the path of hiring to maximize discounted profits

$$\int_0^\infty e^{-\rho t} \pi_i(t) dt$$

Let $J(n,t)$ be the value function of a firm at time t with employment n. Then firms choose hires $\dot{n}$ to maximize the Hamiltonian

$$\rho J(n,t) = \max_{\dot{n}} \left\{ (1-\varphi)\, zn - \vartheta \left(\dot{n} + sn \right) + J_n \dot{n} + J_t \right\}$$

The firm optimality condition is $J_n = \vartheta$, which implies that firms hire workers until the marginal value of an additional worker is equal to the hiring cost.[20] The costate equation of the Hamiltonian is $(\rho + s)\, J_n = (1-\varphi)\, z + \dot{J}_n$. Combining these, we obtain the flexible price job creation condition

$$(\rho + s)\, \vartheta = (1-\varphi)\, z + \dot{\vartheta} \tag{1.23}$$

Equation (1.23) has a straightforward interpretation. When $\dot{\vartheta} = 0$, so that the cost of hiring workers is constant over time, it says that the cost of hiring must equal the present discounted value of profits that the firm will receive from this worker, i.e. $\vartheta = \frac{(1-\varphi)z}{r+s}$. Hiring behavior is altered somewhat when the cost of hiring is expected to change over time. For example, if $\dot{\vartheta} < 0$ so that hiring is becoming less expensive over time, then firms have an incentive to defer hiring until the future when it is less costly. They will do so until they drive down the hiring cost sufficiently that they are indifferent between hiring the marginal worker today, or hiring the worker in the future at a somewhat lower cost.

[20] By assumption, firms are small and do not internalize that their hiring decisions affect the aggregate hiring cost.

Under a few assumptions, (1.23) becomes quite simple. First suppose that the discount rate of firms ρ, the separation rate s, and productivity z are constant over time.[21] Further assume that the economy will eventually reach a steady state equilibrium with interior employment (i.e. $n \in (0,1)$). Finally, suppose that the hiring cost ϑ is a simple increasing function of the job-finding probability of households p. Then (1.23) implies a constant job-finding rate p^* that satisfies $\vartheta(p^*) = \frac{(1-\varphi)z}{\rho+s}$.

If we use our baseline functional form $\vartheta(p) = \psi p^\alpha$, this becomes[22]

$$p^* = \left(\frac{1}{\psi}\frac{(1-\varphi)\,z}{\rho+s}\right)^{1/\alpha} \tag{1.24}$$

We can write (1.24) more intuitively in its implicit form as

$$\underbrace{\psi\,(p^*)^\alpha}_{\text{hiring cost}} = \underbrace{(1-\varphi)}_{\text{profit share}} \times \underbrace{\left(\frac{z}{\rho+s}\right)}_{\text{PDV output}}$$

which makes clear that p^* corresponds to the rate of hiring at which expected discounted profits from the marginal hire equal the hiring cost.

Optimal hiring is given in terms of an optimal *job-finding rate* instead of an optimal *hiring rate* because the cost of hiring a worker is increasing in the tightness of the labor market. The standard definition of tightness depends on the aggregate

[21] If the firm discount rate were not fixed, cyclical variations in firm discount rates could significantly affect hiring incentives. For an analysis of such effects, see Hall (2014). Since the present paper is chiefly concerned with the determination of demand rather than incentives for job creation, I avoid these complications by assuming a fixed discount rate.

[22] This functional form is isomorphic to a fixed vacancy-posting cost combined with a Cobb-Douglas matching function, as in Shimer (2005). If the flow cost of maintaining a vacancy is c, and the matching function is $m = \mu v^b u^{1-b}$, then the cost per worker hired is $\vartheta = c\mu^{-\frac{1}{b}} p^{\frac{1-b}{b}}$.

hiring rate relative to the number of unemployed workers, which is equivalent to the job-finding rate of households. Thus when unemployment is high, the optimal rate of hiring is higher because it is less costly to hire workers (we can interpret this as a high rate of filling posted vacancies in a model with explicit search). This works out to a constant job-finding rate perceived by workers.

1.3.2 Equilibrium under Fixed Prices

Since prices are fixed, there is no inflation and the real interest rate equals the nominal interest rate. Thus the path of interest rates $r(t)$ is set by the central bank. Likewise the paths of non-labor income $e(t)$, labor productivity $z(t)$, and the separation rate $s(t)$ are given.

The law of motion of aggregate employment satisfies

$$\dot{n}(t) = p(t) \cdot (1 - n(t)) - s(t) \cdot n(t)$$

We need a market clearing condition in the asset market. By Walras' law, if the asset market clears, the goods market clears as well. Let $m_e(a,t)$ and $m_u(a,t)$ be the mass of employed and unemployed households with assets a at time t. Thus m_e and m_u are the probability density functions of the asset distribution across households. m_e and m_u are related to employment n by

$$n(t) = \int_a m_e(a,t)da \tag{1.25}$$

$$u(t) = 1 - n(t) = \int_a m_u(a,t)da \tag{1.26}$$

In equilibrium, we need aggregate asset holdings at every time t to equal zero.

We can express this as

$$\int_a a \cdot m_e(a,t)da + \int_a a \cdot m_u(a,t)da = 0 \tag{1.27}$$

Here $\int_a a \cdot m_e(a,t)da$ is total assets held by employed households, and $\int_a a \cdot m_u(a,t)da$ is total assets held by unemployed households.

When the economy is not in its long-term steady state, the asset distribution will evolve over time. The law of motion for the asset distribution is

$$\dot{m}_e = pm_u - sm_e - \frac{d}{da}\left(m_e \dot{a}_e\right) \tag{1.28}$$

$$\dot{m}_u = sm_e - pm_u - \frac{d}{da}\left(m_u \dot{a}_u\right) \tag{1.29}$$

where (1.28) and (1.29) must be consistent with the law of motion of aggregate labor.

Equations (1.28) and (1.29) can be interpreted as *flow equations*. The term $m\dot{a}$ is the mass of households at a point in the asset distribution times their rate of asset accumulation. Since the rate of asset accumulation is the "velocity" of that household along the asset dimension, we can interpret this term as the rate of "flow" of households through a point on the asset dimension, moving from lower to higher assets. Then the rate of change of the mass of households in the neighborhood of a point in the asset distribution is the difference between the rate of flow *into* that neighborhood minus the flow *out* of that neighborhood. This is equivalent to minus the slope of the flow rate along the asset dimension. For example, if $\frac{d}{da}(m\dot{a}) < 0$, so that the rate of flow is decreasing in assets in the neighborhood of a point in the asset distribution, then the flow into that neighborhood is greater than the flow out of that neighborhood, and so the mass of households in that neighborhood is increasing.

Definition 1 (Equilibrium)**.** Given a path of $\{z, e, r, s, D\}$ and initial asset distribution $m_e(a, 0)$ and $m_u(a, 0)$, an equilibrium is a path of $\{m_e, m_u, V, U, c_e, c_u, \dot{a}_e, \dot{a}_u, w, e, \pi, p, \underline{a}, n, u\}$ that satisfies (1.1) - (1.19), (1.22), and (1.25) - (1.29).

1.3.3 Equilibrium Determinacy and Central Bank Policy

Definition 1 specifies an equilibrium for a given path of exogenous variables and initial conditions. However, it leaves open the question of what path of interest rates the central bank sets, and whether a unique equilibrium exists for a given interest rate path. Our benchmark assumption is that the central bank would like to set interest rates in order to replicate the flexible price job-finding rate p^* defined by (1.24). We work out the implications of this assumption for equilibrium determinacy.[23]

Perfectly Flexible Interest Rates. First suppose that there are no restrictions on the path of interest rates that the central bank can select. Then the central bank will always be able to hit its policy target, and we know that the job-finding rate satisfies $p = p^*$. However, without further assumptions this does not specify

[23]This policy rule is equivalent to setting the interest rate equal to the Wicksellian natural rate. This would be optimal if the only friction were sticky prices. However, in this model the flexible-price equilibrium is not Pareto optimal because of the presence of incomplete markets. We may interpret this policy rule as a central bank that limits itself to short-run stabilization, and does not seek to correct inefficiencies arising from long-term structural features of the economy.

an equilibrium: the initial level of employment $n(0)$ can take on any value, which together with constant job-finding rate $p = p^*$ will imply a particular path of employment. These together will imply a path of the interest rate r that is consistent with this path of employment, i.e. in which the path of demand exactly equals the path of output implied by this path of employment.

While many such equilibria exist, and each are consistent with the central bank's policy target, they intuitively correspond to disruptions in the rate of employment by the central bank. If we suppose that the economy starts at the steady state level of employment $n(0) = p^*/(p^* + s)$, then a fall in $n(0)$ would correspond to the central bank raising interest rates and then lowering them over time to maintain a constant job-finding rate. Thus a reasonable equilibrium selection rule is that the central bank will choose to keep employment at its long-run steady state level when it can do so.

Lower Bound on Interest Rates. Now suppose that there is a lower bound on the interest rate $\underline{r}$, so that only policy paths that satisfy $r(t) \geq \underline{r}$ are possible. Then it might not be possible for the central bank to hit its target job-finding rate $p = p^*$. In particular, if the flexible interest rate equilibrium requires the interest rate to fall below $\underline{r}$ at any point, then this equilibrium violates the constraint.

We must now modify the policy rule of the central bank. Instead of assuming a fixed job-finding rate $p(t) = p^*$ and a variable interest rate, suppose that at every point in time *either* the central bank has successfully hit its target $p(t) = p^*$ and the interest rate takes on some value $r \geq \underline{r}$, *or* the interest-rate constraint is binding and the central bank cannot hit its target, meaning that $r = \underline{r}$ and $p(t) \leq p^*$. Note that the central bank can always kill off an excessive boom by raising interest rates, so we don't need to worry about $p > p^*$. This yields the *constrained policy rule*:

$$r(t) \geq \underline{r} \text{ and } p(t) = p^* \quad \text{OR} \quad r(t) = \underline{r} \text{ and } p(t) \leq p^* \tag{1.30}$$

The assumption that the central bank will not allow $p(t) > p^*$ implies a lack of commitment on the part of the central bank. Situations may arise when the lower bound on the interest rate causes a period of low employment and low job-finding, that the central bank could reduce if it could credibly promise to allow $p(t) > p^*$ after the liquidity trap has concluded. In section **??** I will analyze what happens if we allow such forward guidance.[24]

Under (1.30), there are again multiple equilibrium paths. This is analogous to the multiplicity that exists when the interest rate is unconstrained, but in this case it is less reasonable to assume that the central bank can pick any equilibrium that it pleases since its choice of interest rates is constrained.

To fix intuitions, consider the case that the economy experiences an unanticipated adverse demand shock at $t = 0$ that causes $r < \underline{r}$ in the flexible price equilibrium, i.e. the lower bound on the interest rate binds. Suppose further that the economy permanently exits the liquidity trap at some future point T^*.[25] Then

[24]This no commitment assumption is similar to the baseline case considered by Werning (2011), who considers a liquidity trap that ends at a fixed time T, and assumes that the central bank implements the no commitment equilibrium for $t \geq T$.

[25]This formulation is similar to that used in Werning (2011) and related papers in the liquidity trap literature, except that in my model the date of exit from the liquidity trap T^* is endogenous.

from the policy rule (1.30), we know that $\forall t \leq T^*$, $r(t) = \underline{r}$ and $p(t) \leq p^*$, and $\forall t \geq T^*$, $r(t) \geq \underline{r}$ and $p(t) = p^*$. Clearly this implies that $r(T^*) = \underline{r}$ and $p(T^*) = p^*$. The period $t < T^*$ corresponds to the liquidity trap, when hiring is below target and the interest rate is against the zero lower bound. The period $t \geq T^*$ is after the liquidity trap, when the central bank can achieve its hiring target by setting $r \geq \underline{r}$.

If we knew the date of exit from the liquidity trap T^*, this would pin down the equilibrium. However, multiple T^* are possible. In particular, we can always assume that households expect a later T^*, in which case demand will be lower throughout, and the central bank cannot offset this lower demand by lowering interest rates because of the lower bound.

But note that not all T^* can be an equilibrium. For instance, suppose that a very large adverse demand shock occurred at $t = 0$, and suppose that $T^* \to 0$. Then when we enter the post-liquidity trap period, we have $p(t) = p^*$ forever. Then no matter what initial level of employment $n(0)$ is chosen, employment will converge to steady state employment fairly rapidly. If the demand shock has lasted beyond this point (which it could), and since $r(t) \geq \underline{r}$ and so cannot be lowered further to boost demand, this will not be an equilibrium for a sufficiently large and persistent demand shock. Thus for any initial demand shock there exists a continuum of equilibria, corresponding to various dates of exit from the liquidity trap, $T^* \in [\underline{T}, \infty]$. A later exit corresponds to a larger initial fall in initial employment $n(0)$, and a lower path of recovery.

In my baseline experiments, I will assume that the economy follows the equilibrium with the highest path of employment $n(t)$, which in the demand shock cases corresponds to the smallest T^* that is consistent with equilibrium. In section 1.7 I consider the possibility of worse equilibria, and argue that these correspond to pessimistic expectations about the pace of recovery.

1.3.4 Calibration

I use the following parameters in the baseline calibration:

Parameter	Value	Source
z	1	Normalization
e	0.5	Targets income share about 1/3
s	0.1	Shimer (2005)
φ	0.9843	Achieves $p = 1.35$ target
γ	1	Gourinchas and Parker (2002)
ρ	0.39%	Achieves $r^{ss} = 0.25\%$ target
D	3.39	Achieves $\underline{a} = -4.16$ target
$\vartheta(p)$	ψp^{α}	Blanchard and Galí (2010) (isomorphic to Shimer (2005))
α	1	Blanchard and Galí (2010)
ψ	0.115	Matches job-finding cost in Shimer (2005)

Table 1.1: Baseline Calibration

The quarterly separation rate of 0.1 corresponds to the 3.3% monthly separation rate reported as the average job separation rate for the US economy in 1951 – 2003 by Shimer (2005). The target steady state job-finding rate for unemployed households is 1.35, or 0.45 on a monthly basis, which is the 1951 – 2003 average reported by Shimer (2005).[26] The wage share of output φ is set to achieve this target in the flexible price equilibrium. By way of comparison, the wage share of output in the baseline calibration in Shimer (2005) is about 0.9825, so this number is similar.

The coefficient of relative risk aversion $\gamma = 1$ is a standard choice in business cycle models, i.e. in Blanchard and Galí (2010). It is also the middle of the range of estimates reported by Gourinchas and Parker (2002) based on their analysis of household behavior in the Consumer Expenditure Survey and Panel Study of Income Dynamics. Note that this implies a fairly low level of risk aversion compared to many other papers, such as Guerrieri and Lorenzoni (2011), and thus represents a conservative estimate of the degree of risk aversion among households.

The target steady state interest rate is 0.25% on an annual basis. This target was chosen so that the zero lower bound is effectively binding in the steady state. While this calibration is mainly chosen for tractability, such an occurrence is not completely unreasonable, given that Japan has experienced zero short-term policy rates for most of the last 25 years, as have the U.S. and Europe since 2008. Some commentators have suggested that a combination of well-anchored inflation expectations, low population growth rates, and slow technological progress have lowered the natural (flexible price) rate of interest to zero or below, a situation known as "secular stagnation."[27] If they are right, this calibration could be literally correct.

[26]To interpret these numbers, a constant quarterly job-finding flow probability of 1.35 implies a 74.1% total chance of finding a job in a quarter, and a 0.1 separation probability corresponds to a 9.5% probability of losing one's job over the course of a quarter.

[27]The concept of secular stagnation was first proposed by Hansen (1939). Eggertsson and Mehro-

Nevertheless, the qualitative results of the model do not depend on this assumption.

The steady state borrowing constraint $\underline{a}$ in absolute value equals 4.16 times quarterly labor income, as in Guerrieri and Lorenzoni (2011). This produces an analogous asset distribution across households, but shifted to the left. The shift is because I assume that assets are in zero net supply, so that the distribution is centered at zero, whereas Guerrieri and Lorenzoni (2011) assume positive net aggregate asset holdings of households due to a supply of government bonds.[28]

The hiring cost is assumed to be of the form $\vartheta = \psi p^{\alpha}$. I choose $\alpha = 1$, which corresponds to $b = \frac{1}{2}$ in the Cobb-Douglas matching function. In addition to being tractable, this form is consistent with estimates of the matching function according to Blanchard and Galí (2010), who use the same functional form. I set ψ so that the hiring cost in the steady state equals that in Shimer (2005).

1.4 Steady State

For given labor productivity z, nonlabor income e, and separation rate s, the equilibrium conditions, together with the steady state condition $\dot{m}_e = 0$ and $\dot{m}_u = 0$, define steady state equilibrium. In the fixed price case there is a continuum of steady state equilibria, corresponding to different pairs $\{r, p\}$. We may think of each of these as corresponding to different targets of the steady state job-finding rate (and associated employment rate), together with the steady state interest rate that achieves

tra (2014) present a model of secular stagnation.

[28]However, the direct comparison is a bit tricky, because Guerrieri and Lorenzoni (2011) have households that differ in labor productivity.

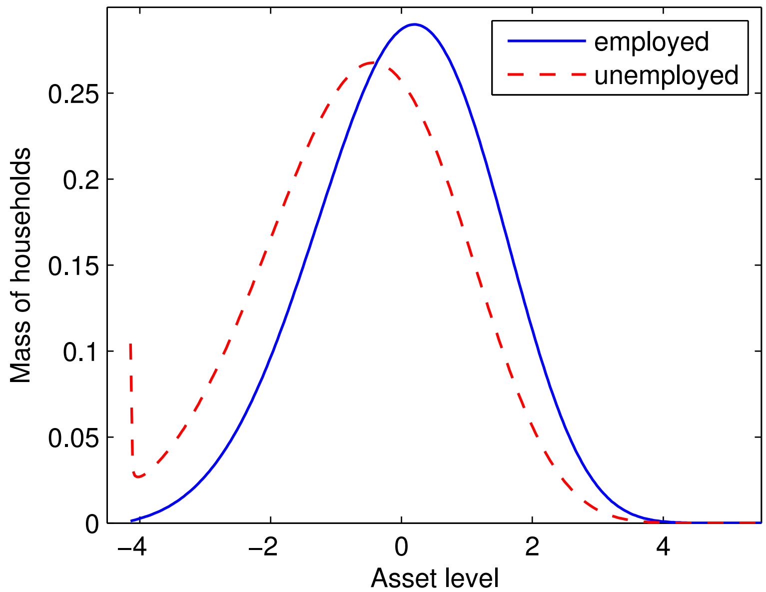

Figure 1.3: Household asset distribution.

this target. One of these equilibria has job-finding rate p^*, which corresponds to the flexible price equilibrium.

For a given job-finding rate p and interest rate r, we can define household saving rules $\dot{a}_e$ and $\dot{a}_u$. Together with $\dot{m}_e = 0$ and $\dot{m}_u = 0$, and the relative mass of employed and unemployed households given by $n = p/(p+s)$, these define a stationary distribution of asset holdings.

The steady state asset distribution for the baseline calibration is depicted in Figure 1.3. The figure depicts the probability mass of households conditional on employment status at each asset level. There is a point mass of unemployed households at the borrowing constraint, but otherwise the asset distribution has a smooth bell shape.

For given (r, p), we can therefore define aggregate desired steady state asset

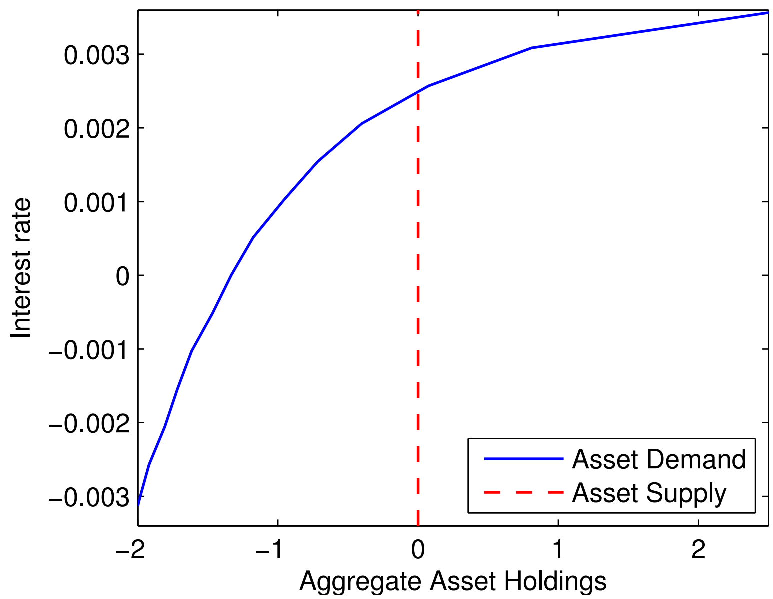

Figure 1.4: Steady state equilibrium.

holdings $A(r, p)$. Since there is no storage or capital in the economy, aggregate total asset holdings must equal zero in equilibrium. Therefore a steady state equilibrium pair $\{r, p\}$ must satisfy $A(r, p) = 0$.

The equilibrium for the flexible price job-finding rate p^* is depicted in figure 1.4. For fixed p, aggregate steady state asset demand is an upward-sloping schedule in r, which we may think of as the steady state equivalent of the supply of savings.[29] A higher interest rate induces households at all wealth levels to save more, shifting out the steady state asset distribution and raising asset demand. The supply of assets is a vertical line at $A = 0$, and equilibrium is the intersection of the two curves.

[29]This is not a supply curve in the traditional sense, because each point corresponds to a different desired steady state asset distribution. A true aggregate saving curve would depict aggregate desired savings for various interest rates for a fixed asset distribution.

What happens to the steady state if there is an increase in the job-finding rate p? For households at every asset level and employment status, higher p implies greater lifetime wealth, and so households increase their consumption, or equivalently *decrease* their savings. However, a higher steady state p also implies a higher steady state employment rate n, which will tend to *increase* desired savings, because employed households save more than unemployed households.

Which effect dominates depends on the current level of n. When n is high, an increase in p will decrease desired savings, while the reverse holds for low n. The reason is that if the majority of households are employed, an increase in p implies a reduction in the aggregate income risk facing households, which will reduce desired precautionary savings. Conversely, when n is low, and in particular when $n < 0.5$, an increase in p will *raise* income risk facing households, leading to an increase in desired savings.[30] Since in all advanced economies the employment rate is well above 50%, the possibility that an increase in p may raise savings is not empirically relevant, and so I will assume for the remainder that aggregate steady state asset demand is strictly decreasing in p.

Figure 1.5 depicts the result of an increase in the steady state job-finding rate p on the steady state equilibrium. The asset demand curve shifts left because higher p implies lower income risk, which depresses desired savings. This shift in asset demand causes the equilibrium interest rate to rise, implying a positive relationship

[30] If this is unintuitive, consider what happens at $p = 0$, so that all households are unemployed in equilibrium, and there is zero income risk. Then an increase in p will clearly raise income risk, since it is now positive.

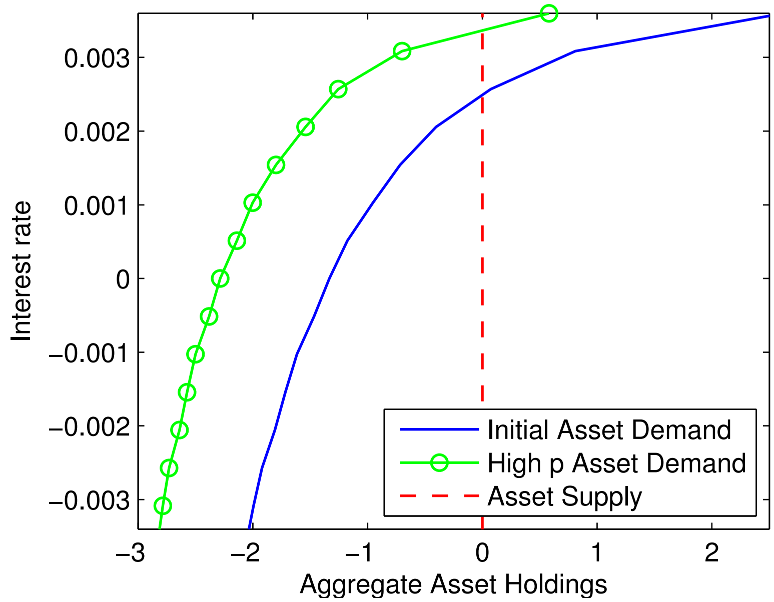

Figure 1.5: Steady state equilibrium with higher p.

between steady state r and steady state n.

We can equivalently interpret this in terms of the aggregate Euler equation (1.14). An increase in p will lower aggregate income volatility, which lowers aggregate consumption volatility $T(\sigma_C^2)$. Then $\dot{C}$ will decrease, and so to achieve $\dot{C} = 0$, which must prevail in steady state, r must rise. Under our earlier definition, this is equivalent to saying that the natural rate of interest has risen.

Since steady state r is increasing in p, the set of steady state equilibria in the fixed price case comprise an upward sloping schedule in (r, p) space.

1.5 Dynamics Following a Credit Shock

Suppose that the economy is in the steady state equilibrium for $t < 0$, and at $t = 0$ experiences an unanticipated shock. In particular, consider a temporary fall in the default penalty D. Since lower D causes borrowing constraints facing consumers to tighten, we can interpret this as a credit shock.[31]

1.5.1 Fixed Price Equilibrium

Consider the dynamic path of the economy under fixed prices in a liquidity trap scenario. The key feature of a liquidity trap is that the interest rate cannot fall below some lower limit. The simplest way to capture this is to suppose that the interest rate is fixed at its steady state level.

Figure 1.6 depicts the path of the economy following a temporary credit shock under the baseline equilibrium selection rule discussed in section 1.3.2. The particular experiment is that the default penalty D falls discretely at $t = 0$, remains at this lower level for 8 quarters, and then steadily recovers to its steady state level over the next 8 quarters. This corresponds to an exogenous discrete worsening of credit conditions for two years, followed by a recovery in credit conditions over the following two years. The initial fall in D is about 1/3 of its steady state value, which was chosen to produce a fall in initial employment of 5 percentage points, equal to the rise in the unemployment rate in the U.S. between January 2008 and October 2009, i.e. from the official beginning of the recession until the trough of the labor market.

[31] Because an unanticipated shock tightens the borrowing constraint instantaneously, some households will violate the constraint. I assume that these households are forced to deleverage very quickly. In particular, I extrapolate these households' saving rules below $\underline{a}$, but place a minimum on their consumption at 0.01.

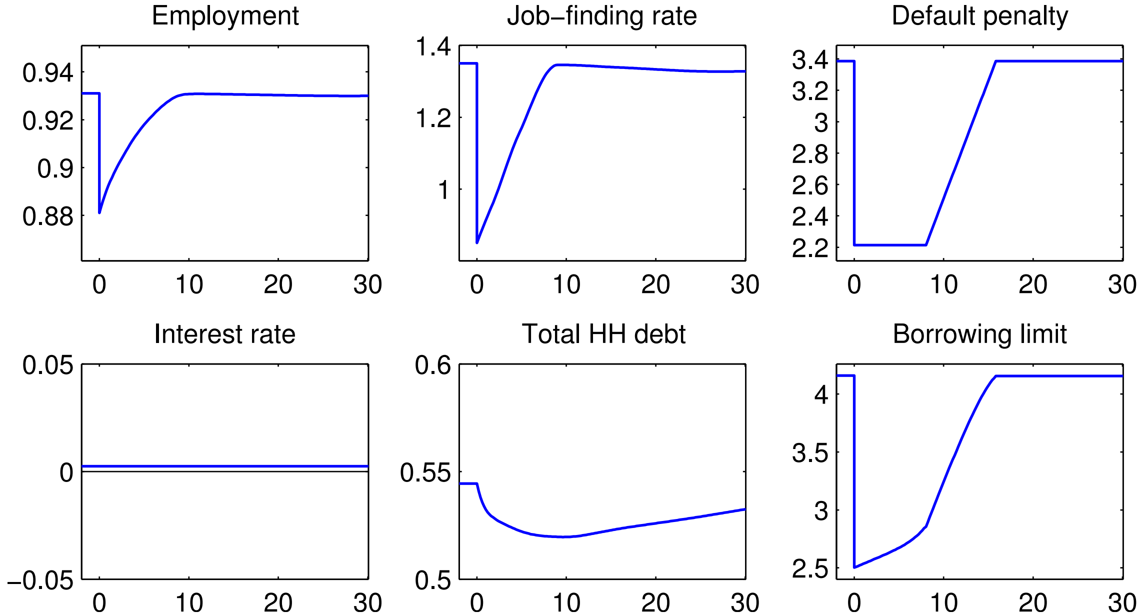

Figure 1.6: Credit shock with fixed interest rate.

Employment initially falls sharply, and then steadily returns to its steady state level after approximately 10 quarters. These dynamics are driven by consumption demand from households. Following the credit shock, demand falls for two reasons. First, there are some households that are at or below the borrowing constraint, and are thus mechanically forced to reduce their spending, what we might call a forced deleveraging effect. Second, households throughout the asset distribution face tighter borrowing constraints that might bind in the future. This reduces households' ability to smooth consumption over unemployment spells, increasing precautionary saving.

The job-finding rate falls and rises in tandem with the employment rate. This occurs because the shock is persistent, and so employment remains depressed for a period of time. This implies a period of weak hiring, and so the job-finding rate falls and remains low during the recovery. As demand recovers, the job-finding rate rises in tandem with employment as firms increase hiring to boost production.

This period of low hiring amplifies the initial fall in demand. This happens both because slow hiring raises income volatility facing households, increasing precautionary savings, and because slow hiring endogenously tightens the borrowing constraint. The latter effect can be seen in the path of the borrowing constraint in Figure 1.6. The borrowing constraint recovers as the hiring rate increases, well before the default penalty begins to increase. Since both of these factors lower demand, they make the initial fall in employment larger than it would otherwise be, i.e. they act to amplify the initial shock.

1.5.2 Precautionary Saving versus Deleveraging

As discussed above, demand falls both because constrained households are forced to reduce spending, a *deleveraging effect*, and because tighter constraints reduce households' ability to smooth consumption, and so raise precautionary saving by unconstrained households, a *precautionary effect*. In this section we compute the relative contributions of these two effects to the total fall in demand and therefore employment.

As discussed in section 1.2.4, the path of demand is governed by the aggregate Euler equation, which can be expressed in terms of the gap between the rate of interest set by the central bank and the natural rate. The fall in employment following the credit shock is therefore equivalent to a fall in the natural rate of

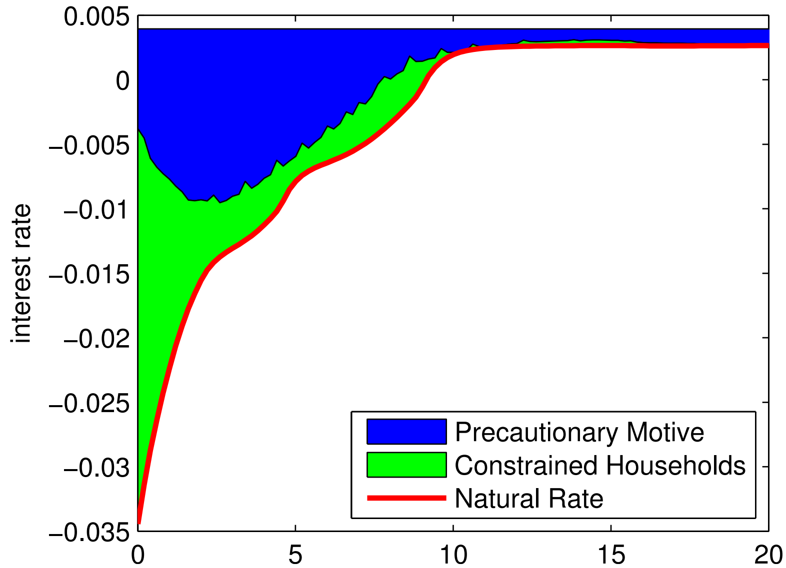

Figure 1.7: Decomposition of natural rate following credit shock.

interest. We can use equation (1.14) to decompose the fall in the natural rate into a component due to deleveraging by constrained households, and a component due to the precautionary response to higher consumption volatility.

Figure 1.7 depicts this decomposition for the baseline credit shock. The natural rate falls sharply to about -3.5% immediately following the shock, and most of this fall is due to a reduction of spending by constrained households. Before the shock, there is an initial mass of households near the constraint (Figure 1.3). When the constraint tightens these households are forced to deleverage, causing an immediate drop in demand. After this initial deleveraging, which lasts about 2 or 3 quarters, demand remains depressed for a time, mainly due to the precautionary effect. The precautionary motive accounts for the majority of the reduction in the natural rate for the remainder of the crisis.

By integrating over the contribution of each term to the fall in the natural rate, we can calculate the contribution of each component to the total fall in demand at time $t = 0$. Doing so reveals that 54% of the initial fall in aggregate demand is due to precautionary saving effects, whereas 46% is due to deleveraging. Thus both deleveraging and precautionary behavior contribute to the fall in employment, but the contribution of the precautionary motive is somewhat larger. This suggests that analyses that focus exclusively on the deleveraging behavior of constrained households may miss a big part of the story.

1.5.3 Measure of Amplification

We now turn to determining the magnitude of the amplification from endogenous unemployment risk. Amplification arises from the interaction of endogenous variation in the job-finding rate with the precautionary behavior of households. Thus a natural comparison case is the complete-markets benchmark — i.e. the baseline model with the additional assumption that households are able to fully insure against employment shocks. Since households are risk-averse, the result will be full risk-pooling, so that all households enjoy the same level of consumption and zero assets. In this case, the aggregate Euler equation is the standard complete markets Euler equation:

$$\frac{\dot{C}}{C} = \gamma^{-1}(r - r^*)$$

and the dynamics of demand, and therefore employment, are determined by the path of $r - r^*$. The model is otherwise identical.

The main obstacle we face in computing a complete markets comparison case is choosing the appropriate shock. The baseline financial shock will not affect the complete markets benchmark at all — since no household wants to borrow, a tighter borrowing constraint makes no difference. The standard way to model a demand shock in this setting is as an exogenous change in the natural rate of interest r^*.[32] Thus we would like to calculate a path of r^* that corresponds to a reduced-form representation of the baseline credit shock in the full incomplete markets model.

To do so, we first compute the flexible price equilibrium following the credit shock. This implies a path of the interest rate that stabilizes output and the job-finding rate. This path corresponds to another notion of the "natural rate" — the interest rate that the central bank would have to set to fully offset the baseline demand shock. We can take this path of the interest rate as r^* in the complete markets benchmark, i.e. as a reduced form representation of the initial demand shock in the absence of feedback from endogenous variation in unemployment risk to demand.

Given this path of the natural rate, we can now compute the complete markets comparison case. The central bank is again constrained by the lower bound on r, but the natural rate r^* now follows the path defined above.[33] The aggregate Euler equation then implies a path of demand and output, and this path of output implies a path of employment. Comparing this to the path of employment from the full model gives a measure of amplification from endogenous employment risk. Further, we can compare the path of r^* to the "natural rate" computed in the previous section,

[32]For instance, this is how demand shocks are modeled in Werning (2011).

[33]Here r^* is equivalent to the households' time discount factor ρ

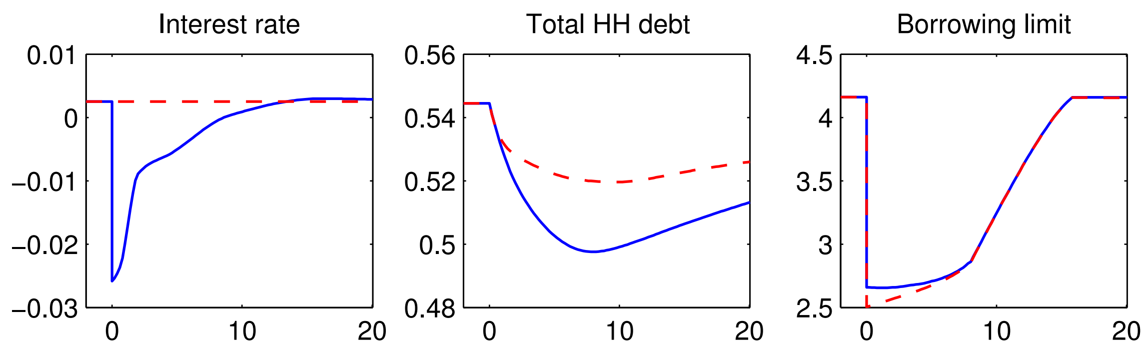

Figure 1.8: Credit shock with flexible prices (solid) vs. fixed prices (dashed).

which allows us to decompose the amplification into deleveraging and precautionary components.

Flexible Price Benchmark. What would happen if the central bank were able to adjust the interest rate in response to the credit shock? In this case the central bank would set r to replicate the flexible price equilibrium. Since the credit shock does not affect the incentives for production, the firm optimality condition (1.23) implies a constant rate of hiring. Thus employment would remain constant at its steady state level, and the interest rate would adjust to produce the conusmption demand necessary to clear the market.

This equilibrium is depicted in Figure 1.8. The resulting path of the interest rate is shown in the left panel. Since the tightening of the borrowing constraint reduces demand, the interest rate must fall to stimulate spending and stabilize output. Since output remains high, households are able to reduce their debt holdings to a much greater degree, as depicted in the middle panel. Finally, the right panel shows that the borrowing constraint is a little looser under flexible prices, because

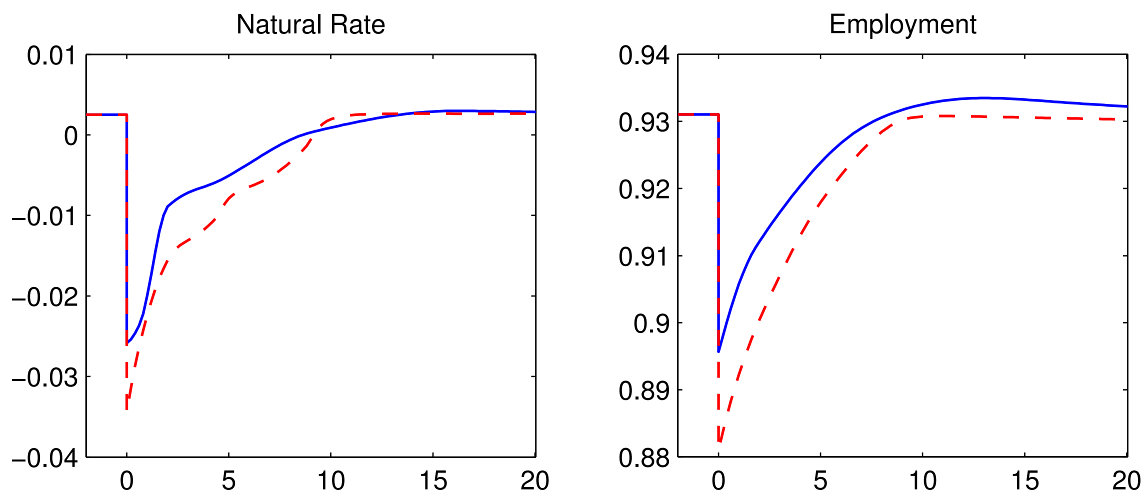

Figure 1.9: Complete markets (solid) vs. incomplete markets (dashed) models.

the improved labor market conditions reduce the endogenous tightening of the constraint.

Comparison to Complete Markets Model. As described above, we can use the path of the interest rate in the flexible price case to obtain a measure of amplification. Treating the flexible price interest rate as a reduced form shock to the natural rate of interest in the complete markets model, we compute the reduction in demand attributable to the exogenous credit shock by itself, with no feedback from labor market conditions. We then compute the resulting path of demand arising from this reduced form shock, and compare it to the path of demand from the full model with endogenous income risk. The difference is a measure of the magnitude of amplification from this channel.

Figure 1.9 presents this comparison. The left panel compares the path of the natural rate of interest in each case. Overall, the natural rate falls farther and remains lower in the incomplete markets model. This reflects amplification arising from endogenous employment risk: since the central bank does not stabilize output, the job-finding rate falls and income risk increases. This lowers demand, which is reflected in a lower natural rate of interest.

The right panel of Figure 1.9 compares the path of output implied by the natural rate paths shown in the left panel. The initial fall in employment in the complete markets model is 3.54 pct. points, compared to 5 pct. points in the incomplete markets model. Thus about 30% of the initial fall in employment in the incomplete markets model is due to the amplification mechanism from endogenous income risk, and the rest is due to the initial exogenous shock. Equivalently, we can say that the employment multiplier from the amplification process is 1.41.

The path of employment overshoots the steady state level somewhat in the incomplete markets model. This is because the forced deleveraging following the financial shock results in high demand after the shock dissipates. This corresponds to a positive reduced form demand shock, leading employment to overshoot.

We can again use the aggregate Euler equation (1.14) to determine the degree to which the amplification process operates through increased deleveraging by constrained households versus increased precautionary saving by constrained households. We first decompose the rate of interest in the flexible price case into deleveraging and precautionary terms, and then subtract these from the terms computed from the whole model.

Figure 1.10 depicts the result of this exercise. It reveals that the majority

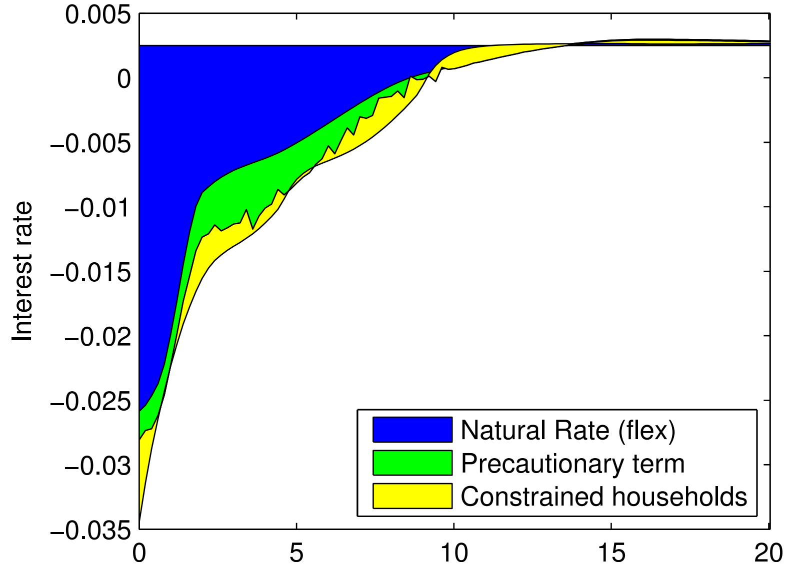

Figure 1.10: Decomposition of amplification term.

(about 70%) of the additional fall in the natural rate due to amplification is due to an increase in the precautionary term, rather than the deleveraging term. This is intuitive, since the credit shock most directly affects households with low wealth who become constrained, whereas a decrease in the job-finding rate increases the income risk facing households throughout the distribution.

1.5.4 Dynamics of the Asset Distribution

Given the emphasis placed on household debt in the wake the 2007 – 2009 recession, it is interesting to consider the dynamics of the asset distribution in the wake of the financial shock. Since assets are in zero net supply in the model, there is by construction no change in aggregate net worth. Instead, we are interested in changes in the *distribution* of assets across households.

The bottom middle panel of Figure 1.6 above shows the path of the *total debt* following the baseline temporary credit shock. Here total debt means the sum of all assets held by households with negative net worth. This statistic is therefore a measure of the *dispersion* of assets held by households.

Total debt falls during the slump because the financial shock causes the borrowing constraint to tighten significantly, forcing low-wealth households to deleverage by reducing their debt holdings. Deleveraging occurs both because households that are highly indebted and close to the constraint are mechanically forced to reduce their debt, and because unconstrained households with low wealth become afraid of hitting the constraint. Because there is a mass of households that become constrained immediately following the shock, there is an initial burst of deleveraging, but after a few quarters the pace of deleveraging slows.

As seen in section 1.2, the saving decision rule of households is convex in assets. Thus an increase in the dispersion of asset holdings will tend to raise desired savings and lower demand, while a decrease in asset dispersion will have the opposite effect. Intuitively, greater asset dispersion implies more poor households and more wealthy households. Wealthy households do not spend much more than medium-wealth households, whereas poor households are close to being constrained and so reduce their consumption greatly relative to medium-wealth households. Thus the lower asset dispersion following the credit shock raises demand and strengthens the recovery.

This occurs because the credit shock forces households to improve their balance sheets, so that after the initial fall in employment their improved asset position allows them to finance a higher level of consumption during the recovery. This suggests that a temporary credit shock may produce a faster recovery than other sorts of demand shocks, because it tends to purge bad balance sheets, enabling a stronger recovery. This is similar to the old view that recessions serve a necessary restorative role by liquidating bad investments, now applied to households rather than firms.[34]

We can approximate the magnitude of the effect of asset dispersion on demand through a partial equilibrium exercise. Let (m_e^*, m_u^*) be the asset distribution at time $t = 0$ after the initial fall in employment. Then we compute the market clearing level of employment $n(t)$ at every point in time using the equilibrium decision rules $(\dot{a}_e(a,t), \dot{a}_u(a,t))$ from our baseline experiment, but leaving the asset distribution fixed. That is, $n(t)$ is implicitly defined by:

$$\int_a \left[m_e^*(a)\dot{a}_e(a,t) + m_u^*(a)\dot{a}_u(a,t) - \left(\frac{n(t)-n(0)}{n(0)}\right) m_e^*(a)\left(\dot{a}_e(a,t) - \dot{a}_u(a,t)\right) \right] da = 0$$

The result is shown in Figure 1.11. The solid line is the baseline path of recovery, identical to that shown in Figure 1.6. The dashed line uses the same decision rules at every point in time, but holds the asset distribution unchanged at the steady state (initial) distribution. This results in a significantly slower recovery in demand than in the baseline case.

[34]This view is most often associated with Austrian writers such as Hayek and Schumpeter. However, the liquidationist view was common before the Keynesian Revolution. See Rognlie, Shleifer, and Simsek (2014), and references therein.

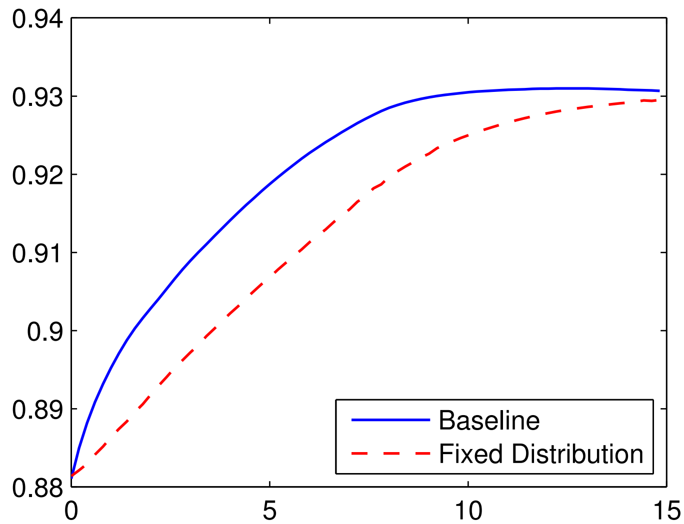

Figure 1.11: Comparison with fixed distribution.

1.6 Recovery from High Debt State

The temporary credit shock examined in section 1.5 does not seem a good description of what happened to the U.S. economy during the 2007 – 2009 recession. As Figure 1.6 reveals, a temporary credit shock produces a temporary decline in household debt, followed by a rapid recovery in hiring when the shock dissipates. In contrast, the 2007 – 2009 recession saw a long decline in household debt with a period of slow hiring.

Many theories of the magnitude of the 2007 – 2009 recession and the duration of the recovery have focused on the role of high levels of household debt in the run up to the crisis. For instance, Mian and Sufi (2010) find that counties with higher household leverage in 2006 experienced larger declines in employment. Similarly, Dynan (2012) argues that household leverage remained high well into the recovery and may be responsible for depressed consumption during the recovery. In light of such arguments, it may make more sense to model the events of 2007 – 2009 as a permanent shock to credit standards, that necessitated a transition to a new credit regime. The present model can provide some insight into this process, since it features a full distribution of household asset holdings, and captures the role of endogenous employment risk during this transition.[35]

In this section, we will consider dynamic equilibrium paths where the economy begins at time $t = 0$ with an initial asset distribution with higher total household debt (i.e. greater wealth dispersion). Since the initial asset distribution is no longer the steady state distribution, we can no longer interpret these experiments as an unanticipated shock that occurs in the steady state. Instead, we can interpret this case as depicting the result of an unmodeled shock that affects the ability of households to carry debt. For instance, perhaps an unmodeled credit boom (such as a housing bubble) enabled households to accumulate excessive debt. Then there was an unanticipated correction to credit conditions (e.g. a collapse of the housing bubble), and we are considering the transition to a new steady state.

1.6.1 Fixed Price Equilibrium

Suppose that the economy enters $t = 0$ with greater dispersion of assets than in steady state. The wider asset distribution acts as a demand shock: it produces an initial fall in employment relative to the steady state, followed by a recovery. We can interpret these dynamics as the result of unwinding a great accumulation of

[35]An important caveat is that this model does not include durable goods or housing, which represent a large fraction of household debt.

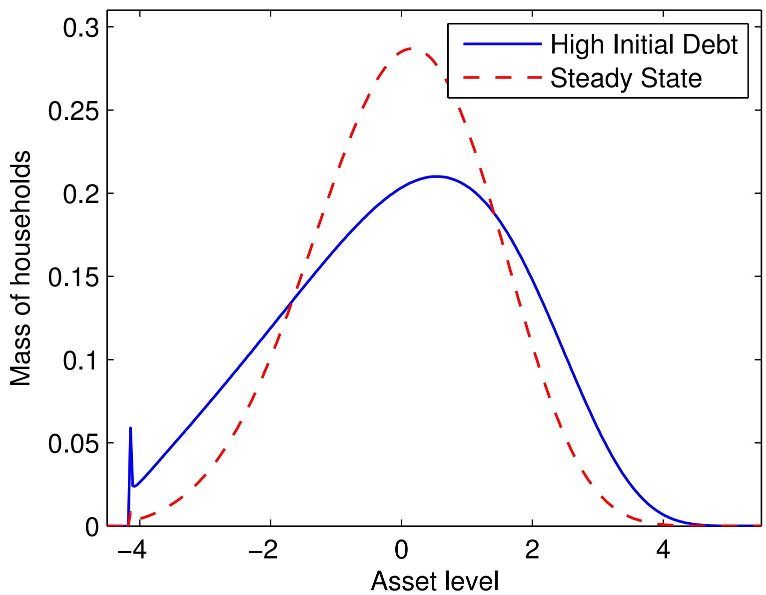

Figure 1.12: Asset distribution with 32.4% higher debt.

debt by some households in some unmodeled past period. This is similar to the circumstances of the U.S. economy in 2009 after the financial crisis, when many more households were highly indebted than during normal times, and a period of deleveraging was necessary.

Suppose the initial asset distribution is a mean-preserving spread of the steady state distribution such that the total debt held by households, i.e. minus the sum of assets held by households with $a < 0$, is 32.4% greater than in steady state. This corresponds to the increase in the household debt to GDP ratio from 2002Q1 to 2008Q1.[36] This asset distribution is depicted in figure 1.12.

We now consider the path of the economy from $t = 0$ onward. This path is de-

[36]Here household debt is taken to be the level of total liabilities held by households and non-profits from the Federal Reserve Financial Accounts of the United States. For details of how the mean-preserving spread of assets is calculated, see the appendix.

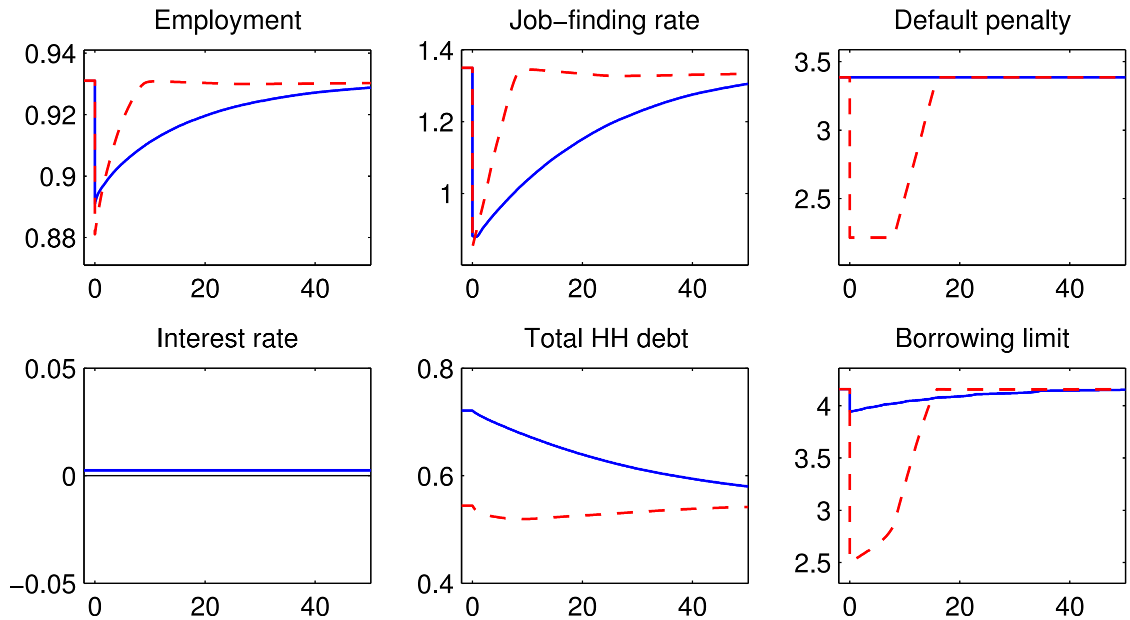

Figure 1.13: Initial high debt (solid) vs. temporary credit shock (dashed).

picted in Figure 1.13, with the baseline credit shock considered in section 1.5 shown for comparison. The result is a 4% fall in initial employment, followed by a very slow recovery. The recovery is much slower than the one following the credit shock. This is because the wide asset dispersion requires a substantial period of deleveraging, during which high levels of household debt depress demand. The sustained period of low employment induces additional debt accumulation by unemployed households, which decreases the rate of deleveraging, further slowing recovery. Moreover, the period of slow hiring induces precautionary saving by employed households, which reduces demand and further increases the variance of asset holdings. The result is that the dispersion of assets falls only very slowly, leading to a sustained period of low employment.

One benefit of the high-debt recovery case relative to the credit shock scenario

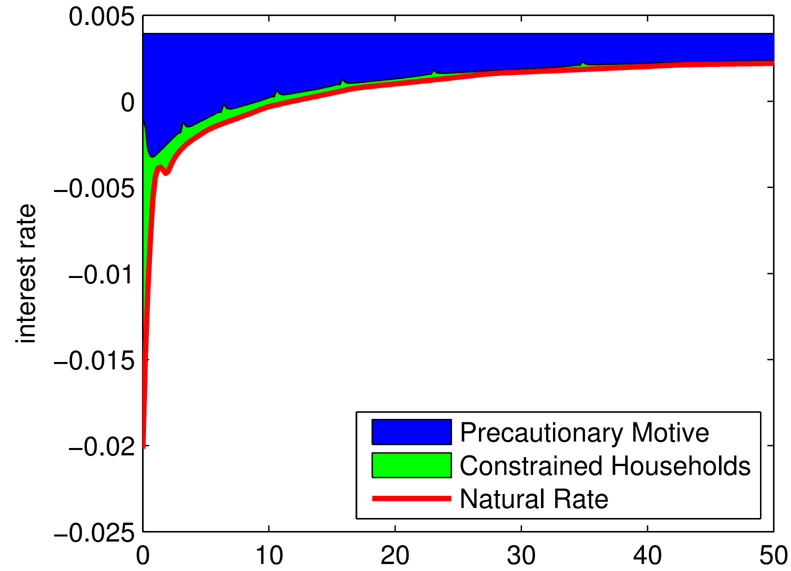

Figure 1.14: Decomposition of natural rate during transition from high-debt state.

is that we can get a sense of the size of the endogenous tightening of borrowing constraints without the complication of an exogenous tightening. Figure 1.13 shows that the borrowing limit tightens by about 9% relative to the steady state. This is a significant tightening of the borrowing constraint, though much less than what is seen in the credit shock case.

Decomposition of Natural Rate. As in the temporary credit shock case, we can represent the period of weak demand resulting from the initially wide asset distribution as a reduction in the natural rate of interest. We can then decompose the fall in the natural rate into components deriving from deleveraging of constrained households and precautionary behavior by unconstrained households.

Figure 1.14 depicts this decomposition. It reveals that the great majority of the reduction in the natural rate during the transition from the high debt state is due to the increase in the precautionary motive. The only exception is at the very beginning of the transition, when constrained households are forced to deleverage following the endogenous tightening of the constraint. Overall, 77% of the total reduction in initial consumption demand is attributable to the precautionary motive.

1.6.2 Flexible Price Benchmark

The flexible price equilibrium again provides a useful point of comparison. Figure 1.15 depicts the equilibirum paths in both the flexible and fixed price models. The left panel depicts the path of the interest rate. In the flexible price model, the interest rate falls initially by about 60 bp to accommodate an initial burst of deleveraging by highly indebted households. The interest rate then quickly recovers to about 10 bp below its steady state value after 5 quarters. It remains slightly depressed for some time, as the economy slowly transitions to its new stationary asset distribution. This indicates that the total reduction in demand due to the necessary deleveraging process is not so great, since it requires only a small fall in the interest rate to accommodate.

As the center panel of Figure 1.15 makes clear, the reduction in total household debt occurs somewhat faster under flexible prices than in the fixed price case. In particular, after 20 quarters excess household debt has fallen by 60% in the flexible price case compared to 47% in the fixed price case. Likewise, after 40 quarters, excess household debt has fallen by 86% in the flexible price case compared to 74% in the fixed price case. This happens because labor market conditions are stabilized

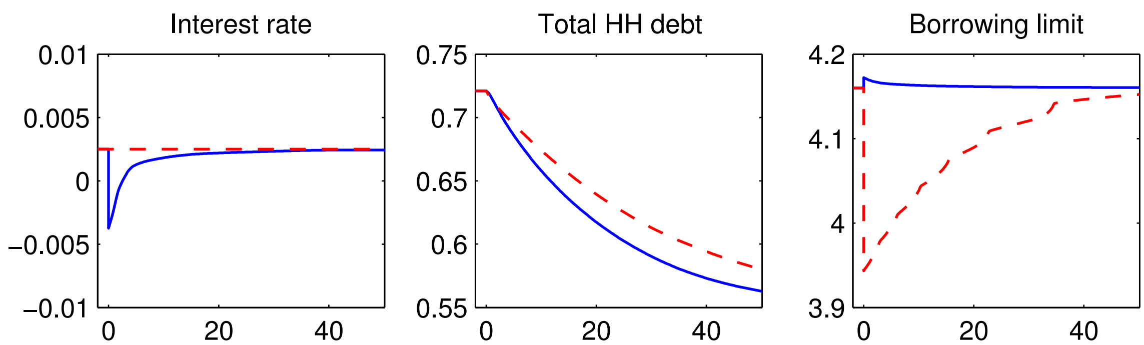

Figure 1.15: Recovery from high debt state with flexible prices (solid) vs. fixed prices (dashed).

under flexible prices, and so households are in a better position to pay down their debts. As shown in Figure 1.13, under fixed prices employment falls and remains low throughout the transition. Thus a higher fraction of highly indebted households are unemployed. These households do not reduce their debts while they are unemployed, and so high unemployment slows down the pace of deleveraging.

Finally, as the right panel indicates, borrowing constraints do not tighten in the flexible price case. Since there is no exogenous credit shock in this case, the tightening of the borrowing constraint in the fixed price case is entirely due to worse labor market conditions. Since output and hiring are stabilized in the flexible price case, there is no endogenous feedback to credit conditions, and so borrowing constraints fail to tighten. In fact, they loosen slightly, because lower interest rates reduce the cost of repayment.

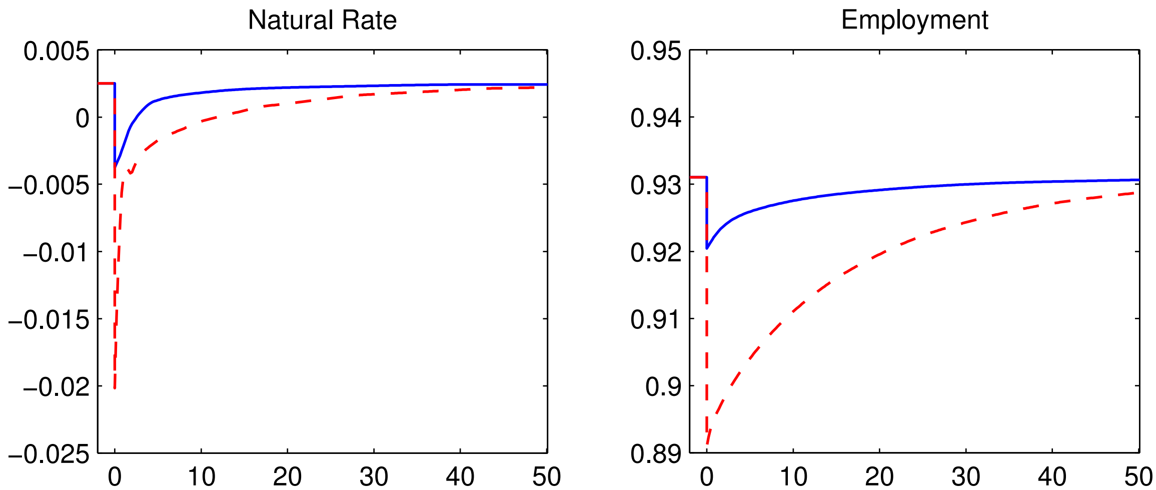

Figure 1.16: Reduced form shock (solid) vs. endogenous income risk (dashed).

1.6.3 Measure of Amplification

We can again take the flexible price interest rate as a reduced form representation of the initial demand shock, absent the feedback from endogenous income risk. The difference between the flexible interest rate path and the natural rate in the full model with fixed prices is a measure of the amplification from this feedback. The left panel of Figure 1.16 presents both interest rate paths, while the right panel presents the implied paths of employment from the standard aggregate Euler equation (1.17).

Overall, the natural rate of interest as defined by equation (1.18) in the model with endogenous income risk falls substantially more than the flexible price path of interest. This is particularly true at $t = 0$, when the natural rate plungs 225 bp, but remains true throughout the recovery. This larger fall in the natural rate implies a similarly larger fall in employment. Initial employment falls nearly four times as much in the model with endogenous employment risk than in the complete markets model with the reduced form demand shock. This implies that 70% of the intial fall in aggregate consumption, and 69% of the initial fall in employment, is due to amplification from endogenous income risk. This implies a "multiplier" from the amplification mechanism of 3.38.

In the case of the temporary credit shock in section 1.5, only 30% of the fall of consumption was due to amplification, and the multiplier was 1.41. Why is the amplification process so much stronger in the deleveraging case than in the case of a credit shock? The answer is that while the direct reduction in initial demand implied by deleveraging is less than in the case of the credit shock, the persistence of the deleveraging shock is much greater. We can see this by comparing the path of employment implied by setting the natural rate of interest equal to the flexible price interest rate in each case. These are the solid lines in Figures 1.16 and 1.9. While the *initial* fall in employment in the temporary credit shock case is 2.2 pct. points greater than in the deleveraging case, the cumulative loss of employment over the entire path of recovery is nearly 5 pct. points greater in the deleveraging case. Thus the deleveraging case implies a greater total total loss of employment, and therefore prompts a greater precautionary saving response from forward-looking households.

One implication of this result is that it is better to have a sharp and short demand shock, rather than a long and slow shock. Whereas the total initial *exogenous* reduction in demand may be the same in each case, they imply very different paths of future hiring, and therefore different income risk perceived by households. A short and sharp demand shock implies a large initial fall in employment followed by a rapid recovery. Households anticipate this rapid recovery, and so have limited

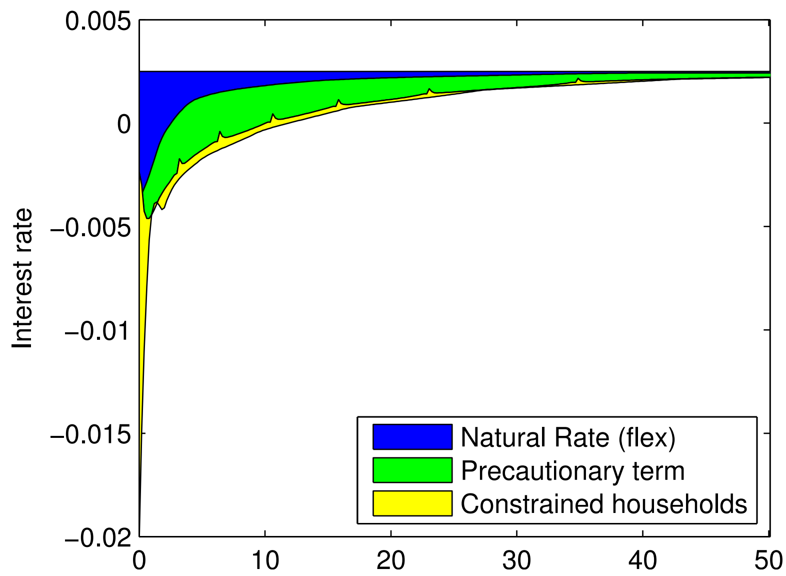

Figure 1.17: Decomposition of amplification term.

desire to engage in precautionary saving, leading to limited amplification through endogenous income risk. By contrast, a long and shallow demand shock implies a lengthy period of slow hiring, with higher income risk facing households. This increases precautionary saving by households, and amplifies the fall in demand.

Decomposition of Amplification. We can again use equation (1.14) to decompose the amplification term into a component corresponding to forced deleveraging by constrained households, and a component arising from the precautionary saving behavior of unconstrained households. The result of this exercise is depicted in Figure 1.17. Unsurprisingly given the discussion above, most (76%) of the amplification is due to precautionary saving by unconstrained households.

1.6.4 Credit Shock in High Debt Initial State

The analysis above modeled the events of 2007 – 2009 as a high debt initial state that required a period of deleveraging. However, it is reasonable to think that the financial crisis that occurred in the Fall of 2008 represented a qualitatively different phenomenon than simply high levels of household debt. Arguably this was an exogenous credit shock, with subsequent recovery slowed by high levels of household debt and the lingering effects of the credit shock.

We can model this as an initial high level of household debt *together with* a temporary credit shock. In particular, suppose that initial household assets are as shown in Figure 1.12, and that the economy experiences a temporary credit shock as analyzed in section 1.5. Then the equilibrium path reflects *both* the effects of a high level of initial debt that requires a period of deleveraging, and a temporary credit crunch facing households.

The resulting equilibrium path is shown in Figure 1.18, with the recovery from the high debt state with no exogenous credit shock shown for comparison. The initial fall in employment is 10.7 pct. points, compared to a fall of 4 in the high debt case without a credit shock, and a fall of 5 in the case of a credit shock that hits an economy already at the stationary distribution of assets. Thus the total initial fall in employment is greater than the sum of its constituent parts, implying an interaction between high initial levels of debt and the credit shock. One way to interpret these results is that greater initial dispersion in asset holdings increases the

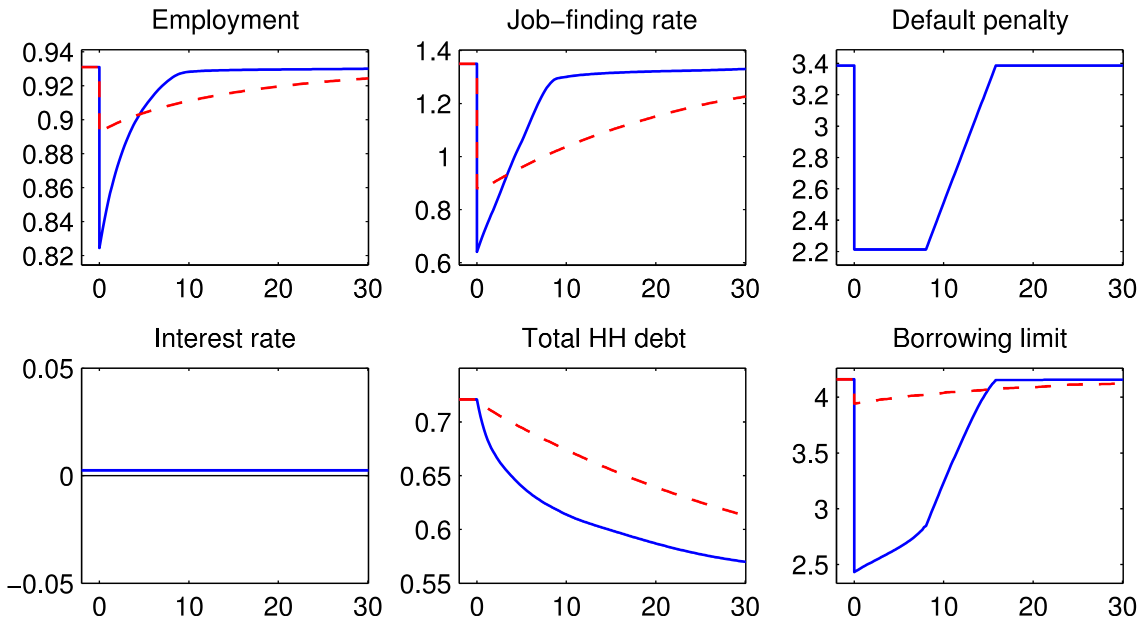

Figure 1.18: High debt distribution with credit shock (solid) vs. no shock (dashed).

sensitivity of the economy to financial shocks.[37] This suggests that higher moments of the asset distribution are important variables for central banks to monitor, and that policies that aim to reduce excessive levels of debt held by households can be welfare-improving. Such policies are an example of macroprudential policies.[38]

This interaction occurs because the high debt initial asset distribution has many more households close to the borrowing constraint than the stationary distribution, as can be seen in Figure 1.11. Thus when the borrowing constraint tightens substantially following the shock, the initial fall in demand is much greater than

[37]Heathcote and Perri (2014) obtain a similar result by showing how a low level of household wealth can produce greater volatility by allowing low-employment equilibria to become feasible.

[38]For an analysis of macroprudential policies in a model with a demand externality from household debt, see Korinek and Simsek (2014). For a general analysis of macroprudential policies in the presence of nominal rigidities, see Farhi and Werning (2013).

with the stationary distribution. This can be illustrated by a decomposition of the natural rate using equation (1.14), which reveals that 56% of the cumulative fall of the natural rate in the credit shock case is due to deleveraging by constrained households, compared to 23% in the high debt path with no credit shock.

Although the initial fall in employment is much greater with the credit shock than without, the *persistence* of the fall in employment is much less. The reason is that deleveraging happens much faster in the credit shock case, as can be seen directly from the bottom middle panel of Figure 1.11. After 20 quarters, excess household debt has been reduced by 76% in the financial shock case, compared with 47% in the no shock case. Likewise, after 40 quarters 90% of deleveraging has been accomplished in the financial shock case, compared with 74% in the no shock case.

The faster pace of deleveraging implies a much faster recovery in employment, as can be seen from the first panel in Figure 1.11. Although initial employment falls more than twice as far following the financial shock than in its absence, the employment path following the financial shock surpasses the no shock employment path after just 5 quarters, and is nearly fully recovered after 10. This leads to the surprising result that the cumulative loss of employment and therefore output is nearly twice as large in the *absence* of the credit shock! Cumulative employment loss in the credit shock case comes to 37.8 pct. points of employment, compared with 63.5 points in the no shock case. Thus the total decline in output is less in an economy with an additional adverse demand shock, if that shock speeds up the pace of deleveraging.

This is particularly surprising given that the initial demand shock is so much greater in the credit shock case. Taking the magnitude of the initial demand shock to be the initial fall in employment implied by the flexible price path of the interest rate, the credit shock generates a 6 times greater fall in initial employment than that due to the high debt distribution alone. Yet despite a six-fold greater initial shock, the faster recovery in the financial shock case produces significantly less cumulative employment loss. Part of the reason for this is that the faster employment recovery produces less amplification compared with the no financial shock case: 42% of the initial fall in aggregate consumption is due to amplification in the financial shock case, compared with 70% in the no shock case. This implies an initial demand multiplier of 1.73 instead of 3.38.

1.7 Expectations and Recovery

Thus far we have assumed that the economy follows the path with the highest initial level of employment subject to the central bank's no-commitment policy rule. However, there are many other equilibrium paths that are consistent with the equilibrium conditions of the model. Which one the economy follows depends on the expectations of households, which may be influenced by the policy rule adopted by the central bank.

1.7.1 Forward Guidance

The analysis in sections 1.5 and 1.6 assumed that the central bank followed the no commitment policy rule after the economy exits the liquidity trap. However, if the central bank can credibly commit to allowing $p > p^*$ after the recovery, then an equilibrium path with higher employment is possible. If the central bank has a loss function that penalizes deviations from the target job-finding rate in both directions, then it would like to commit to producing such a hiring boom. In this case, the central bank commits to allowing excessive hiring after the trap without raising interest rates.

This is a form of forward guidance, analogous to committing to keeping interest rates low despite high inflation in a standard New Keynesian model. Here the cost of this policy is not excessive inflation, but an inefficiently high rate of hiring, which is costly due to the hiring cost ϑ.The mechanism for raising demand is also somewhat different: in the standard New Keynesian model, the expectation of a future consumption boom raises current demand due to consumption smoothing, whereas in this paper an expected hiring boom also raises demand by reducing unemployment risk, which reduces precautionary saving.

Figure 1.19 shows the dynamics of the economy in response to the credit shock analyzed in section 1.5. The solid line shows the path of recovery when the central bank lacks commitment (identical to Figure 1.6), whereas the dashed line shows the recovery path under forward guidance. The forward guidance equilibrium allows a hiring boom with the job-finding rate topping out at 1.55, 15% higher than the central bank's preferred level. This policy rule was chosen to minimize the loss function

$$L = \int_0^\infty e^{-\rho t}[zn(t) - zn^*]^2 dt$$

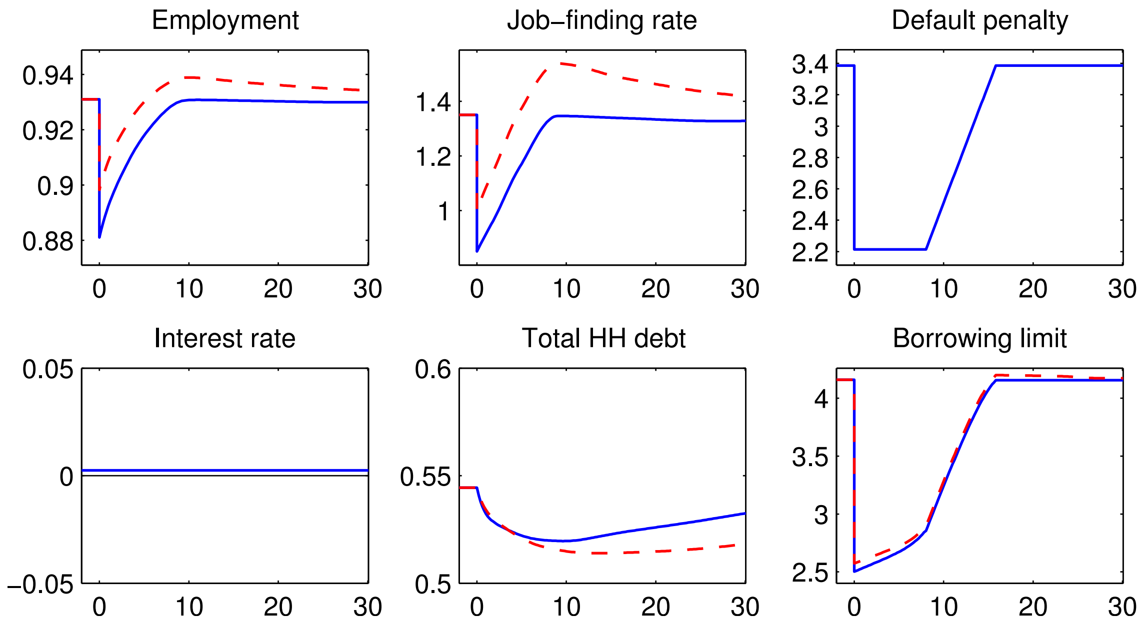

Figure 1.19: Hiring boom (dashed) vs. no commitment (solid).

where n^* is the steady state level of employment. That is, the central bank seeks to minimize the discounted quadratic output gap.[39]

The equilibrium path under forward guidance sees an initial fall of employment of just over 3 percentage points, about 40% smaller than in the no-commitment case. Employment recovers to its steady state level after just 5 quarters (well before the shock begins to dissipate) and then overshoots the steady state, reaching a maximum level of 0.8 percentage points above steady state after 10 quarters. Overall, the average level of employment over the entire period is about equal to steady state

[39] A discounted quadratic loss function in the output gap and inflation is common in the New Keynesian literature because this is the linear approximation to the optimal policy rule around the efficient zero-inflation steady state. This rule is not optimal in the present model because of the presence of incomplete markets, but offers a simple benchmark to illustrate the power of forward guidance in this setting.

employment (0.14 percentage points higher).

The initial fall in the job-finding rate is likewise about half of what it is in the baseline case (a fall of 0.27 instead of 0.56), and the job-finding rate likewise overshoots the steady state rate by about 0.20. This higher job-finding rate both mechanically produces the higher growth path of employment, and creates the demand that enables it by reducing precautionary saving and by relaxing the borrowing constraint relative to the no-commitment path.

1.7.2 Pessimistic Expectations Following Credit Shock

The discussion above assumes that the economy always follows the highest path of employment consistent with the central bank's policy rule. However, there is no reason that this must be true. While the central bank has the means to rule out any path with excessive hiring by threatening to raise the interest rate, it cannot rule out paths with lower hiring because it is constrained by the lower bound on the interest rate.

Figure 1.20 shows a pessimistic path of recovery in response to the credit shock analyzed in section 1.5. This produces an initial fall in employment that is 2 percentage points greater than the highest no-commitment path. The resulting path of employment remains about 1 percentage point below the steady state for a significant period after the credit shock has dissipated.

Simultaneous to the slow recovery in employment is a period of slow hiring, which lowers demand and tightens the borrowing constraint. The job-finding rate

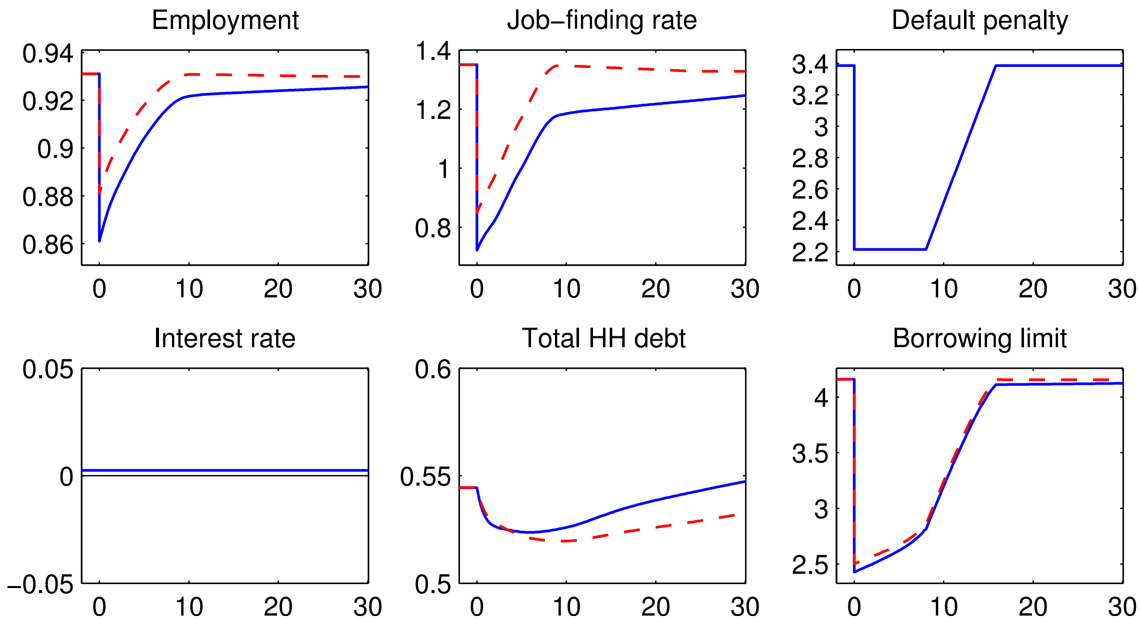

Figure 1.20: Pessimistic recovery from credit shock (solid) vs. baseline (dashed).

remains well below the steady state level for the entire period shown, whereas in the baseline case it returns to the steady state level after 12 quarters and remains close to this level thereafter. The larger initial fall in the job-finding rate is by construction — this is exactly what it means in this model for households to be pessimistic about the rate of recovery, and this is the proximate cause of the larger initial fall in employment. However, a further effect arises from the dynamics of the asset distribution. The lower rate of job-finding raises precautionary savings, while the lower employment rate increases the number of borrowers in the economy. This increases the dispersion of asset holdings relative to the baseline case, raising the total debt held by net borrowers as seen in the bottom middle panel of Figure 1.20. This greater dispersion of asset holdings weakens demand, further lowering the path of employment.

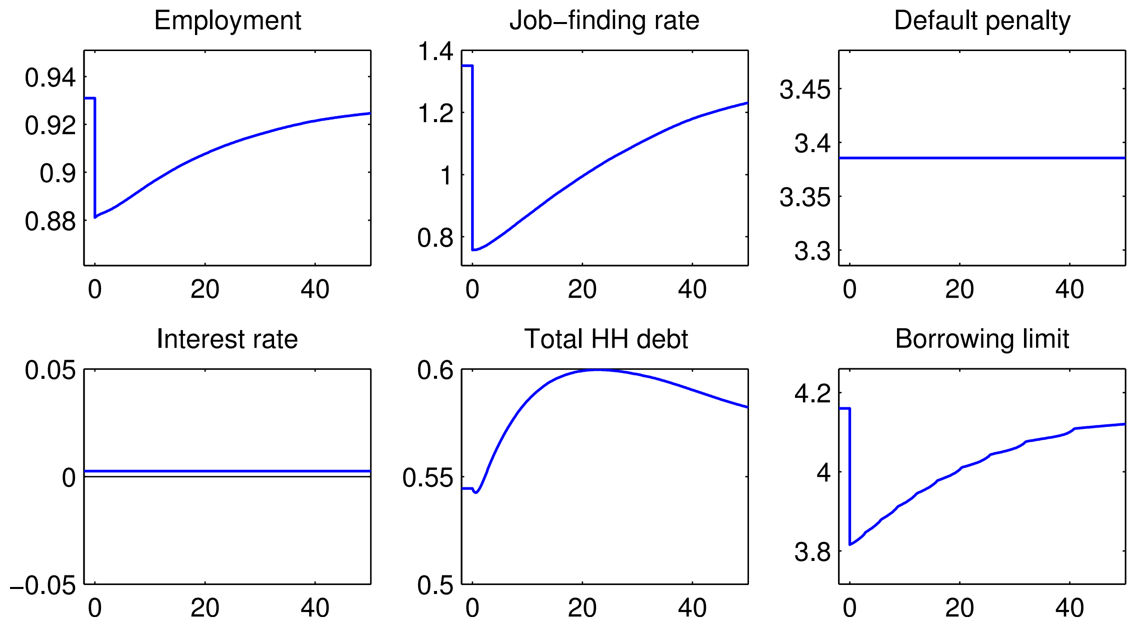

Figure 1.21: Recovery following expectational shock.

1.7.3 Pure Expectational Shock

Given the significant effects of pessimistic expectations following a shock to credit market fundamentals found above, a natural question is what effect a pure expectational shock can produce. Suppose that the economy is at the steady state for $t < 0$, and at time 0 households come to expect a period of slow hiring. This will lead them to reduce their desired spending, decreasing demand and causing $n(0)$ to fall. The central bank cannot rule out this equilibrium because it cannot lower interest rates, and so it is powerless to prevent this equilibrium.

Under our maintained assumption that the economy converges to the steady state in the long-run, there is a unique path corresponding to each level of $n(0) < n^*$, with lower $n(0)$ corresponding to a larger shock to expectations. Figure 1.21 shows the path corresponding to a 5 percentage point drop in initial employment relative to the steady state. The first thing to note is how persistent the fall in employment is. The time horizon shown in Figure 1.21 extends to 50 quarters, 12.5 years after the initial shock. By the end of this time, employment has still not fully recovered to its steady state level.

This persistence is driven by the increase in asset dispersion during the recovery. As discussed above, a low job-finding rate increases asset dispersion because it stimulates greater saving from employed households while also increasing the total number of households that are borrowers (i.e. unemployed). This effect is sufficiently great that it offsets the narrowing of the asset distribution due to the tightening of the borrowing constraint, and significantly raises total debt held by borrowers over the course of the recovery, as shown in the bottom middle panel of Figure 1.21. This greater asset dispersion generates persistence of low employment by lowering demand. Intuitively, high unemployment worsens household balance sheets, which lowers demand and perpetuates high unemployment. Since the asset distribution is a relatively slow-moving variable (excepting initial deleveraging), this generates substantial persistence.

1.8 Conclusion

Unemployment risk has significant implications for the dynamics of recovery from a liquidity trap. Unemployment risk both amplifies and increases the persistence of demand shocks. Amplification occurs due to the feedback from slow hiring to weak demand, through precautionary savings and endogenous borrowing constraints. Persistence arises from the dynamics of the asset distribution, since a period of high unemployment raises the burden of debt, reducing demand during the recovery.

One important conclusion of this paper is that the distribution of debt is a critical variable in determining the dynamics of the economy in response to demand shocks. From a positive perspective, I find that persistence of demand shocks is driven in part by the evolution of the asset distribution, and that high initial asset dispersion increases the sensitivity of the economy to shocks. From a normative perspective, these results suggest that central banks should monitor the distribution of debt, and take steps to reduce high household leverage. These results are particularly relevant given recent trends toward increased credit access and greater wealth inequality in the developed world.

This paper also highlights the significant role of expectations in determining the path of recovery. These results suggest that if the central bank can credibly manage expectations about future policy, i.e. engage in forward guidance, it can significantly mitigate negative shocks. More ominously, these results also suggest the possibility of self-fulfilling negative expectations about the pace of recovery. Given the zero lower bound on interest rates, the central bank may find itself powerless to prevent such outcomes. In light of slowing growth and demographic transitions around the world that depress the natural rate of interest, these possibilities are highly relevant to policymakers today.

I leave many questions to future research. This paper does not explicitly model durable goods and housing, although these constitute a large fraction of household credit use. Particularly given the centrality of mortgage debt and house prices in the 2007 – 2009 recession, explicitly considering such forms of debt could be a fruitful avenue of further research. In addition, this paper assumes that prices are fixed, which eliminates the role of inflation in equilibrium dynamics. Allowing for partially flexible prices could generate interesting interactions between inflation, unemployment risk, and the asset distribution.

Overall, my results suggest that a precautionary saving channel in response to high unemployment may have played a role in the 2007 – 2009 recession and the subsequent recovery. Such a channel could help explain the size of the initial fall in demand, as increased precautionary saving led households to decrease spending in the Fall of 2008. Further, this channel could also explain the slow recovery, as high unemployment increased the burden of debt and slowed the process of deleveraging.

Chapter 2: Financial Intermediaries and the Supply of Liquidity

2.1 Introduction

An important function of the financial system is liquidity provision, i.e. the creation of financial instruments for firms and households to store wealth while maintaining access to these funds to meet unpredictable financing needs. Though liquidity shortages played a prominent role in the 2008 financial crisis, relatively little macroeconomic analysis has focused on the role of the financial sector in providing liquidity to the productive sector. Such analysis is particularly relevant to policy, since government liabilities comprise a significant fraction of the economy's aggregate supply of liquidity. This paper analyzes the role of financial intermediaries in liquidity provision. I investigate two questions: What effect do fluctuations in the value of intermediaries' balance sheets have on the supply of liquidity to the real sector? What is the role of public liquidity provision in response to such fluctuations?

My analysis builds on the model of Holmstrom and Tirole (1998), in which productive projects are subject to ongoing uncertain financing needs. Agency costs may prevent firms' raising these funds by new borrowing. Thus firms purchase and hold liquid assets to cover these expenses. When liquid assets are in short supply, financial intermediaries can offer financial instruments such as lines of credit that substitute for liquid assets.

I depart from Holmstrom and Tirole (1998) by assuming that financial intermediaries are also subject to agency costs. I further assume that these agency costs worsen when intermediaries' balance sheets deteriorate, so that aggregate frictions in bank liquidity supply are decreasing in the level of bank equity. Under this assumption, variations in financial sector equity affect the aggregate supply of liquidity. I find that the aggregate scarcity of liquidity can be summarized by the liquidity premium: the difference between the price of liquid assets and their fundamental value based on their expected return if held to maturity. When bank equity falls below a certain threshold, there is a positive equilibrium liquidity premium, and investment and production fall below their constrained optimum quantities. In such liquidity-constrained states, it is optimal for the government to increase issuance of debt for its liquidity properties, even if funds can only be raised by costly taxation.

As this paper is concerned with the supply of liquidity, we should begin by precisely defining the term. By liquidity (or equivalently, liquid assets), I mean assets held by firms and households that serve as a store of value, and that can be sold (liquidated) for their full value at any time before maturity. Examples of liquidity include government bonds, some corporate bonds, commercial paper, money market funds, and demand deposits issued by banks. The key feature of liquid assets is information insensitivity — since all agents know these assets' worth, they can be sold for their full price without discounts due to adverse selection.[1]

[1]The importance of adverse selection in generating market illiquidity was first pointed out by Akerlof (1970). The importance of information insensitivity was emphasized by Gorton and

Following Holmstrom and Tirole (1998), I model the demand for liquidity from the business sector by supposing that investment projects are subject to stochastic needs for funds. Firms can raise these funds by new borrowing, but due to agency costs are unable to pledge the full value of their projects. If required funds exceed the pledgeable amount, the firm will be borrowing constrained. Anticipating this possibility, firms undertaking a new project can borrow more than required for the initial investment, and hold the excess as liquid assets. Should a need for additional funds arise, firms use these assets as collateral to enable additional borrowing. If liquid assets yield a return equal to the firm's cost of borrowing, firms can hold the optimal amount of liquid assets given their initial borrowing capacity. However, when liquidity is scarce liquid assets command a premium over illiquid assets. Then firms must use some of their limited pledgeable funds to pay this premium, and reduce investment below the constrained efficient level.

Firms obtain these funds in two ways. First, firms may borrow funds when they are not borrowing constrained and purchase assets that can be liquidated in the event of a liquidity shock. If these assets are not liquidated, their proceeds can be returned to the firm's investors. This arrangement is equivalent to a sort of insurance contract between firms and original investors, in which investors agree to provide additional financing to firms that require it, even when the firm is unable to raise new financing on private markets. The liquid assets held by firms serve as collateral on this promise, since otherwise investors might not want to provide these

Pennacchi (1990), and more recently by DeMarzo and Duffie (1999). For a summary of the various sources of illiquidity, see Tirole (2011) and Holmstrom and Tirole (2011).

funds. Such assets include corporate and government bonds and shares of money market funds. Since it doesn't matter whether the firm liquidates the asset or uses it as collateral to obtain a loan, this category also includes repos and other securitized borrowing.

The second way firms obtain liquid funds is by locking in financing at convenient terms by purchasing a credit line from a financial intermediary. Typically the firm pays a premium to an intermediary to provide financing at specified terms up to some limit, even if such financing becomes unprofitable in expectation. This category primarily consists of credit lines issued by banks. Note that this is analogous to the arrangement above, except that banks are presumed to be able to credibly provide the promised financing without requiring collateral. Thus in equilibrium the premium paid on bank credit lines will equal the liquidity premium on liquid assets.

I refer to assets that can be easily sold to meet liquidity shocks as liquid assets; I refer to the market for liquid assets as the market for liquidity; and I refer to the availability of liquidity on good terms as the supply of liquidity.

I depart from Holmstrom and Tirole (1998) by supposing that intermediaries are also subject to agency costs. When lending is subject to agency costs, the supply of liquidity provided by intermediaries depends on their net worth.[2] When banks

[2] Several papers use agency costs in models of banks. One notable example is Mattesini, Monnet, and Wright (2009), which argues that agents with large stakes in the continuation of the economy are better suited to serve as banks. This is analogous to my model, in which bank agency costs depend on their asset holdings.

have low net worth, they have less collateral to be seized in the event of bankruptcy and therefore have more incentive to engage in fraud, for instance by failing to exert effort in screening loan applicants, or by taking on excessively risky loans since they do not bear downside risk. When these agency constraints are binding, the economy will not be able to achieve perfect risk pooling of idiosyncratic liquidity risk even in the presence of intermediaries. Moreover, if the average agency cost is increasing in the quantity of funds intermediated, then the liquidity premium will be increasing in the quantity of bank financing and decreasing in bank net worth.

Why might liquidity be scarce? Why can't firms simply buy assets that pay at the prevailing interest rate, sell them to raise funds in the event of a liquidity shock, and if no shock is realized hold them to maturity? If there were a sufficient quantity of such assets then this would indeed be possible. However, such assets may be in limited supply. This is somewhat removed from everyday experience: we are used to being able to save (buy assets) easily, while borrowing (issuing liabilities) may be difficult. But while we can always buy assets to save, if we want to buy a *liquid* asset we must pay a premium in the form of a lower expected return relative to illiquid assets. This indicates that liquid assets are not in unlimited supply.

Why is this? It derives from asymmetric information. If pricing an asset requires some specialized knowledge, then sellers may not be able to obtain the fundamental price of the asset in a sale — the very act of selling signals to potential buyers that the asset is not as valuable as the owner knows it to be. This phenomenon affects all assets to some degree, but is especially acute for assets that require specialized knowledge to price correctly. By contrast, safe government debt and money are much less sensitive to information from selling decisions, and thus are highly liquid.

If the price q of any asset is higher than its fundamental value based on its expected return, $\hat{q}$, then that asset has a positive liquidity premium, $q - \hat{q}$. [3] There is ample evidence of positive liquidity premia in reality. The simplest example is money, which commands a positive liquidity premium relative to safe government bonds. This is intuitive because money has some utility in transactions. Likewise, the interest rate on safe government debt is lower than the expected returns on other assets in the economy. This suggests a positive liquidity premium on government debt.

During times of crisis, interest rates on many forms of safe assets fall, a phenomenon often called a *flight to quality*. This may reflect falling returns to capital due to productivity shocks, but since interest rate *spreads* rise, it plausibly indicates an increase in the liquidity premium due to a decrease in the supply of liquidity, which makes liquid government bonds more desirable than longer-term investments that cannot be easily liquidated. In general, a positive liquidity premium indicates that liquid assets are *scarce*, i.e. if liquid assets yielded a return equal to their fundamental value, demand for these assets would exceed supply. Thus a decrease in the supply of liquid assets will imply a rise in the liquidity premium in order to clear

[3] Since the price of a bond with face value one and interest rate r is $q = 1/(1+r)$, a positive liquidity premium on an asset implies a *lower* interest rate on this asset than the long-run expected return of the asset when held to maturity, i.e. $r < r^*$, where r^* is the return on illiquid assets.

the market for liquidity.

Why might an economic downturn cause a fall in the supply of liquidity? I consider two mechanisms. First, a recession leads to losses on projects that underlie liquid assets in the economy. For instance, the 2008 financial crisis was preceded by a drop in housing prices, which led to an increase in mortgage delinquencies. Since mortgage-backed securities made up a significant fraction of asset-backed commercial paper in 2008,[4] this rise in mortgage delinquencies represented the destruction of a fraction of the economy's stock of liquid assets. In a world with scarce liquidity, this should raise the cost of investment and raise the prices of liquid assets, for instance by reducing interest rates on government bonds. I label this mechanism the *asset destruction channel*.

In addition to the direct destruction of assets held by agents to meet liquidity shocks, a recession will generally involve falling asset prices and losses on bank loans. These reduce bank net worth, causing banks to reduce credit supply, including liquid instruments such as lines of credit held by firms. This mechanism is stronger if the financial system is poorly capitalized. Leading up to the 2008 crisis, banks were highly leveraged, and so intermediaries suffered large losses due to the housing crash. This was a major source of worsening credit conditions and falling investment. I refer to this second mechanism as the *bank channel*.

I can distinguish between these channels in my model. The asset destruction channel corresponds to a fall in outside liquidity in my model. The bank channel

[4] As documented by Brunnermeier (2009).

is modeled as a fall in bank net worth, which decreases the supply of liquidity by banks. These two channels produce similar macroeconomic results, but imply different policy responses.

If the economy is initially in a region with a positive liquidity premium, then a fall in outside liquidity reduces investment and increases the liquidity premium. I find that if the economy is initially in an equilibrium with positive bank financing, then a decrease in bank net worth will tighten the supply of liquidity, also leading to a fall in investment and an increase in the liquidity premium. So the results in each case are analogous.

The final step is to consider optimal government liquidity provision. As described above, when there is a positive liquidity premium in equilibrium, the government can improve the allocation by issuing liquid liabilities such as government debt. If issuing bonds were costless, then the optimal policy would always be for the government to issue bonds until liquidity is no longer scarce. However, realistically there are some costs of government debt. One cost is that debt must be repaid, and the government must levy distortionary taxes to do so. Thus the liquidity benefits of issuing additional debt must be weighed against the costs from repaying this debt.

I analyze optimal government policy within the context of this model. I find that when there is a positive liquidity premium at an interior equilibrium, it will always be optimal for the government to provide some public liquidity. This result is quite interesting by itself. It implies that in an economy with a positive liquidity premium there is positive value of issuing government debt. If total government debt is sufficiently low, the issuance of additional debt will crowd in investment, and these benefits will outweigh the costs of higher debt.

I also analyze how the optimal supply of government debt *changes* when the economy's initial stock of liquid assets changes. I find that when the stock of bank capital falls, the optimal supply of public liquidity increases. This result can be interpreted in several ways. One interpretation is simply that the government should issue more debt when the banking sector is poorly capitalized. Since this is likely true during recessions, this result can be taken as a justification for procyclical budget deficits. The result can also be interpreted as a need for public liquidity provision more broadly. The government can provide liquidity by lending directly to firms or financial intermediaries in distress. Thus during the 2008 crisis, the Federal Reserve established a number of liquidity facilities to provide funds to intermediaries responsible for a large portion of the economy's supply of liquidity.

By contrast, when the economy's stock of private outside liquidity falls, optimal public liquidity provision *decreases*. This result is driven by the elasticity of private liquidity supply. An increase in government bonds crowds out private liquidity, meaning that government bonds raise total liquidity less than one-for-one. The degree of crowding out depends on the elasticity of private liquidity supply. A fall in (inelastic) *outside* liquidity implies that (elastic) *bank* liquidity is a larger share of total liquidity. Thus the elasticity of private liquidity supply is higher, and crowding out is greater. This decreases the effectiveness of public liquidity provision, leading the government to reduce liquidity supply.

Related Literature. This paper is part of the literature exploring the role of public liabilities in providing liquidity by serving as stores of value. Samuelson (1958) shows that a government bond can enable intergenerational trades that would not otherwise occur. Woodford (1990) shows that when income and investment opportunities are not synchronized, investment and liquid assets are complements, so that under some conditions the issuance of government debt will "crowd in" investment. Kiyotaki and Moore (2008) consider both public and private liquidity in a similar framework, and Farhi and Tirole (2012) introduce bubbles as a store of value to explore the interplay between public, private, and bubble liquidity. In all of these models, agents would like to transfer funds forward in time, but are unable to do so due to a market incompleteness. I build on these models by introducing financial intermediaries that supply liquidity.

Several papers have examined the role of banks in creating liquid assets. In Diamond and Dybvig (1983), banks offer liquidity insurance by pooling claims to investment projects and selling demand deposits to households. Holmstrom and Tirole (1998) model liquidity in a framework identical to my own, and show that banks can perfectly insure against idiosyncratic liquidity shocks but not aggregate liquidity shocks. Brunnermeier and Sannikov (2010) introduce intermediaries to a model similar to Kiyotaki and Moore (2008), and show that a fall in bank capital will decrease the supply of inside liquidity.[5] I depart from these models by making bank lending subject to an agency cost that depends on bank capital.

[5] Other papers that discuss liquid asset creation by banks in comparison to government liabilities include Stein (2012) and Greenwood, Hanson, and Stein (2010).

The idea that bank lending is subject to agency costs has several precedents in the banking literature. Calomiris and Kahn (1991) argue that demandable debt is a mechanism to ensure the cooperation of banks. They argue that agency problems are paramount in banking, writing "...studies of banking failures give fraud a prominent place in the list of causes. Studies of 19th- and 20th-century banking indicate that fraud and conflicts of interest characterize the vast majority of bank failures for state and nationally chartered banks." Likewise, Diamond and Rajan (2001) theorize that banks adopt fragile asset structures as a commitment device. Mattesini, Monnet, and Wright (2009) examine which agents will serve as banks in a mechanism design framework, and find that agents who have a larger stake in the system serve as banks, since exclusion from future trades is more costly for such agents. Such arguments provide a microfoundation for my assumption that agency costs are decreasing in bank capital, since bank capital serves as collateral and well-capitalized banks have more to lose if their reputations are damaged.

This paper is also related to the recent literature exploring the connection between the 2008 financial crisis and the market for liquidity. Pozsar et al. (2010) analyzes the 2008 financial crisis as a run on the "shadow banking" system, and discusses Fed policy in response to these events. Pozsar (2011) argues that the rise of shadow banking was driven by high demand for safe and liquid secured assets similar to Treasury debt, primarily driven by institutional cash pools. In the years immediately preceding the crisis, the demand for safe liquid assets exceeded the supply of government liabilities by at least $1.5 trillion, and the shadow banking sector developed to fill this need. Pozsar (2011) recommends that policy makers consider issuing a greater volume of Treasury bills to fill this demand for liquidity.[6]

There have been several empirical studies of the role of government liquidity. Krishnamurthy and Vissing-Jorgensen (2012) find that government bonds hold a liquidity premium over corporate bonds, with about half of the 100 bps average spread explained by the superior liquidity of Treasury bonds. Kashyap and Stein (2000) find evidence of a credit channel of monetary policy.

Finally, there is a large literature exploring the linkages between the financial sector and the real economy.[7] Most of this literature focuses on the role of banks in lending funds to firms for investment. I depart from this literature by considering the role of the financial sector in providing liquid assets as means of saving, not providing loans. Banks provide both loan services to borrowers and liqudity services to depositors, and I focus on the latter in this paper.

2.2 Model Without Banks

I first consider the model without banks. This case is a simplified version of Holmstrom and Tirole (1998). I give the derivation in detail because it will serve as the basis for the model with banks that follows.

[6]Bernanke et al. (2011), Caballero (2010), and Acharya and Schnabl (2010) also discuss the role of high demand for safe assets in the run-up to the 2008 crisis.

[7]Prominent examples include Bernanke and Gertler (1989) and Kiyotaki and Moore (1997).

2.2.1 Preferences, Endowment, and Technology

I consider an economy containing three periods, labeled $t = 0, 1, 2$. The economy is populated by a unit measure of two types of agents: households and firms. There is a single good used for both consumption and investment, which is not storable between periods. All agents have linear utility over consumption across all three periods, i.e. they have utility functions $u(c_0, c_1, c_2) = c_0 + c_1 + c_2$.

Households have an endowment of the good in each period that is sufficiently large that their ability to lend to firms is never limited in the constrained case. I denote their period t endowment by H_t.[8] Households also enter period 0 holding a stock $\bar{\ell}$ of trees that yield a unit return of the good in period 2. These trees are the economy's stock of liquid assets, which I will refer to as outside liquidity. Firms have an initial endowment of A units of the good in period 0, and no endowment in any other period.

Firms operate a linear production technology which yields a gross return of $\rho_1 > 1$ between period 0 and period 2. Thus if a firm invests I units of the good in period 0, the project will produce $\rho_1 I$ in period 2 if the project is completed. During period 1, each firm receives an idiosyncratic shock ρ. A firm that suffers shock ρ must supply an additional ρI units of the good to the project to continue

[8] At the unconstrained optimum defined below, this endowment will limit investment. The assumption here is that H_t is sufficiently large relative to firm assets that for a non-trivial leverage constraint, firm borrowing will not be limited by available funds, but only by the firm's ability to commit to repaying their investors.

its operation. If these funds are not provided, the project will produce nothing. ρ is drawn from $\{\rho_L, \rho_H\}$, with $\rho_L < \rho_H < \rho_1$, and takes on the value ρ_L with probability p, and ρ_H with probability $1-p$. Therefore in period 1 a measure p of firms will suffer the low shock ρ_L, and a measure $1-p$ will suffer the high shock ρ_H.

We can summarize the production plan of a firm by the initial investment I and a continuation policy rule $\lambda_s \in \{0,1\}$ for $s \in \{H, L\}$, where $\lambda_s = 1$ means that the project continues given liquidity shock s.

2.2.2 Unconstrained Optimum

I now characterize the unconstrained optimum. Consider a consumption plan $\{C_1, C_2, C_3\}$, where C_t is total consumption of firms and households in period t. The optimal production plan maximizes total consumption subject to the economy's resource constraints, which are

$$C_0 + I \leq A + H_0 \tag{2.1}$$

$$C_1 + p\lambda_L\rho_L I + (1-p)\lambda_H\rho_H I \leq H_1 \tag{2.2}$$

$$H_2 + p\lambda_L\rho_1 I + (1-p)\lambda_H\rho_1 I + \bar{\ell} \geq C_2 \tag{2.3}$$

In period 0, households have endowment H_0 and firms have endowment A. These funds are spent on consumption $C_0 \geq 0$ and investment $I \geq 0$. In period 1, households have endowment H_1. Funds $C_1 \geq 0$ are used for consumption, funds $p\lambda_L\rho_L I$ are used to meet low liquidity shocks, and funds $(1-p)\lambda_H\rho_H I$ are used to meet high liquidity shocks. In period 2, households have endowment H_2 and earn return $\bar{\ell}$ from their holdings of outside liquidity, and firms produce $p\lambda_L\rho_1 I + (1-p)\lambda_H\rho_1 I$.

These funds are spent on consumption $C_2 \geq 0$.

The unconstrained optimal production plan is the solution to

$$\max_{\lambda, I} \{C_0 + C_1 + C_2\} \tag{2.4}$$

$$\text{s.t.} \quad (2.1) \text{ - } (2.3)$$

$$C_i \geq 0,\ I \geq 0,\ \lambda_L \in \{0,1\},\ \lambda_H \in \{0,1\}$$

Proposition 1 (Optimal Production Plan)**.** The unconstrained optimal production plan is $\lambda_H = 1$, $\lambda_L = 1$, and

$$I = \begin{cases} A + H_0 & R_1 \geq 0 \\ \\ 0 & R_1 < 0 \end{cases} \tag{2.5}$$

where $R_1 = p(\rho_1 - \rho_L) + (1-p)(\rho_1 - \rho_H) - 1$.

Proof. See appendix B.1.[9] □

Since utility is linear with no discounting, any distribution of consumption between agents is Pareto efficient as long as total consumption is maximized. Here R_1 is the net expected return on investment. Since agents are indifferent between consuming in periods 0, 1, or 2, as long as there is a positive expected return to investment, the optimal production plan is to invest all available resources in period 0. Since both liquidity shocks are smaller than the final output of the project ρ_1, if the project has been undertaken it is profitable to meet any liquidity shock that occurs and bring the project to completion.

[9] This proposition assumes that H_1 is sufficiently large relative to H_0 that the limiting factor on investment is the supply of funds in period 0, not funds available to meet the liquidity shocks in period 1. This is equivalent to the condition $H_1 \geq [p\rho_L + (1-p)\rho_H](A + H_0)$.

I assume for the rest of the paper that the project is profitable even if only the low shock triggers continuation, meaning

$$p\left(\rho_1 - \rho_L\right) > 1 \tag{2.6}$$

Condition (2.6) together with $\rho_1 > \rho_H$ implies $R_1 > 0$, so investment yields a positive return in expectation. Thus the unconstrained optimal production plan is to invest all available resources as defined in (2.5).

2.2.3 Limited Pledgeability

I assume that firm borrowing is subject to a limited pledgeability constraint that makes it impossible to implement the first-best production plan. This constraint arises from moral hazard. At the end of period 1, each firm with a functioning project is presented with an alternative opportunity. If a firm shirks by pursuing this opportunity, the firm's project fails and the firm earns private benefit BI, where $B > 0$. There is no legal recourse for investors to seize any portion of this private benefit. Therefore in every state in which a firm's project is successful, the firm must receive at least a share BI of the output in order to cooperate.[10]

Let $\rho_0 = \rho_1 - B$ be the return on investment net the firm's outside opportunity. Since a successful project produces $\rho_1 I$ output in period 2, and since the firm must receive BI in order to operate the project, external investors may receive no more than $\rho_1 I - BI = \rho_0 I$.[11] Therefore $\rho_0 I$ is the portion of a project's final output that

[10] There are many alternate motivations for limited pledgeability. This description corresponds to the microfoundation provided in Holmstrom and Tirole (1998).

[11] It is also possible to make firms' other asset holdings subject to an agency cost. Then firms

a firm can credibly promise to repay investors. I refer to ρ_0 as firm pledgeability, and assume it satisfies

$$\rho_0 < \rho_H \tag{2.7}$$

$$\rho_0 < 1 + p\rho_L + (1-p)\rho_H \tag{2.8}$$

$$p\rho_0 < 1 + p\rho_L \tag{2.9}$$

$$\rho_0 > p\rho_L + (1-p)\,\rho_H \tag{2.10}$$

Condition (2.7) says that pledgeable funds are insufficient to meet some liquidity shocks. Condition (2.8) and (2.9) say that expected funds required per unit of investment exceed the pledgeable portion of the return, no matter which shocks are met, so that projects cannot be financed solely with external funds. Thus limited pledgeability is sufficiently severe that it implies a "skin in the game" constraint that requires firms to put up some of their own capital. If either of these conditions failed to hold, the scale of investment would not be limited by pledgeability. Finally, condition (2.10) says that pledgeable funds are sufficient to meet the expected size of the liquidity shock, so that there are enough pledgeable funds owned by firms in the aggregate to meet all liquidity shocks in period 1. Thus liquidity will only be scarce when there is imperfect pooling of these funds. Note that (2.10) and (2.7) together imply that $\rho_0 > \rho_L$.

would have to keep some of the return on these assets, and only a fraction could be used to repay the original investors. This is analogous to a tax on liquid assets purchased by firms, since firms would have to buy $(1+\tau)\,\ell$ assets in order to increase their pledgeable funds by ℓ. From the perspective of the firm, this is analogous to an increase in the liquidity premium.

Households cannot borrow at all. Since they cannot offer any collateral, they will renege on any promise they make.[12]

2.2.4 Households

Households supply funds by purchasing state-contingent assets from firms. Given linear utility and no discounting, households are willing to purchase any asset that promises an expected return of at least 1. Therefore the supply of loanable funds in the economy is perfectly elastic at $r = 1$. Households consume all of their income that is not used to purchase assets.

Let q be the price of trees in period 0. Since trees yield a unit return, households will demand an infinite quantity of trees if $q < 1$ and no trees if $q > 1$. If $q = 1$ households are indifferent between any quantity of trees. Therefore household asset demand is perfectly elastic at $q = 1$.

2.2.5 Firms

Firms raise funds from households in period 0 by selling a contract offering a state-contingent return in period 2. The contract specifies payments to initial investors of $R_s^I \geq 0$ in the case of shock $s \in \{L, H\}$. These payments must be positive because households cannot commit to providing future funds. Firms also purchase ℓ trees from households at a price of q. In period 1, firms experience liquidity shocks. Firms raise funds to meet these shocks by selling new claims to households to be paid in period 2. Since households know the shock experienced by a firm, they will require

[12] We can interpret household's holdings of outside liquidity $\bar{\ell}$ as their pledgeable funds.

exact compensation for funds provided. I denote by R_s^1 the funds repaid in period 2 to period 1 investors by a firm that experiences shock $s \in \{L, H\}$.

Any equilibrium contract must satisfy a number of constraints arising from incentive compatibility. First, the initial investors must be compensated in expectation for the funds they provide. The firm requires I funds in period 0 for its initial investment, plus $q\ell$ funds to purchase trees. Firms have initial assets A, and the rest must be raised from outside investors. In order for initial investors to buy this contract, their expected payments $pR_L^I + (1-p)R_H^I$ must satisfy

$$pR_L^I + (1-p)R_H^I \geq I + q\ell - A \tag{2.11}$$

In period 1, firms receive liquidity shocks, and those that continue production meet these shocks by issuing new liabilities to investors. I denote by R_s^1 the repayment in period 2 of a firm that experiences liquidity shock $s \in \{L, H\}$ to its period 1 investors. These funds must be sufficient to cover liquidity shocks, which implies constraints

$$\lambda_H \rho_H I \leq R_H^1 \tag{2.12}$$

$$\lambda_L \rho_L I \leq R_L^1 \tag{2.13}$$

Firms need a sufficient share of profits to cooperate. They repay outside investors from the pledgeable portion of their output and from the return on their asset holdings ℓ. Total repayments to outside investors may not exceed

$$R_L^I + R_L^1 \leq \lambda_L \rho_0 I + \ell \tag{2.14}$$

$$R_H^I + R_H^1 \leq \lambda_H \rho_0 I + \ell \tag{2.15}$$

Note that firms do not use outside assets ℓ to pay for liquidity shocks directly. Rather, buying assets in period 0 allows firms to increase their pledgeability by the amount ℓ, and therefore borrow more from period 1 investors. These two specifications are analogous, but writing the constraints in this manner will simplify the exposition.

How can liquid assets improve the allocation? We can think of the fundamental problem as a lack of commitment on the part of households. Since projects that experience a high shock are profitable to continue ($\rho_1 > \rho_H$), it is optimal for households to provide funds to firms to meet high liquidity shocks. However, once a high liquidity shock is realized firms do not have sufficient pledgeable funds to meet them, since they can only promise a fraction $\rho_0 < \rho_H$ of their output. Ex ante, original investors would like to promise to provide funds at a loss in this event, because this would allow higher investment and therefore higher payments in good states. Since $\rho_0 > p\rho_L + (1-p)\rho_H$, the higher payments in good states can be large enough to compensate households for the negative returns in bad states so that households come out ahead on average. However, by assumption households are not able to commit to providing funds at a loss once the liquidity shock is realized.

Liquid assets provide a mechanism for circumventing this commitment problem. Rather than promising to provide funds in the future, households can provide the funds up front, and the firm can use them to purchase assets ℓ. These assets can then be used by the firm as collateral to raise additional funds in the event of a high liquidity shock. Effectively, households are providing collateral for their promise to pay firms in period 1 to meet liquidity shocks. Therefore liquid assets serve as a social commitment mechanism.

2.2.6 Equilibrium Contract

Firms choose the profit-maximizing contract

$$\max_{R,\lambda,I,\ell} \left\{ p\left(\lambda_L \rho_1 I - R_L^I - R_L^1 + \ell\right) + (1-p)\left(\lambda_H \rho_1 I - R_H^I - R_H^1 + \ell\right) \right\}$$

$$\text{s.t.} \quad (2.11) \text{ - } (2.15),\ R_L^I \geq 0,\ R_H^I \geq 0,\ R_L^1 \geq 0,\ R_H^1 \geq 0,\ I \geq 0$$

where $R = \left\{R_L^I, R_H^I, R_L^1, R_H^1\right\}$ and $\lambda = \{\lambda_L, \lambda_H\}$. Non-negativity constraints on payments to investors arise from the assumption of limited commitment by households.

Lemma 1. Under the equilibrium contract, the firm always meets the low shock ($\lambda_L = 1$), and constraints (2.11) - (2.14) hold with equality.

Proof. See appendix B.1. □

Intuitively, since our assumptions on parameters imply a positive return to investment, firms will invest until the pledgeability constraint is binding. Moreover, since there are sufficient funds available to finance all desired investment at a unit return, firms will exactly compensate all investors.

Lemma 1 greatly simplifies the statement of the problem. Since constraints (2.12) - (2.13) and (2.14) - (2.15) hold with equality, we can substitute them directly into the various expressions in the problem. Substituting constraints (2.14) and (2.15) into the objective function, we find that firm payoffs are $p\,(\rho_1 - \rho_0)\, I$ if the firm does not meet the high shock, and $(\rho_1 - \rho_0)\, I$ if the firm does. Since the pledgeability constraint binds, firms receive exactly the amount $(\rho_1 - \rho_0)\, I$ necessary for them to cooperate in equilibrium. Also, these equations allow us to derive exact expressions for R_s^1 and R_s^I. These are

$$R_H^1 = \lambda_H \rho_H I$$

$$R_L^1 = \rho_L I$$

$$R_L^I = (\rho_0 - \rho_L)\, I + \ell$$

$$R_H^I = \lambda_H\, (\rho_0 - \rho_H)\, I + \ell$$

Together with (2.11), these imply leverage constraint

$$I \leq \frac{A - (q-1)\,\ell}{1 - p\,(\rho_0 - \rho_L) - \lambda_H (1-p)\,(\rho_0 - \rho_H)} \tag{2.16}$$

which will hold with equality at the optimum. Intuitively, the pledgeability constraint is sufficiently severe that firms must supply a fraction of the capital for investment (a skin-in-the-game constraint). This implies that investment is limited by firms' initial assets, which can be expressed as a limit on leverage, as given in (2.16). When $q > 1$, firms must put up some of their own funds to purchase additional liquidity, reducing available funds and thus equilibrium investment.

The non-negativity constraints on R_H^1, R_L^1, and R_L^I are trivially satisfied. The non-negativity constraint on R_H^I can be expressed as $\ell \geq \lambda_H\, (\rho_H - \rho_0)\, I$, which implies that holdings of outside liquidity ℓ must be sufficient to finance the high liquidity shock.

We can now express the optimal contracting problem of the firm as

$$\max_{\lambda_H, I, \ell} \{(\rho_1 - \rho_0)\,[p + (1-p)\lambda_H]\,I\}$$

$$\text{s.t.} \quad (2.16),\ \lambda_H\,(\rho_H - \rho_0)\,I \leq \ell$$

plus non-negativity constraints on I and ℓ. The following proposition characterizes the solution:

Proposition 2. The optimal production plan of firms is to meet the high shock ($\lambda_H = 1$) if and only if

$$q - 1 \leq \frac{(1-p)\,[1 - p\,(\rho_H - \rho_L)]}{p\,(\rho_H - \rho_0)} \tag{2.17}$$

If (2.17) holds with equality, the firm is indifferent between $\lambda_H = 1$ and $\lambda_H = 0$. The optimal choice of I is

$$I = \frac{A}{1 - p\,(\rho_0 - \rho_L) - \lambda_H\,(q-p)\,(\rho_0 - \rho_H)}$$

and the optimal choice of ℓ is $\ell = \lambda_H\,(\rho_H - \rho_0)\,I$.

Proof. See appendix B.1. □

Since $q - 1$ is always non-negative, (2.17) will never hold if the parameters satisfy $p\,(\rho_H - \rho_L) > 1$. This condition is a simplification of

$$\frac{pA}{1 - p\,(p_0 - \rho_L)} > \frac{A}{1 - p\,(\rho_0 - \rho_L) - (1-p)\,(\rho_0 - \rho_H)}$$

where the left-hand side is output when firms meet only the low shock, and the right-hand side is output when firms meet both shocks. When $p\,(\rho_H - \rho_L) \leq 1$, (2.17) defines a cutoff level of q above which it is optimal to meet only the low shock. Intuitively, firms allocate limited pledgeable funds between financing the initial investment and meeting liquidity shocks. If liquidity is sufficiently expensive, firms substitute away from the more expensive input to the production process by increasing the scale of the initial investment, and reducing the fraction of liquidity shocks they meet.

2.2.7 Equilibrium

Proposition 2 defines a demand for liquid assets ℓ, which is a decreasing function of q and drops to zero at the cutoff defined by (2.17). To determine equilibrium q, we impose clearing in the liquidity market. Since households will only hold assets that yield a unit return or better, if $q > 1$ aggregate liquidity demand is exactly equal to demand from firms. Equilibrium is depicted in Figure 2.1 and characterized in Proposition 3.

Proposition 3. Let $I_0 = A/\chi_0$ and $I_1(q-1) = A/(\chi_1 + (q-1)(\rho_H - \rho_0))$, where $\chi_0 = 1 - p\,(\rho_0 - \rho_L)$ and $\chi_1 = 1 - \rho_0 + p\rho_L + (1-p)\,\rho_H$. Then,

(i) If $p\,(\rho_H - \rho_L) > 1$, equilibrium $q - 1 = 0$ and all firms choose $\lambda_H = 0$, $\ell = 0$, and $I = I_0$.

(ii) If $p\,(\rho_H - \rho_L) \leq 1$ and $\bar{\ell} \geq (\rho_H - \rho_0)\, I_1(0)$, equilibrium $q - 1 = 0$ and firms choose $I = I_1(0)$, $\ell = (\rho_H - \rho_0)\, I_1(0)$, and $\lambda_H = 1$.

(iii) If $p\,(\rho_H - \rho_L) \leq 1$ and $p\,(\rho_H - \rho_0)\, I_0 \leq \bar{\ell} \leq (\rho_H - \rho_0)\, I_1(0)$, equilibrium $q - 1$ is

$$q - 1 = \frac{A}{\bar{\ell}} - \frac{\chi_1}{\rho_H - \rho_0}$$

and firms choose $\lambda_H = 1$, $\ell = \bar{\ell}$, and $I = I_1(q-1) = \bar{\ell}/(\rho_H - \rho_0)$.

(iv) If $p(\rho_H - \rho_L) \leq 1$ and $\bar{\ell} < p(\rho_H - \rho_0) I_0$, equilibrium q is

$$q - 1 = \frac{\chi_0 - p\chi_1}{p(\rho_H - \rho_0)}$$

and firms are indifferent between $\lambda_H = 0$ and $\lambda_H = 1$. A fraction ζ of firms choose $\lambda_H = 1$, $I = I_1(q-1) = pI_0$, and $\ell = (\rho_H - \rho_0) pI_0$, where $\zeta = \bar{\ell}/[(\rho_H - \rho_0) pI_0]$. The remaining fraction $1 - \zeta$ choose $\lambda_H = 0$, $\ell = 0$, and $I = I_0$.

Proof. See appendix B.1. $\square$

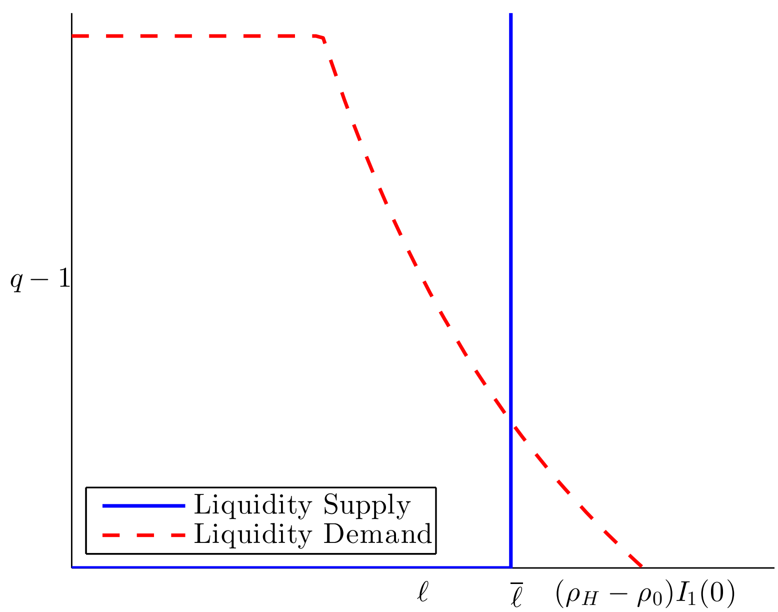

Figure 2.1: Equilibrium without banks.

Assuming $p(\rho_H - \rho_L) \leq 1$ so that it is potentially optimal to meet the high shock, Proposition 3 defines two cutoff levels of $\bar{\ell}$.[13] If outside liquidity $\bar{\ell}$ is above

[13]Note that $p(\rho_H - \rho_0) I_0 \leq (\rho_H - \rho_0) I_1(0)$ if and only if $p(\rho_H - \rho_L) \leq 1$, and so the cutoffs satisfy the assumed ordering.

the higher cutoff, then all firms meet both liquidity shocks and there is no liquidity premium ($q-1=0$). In this case, the economy still fails to achieve the unconstrained optimum due to the usual effect of credit constraints restricting initial investment, but there are no additional effects due to limited liquidity.

I_0 is the optimal level of investment when $\lambda_H = 0$, and $I_1(q-1)$ is the optimal level of investment when $\lambda_H = 1$, which is strictly decreasing in $q-1$. I refer to the equilibrium allocation with $q = 1$ as the *constrained optimum*. The constrained optimum achieves the highest level of investment that respects the aggregate limited pledgeability constraint. When both shocks are met, this level of investment is $I_1(0) = A/\chi_1$, meaning that firms must provide χ_1 of their own funds for every unit of investment, and so $1/\chi_1$ is the leverage ratio in the economy. $\chi_1 = 1 + p\rho_L + (1-p)\rho_H - \rho_0$ is the amount by which expected required funds per unit of investment exceed pledgeable funds. Our assumptions on ρ_0 imply $\chi_1 \in (0,1)$, and so $I \in (A, \infty)$. When ρ_0 is large relative to $1 + p\rho_L + (1-p)\rho_H$, most investment is financed using external funds and leverage is high. When ρ_0 is small relative to $1 + p\rho_L + (1-p)\rho_H$, external financing is limited and leverage is low.

If $\bar{\ell}$ lies below the higher cutoff, then the economy is liquidity constrained, in the sense that there is insufficient liquidity for all firms to meet the high shock at the constrained optimal level of investment. Therefore equilibrium q must rise until the quantity of liquidity demanded by firms equals available liquidity $\bar{\ell}$. Liquidity demand is downward sloping because firms' pledgeable funds are used to pay the higher price of liquidity, and so funds available to finance initial investment fall, tightening the leverage constraint.

If $\bar{\ell}$ lies below the lower cutoff, then there is insufficient liquidity available for all firms to meet the high shock. Intuitively, firms are indifferent between meeting and not meeting the high shock when $I_1 = pI_0$. Thus the lower cutoff is the level of $\bar{\ell}$ that is just sufficient for all firms to meet the high liquidity shock given investment $I_1 = pI_0$. The corresponding liquidity premium $q-1$ is implicitly defined by $I(q-1) = pI_0$. For $\bar{\ell}$ below this level, some firms will meet the high shock and some will not, with the fraction determined by available $\bar{\ell}$.

I refer to an equilibrium that lies strictly between these two cutoffs as an interior equilibrium. This is the most interesting case, because here a marginal change in available liquidity will change both the equilibrium investment and the equilibrium liquidity premium q. Such a shift is depicted in Figure 2.2, which shows that a fall in outside liquidity $\bar{\ell}$ will increase the equilibrium liquidity premium $q-1$ and decrease equilibrium investment I.

In the remainder of the paper, I assume that $p(\rho_H - \rho_L) < 1$, so that firms strictly prefer to meet the high shock if there is no liquidity premium.

2.3 Model With Banks

In the previous section, I showed that there may be insufficient liquidity in the economy, resulting in lower investment and, in some cases, a lower share of firms meeting high liquidity shocks. But by assumption, there are sufficient pledgeable funds in the aggregate to meet all liquidity shocks, since firms that experience a low shock have more pledgeable funds available than they need, whereas firms that

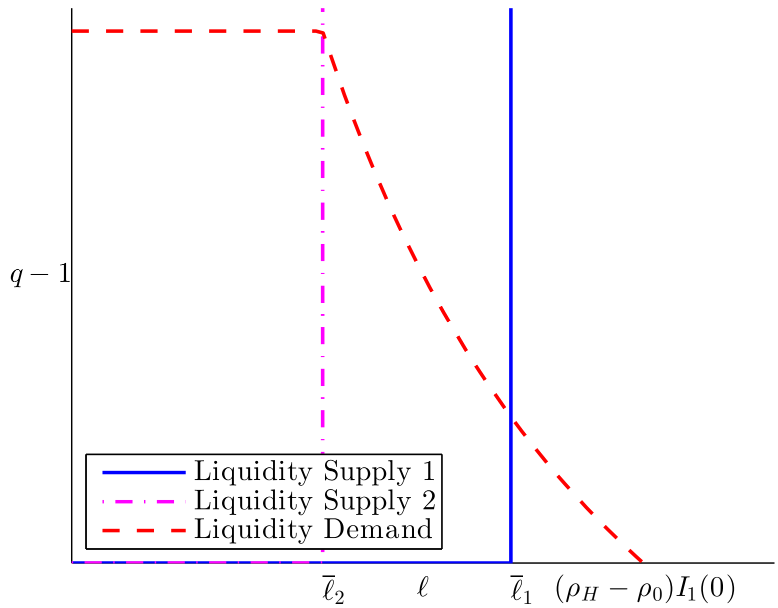

Figure 2.2: Fall in outside liquidity without banks.

experience a high shock are short pledgeable funds. Thus firms can achieve the constrained optimum if they are able to pool their pledgeable funds by arranging for transfers from firms with excess funds to firms that receive a high shock and need additional funds.

To implement this liquidity pooling, I introduce financial intermediaries (banks) into the model.[14] Banks sell financial contracts in period 0 to firms in exchange for shares in the firm. These contracts obligate banks to provide M funds in period 1 to firms that experience a high liquidity shock. Banks obtain these period 1 funds from households by borrowing against the future returns on their portfolio. Since firms that experience a high liquidity shock will dilute their period 0 shares with new borrowing, banks take losses on these credit lines. To make up for this, banks

[14] I assume that firms cannot implement liquidity pooling by themselves.

Period 0	Period 1	Period 2
• Firms buy credit lines from banks with maximum M. • Banks receive firm shares worth $(1-p)M+\pi$.	• Measure $1-p$ of firms receive high shock, receive M each from banks. • Banks raise $(1-p)M$ from households, pay to firms. • Households receive bank claims worth $(1-p)M$.	• Measure $1-p$ of firms pay $R_H^B = 0$ each to banks. • Measure p of firms pay $R_L^B = ((1-p)M+\pi)/p$ each to banks. • Banks pay $(1-p)M$ to households, keep π.

Figure 2.3: Timeline of bank credit lines.

require higher payments from firms that experience the low shock. The resulting pattern of payments between banks, firms, and households is shown in Figure 2.3.

I call these arrangements *credit lines*, since they are analogous to the credit lines in Holmstrom and Tirole (1998). Their key feature is that they pool pledgeable funds via effective transfers from firms that experience low liquidity shocks to firms that experience high shocks, intermediated by banks. The most straightforward implementation of this arrangement is a system of direct contingent transfers, as described above. This arrangement can be interpreted as an insurance contract.

Alternatively, the same set of net payments can be implemented using demand deposits, as discussed in more detail in section 2.3.5. In this case, firms raise additional initial funds and deposit them with banks, and banks invest these funds by purchasing shares in firms. In period 1, firms withdraw funds from their accounts to meet liquidity shocks, and banks raise new funds from households to finance these withdrawals. Another implementation is a contract that resembles a corporate credit line. Under this arrangement, firms pay an upfront fee in exchange for access to a line of credit at a specified rate, together with a fee on unused credit.[15] In period 1, firms draw on their lines of credit to meet liquidity shocks, and firms that experience a high shock default in period 2. The upfront fee covers the losses on defaulting credit lines. In effect, firms pay a premium to lock in future credit at good terms, with this premium compensating banks for the commitment to provide credit at an expected loss to some firms. Thus the described contract is equivalent to several forms of liquidity services actually provided to firms by financial intermediaries.

If banks are not subject to any commitment frictions, they will be able to perfectly insure against idiosyncratic liquidity shocks, and the economy will achieve the constrained optimal level of investment, even when there is no outside liquidity ($\bar{\ell} = 0$). However, it is reasonable to think that banks are subject to agency costs of a similar nature to firms. Banks act as agents that invest on behalf of their depositors, in this case households that provide them with funds in period 1. In the process of lending, banks may need to exert effort in screening firms, verifying the liquidity shock, and collecting funds. Banks also have opportunities to defraud their depositors by withholding funds. If banks do not receive sufficient profits relative to the size of their portfolios, they will not exert full effort screening or collecting loans, and may have an incentive for fraud.

Let bank agency costs be represented by the function $C(\varphi M, K)$, where M is the payment in period 1 to firms that experience the high liquidity shock, φ is the measure of credit lines sold by the bank, and K is bank equity at the start of period 2. I assume that K is exogenous and known to all agents at the beginning of period

[15] This is a common arrangement for corporate credit lines, as described by Sufi (2009).

0. In order to commit to properly intermediating a measure φ of credit lines of size M, a bank with equity K must receive at least $C(\varphi M, K)$ in profits, i.e.

$$\varphi\pi \geq C\left(\varphi M, K\right)$$

where π is profits received per credit line sold. This is the bank's incentive compatibility constraint.

Letting $D = \varphi M$ denote total funds intermediated, I assume that the agency cost function $C(\cdot)$ is twice continuously differentiable at all points (D, K), with $D \geq 0$ and $K > 0$. I assume that agency costs are increasing and convex in funds intermediated, $C_1(\cdot) > 0$ and $C_{11}(\cdot) > 0$, and decreasing in bank equity, $C_2(\cdot) < 0$. Finally, I assume that $\lim_{D\to 0} C_1\left(D, K\right) = 0$, $\lim_{D\to\infty} C_1\left(D, K\right) = \infty$, $C\left(0, K\right) = 0$ for $K \geq 0$, and $C(D, 0) = \infty$ for $D > 0$.

2.3.1 Preferences, Endowment, and Technology

Preferences, endowments, and technology are the same as in section 2.2, except for the inclusion of banks. Like firms and households, there is a unit measure of banks, and banks have linear utility over consumption across all three periods, i.e. $u(c_0, c_1, c_2) = c_0 + c_1 + c_2$. As described above, banks have equity K entering period 2.

2.3.2 Households

As before, households supply funds perfectly elastically at an expected return of 1. Therefore every asset will be priced at its expected return, or households will not hold that asset. Households are willing to sell their supply of trees at a price $q \geq 1$, and will not buy any trees if $q > 1$.

2.3.3 Firms

In period 0, firms choose production plan $\{I, \lambda\}$, with $I \geq 0$ investment in period 0, and project continuation policy rule λ_s for state $s \in \{L, H\}$. Firms raise funds in period 0 by selling shares offering a state-contingent return in period 2. The contract specifies payments to initial investors of R_s^I in the case of shock $s \in \{L, H\}$. Firms also purchase ℓ trees from households at a price of q.

Altogether, the firm requires total funds $I + q\ell$ in period 0. It has its own funds A, and raises the remainder from outside investors. Since there is a unit cost to external funds, expected payments to initial investors must satisfy

$$pR_L^I + (1-p)R_H^I \geq I + q\ell - A \tag{2.18}$$

In period 0, firms also choose whether to purchase credit lines from banks. I represent this decision by $\lambda_B \in \{0, 1\}$. The equilibrium credit line consists of a pair (π, M), where M is the maximum amount that can be drawn on in period 1, and π is the expected profits received by the bank. The firm pays for the credit line with shares that yield net payments of R_s^B in period 2 in the event of shock s. These payoffs must satisfy

$$pR_L^B + (1-p)R_H^B \geq \pi + (1-p)M \tag{2.19}$$

In period 1, firms experience liquidity shocks. Firms raise funds to meet these shocks by issuing new shares offering a return in period 2. I denote by R_s^1 the funds repaid in period 2 to period 1 investors by a firm that experiences shock s. Firms that experience a low shock raise funds ρ_L, whereas firms that experience a high shock raise new funds $\rho_H - \lambda_B M$ and draw on their credit lines for the remainder. Since there is a unit cost of funds in period 1, these repayments must satisfy

$$R_L^1 \geq \lambda_L \rho_L I \tag{2.20}$$

$$R_H^1 \geq \lambda_H \rho_H I - \lambda_B M \tag{2.21}$$

Firms repay their outside investors from the pledgeable portion of their income and from the return on their asset holdings ℓ. Thus total repayments in each state cannot exceed available pledgeable funds

$$R_L^I + R_L^1 + \lambda_B R_L^B \leq \lambda_L \rho_0 I + \ell \tag{2.22}$$

$$R_H^I + R_H^1 + \lambda_B R_H^B \leq \lambda_H \rho_0 I + \ell \tag{2.23}$$

Firms choose production to solve

$$\max_{R,\lambda,I,\ell} \left\{ p \left(\lambda_L \rho_1 I - R_L^I - R_L^1 - \lambda_B R_L^B + \ell \right) + (1-p) \left(\lambda_H \rho_1 I - R_H^I - R_H^1 - \lambda_B R_H^B + \ell \right) \right\}$$

$$\text{s.t.} \quad (2.18)\ \text{-}\ (2.23),\ \{R, I\} \geq 0$$

where $R = \left\{ R_L^I, R_H^I, R_L^1, R_H^1, R_L^B, R_H^B \right\}$. Non-negativity constraints on payments to households arise from the assumption of limited commitment by households.

Lemma 2. Under optimal firm behavior, $\lambda_L = 1$ and constraints (2.18) - (2.23) hold with equality.

Proof. See appendix B.2. □

The intuition for this result is similar to Lemma 1. Since investment yields a positive return, firms will borrow until the pledgeability constraints are binding. The other constraints bind because this minimizes payments to creditors.

Lemma 2 greatly simplifies the statement of the problem. Since (2.22) - (2.23) hold with equality, we can substitute them into the objective function to find that firm payoffs are $p\,(\rho_1 - \rho_0)\,I$ if $\lambda_H = 0$, and $(\rho_1 - \rho_0)\,I$ if $\lambda_H = 1$. Since the pledgeability constraint binds, firms receive exactly the amount $(\rho_1 - \rho_0)\,I$ necessary for them to cooperate in equilibrium.

Taking a weighted sum of (2.22) and (2.23), we have:

$$p\left(R_L^I + R_L^1 + \lambda_B R_L^B\right) + (1-p)\left(R_H^I + R_H^1 + \lambda_B R_H^B\right) = p\,(\rho_0 I + \ell) + (1-p)\,(\lambda_H \rho_0 I + \ell)$$

Substituting in (2.18) - (2.21), we obtain the leverage constraint

$$I = \frac{A - (q-1)\,\ell - \lambda_B \pi}{1 - p\,(\rho_0 - \rho_L) - \lambda_H (1-p)\,(\rho_0 - \rho_H)} \tag{2.24}$$

Intuitively, the pledgeability constraint is sufficiently severe that firms must supply a fraction of the capital for investment (a skin-in-the-game constraint). This implies that investment is limited by firms' initial assets, which can be expressed as a constraint on leverage.

We can now rewrite the other constraints in the problem. The non-negativity constraints on R_L^1, R_H^1, R_L^I and R_L^B are trivially satisfied, given $M \geq 0$. The non-negativity constraints on R_H^I and R_H^B imply

$$\lambda_H (\rho_H - \rho_0) I \leq \lambda_B M + \ell \tag{2.25}$$

Thus the firm problem reduces to

$$\max_{\lambda_H,\lambda_B,I,\ell} \{(\rho_1-\rho_0)\,[p+(1-p)\lambda_H]\,I\}$$

$$\text{s.t.} \quad (2.24) \text{ and } (2.25)$$

Optimal firm behavior is characterized by Proposition 4.

Proposition 4. Let $q \geq 1$ and $(M,\pi) \geq 0$ with $M\chi_1 \leq (\rho_H-\rho_0)(A-\pi)$ be given.[16] Let $\chi_1 = 1-\rho_0+p\rho_L+(1-p)\,\rho_H$ and $\chi_0 = 1-p(\rho_0-\rho_L)$. Then,

(i) Firms will choose $\lambda_B = 1$ if and only if

$$\frac{\pi}{M} \leq q-1 \tag{2.26}$$

and will be indifferent between $\lambda_B = 0$ and $\lambda_B = 1$ if (2.26) holds with equality.

(ii) Firms will choose $\lambda_H = 1$ if and only if

$$\frac{A-\lambda_B(\pi-(q-1)M)}{\chi_1+(q-1)(\rho_H-\rho_0)} \geq \frac{pA}{\chi_0} \tag{2.27}$$

and will be indifferent between $\lambda_H = 0$ and $\lambda_H = 1$ if (2.27) holds with equality.

(iii) Firms will invest

$$I = \frac{A-\lambda_H\lambda_B\,[\pi-(q-1)M]}{1-p(\rho_0-\rho_L)-\lambda_H(q-p)(\rho_0-\rho_H)}$$

and will purchase outside liquidity

$$\ell = \lambda_H\left((\rho_H-\rho_0)I-\lambda_B M\right)$$

[16] The condition $M\chi_1 \leq (\rho_H-\rho_0)(A-\pi)$ rules out cases where M is larger than the amount of liquidity any firm seeks to hold.

Proof. See appendix B.2. □

Since by Lemma 2 all constraints bind and firms always meet the low liquidity shock, the only decisions that remain are whether to meet the high liquidity shock, and whether to purchase a credit line from banks. Proposition 4 says that firms hold credit lines as long as they are cheaper than outside liquidity. Since the unit cost of outside liquidity is q, whereas the unit cost of credit lines is $1 + \pi/M$, this implies condition (2.26).

Optimal investment when firms do not meet the high shock is $I_0 = A/\chi_0$, and when firms meet the high shock is

$$I_1(q-1) = \frac{A - \lambda_B[\pi - (q-1)M]}{\chi_1 + (q-1)(\rho_H - \rho_0)}$$

Firm profits are proportional to $(p + \lambda_H(1-p))I$, and so firms will meet the high shock as long as $I_1 > pI_0$, which is condition (2.27). Since $\rho_H > \rho_0$, meeting the high shock means there are less pledgeable funds available per unit invested, and so the leverage constraint (2.24) implies $I_1 < I_0$. Moreover, a higher price of liquidity will raise the costs of meeting the high shock, further reducing I_1.

2.3.4 Banks

There is a unit measure of banks that sell credit lines to firms. I assume that bank issuance of credit lines is perfectly diversified across firms. Thus although there is a unit measure of each agent, in a sense there are "more" firms than banks.[17]

[17] We can formalize this notion by supposing there are N banks of measure $1/N$, and N^2 firms of measure $1/N^2$, so that each bank has N firms as customers, and letting $N \to \infty$.

Banks sell credit lines to firms, which are contracts represented by $\left(M, R_L^B, R_H^B\right)$. M is the credit maximum, which are the funds the firm may draw on in period 1 to meet a high liquidity shock. R_L^B and R_H^B are net payments from the firm to the bank in period 2 in case of a low shock and high shock respectively.

If all firms purchase credit lines and both shocks are met, then in period 1 a total of $(1-p)M$ funds will be supplied by banks to firms that received a high shock. Banks obtain these funds by selling claims to households, which they repay in period 2. Each bank receives payments $pR_L^B + (1-p)\, R_H^B$ from firms in period 2, and it must compensate its depositors by repaying the $(1-p)\, M$ funds raised in period 1. For notational convenience, I suppose that firms that did not draw on their credit lines in period 1 receive the funds M from the bank in period 2. Since R_H^B was defined as the net payments to banks, bank expected profits for each credit line are

$$\pi = p(R_L^B - M) + (1-p)\, R_H^B - M$$

As described in the beginning of this section, banks must be sufficiently compensated to cooperate. For each credit line issued, banks repay funds M and receive profits π, which must satisfy

$$\varphi\pi \geq C\left(\varphi M, K\right) \tag{2.28}$$

where φ is the measure of firms to which a particular bank issues credit lines.

2.3.5 Credit Lines and Demand Deposits

Throughout this paper, I refer to the financing arrangement offered by banks as a credit line. This terminology is consistent with Holmstrom and Tirole (1998), but the arrangement described can also be thought of as a demand deposit account. Under this interpretation, we suppose that banks offer firms an account that requires the deposit of $M+\pi$ funds in period 0. In return, the firm receives a demand account in the amount of M, that can be drawn in either period 1 or 2. The bank then uses the $\pi + M$ deposits to buy shares in firms as an initial investor. Firms that receive a high liquidity shock draw on their accounts in period 1, and firms that receive a low liquidity shock wait until period 2. In either case, banks raise the requisite M funds from households to meet these withdrawals. Then in period 2, banks receive a return of $M + \pi$ from the shares they hold, and repay M to the households from whom they raised the funds to meet the firm's withdrawals.

It is clear that the described arrangement is isomorphic to the credit lines derived above. The only difference is that in our description of credit lines banks accepted as payment a stake in the firm they were lending to, whereas in the description above banks receive funds, and then use these funds to purchase shares in firms. However, given our assumption of complete diversification in each instance, these amount to the same thing.

Note that in each case, banks are offering an insurance contract. Banks take a loss on firms that experience a high liquidity shock, which is compensated by profits on firms that experience a low shock. This is consistent with the notion of a committed credit line, which can be interpreted as insurance against the loss of credit access. The demand deposit interpretation shows that such arrangements are substitutes for liquid assets that can be sold to meet stochastic financing needs.

2.3.6 The Equilibrium Credit Line

What will be the equilibrium credit line? Suppose that rather than selling credit lines of a particular price and size (π, M), banks set a unit price ϑ, and sell any credit line with $(\vartheta M, M)$. Thus $\vartheta = \pi/M$ is equal to profits received by banks per unit of liquidity provided. Let $D = \varphi M$ be total funds intermediated by the bank, and let $\Pi = \varphi\pi = \vartheta D$ be profits from selling this quantity of credit lines.

From Proposition 4 we know that firms will never purchase a credit line at price $\pi/M = \vartheta > q - 1$. If credit lines are offered at different prices, firms will naturally prefer to purchase the one with the lowest price ϑ. However, firms will not purchase a credit line from a bank whose agency constraint is violated, since such a bank would not honor its obligations. This implies that a bank's total sale of credit lines D must satisfy $\vartheta D \geq C(D, K)$.

Let $\tilde{D}(\vartheta)$ denote the largest value of D for which $\vartheta D \geq C(D, K)$. Given our assumptions $C_{11} > 0$, $C(0, K) = 0$, and $C(D, K) \to \infty$ as $t \to \infty$, $\vartheta\tilde{D}(\vartheta) = C(\tilde{D}(\vartheta), K)$ will hold with equality, and $\tilde{D}(\cdot)$ satisfies $\tilde{D}(0) = 0$ and $\tilde{D}'(\vartheta) > 0$. Then we may write the constraint $\vartheta D \geq C(D, K)$ as $D \leq \tilde{D}(\vartheta)$.

Even if banks set a price for credit lines $\vartheta < q - 1$, firms may not purchase all the way up to $\tilde{D}(\vartheta)$. This is because there is a maximum amount of liquidity firms desire to hold. When the price of liquidity is such that all firms meet both liquidity shocks, maximum desired liquidity satisfies

$$\bar{D}(\vartheta) = \frac{A(\rho_H - \rho_0)}{\chi_1 + \vartheta(\rho_H - \rho_0)}$$

Now we can define the demand for credit lines facing a particular bank i, which we denote by $D^d(\vartheta_i)$. Suppose that the price of outside credit is q and the price of credit lines set by other banks is ϑ. Then demand for credit lines satisfies

Definition 2 (Demand for Credit Lines)**.** Demand for credit lines $D^d(\vartheta_i)$ satisfies

(i) $D^d(\vartheta_i) = 0$ if $\vartheta > q - 1$, or $\vartheta_i > \vartheta$ and $\tilde{D}(\vartheta) \geq \bar{D}(\vartheta)$.

(ii) $D^d(\vartheta_i) = \tilde{D}(\vartheta_i)$ if $\vartheta \leq q - 1$ and either $\vartheta_i < \vartheta$, or $\vartheta_i \geq \vartheta$ and $\tilde{D}(\vartheta) < \bar{D}(\vartheta)$.

(iii) $D^d(\vartheta_i) = \bar{D}(\vartheta_i)$ if $\vartheta_i \leq q - 1$ and $\vartheta_i = \vartheta$ and $\tilde{D}(\vartheta) = \bar{D}(\vartheta)$.

Clearly no firm will purchase a credit line with a cost above $q-1$. If $\vartheta_i \leq q-1$ we may distinguish three cases. First if $\vartheta_i < \vartheta$, then the bank underprices its competition and sells up to its agency costs $D^d(\vartheta_i) = \tilde{D}(\vartheta^i)$. Second, if $\vartheta_i = \vartheta$, then a bank will sell up to its agency costs if there is enough total demand to go around, but if not it will receive an equal share of firm's desired credit lines, i.e. $D^d(\vartheta_i) = \min\left(\tilde{D}(\vartheta_i), \bar{D}(\vartheta)\right)$. Finally, if a bank chooses $\vartheta_i > \vartheta$, then it will receive any residual demand after firms have purchased from all other banks. Thus if $\bar{D}(\vartheta) > \tilde{D}(\vartheta)$, then $D^d(\vartheta_i) = \tilde{D}(\vartheta)$, and if not then $D^d(\vartheta_i) = 0$.

We can write the bank pricing problem as

$$\max_{\vartheta_i} \vartheta_i D^d(\vartheta_i)$$

Proposition 5 describes the equilibrium credit line.

Proposition 5 (Equilibrium Credit Line)**.** Given price of outside liquidity q, the equilibrium credit line satisfies

(i) If $\tilde{D}(q-1) \leq \bar{D}(q-1)$, then $\vartheta = q-1$ and firms purchase $D = \tilde{D}(q-1)$ credit lines.

(ii) If $\tilde{D}(q-1) > \bar{D}(q-1)$, then ϑ satisfies $\tilde{D}(\vartheta) = \bar{D}(\vartheta)$, and firms purchase $D = \bar{D}(\vartheta)$ credit lines.

Proof. See appendix B.2. □

2.3.7 Equilibrium

We are now able to define equilibrium. Given price of outside liquidity q, Proposition 5 defines the optimal credit line (π, M), and Proposition 4 defines optimal firm behavior, including holdings of outside liquidity $\ell(q)$.

To define equilibrium, we need one further condition to determine q, for which we use the outside liquidity market clearing condition. Let ζ be the fraction of firms that meet the high shock, and suppose that each of these firms purchase credit lines (π, M) and hold liquidity ℓ. Since they have no need of liquidity, households only desire to hold outside liquidity when $q = 1$. Therefore either $q = 1$ and firms may hold any ℓ such that $\zeta\ell \leq \bar{\ell}$, or else $q > 1$ and firms must hold $\zeta\ell = \bar{\ell}$. We can express this market clearing condition as

$$(q-1)\left(\bar{\ell} - \zeta\ell\right) = 0 \tag{2.29}$$

This equilibrium is described in Proposition 6.

Proposition 6. Let $K > 0$ and $\bar{\ell} \geq 0$ be given, and let $\chi_1 = 1 - \rho_0 + p\rho_L + (1-p)\,\rho_H$ and $\chi_0 = 1 - p(\rho_0 - \rho_L)$. Let $I_0 = A/\chi_0$ and $I_1(q-1) = A/(\chi_1 + (q-1)(\rho_H - \rho_0))$. Let M_1 be defined implicitly by $(\rho_H - \rho_0) I_1(\frac{C(M_1,K)}{M_1}) = M_1 + \bar{\ell}$, and let M_2 be defined by $I_1(\frac{C(M_2,K)}{M_2}) = pI_0$. Then,

(i) If $\bar{\ell} \geq (\rho_H - \rho_0)I_1(0)$, then $\zeta = 1$, $I = I_1(0)$, $q = 1$, $M = \pi = 0$, and $\ell = (\rho_H - \rho_0)I_1(0)$.

(ii) If $\bar{\ell} < (\rho_H - \rho_0)I_1(0)$, then M_1 and M_2 are uniquely defined, and if $M_1 \leq M_2$, equilibrium satisfies $M = M_1$, $q - 1 = \frac{C(M,K)}{M}$, $\ell = \bar{\ell}$, $\zeta = 1$, $\pi = (q-1)M$, and $I = I_1(0)$.

(iii) If $\bar{\ell} < (\rho_H - \rho_0)I_1(0)$ and $M_1 > M_2$, then equilibrium satisfies $q - 1 = \frac{C(M_2,K)}{M_2}$, $I = I_1(q-1) = pI_0$, $\zeta = \frac{M_2 + \bar{\ell}}{(\rho_H - \rho_0)I}$, $\ell = \bar{\ell}/\zeta$, $M = M_2/\zeta$, and $\pi = (q-1)M$.

Proof. See appendix B.2. □

Once again, I_0 is optimal investment when $\lambda_H = 0$, and $I_1(q-1)$ is optimal investment when $\lambda_H = 1$, which is decreasing in $q - 1$.

If there is sufficient outside liquidity to meet all liquidity needs at the constrained optimal level of investment $I_1(0) = A/\chi_1$, then there is no need for banks to supply credit lines and there is no liquidity premium.

If there is insufficient outside liquidity to finance all shocks, then banks sell credit lines. Banks will require some profits to cooperate, and arbitrage between bank lines of credit and outside liquidity will set $q - 1$ equal to average bank agency costs. Thus there is a positive liquidity premium. Since households do not value

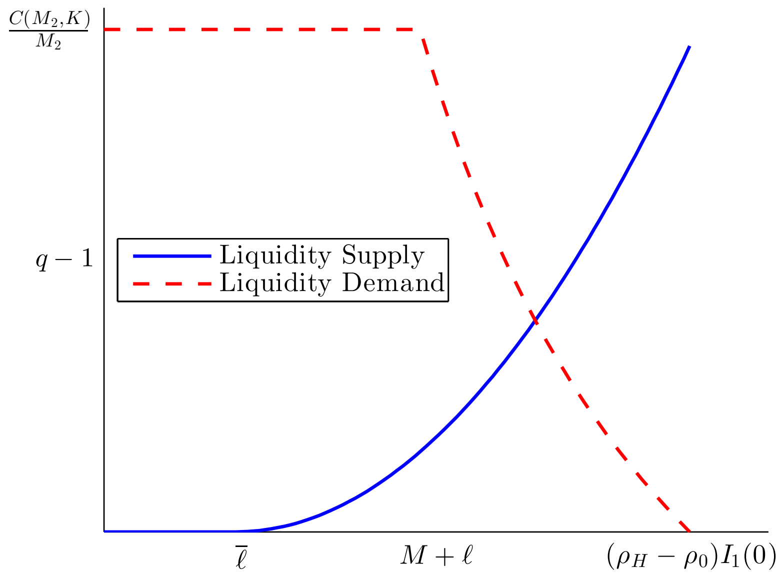

Figure 2.4: Equilibrium with banks.

outside liquidity for its liquidity properties, they are unwilling to hold it if there is a positive liquidity premium, and so all outside liquidity will be held by firms.

When there is a positive liquidity premium, it may be the case that all firms meet the high liqudity shock, or that only a fraction of firms meet it. To distinguish these cases, we compute two levels of M. M_1 corresponds to the equilibrium credit line that would prevail if all firms met the high liquidity shock. Firms will choose to meet the high liquidity shock as long as investment $I_1(q-1) \geq pI_0$. M_2 corresponds to the highest level of M for which this expression holds when all firms meet the high shock.

We can illustrate equilibrium by means of a supply and demand diagram in the market for liquidity, as depicted in Figure 2.4. Total liquidity is $M+\bar{\ell}$. Since households are willing to sell their liquidity holdings at any price $q \geq 1$, the supply of liquidity is horizontal at $q=1$ up to $\bar{\ell}$. Liquidity above this level is provided by banks, who issue credit lines that provide total liquidity $D(q)$, which is implicitly defined by $q - 1 = C(D, K)/D$.

Demand for liquidity equals $\zeta(q)(\rho_H - \rho_0)I_1(q)$. This is downward sloping with x-intercept at $(\rho_H - \rho_0)I_1(0)$ up to the value of q at which $I_1(q) = pI_0$. At this q firms are indifferent with respect to meeting the high liquidity shock, and so demand for liquidity is horizontal.

Proposition 6 shows that if there is insufficient outside liquidity $\bar{\ell}$, there will be a positive liquidity premium and lower investment in equilibrium compared to the constrained optimum. Thus in the presence of bank agency costs, the addition of banks to the model does not make outside liquidity unnecessary. This is worth emphasizing because it is the point of departure from Holmstrom and Tirole (1998). Holmstrom and Tirole (1998) do not have bank agency costs, and thus the introduction of banks allows the economy to achieve the constrained optimum. In order to allow for a non-trivial discussion of liquidity and optimal policy in the presence of banks, Holmstrom and Tirole (1998) introduce an aggregate liquidity shock. By including bank agency costs, I find that liquidity can be scarce in the presence of banks even when there is no aggregate liquidity shock.

Moreover, equilibrium investment is determined by the liquidity premium, which in equilibrium equals the average agency costs of banks. Since firms can meet liquidity shocks either using outside liquidity or credit lines, these assets must have the same price by arbitrage. Intuitively, we can think of banks as a sector of the economy that produces liquid assets, and agency costs define the production function of this sector. The precise nature of agency costs in the banking sector determines the liquidity premium, which is a component of the overall cost of the asset, along with its riskiness and return. Thus anything that affects agency costs in the banking sector will affect the capacity of the economy to commit funds, which will in turn affect investment and output.

Note that when $K = 0$, the equilibrium is as in section 2.2 without banks. Thus the specification with banks nests the specification without banks. In this case the supply of liquidity would be vertical at $\bar{\ell}$, so that $M(q) = 0$ for all q.

2.4 Comparative Statics

I now explore the properties of the equilibrium defined in Proposition 6. Given our discussion of the role of liquidity, we are interested in how changes in the supply of liquidity affect equilibrium values, notably investment I and the liquidity premium $q - 1$. I explore these questions by deriving and discussing the comparative statics of the equilibrium with respect to changes in outside liquidity $\bar{\ell}$ and bank capital K, which are the two determinants of private liquidity.

As it is the most interesting case, we restrict the following discussion to the case with $\bar{\ell} < \frac{(\rho_H - \rho_0)A}{\chi_1}$ and $M_1 > M_2$, so that liquidity is scarce in equilibrium and all firms meet the high liquidity shock. I refer to such a point as an interior equilibrium. Then all firms meet the high liquidity shock, $\zeta = 1$, and the liquidity premium satisfies $q - 1 = C(M_1, K)/M_1$. Thus investment satisfies $I(M_1) = I_1(C(M_1, K)/M_1)$.

Since we have simple expressions for $q(M)$ and $I(q)$, the key determinant of equilibrium is $M = M_1$. We know from Proposition 6 that M_1 satisfies

$$(\rho_H - \rho_0) I_1 \left(\frac{C(M_1, K)}{M_1} \right) = M_1 + \bar{\ell}$$

which allows us to implicitly define a function $M_1(K, \bar{\ell})$.

There is no analytic expression for $M_1(\cdot)$ for general agency costs, so we use the implicit function theorem to derive expressions for the marginal change in equilibrium M_1 from a change in outside liquidity $\bar{\ell}$ or bank capital K. Once we have the derivatives of M with respect to exogenous variables, we can easily compute changes in I and q using $q(M, K) = C(M, K)/K$ and $I(q)$.

2.4.1 Variation in Outside Liquidity

We first consider the effect of variations in outside liquidity $\bar{\ell}$. Following the approach described in the previous section, we compute the following changes in equilibrium variables in response to a marginal change in $\bar{\ell}$.

Proposition 7. At an interior equilibrium the effect of a marginal change in $\bar{\ell}$ on equilibrium variables is

$$\begin{aligned}
\frac{dM}{d\bar{\ell}} &= -\left(1 + \epsilon/\eta\right)^{-1} \in (-1, 0) \\
\frac{dq}{d\bar{\ell}} &= -\left(\frac{C_M - C/M}{M} \right) \left(1 + \epsilon/\eta\right)^{-1} < 0 \\
\frac{dI}{d\bar{\ell}} &= \left(\rho_H - \rho_0\right)^{-1} \left(1 + \eta/\epsilon\right)^{-1} > 1
\end{aligned}$$

where $\eta = \left(\frac{C/M}{C_M - C/M} \right) \left(\frac{M}{M + \bar{\ell}} \right)$ is the elasticity of liquidity supply, and $\epsilon = \left(M + \bar{\ell} \right) \frac{C/M}{A}$ is the elasticity of liquidity demand.

Proof. See appendix B.3. □

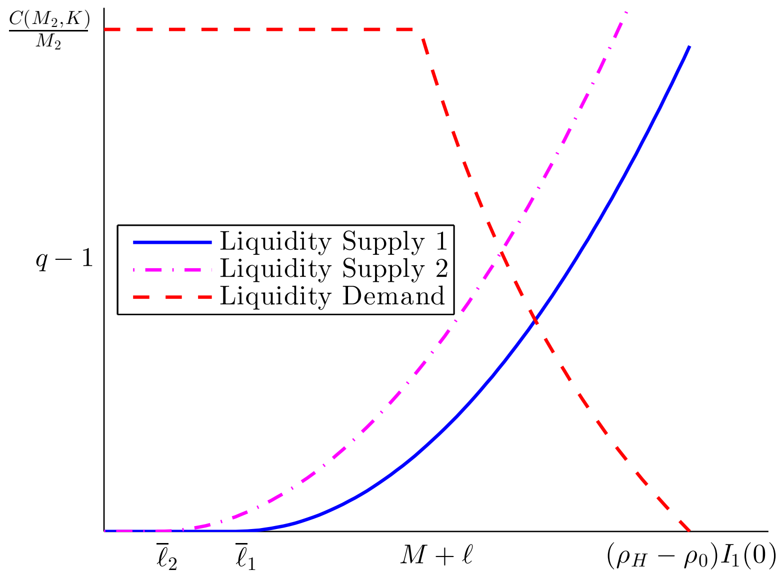

Figure 2.5: Fall in outside liquidity with banks.

The signs given follow from our assumptions that $C_M - C/M > 0$ so that $\eta > 0$. Naturally an increase in the supply of outside liquidity will reduce its equilibrium price, so that $dq/d\bar{\ell} < 0$. This is just the typical result that an increase in supply for a good lowers its price. It is equally intuitive that the lower price of liquidity leads to higher investment $dI/d\bar{\ell} > 0$, since liquidity is an input into production.

The result $dM/d\bar{\ell} \in (-1, 0)$ is also quite intuitive. M and $\bar{\ell}$ are both forms of liquidity, and so are substitutes. An increase in the supply of one reduces the demand for the other. But the reduction in credit lines must be less than one-for-one, because lower q induces higher investment from firms. Total demand for liquidity $M + \ell$ is proportional to investment, and so must rise also. This implies $dM/d\bar{\ell} > -1$.

Note that since $\partial q/\partial M = (C_M - C/M)/M$, the expression for $dq/d\bar{\ell}$ is equivalent to $\frac{\partial q}{\partial M}\frac{dM}{d\bar{\ell}}$, which is just an application of the chain rule. Also note that $dI/d\bar{\ell}$

is equal to $(\rho_H - \rho_0)^{-1}(1 + \frac{dM}{d\bar{\ell}})$, as we would expect given $(\rho_H - \rho_0)I = M + \bar{\ell}$ in equilibrium.

A fall in outside liquidity is depicted in figure 2.5. This corresponds to a leftward translation of the liquidity supply curve, which raises the liquidity premium and lowers investment.

2.4.2 Variation in Bank Capital

We next consider the effect of a change in bank capital K. We can compute changes in equilibrium variables using the same approach as above. The results are given in Proposition 8.

Proposition 8. At an interior equilibrium the effects of a marginal change in K on equilibrium variables are

$$\frac{dM}{dK} = -\left(\frac{M}{C_M - C/M}\right)(1 + \eta/\epsilon)^{-1}\frac{C_K}{M} > 0$$
$$\frac{dq}{dK} = (1 + \epsilon/\eta)^{-1}\frac{C_K}{M} < 0$$
$$\frac{dI}{dK} = (\rho_H - \rho_0)^{-1}\frac{dM}{dK} > 0$$

where $\eta = \left(\frac{C/M}{C_M - C/M}\right)\left(\frac{M}{M+\bar{\ell}}\right)$ is the elasticity of liquidity supply, and $\epsilon = \left(M + \bar{\ell}\right)\frac{C/M}{A}$ is the elasticity of liquidity demand.

Proof. See appendix B.3. $\square$

The given signs follow from the assumptions $C_M > C/M$ and $C_K < 0$. Intuitively, an increase in bank capital lowers the profits necessary to induce banks'

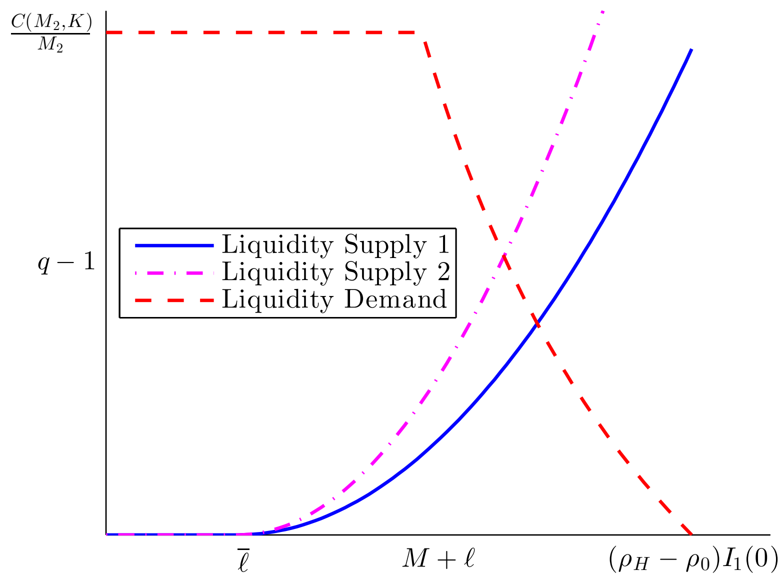

Figure 2.6: Fall in bank capital.

cooperation, which raises the effective supply of liquidity. Bank creation of credit lines M increases, the liquidity premium $q-1$ falls, and investment rises.

Comparing these expressions to those for changes in $\bar{\ell}$, we see that the expression for dq/dK is the same as the expression for $dq/d\bar{\ell}$ times the term C_K/M. The latter is the first-order change in q due to the change in bank capital, and thus corresponds to the upward shift in the liquidity supply curve. The reason these expressions are otherwise identical is because the change in quantity due to a rightward shift of a curve is the same as the change in price due to a downward shift of the same curve, as can be easily verified geometrically.

Figure 2.6 depicts a fall in bank capital. This again shifts the liquidity supply curve to the left. However, rather than being a translation in the curve, it rotates the curve about the point $\bar{\ell}$. The larger the initial M, the larger is the rotation in the supply curve, since larger bank financing implies a larger effect from a fall in bank capital.

2.4.3 Ratio of Elasticities

The ratio of the elasticities of liquidity supply and demand η/ϵ is a key term in the expressions derived above. In each case, changes in investment are proportional to $(1+\eta/\epsilon)^{-1}$, although in the case of changes in K there is the additional term $(\partial q/\partial M)^{-1}(\partial q/\partial K)$. Since η/ϵ appears in the denominator, higher η/ϵ implies that investment is less responsive to changes in liquidity.

Intuitively, if liquidity supply is very inelastic (η small), a change in the economy's stock of liquidity will result in a large price response. This will produce a large change in equilibrium liquidity and investment, unless liquidity demand is also very inelastic (ϵ small). Conversely, if liquidity supply is very elastic, a change in the stock of liquidity will produce a relatively small price response, because other sources of liquidity adjust. This results in relatively small movements in equilibrium liquidity and investment, unless liquidity demand is very inelastic.

Using the expressions for η and ϵ derived above, we can write η/ϵ as

$$\eta/\epsilon = \eta_B \left(\frac{M}{M+\bar{\ell}}\right) \cdot \frac{A}{\frac{C}{M}(M+\bar{\ell})} \tag{2.30}$$

where $\eta_B = \partial M/\partial q \cdot (q-1)/M$ is the supply elasticity of *bank liquidity*. If we assume that η_B is close to constant, i.e. an isoelastic supply of bank liquidity, then the relative elasticity is driven primarily by the ratio $M/(M+\bar{\ell})$. Intuitively, the supply of liquidity is composed of two parts — elastically supplied bank liquidity, and inelastically supplied outside liquidity. The higher is the fraction of bank liquidity,

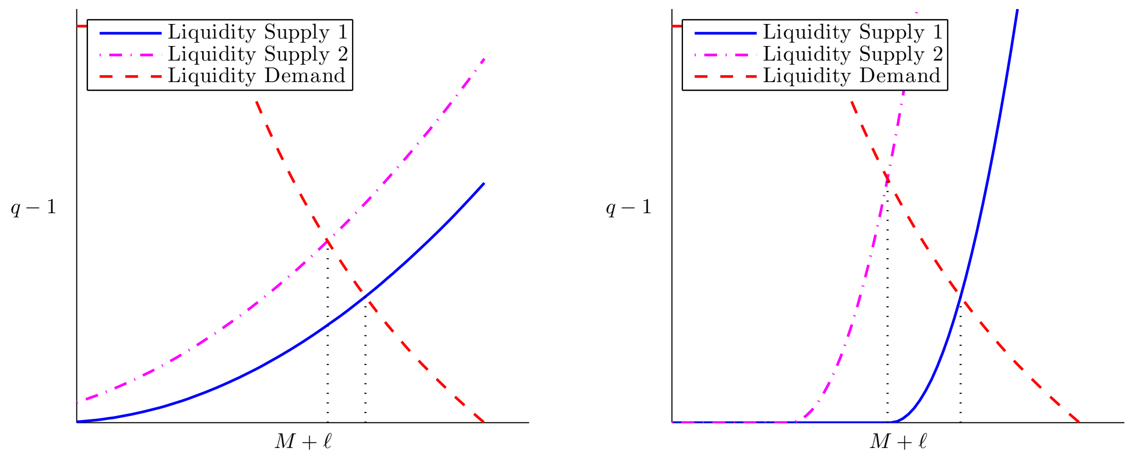

Figure 2.7: Fall in outside liquidity with high (left) and low (right) elasticity ratios.

the higher is the overall elasticity of liquidity supply, and therefore the smaller is the sensitivity of investment due to changes in the economy's stock of liquidity.

This difference is depicted in Figure 2.7. The left panel depicts a case with low $\bar{\ell}$ and high K, so that a relatively high share of equilibrium liquidity is provided by banks. The right panel depicts a case with high $\bar{\ell}$ and low K, so that little liquidity is provided by banks. In each case, initial equilibrium liquidity and investment are the same. Both panels depict an identical absolute fall in $\bar{\ell}$, which produces a leftward translation of each liquidity supply curve by the same amount. This results in a signficantly larger fall in investment in the right panel, reflecting the lower elasticity ratio η/ϵ in an economy with lower bank liquidity.

2.5 Welfare

We now turn to the welfare properties of the equilibrium described in section 2.3. Clearly equilibrium will fail to attain the unconstrained optimal level of investment described in Proposition 1, and so the equilibrium is not Pareto optimal. A more interesting question is whether the equilibrium corresponds to the choice of a *constrained planner*, i.e. a planner with access to a limited set of policy instruments.

Since this notion of optimality is contingent on the instruments available to the constrained planner, the analysis of welfare and policy are closely related. I proceed by first defining a notion of welfare, and comparing it to the equilibrium allocation. I then consider what instruments would enable a planner to implement this allocation. I show that when liquidity is scarce, the equilibrium allocation is generally suboptimal relative to the choices of a constrained planner.

Since utility is linear with no discounting, total welfare is the sum of all agents' consumption. Normalizing zero-investment welfare to 0, total welfare is

$$W = \int_i \left[p(\rho_1 - \rho_L) + \lambda_H^i (1-p)(\rho_1 - \rho_H) - 1\right] I_i$$

where (I_i, λ_H^i) is the production plan of firm i.

In the following analysis, I assume that the constrained planner maximizes total welfare W. The result may not be a Pareto-improvement relative to equilibrium as the distribution of consumption across agents may change. However, if the planner has access to lump-sum taxes and transfers at the end of period 2, any policy that increases W can be made Pareto improving through suitable uncontingent transfers between agents. Such transfers will not affect agents' behavior given linear utility.

2.5.1 Unconstrained Optimum

First consider the unconstrained optimal allocation described in Proposition 1. This allocation involves investing all available period 0 resources $I = A + H_0$, and meeting the high liquidity shock $\lambda_H = 1$. Since the maximum level of investment in equilibrium with $\lambda_H = 1$ is $I = A/\chi_1 < A + H_0$, the equilibrium is not Pareto optimal.

A planner that has access to lump-sum taxes and transfers in all periods can implement the unconstrained optimum by transferring sufficient funds from households to firms in periods 0 and 1, as described in Proposition 9.

Proposition 9 (Planner's Problem with Unlimited Transfers)**.** A planner that has access to unlimited lump-sum taxes and transfers can achieve the allocation described in Proposition 1 by transferring H_0 from households to firms in period 0, and transferring $(\rho_H - \rho_0)(A + H_0)$ from households to firms that experience a high liquidity shock in period 1.

Proof. See appendix B.4. □

Since the only constraint on investing all available period 0 funds is limited access to financing in periods 0 and 1 arising from limited pledgeability, a planner can achieve the optimal production plan by simply transferring the necessary funds to firms. Intuitively, if a planner can identify investment opportunities that cannot be financed by the private sector via credit markets, and has access to nondistortionary tax instruments, then the planner can improve welfare by transferring funds directly to those agents with investment opportunities. Thus the planner wholly supplants private credit markets. However, such policies are rarely observed in practice since it is rare that governments hold the necessary informational advantage over the private sector.

2.5.2 Liquidity-Constrained Optimum

In Sections 2.2 and 2.3 I defined the constrained optimum as the equilibrium when liquidity is abundant, i.e. when $\bar{\ell}$ is sufficiently large that $q = 1$. For the discussion of welfare, another notion of constrained optimality is useful: the optimal allocation given fixed liquidity supply $\bar{\ell}$ and K, and subject to agents' participation and incentive compatability constraints. To distinguish this from the previously defined notion of constrained optimum, I refer to this as the *liquidity-constrained optimum.*

Here I define an allocation as $\{I_i, \ell_i, \pi_i, M_i\}_i$ for firms $i \in [0, 1]$, together with price of liquidity q. The participation constraint for suppliers of outside liquidity is $q \geq 1$. Investors receive compensation out of the pledgeable funds of firms, so investment satisfies

$$I_i \leq \frac{A - \pi_i - (q-1)\ell_i}{1 - p(\rho_0 - \rho_L) - \lambda_H^i (1-p)(\rho_0 - \rho_H)} \tag{2.31}$$

Each firm that meets the high shock must have sufficient liquidity to meet the shock

$$\lambda_H^i I_i (\rho_H - \rho_0) \leq M_i + \ell_i \tag{2.32}$$

Banks must receive sufficient compensation for providing liquidity

$$\int_i \pi_i \geq C\left(\int_i M_i, K\right) \tag{2.33}$$

Finally, firms will not pay for liquidity if they are not using it, so $\pi_i = 0$ if $\lambda_H^i = 0$, and outside liquidity use is limited by the available supply $\int_i \ell_i \leq \bar{\ell}$.

Definition 3. The *liquidity constrained optimum* (LCO) is the solution to

$$\max_{I_i, \lambda_H^i, M_i, \ell_i, \pi_i, q} \int_i \left[p(\rho_1 - \rho_L) + \lambda_H^i (1-p)(\rho_1 - \rho_H) - 1\right] I_i$$

$$\text{s.t.} \quad (2.31),\ (2.32),\ (2.33),\ \textstyle\int_i \ell_i \leq \bar{\ell},\ q \geq 1$$

Proposition 10 (Liquidity Constrained Optimum)**.** Let $\chi_1 = 1 - \rho_0 + p\rho_L + (1-p)\rho_H$ and $\chi_0 = 1 - p(\rho_0 - \rho_L)$. Let ζ designate the fraction of firms that meet the high liquidity shock. Let $R_1 = \rho_1 - p\rho_L - (1-p)\rho_H - 1$ and $R_0 = p(\rho_1 - \rho_L) - 1$. Let $I(\zeta)$ be defined by $\chi_1 I = A - C\left(\max(\zeta(\rho_H - \rho_0)I - \bar{\ell}, 0), K\right)/\zeta$, and let $\hat{\ell}$ be defined by $R_1\chi_0 = R_0\left[\chi_1 + (\rho_H - \rho_0)C_1\left((\rho_H - \rho_0)I(1) - \hat{\ell}, K\right)\right]$. Then the liquidity constrained optimum is defined as follows:

(i) If $\bar{\ell} \geq (\rho_H - \rho_0)\frac{A}{\chi_1}$, then all firms choose $\lambda_H = 1$, $I = \frac{A}{\chi_1}$, $\ell = (\rho_H - \rho_0)\frac{A}{\chi_1}$, and $M_i = \pi_i = 0$.

(ii) If $\hat{\ell} \leq \bar{\ell} < (\rho_H - \rho_0)\frac{A}{\chi_1}$, then all firms choose $\lambda_H = 1$, $I = I(1)$, $\ell = \bar{\ell}$, $M = (\rho_H - \rho_0)I - \bar{\ell}$, and $\pi = C(M, K)$.

(iii) If $\bar{\ell} < \hat{\ell}$, then a fraction ζ of firms choose $\lambda_H = 1$, where ζ is implicitly defined by

$$R_1\chi_0 = R_0\left[\chi_1 + (\rho_H - \rho_0)C_1\left(\zeta(\rho_H - \rho_0)I(\zeta) - \bar{\ell}, K\right)\right]$$

These firms choose $I = I(\zeta)$, $\ell = \bar{\ell}/\zeta$, $M = (\rho_H - \rho_0)I(\zeta) - \ell$, and $\pi = C(\zeta M, K)/\zeta$. The remaining fraction $1 - \zeta$ choose $\lambda_H = 0$, $I = A/\chi_0$, $\ell = M = \pi = 0$.

$q = 1$ in all cases.

Proof. See appendix B.4. $\square$

The differences between the equilibrium and the liquidity constrained optimum arise because firms do not consider the effects of their demand for liquidity on the prices of liquidity q and π, whereas the planner internalizes these pecuniary effects. Since payments for liquidity reduce pledgeable funds, a higher price of liquidity reduces possible investment.

When liquidity is scarce in equilibrium, firms bid up the price of outside liquidity to some $q > 1$. Since higher q transfers period 0 funds from firms to households who hold outside liquidity, it inefficiently decreases investment. Since the supply of outside liquidity is inelastic, there is no gain from increasing its price. Thus the constrained planner sets the price of ℓ as low as possible, resulting in higher investment.

We can see this by comparing the expressions for I in the cases of equilibrium and the liquidity constrained optimum. For simplicity consider the case that $\zeta = 1$, so that all firms meet the liquidity shocks. Then equilibrium investment satisfies

$$I = \frac{A - C(M,K) - \frac{C(M,K)}{M} \cdot \bar{\ell}}{\chi_1}$$

where $M = I/(\rho_H - \rho_0) - \bar{\ell}$. By comparison, the liquidity-constrained level of investment satisfies

$$I = \frac{A - C\left(M,K\right)}{\chi_1}$$

where again $M = I/(\rho_H - \rho_0) - \bar{\ell}$.

The difference is that in equilibrium, pledgeable funds are reduced by a further amount $C/M \cdot \bar{\ell}$. This occurs because arbitrage pushes up the price of outside liquidity to equal the cost of credit lines provided by banks. Since outside liquidity is in fixed supply, this increase in price does not raise the supply of outside liquidity. Instead, owners of outside liquidity enjoy an excess return that comes at a steep social cost, since it reduces pledgeable funds and therefore investment. A planner would prefer to push down the price of liquid assets in order to increase investment, which implies a wedge between the price of outside liquidity and the premium paid to financial intermediaries. The former fulfills no social purpose, while the latter is necessary to overcome the agency costs of banks.

Implementation. Given the analysis above, a constrained planner should try to lower the price of outside liquidity in order to increase welfare. The simplest way to do this is to put a price ceiling on the price of outside liquidity at $q = 1$. Then firms would first buy outside liquidity until it was sold out, and then substitute to the (more expensive) bank credit lines.

2.6 Public Liquidity Provision

The previous sections assumed that the supply of liquid assets in the economy is fixed. However, many government policies directly affect this supply. When central banks engage in conventional monetary policy, they do so by buying and selling assets of varying liquidity. Likewise, when governments issue new bonds, they increase the economy's store of liquid assets. We can capture these activities in our model as the issuance of government bonds that are perfect substitutes for outside liquidity.

Throughout this section, we will focus on the interior equilibrium where both liquidity shocks are met and the liquidity premium is strictly positive.

2.6.1 Public Bond Issuance

Suppose that the government issues x bonds with a face value of 1 in period 0. The funds from the sale of these bonds are returned to households via a lump sum transfer and consumed immediately. In period 2, the government levies a tax on the households in order to raise funds that it uses to repay the bonds.

Since government bonds are perfect substitutes for outside liquid assets $\bar{\ell}$, government bonds will sell at the same equilibrium price q. From the perspective of firms and households, it is as though the stock of outside liquidity had increased from $\bar{\ell}$ to $\bar{\ell}+x$, and so equilibrium will be exactly the same as given in Proposition 6, except that $\bar{\ell}$ is replaced by $\bar{\ell}+x$. Likewise the comparative statics of real variables with respect to changes in government debt x will be as given in Proposition 7, with $\bar{\ell}$ replaced by $\bar{\ell}+x$, and $d\bar{\ell}$ replaced by dx.

By assumption, the government has perfect credibility, and therefore is able to commit to repaying its debt. Since the government can raise taxes, it can promise a sufficiently large quantity of pledgeable funds in period 1 to cover any potential liquidity shock. By Proposition 6, there is a level of outside liquidity that achieves the constrained optimum, and so the government can issue bonds in order to achieve the constrained optimum level of investment. If there were no costs to taxation, then this would be the optimal policy, and the government could achieve the constrained optimum. However, I assume that the government can only raise funds using a distortionary tax, although it can disburse funds to households in a lump-sum transfer. Suppose that the deadweight loss from raising x funds is given by the function $D(x)$, which I assume is increasing and convex in x, and satisfies $D'(0) = 0$ and $D'(x) \to \infty$ as $x \to \infty$.[18]

2.6.2 Results with General Agency Costs

The total social surplus if both shocks are met is

$$RI - D(x)$$

where $R = \rho_1 - p\rho_L - (1-p)\rho_H - 1$, since all other aspects of the allocation represent transfers.[19] Then at an interior solution the necessary condition for optimality is

$$R\frac{dI}{dx} = D'(x) \tag{2.34}$$

which will imply a unique optimum if $\frac{d^2I}{dx^2} < 0$ for all x. However, given the expression for $\frac{dI}{d\ell}$ derived previously in Proposition 7, the sign of $\frac{d^2I}{dx^2}$ will depend on how the ratio of elasticities η/ϵ vary as outside liquidity is changed. Thus the sign will be ambiguous, and without further assumptions we cannot say whether $\frac{dI}{dx}$ is increasing

[18] I assume here that the government must raise funds by taxing activity in some unmodeled market.

[19] This is with the no-production case normalized to 0. Additional liquidity has no value unless both shocks are met, so we will restrict attention to this case.

or decreasing in x, and we cannot give an expression that uniquely defines optimal policy. We can, however, say a few things about optimal policy in general.

Proposition 11. Suppose that in the absence of government policy the economy is initially at an interior equilibrium. Then

(i) The optimal supply of government liquid assets x^* will be positive.

(ii) If optimal x^* implies that the economy is still at an interior equilibrium, and if the current choice of x^* is unique, then a marginal change in $\bar{\ell}$ or K will shift the optimal point according to

$$\frac{dx^*}{d\bar{\ell}} = -\frac{R\frac{\partial^2 I}{\partial x^2}}{R\frac{\partial^2 I}{\partial x^2} - D''(x)}$$
$$\frac{dx^*}{dK} = -\frac{R\frac{\partial^2 I}{\partial x \partial K}}{R\frac{\partial^2 I}{\partial x^2} - D''(x)}$$

(iii) At an interior optimal point x^*, we have $R\frac{\partial^2 I}{\partial x^2} < D''(x)$.

(iv) We have $\frac{dx}{d\bar{\ell}} > 0 \iff \frac{\partial^2 I}{\partial x^2} > 0$ and $\frac{dx}{dK} > 0 \iff \frac{\partial^2 I}{\partial x \partial K} > 0$.

Proof. See appendix B.5. □

The first result in Proposition 11 is that if the economy is at an interior equilibrium in the absence of any policy, it will always be optimal to choose $x^* > 0$. Thus when liquidity is scarce it will in general be optimal for the government to issue bonds solely for their value in providing liquidity, even though taxation is distortionary. The intuition for this result is that at an interior solution with a positive liquidity premium there is a positive marginal value of issuing bonds because they serve as liquid assets. Since the marginal cost of taxation is zero at $x = 0$, the marginal value of issuing bonds at $x = 0$ is greater than the marginal cost, and therefore it will be always be optimal to issue some positive quantity of bonds.[20]

This intuition is depicted in Figure 2.8. The line that passes through the origin is the marginal cost of raising x in funds via taxation, and the other line is the marginal value of increasing the supply of liquid assets $R\frac{dI}{dx}$. Moreover, there is a point $\bar{x} = \frac{A(\rho_H - \rho_0)}{\chi_1} - \bar{\ell}$, such that if the government chooses $x \geq \bar{x}$ the economy achieves the constrained optimum, and the marginal value of increasing x above $\bar{x}$ drops to zero. Therefore there will be some point at which the two curves intersect, and this will be at some $x > 0$.

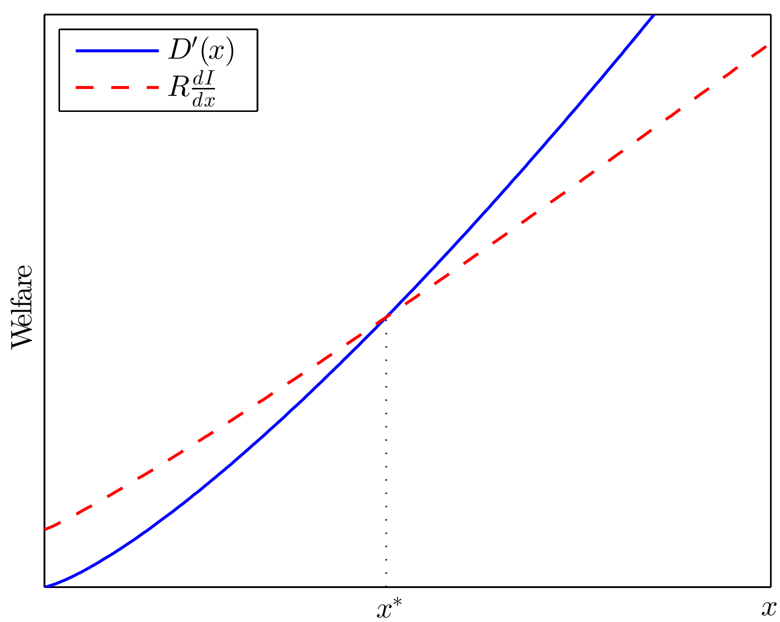

Figure 2.8: Optimal policy.

Results (ii) - (iv) in Proposition 11 have a similarly straightforward intuition.

[20]This assumes that the cost of the first dollar raised is zero. If $D'(0) > 0$, as will be the case if there are distortionary taxes already in place, optimal public debt issuance may be zero.

Suppose that there is an interior unique solution at the point x^*. Then at x^* the marginal cost and marginal benefit curves intersect, as depicted in Figure 2.8. Since x^* is a local maximum, the marginal benefit curve crosses the marginal cost curve from above, as claimed in statement (iii). A change in $\bar{\ell}$ or K will not shift the marginal cost curve. Since the equation for the marginal benefit curve is $R\frac{dI}{dx}$, an increase in $\bar{\ell}$ will shift the marginal benefit curve upwards if $\frac{\partial^2 I}{\partial x^2} > 0$, which will increase the point of intersection x^*. Similarly if $\frac{\partial^2 I}{\partial x \partial K} > 0$, an increase in K will shift the marginal benefit curve upwards, which will again increase the point of intersection x^*, and the reverse. Such a shift is depicted in Figure 2.9.

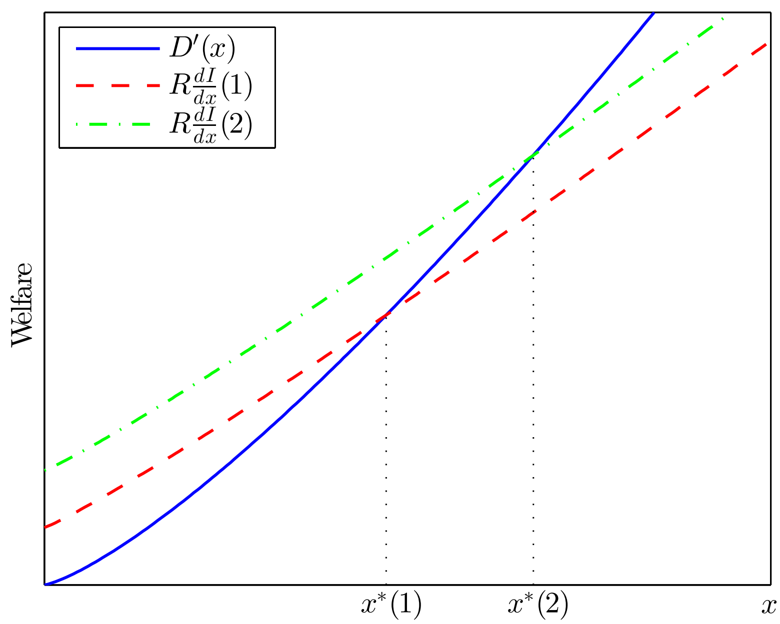

Figure 2.9: Optimal policy after increase in outside liquidity.

2.6.3 Optimal Policy and the Elasticity of Bank Liquidity Supply

The previous section made no assumptions about the bank agency cost functions beyond what we have already assumed, and so was not able to say very much about either the level of optimal bond supply x^*, or how x^* varies with private outside liquidity or bank capital. However, under a few reasonable assumptions about the agency cost function, we can say a lot more.

For this section I change the terminology slightly for ease of exposition. We continue to only consider the case that firms meet both liquidity shocks. Let the total liquidity used be L, so that $L = M + x + \bar{\ell}$; let the liquidity premium be $\vartheta = q - 1$; let total outside liquidity, including both government bonds and private outside liquidity, be $z = x + \bar{\ell}$; let $\vartheta(M, K) = C(M, K)/M$ be the *average* agency cost function; finally, let $\phi = (\rho_H - \rho_0)^{-1}$ be the liquidity multiplier, so that $I = \phi L$.

Using this terminology, the elasticity of liquidity demand is

$$\epsilon = \frac{\vartheta}{\vartheta + \phi\chi} = \frac{L\vartheta}{A}$$

and the elasticity of liquidity supply is

$$\eta = \frac{\vartheta}{\vartheta_M L} = \eta_M \left(\frac{L - z}{L}\right)$$

where $\eta_M = \frac{\vartheta}{\vartheta_M M}$ is the elasticity of the supply of bank credit lines. This decomposition reflects that the economy's liquidity is the sum of two components: outside liquidity that is supplied inelastically, and bank liquidity that is supplied with elasticity η_M. Thus total elasticity of liquidity supply is a function of η_M, and the fraction of total liquidity that is supplied by banks.

Using these expressions, we can prove the following about the optimal supply of government bonds x^*:

Proposition 12. Suppose that under optimal government bond issuance x^* the economy is at an interior equilibrium. Then

(i) x^* is decreasing in the ratio of the elasticity of liquidity supply to the elasticity of liquidity demand η/ϵ.

(ii) x^* is *increasing* in private outside liquidity $\bar{\ell}$ iff

$$\frac{\frac{d}{dz}\left(\eta_M\right)}{\eta_M} + \frac{1 - 2\epsilon - \eta_M}{L\left(\epsilon + \eta\right)} < 0 \tag{2.35}$$

(iii) A sufficient condition for (ii) to hold is that the supply of credit lines by bank be isoelastic with elasticity $\eta_M \geq 1$.

(iv) x^* is *decreasing* in bank capital K iff

$$\frac{L_{zz}}{L_z} > \frac{d}{dM}\left(\frac{\vartheta_M}{\vartheta_K}\right) \cdot \frac{\vartheta_M}{\vartheta_K} M_z \tag{2.36}$$

Proof. See appendix B.5. □

These results suggest that for a reasonable parameterization of the bank agency cost, the optimal supply of government bonds is *increasing* in private outside liquidity $\bar{\ell}$, and *decreasing* in bank equity K. In other words, the government should provide additional public liquidity when private liquidity from banks decreases, but should provide *less* liquidity when private outside liquidity decreases. I discuss the intuition for each of these results below. The key mechanism is the response of bank liquidity supply to government liquidity provision, since government liquidity crowds out (elastically supplied) bank liquidity, but does not crowd out (inelastically supplied) private outside liquidity.

Result (i) is immediate from the expression for L_x:

$$L_x = 1 + M_x = \frac{\epsilon/\eta}{1 + \epsilon/\eta}$$

together with the optimality expression for government debt (2.34). Intuitively, the ratio η/ϵ tells us to what extent an increase in outside liquidity will be offset by a reduction in bank liquidity. η tells us how much bank liquidity will fall for a given fall in the price of liquidity ϑ, and $1/\epsilon$ tells us how much the price of liquidity will fall to absorb a given increase in liquidity. If η is large, issuance of government liquidity will cause a large reduction in bank liquidity, and so the benefits of public liquidity provision are attenuated. Likewise, if ϵ is small, demand for liquidity is very inelastic, which suggests a large price swing from public liquidity provision, which will prompt a greater crowding out of bank liquidity. Figure 2.10 depicts the optimal supply of government liquidity x^* for high and low η/ϵ.

Result (ii) and (iii) give conditions under which x^* decreases given a decrease in private outside liquidity $\bar{\ell}$. This will hold as long as bank liquidity is reasonably elastic, and its elasticity is not falling too rapidly in the price of liquidity. In particular, this will hold if bank liquidity is isoelastic, with an elasticity of at least 1. This result might seem somewhat surprising — one might intuitively expect that the government should provide outside liquidity when it is scarce. However, the normal intuition rests on the idea of decreasing marginal returns on investment, with liquidity a scarce input. That does not hold here, since we assumed R was

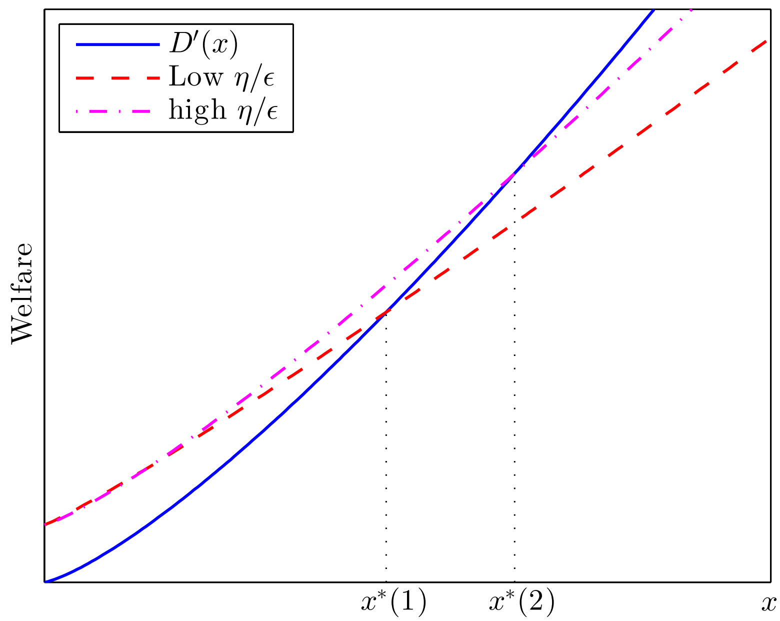

Figure 2.10: Optimal policy for high and low elasticity ratios.

fixed. Instead, an increase in private outside liquidity $\bar{\ell}$ makes the overall supply of liquidity more inelastic, which decreases crowding out from government supply of liquidity. This is what drives the sign.

Result (iv) gives a condition under which x^* *increases* given a decrease in bank capital K. This is close to what we would expect intuitively, but in this case it is again driven by changes in the elasticity of liquidity supply. The condition that establishes this result requires that the cross elasticity of the agency cost function not be too great. If this failed, the increase in K might end up *decreasing* the elasticity of liquidity supply: although it would increase the share of (elastic) bank liquidity in the total, it would simultaneously reduce the elasticity of bank liquidity supply enough that the net result would be to decrease the elasticity of liquidity supply overall.

It is worth emphasizing that these results do not assume any sort of aggregate liquidity shock as in Holmstrom and Tirole (1998). Here the provision of public liquidity acts as a substitute to bank financing, and therefore decreases the cost of liquidity in the economy, which is positive because of bank agency costs. Thus while Holmstrom and Tirole (1998) discusses countercyclical policy, the cycle in question is variation in an aggregate liquidity shock. By contrast, my discussion applies to variations in bank assets and the stock of liquid assets in the economy. My model exhibits cyclical variation in the economy's capacity to meet liquidity shocks, rather than cyclical variation in liquidity shocks themselves.

The benefit of my approach is that variations in bank net worth can be observed and are procyclical in nature. Thus my model provides both a mechanism for the propagation of shocks and a justification for countercyclical public liquidity provision.

2.7 Conclusion

I analyze the supply and demand for liquidity in the economy using a simple model of investment and liquidity provision. I find that if there is insufficient liquidity, equilibrium investment will be below the constrained optimum. This effect can be reduced by introducing intermediaries that pool liquidity and insure against shocks. When banks are subject to agency costs that are convex in funds intermediated and decreasing in bank assets, firms will face a tradeoff between financing their liquidity shocks by holding scarce liquidity, and financing by holding credit lines issued by banks. At an interior equilibrium, the liquidity premium will equal the average agency costs of banks, liquid assets will sell at a premium over their expected return because of their scarcity, and investment will be lower than when liquidity is abundant.

The government can provide liquidity by issuing public liabilities such as bonds. If public liabilities were costless to issue, the government could achieve the constrained optimum. However, if funds can only be raised through distortionary taxation, the government may choose an optimal supply of liquidity that does not fully eliminate the liquidity premium. Nevertheless, if the economy is initially at an interior equilibrium with scarce liquidity, the government will find it optimal to issue a strictly positive quantity of bonds. Moreover, the optimal issuance of government debt is higher when the elasticity of private liquidity supply is low, because then public liquidity crowds out less private liquidity, or when the elasticity of liquidity demand is high.

I consider two comparative static exercises: a fall in bank assets, and a fall in outside liquidity. We can think of these as two channels by which shocks propagate through the financial system via the supply of liquidity. A fall in bank assets could represent losses to the economy's productive factors, or simply accounting losses due to a credit bust. A fall in outside liquidity could represent asset destruction (such as a credit freeze) or a sharp drop in liquid assets due to capital flight or a decision by the government to decrease its supply of liquidity.

I find that a fall in bank assets will worsen agency costs and reduce the supply of liquidity from financial intermediaries. This raises the equilibrium liquidity premium, making investment more costly, and thus reduces equilibrium investment. The absolute reduction in investment will depend only on the sensitivity of average agency costs to bank equity, but the percentage effects will depend on the share of bank liquidity in the economy's total supply of liquidity. The optimal policy response is to increase provision of public liquidity, because the decrease in bank liquidity provision reduces the elasticity of private liquidity supply, and thus makes public liquidity issuance more effective at the margin, since it crowds out less private liquidity.

I find that a fall in private outside liquidity will decrease investment and raise the liquidity premium. The optimal policy response is to *decrease* provision of public liquidity, because the reduction in private outside liquidity shifts the economy's total stock of liquidity away from inelastically supplied outside liquidity and towards elastically supplied bank liquidity. This implies higher overall elasticity of liquidity supply, and therefore greater crowding out from public liquidity issuance.

Chapter 3: The Redistributive Effects of Financial Deregulation[1]

3.1 Introduction

Financial regulation is often framed as a question of economic efficiency. However, the intense political debate on the topic suggests that redistributive questions are front and center in setting financial regulation. In the aftermath of the financial crisis of 2008/09, for example, consumer organizations, labor unions and political parties championing worker interests have strongly advocated a tightening of financial regulation, whereas financial institutions and their representatives have argued the opposite case and have issued dire warnings of the dangers and costs of tighter regulation.

This paper makes the case that there is a distributive conflict over the level of risk-taking in the financial sector, and by extension over the tightness of financial regulation. Financial institutions prefer more risk-taking than what is optimal for the rest of society because risk-taking delivers higher expected returns. However, it also comes with a greater incidence of large losses that lead to credit crunches and negative externalities on the real economy. This link between financial regulation

[1] The work appearing in this chapter was coauthored with Professor Anton Korinek, and was published as Korinek and Kreamer (2014)

and volatility in the real economy has been documented e.g. by Reinhart and Rogoff (2009).

We develop a formal model to analyze the distributive conflict inherent in regulating risk-taking in the financial sector. The financial sector plays a special role in the economy as the only sector that can engage in financial intermediation and channel capital into productive investments. This assumption applies to the financial sector in a broad sense, including broker-dealers, the shadow financial system and all other actors that engage in financial intermediation. For simplicity, we will refer to all actors in the financial sector broadly defined as "bankers."

There are two types of financial imperfections. First, bankers suffer from a commitment problem and need to have sufficient capital in order to engage in financial intermediation. This captures the standard notion that bankers need to have "skin in the game" to ensure proper incentives. Secondly, insurance markets between bankers and the rest of society are incomplete, and bank equity is concentrated in the hands of bankers.

Because of the "skin in the game"-constraint, a well-capitalized financial sector is essential for the rest of the economy. In particular, the financial sector needs to hold a certain minimum level of capital to intermediate the first-best level of credit and achieve the optimal level of output. If aggregate bank capital declines below this threshold, binding financial constraints force bankers to cut back on credit to the rest of the economy. The resulting credit crunch causes output to contract, wages to decline and lending spreads to increase. At a technical level, these price movements constitute pecuniary externalities that hurt the real economy but benefit bankers.

When financial institutions decide how much risk to take on, they trade off the benefits of risk-taking in terms of higher expected return with the risk of becoming constrained. They always find it optimal to choose a positive level of risk-taking. By contrast, workers are averse to fluctuations in bank capital. They prefer less financial risk-taking and a stable supply of credit to the real economy. This generates a Pareto frontier along which higher levels of risk-taking correspond to higher levels of welfare for bankers and lower levels of welfare for workers. Financial regulation imposes constraints on risk-taking, which move the economy along this Pareto frontier. Financial regulators have to trade off greater efficiency in the financial sector, which relies on risk-taking, against greater efficiency in the real economy, which requires a stable supply of credit.[2]

The distributive conflict over risk-taking and regulation is the result of both financial imperfections in our model. If bankers weren't financially constrained, then Fisherian separation would hold: they could always intermediate the optimal amount of capital, and their risk-taking would not affect the real economy. Similarly,

[2]Our findings are consistent with the experience of a large number of countries in recent decades: deregulation allowed for record profits in the financial sector, which benefited largely the financial elite (see e.g. Philippon and Reshef (2013)). Simultaneously, most countries also experienced a decline in their labor share (Karabarbounis and Neiman (2014)). When crisis struck, e.g. during the financial crisis of 2008/09, economies experienced a sharp decline in financial intermediation and real capital investment, with substantial negative externalities on workers and the rest of the economy. Such occasionally binding financial constraints are generally viewed as the main driving force behind financial crises in the quantitative macro literature (see e.g. Korinek and Mendoza (2014)).

if risk markets were complete, then bankers and the rest of the economy would share not only the downside but also the benefits of financial risk-taking. In both cases, the distributive conflict would disappear.

Drawing an analogy to more traditional forms of externalities, financial deregulation is similar to relaxing safety rules on nuclear power plants: such a relaxation will reduce costs, which increases the profits of the nuclear industry and may even benefit the rest of society via reduced electricity rates in good states of nature. However, it comes at a heightened risk of nuclear meltdowns that impose massive negative externalities on the rest of society. In expectation, relaxing safety rules below their optimum level increases the profits of the nuclear sector at the expense of the rest of society.

We analyze a number of extensions to study how risk-taking in the financial sector interacts with the distribution of resources in our model economy. When bank managers receive asymmetric compensation packages — for instance, that grant high bonuses following good performance, but do not reduce earnings an equivalent amount following losses — managers will take on greater risk and expose the economy to larger negative externalities. If bankers have market power, their precautionary incentives are reduced and they take on more risk which hurts workers, highlighting a new dimension of welfare losses from concentrated banking systems. Financial innovation that expands the set of available assets allows the financial sector to take on more risk, and in some cases can make workers unambiguously worse off. Finally, greater risk-taking induced by bailouts likely leads to a significantly larger redistribution of surplus than the explicit transfers that financial institutions receive during bailouts. These extensions suggest that the externalities from credit crunches may easily represent the most significant social cost of distortions in the financial sector.

Our analytic findings suggest a number of policy interventions in the real world that regulators could implement if their main concern is a stable supply of credit to the real economy: they could (i) separate risky activities, such as proprietary trading, from traditional financial intermediation, (ii) impose higher capital requirements on risky activities, in particular on those that do not directly contribute to lending to the real economy, (iii) limit payouts if they endanger a sufficient level of capitalization in the financial sector, (iv) use structural policies that reduce incentives for risk-taking, and (v) force recapitalizations when necessary, even if they impose private costs on bankers.

A Pareto-improvement could only be achieved if deregulation was coupled with measures that increase risk-sharing between bankers and the rest of the economy so that the upside of risk-taking also benefits workers. Even if formal risk-sharing markets are absent, redistributive policies such as higher taxes on financial sector profits that are used to strengthen the social safety net for the rest of the economy would constitute such a mechanism.

Literature. This paper is related to a growing literature on the effects of financial imperfections in macroeconomics (see e.g. Gertler, Kiyotaki et al. (2010) for an overview). Most of this literature describes how binding financial constraints may amplify and propagate shocks (see e.g. Bernanke and Gertler (1989); Kiyotaki and

Moore (1997)) and lead to significant macroeconomic fluctuations that affect output, employment and interest rates (see e.g. Gertler and Karadi (2011)). The main contribution of our paper is to focus on the redistributive effects of such fluctuations.

Our paper is also related to a growing literature on financial regulation (see e.g. Freixas and Rochet (2008) for a comprehensive review), but puts the distributive implications of such financial policies center stage. One recent strand of this literature argues that financial regulation should be designed to internalize pecuniary externalities in the presence of incomplete markets.[3] Our paper is based on pecuniary externalities from bank capital to wage earners and studies the redistributive implications.

In the discussion of optimal capital standards for financial institutions, Admati et al. (2013) and Miles, Yang, and Marcheggiano (2013) have argued that society at large would benefit from imposing higher capital standards. They focus on the direct social cost of risk-shifting by banks on governments, whereas this paper highlights an additional indirect social cost from the increased incidence of costly credit crunches. Estimates for the financial crisis of 2008/09 suggest that in most coun-

[3] See e.g. Lorenzoni (2008), Jeanne and Korinek (2010a), Jeanne and Korinek (2010b), Jeanne and Korinek (2013), Bianchi and Mendoza (2010), Korinek (2011), Benigno et al. (2012) and Gersbach and Rochet (2012) for papers on financial regulation motivated from asset price externalities, or Caballero and Lorenzoni (2014) for a paper on currency intervention based on wage externalities in an emerging economy. Campbell and Hercowitz (2009) study pecuniary externalities on the interest rate that arise in the transition from an equilibrium with low household debt to an equilibrium with high household debt. Kreamer (2014) studies the effect of household debt on the speed of recovery following a period of high unemployment.

tries, including the US, the indirect social cost of the credit crunch far outweighed the direct monetary costs of crisis-related bailouts (see e.g. Haldane (2010)).

In the empirical literature, Kaplan and Rauh (2010) and Philippon and Reshef (2012) provide evidence that the surplus created during booms accrued in large part to insiders in the financial sector, i.e. bankers in our framework. Larrain (2013) provides evidence on adverse effects of financial liberalization on wage inequality.

The rest of the paper proceeds as follows: The ensuing section develops an analytical model in which bankers intermediate capital to the real economy. Section 3.3 analyzes the determination of equilibrium and how changes in bank capital differentially affect the banking sector and the real economy. Section 3.4 describes the redistributive conflict over risk-taking between bankers and the real economy. Section 3.5 analyzes the impact of factors such as market power, agency problems, financial innovation, and bailouts on this conflict. All proofs are collected in Appendix C.

3.2 Model of Bank Capital and Workers

Consider an economy with three time periods, $t = 0, 1, 2$, and a unit mass each of two types of agents: bankers and workers. Furthermore, there is a single good that serves both as consumption good and capital.

Bankers. In period 0, bankers are born with one unit $e_0 = 1$ of the consumption good. They invest a fraction $x \in [0, 1]$ of it in a project that delivers a risky payoff

$\tilde{A}$ in period 1 with a continuously differentiable distribution function $G(\tilde{A})$ over the domain $[0, \infty)$, a density function $g(\tilde{A})$ and an expected value $E[\tilde{A}] > 1$. The realization of $\tilde{A}$ is common across all bankers.[4] Bankers hold the remainder $(1 - x)$ in a storage technology with gross return 1.

After the realization of the risky payoff $\tilde{A}$ in period 1, the resulting equity level of bankers is

$$e = x\tilde{A} + (1 - x) \tag{3.1}$$

Consistent with the banking literature, we use the term "bank capital" to refer to bank equity e in the following. However, this is a pure naming convention; bank capital is distinct from physical capital.

In period 1, bankers raise d deposits at a gross deposit rate of r and lend $k \leq d + e$ to the productive sector of the economy at a gross interest rate R. In period 2, bankers are repaid and value total profits in period 2 according to a linear utility function $\pi = Rk - rd$.

Workers. Workers are born in period 1 with a large endowment m of consumption goods. They lend an amount d of deposits to bankers at a deposit rate of r and hold the remainder in a storage technology with gross return 1. No arbitrage implies that the deposit rate satisfies $r = 1$.[5]

[4] This structure implies that banks are able to perfectly diversify their investment portfolios in order to eliminate idiosyncratic risk across projects. Then the only way to increase expected return is to increase exposure to aggregate risk, which is captured by the choice of x.

[5] Workers' endowments m are assumed to be sufficiently large to rule out cases where deposits are scarce and banks compete for funds by raising deposit rates above $r = 1$.

In period 2, workers inelastically supply one unit of labor $\ell = 1$ at the prevailing market wage w. Worker utility depends only on their total consumption, which for notational simplicity is normalized by subtracting the constant m so that $u = w\ell$.

In the described framework, risk markets between bankers and workers are incomplete since workers are born in period 1 after the technology shock $\tilde{A}$ is realized and cannot enter into risk-sharing contracts with bankers in period 0. All the risk $x\tilde{A}$ from investing in the risky technology therefore needs to be borne by bankers. An alternative microfoundation for this market incompleteness would be that obtaining the distribution function $G(\tilde{A})$ requires that bankers exert an unobservable private effort, and insuring against fluctuations in $\tilde{A}$ would destroy their incentives to exert this effort. In practice, bank capital is subject to significant fluctuations, and a large fraction of this risk is not shared with the rest of society.[6] Section 3.4.1 investigates the implications of reducing this market incompleteness.

Firms. Workers collectively own firms, which are neoclassical and competitive. Firms borrow k units of good from bankers at interest rate R at the end of period 1,

[6]For example, Wall Street banks routinely pay out up to half of their revenue as employee compensation in the form of largely performance-dependent bonuses, constituting an implicit equity stake by insiders in their firms. A considerable fraction of remaining explicit bank equity is also held by insiders. Furthermore, only 17.9% of US households hold direct stock investments, and another 33.2% hold equity investments indirectly, e.g. via retirement funds or other mutual funds. And this equity ownership is heavily skewed towards the high end of the income distribution (see e.g. Table A2a in Kennickell (2009)).

Period 0	Period 1	Period 2
• Banks enter with initial endowment 1 • Banks choose risky investment $x \in [0,1]$	• Shock $\tilde{A}$ realized • Bank equity $e = (1-x) + \tilde{A}x$ • Households enter and deposit d at rate r in banks • Bankers supply capital $k \leq d + e$ to firms	• Households supply labor $\ell = 1$ • Firms produce $F(k, \ell)$ • Banks receive return Rk, households obtain $w\ell$ • Banks pay households rd

Figure 3.1: Timeline

which they invest as physical capital. They hire labor ℓ from workers at wage w in period 2. They combine the two factors to produce output in period 2 according to the production function $F(k, \ell) = Ak^{\alpha}\ell^{1-\alpha}$ with $\alpha \in (0, 1)$. There is no uncertainty in firms' production. Firms maximize profits $F(k, \ell) - w\ell - Rk$ and find it optimal to equate the marginal product of each factor to its price, $F_k = R$ and $F_\ell = w$. In equilibrium they earn zero profits.

A timeline that summarizes our setup is presented in Figure 3.1.

Remark 1: In the described setup, the risk-taking decision x of bankers is separate from the financial intermediation function k, since they occur in separate time periods. This simplifies the analysis and sharpens our focus on the asymmetric costs of credit crunches, but implies that there is no direct contemporaneous benefit to workers if bankers invest more in the risky payoff with higher expected return. Appendix D.2 shows that our results continue to hold if the risk-taking and financial intermediation functions of bankers are intertwined. It considers an aggregate production function in periods 1 and 2 of $[\tilde{A}_t x_t + 1 - x_t]F(k_t, \ell_t)$, so that workers benefit immediately from risk-taking x_t through higher wages in period t.[7]

Remark 2: The model setup assumes for simplicity that the endowments of labor and savings as well as the firms are owned by workers. The results would be unchanged if these ownership claims were assigned to separate types of agents, since savers earn zero net returns and firms earn zero profits in equilibrium. For example, there could be an additional type of agent called capital owners who own all the savings and firms of the economy. Furthermore, our main insights are unchanged if labor supply is elastic.

3.2.1 First-Best Allocation

A planner who implements the first-best maximizes aggregate surplus in the economy

$$\max_{x,e,k,\ell} E\left[F(k,\ell) + e + m - k\right] \quad \text{s.t.} \quad e = x\tilde{A} + (1-x) \tag{3.2}$$

where $x \in [0,1]$ and $\ell \in [0,1]$. In period 2, the optimal labor input is $\ell^* = 1$, and the optimal level of capital investment satisfies $k^* = (\alpha A)^{\frac{1}{1-\alpha}}$, i.e. it equates the marginal return to investment to the return on the storage technology, $R^* = F_k(k^*, 1) = 1$. As discussed earlier, m is large enough that the resource constraint $k \leq e + m$ can be omitted, i.e. there are always sufficient funds available in the economy to invest k^* in the absence of market frictions. The marginal product of labor at the first-best level of capital is $w^* = F_\ell(k^*, 1)$.

[7] In a similar vein, it can be argued that risky borrowers (e.g. in the subprime segment) benefited from greater bank risk-taking because they obtained more and cheaper loans.

In period 0, the first-best planner chooses the portfolio allocation that maximizes expected bank capital $E[e]$. Since $E[\tilde{A}] > 1$, she will pick the corner solution $x = 1$. Since a fraction $\alpha F(k^*, 1)$ of production is spent on investment, the net social surplus generated in the first-best is $S^* = (1-\alpha) F(k^*, 1) + E[\tilde{A}]$.

3.2.2 Financial Constraint

We assume that bankers are subject to a commitment problem to capture the notion that bank capital matters. Specifically, bankers have access to a technology that allows them to divert a fraction $(1-\phi)$ of their gross revenue in period 2, where $\phi \in [0, 1]$. By implication depositors can receive repayments on their deposits that constitute at most a fraction ϕ of the gross revenue of bankers. Anticipating this commitment problem, depositors restrict their supply of deposits to satisfy the constraint

$$rd \leq \phi R k \tag{3.3}$$

An alternative interpretation of this financial constraint follows the spirit of Holmstrom and Tirole (1998): Suppose that bankers in period 1 can shirk in their monitoring effort, which yields a private benefit of B per unit of period 2 revenue but creates the risk of a bank failure that may occur with probability Δ and that results in a complete loss. Bankers will refrain from shirking as long as the benefits are less than the costs, or $BRk \leq \Delta[Rk - rd]$. If depositors impose the constraint above for $\phi = 1 - \frac{B}{\Delta}$, they can ensure that bankers avoid shirking and the associated risk of bankruptcy.[8]

[8] If the equilibrium interest rate is sufficiently large that $R > \frac{1}{1-\Delta+B}$, banks would prefer to offer

Remark: Our model assumes that all borrowing is used to finance capital investment, so that binding constraints directly reduce supply in the economy. An alternative and complementary assumption would be that credit is required to finance (durable) consumption so that binding constraints reduce demand. In both setups, binding financial constraints hurt the real economy, with similar redistributive implications.[9]

3.3 Laissez-Faire Equilibrium

The laissez-faire equilibrium of the economy is defined as the set of prices $\{r, R, w\}$ and allocation $\{x, e, d, k\}$, with all variables except x contingent on $\tilde{A}$, such that the decisions of bankers, workers, and firms are optimal given their constraints, and the markets for capital, labor and deposits clear.

We solve for the laissez-faire equilibrium in the economy with the financial constraint using backward induction, i.e. we first solve for the optimal period 1 equilibrium of bankers, firms and workers as a function of a given level of bank capital e. Then we analyze the optimal portfolio choice of bankers in period 0,

depositors a rate $r = \frac{1}{1-\Delta}$ and shirk in their monitoring, incurring the default risk Δ. However, this outcome is unlikely to occur in practice because such high interest rates would likely prompt a bailout, as discussed in Section 3.5.4

[9] Note that the benchmark model does not account for the procyclicality of financial leverage, which is documented e.g. in Brunnermeier and Pedersen (2009). However, this could easily be corrected by making the parameter ϕ vary with the state of nature so that $\phi(\tilde{A})$ is an increasing function.

which determines e.

3.3.1 Period 1 Equilibrium

Employment is always at its optimum level $\ell = 1$, since wages are flexible. The financial constraint is loose if bank capital is sufficiently high that bankers can intermediate the first-best amount of capital, $e \geq e^* = (1-\phi)k^*$. In this case, the deposit and lending rates satisfy $r = R = 1$ and bankers earn zero returns on lending. The wage level is $w^* = (1-\alpha) F(k^*, 1)$. This situation corresponds to "normal times."

If bank capital is below the threshold $e < e^*$ then the financial constraint binds and the financial sector cannot intermediate the first-best level of physical capital. This corresponds to a "credit crunch " or "financial crisis" since the binding financial constraints reduce output below its first-best level. Workers provide deposits up to the constraint $d = \phi Rk/r$, the deposit rate is $r = 1$, and the lending rate is $R = F_k(k, 1)$. Equilibrium capital investment in the constrained region, denoted by $\hat{k}(e)$, is implicitly defined by the equation

$$k = e + \phi k F_k(k, 1) \tag{3.4}$$

which has a unique positive solution for any $e \geq 0$. Overall, capital investment is given by the expression

$$k(e) = \min\left\{\hat{k}(e), k^*\right\} \tag{3.5}$$

Equilibrium $k(e)$ is strictly positive, strictly increasing in e over the domain $e \in [0, e^*)$ and constant at k^* for $e \geq e^*$. The equilibrium lending rate is then $R(e) =$

$\alpha F(k(e),1)/k(e)$. Equilibrium profits of the banking sector and worker utility are

$$\pi(e) = e + \alpha F(k(e),1) - k(e) \tag{3.6}$$

$$w(e) = (1-\alpha) F(k(e),1) \tag{3.7}$$

and total utilitarian surplus in the economy is $s(e) = w(e) + \pi(e)$.

Focusing on the decisions of an individual banker i, it is useful to distinguish individual bank capital e^i, which is a choice variable, from aggregate bank capital e, which is exogenous from an individual perspective. Then the level of physical capital intermediated by banker i and the resulting profits are respectively [10]

$$k\left(e^i, e\right) = \min\left\{k^*, \frac{e^i}{1-\phi R(e)}\right\} \tag{3.8}$$

$$\pi\left(e^i, e\right) = e^i + [R(e) - 1] \cdot k\left(e^i, e\right) \tag{3.9}$$

In equilibrium, $e^i = e$ will hold.

Panel 1 of Figure 3.2 depicts the payoffs of bankers and workers as a function of aggregate bank capital e. As long as $e < e^*$, physical capital investment falls short of the first best level. In this region, the welfare of workers and of bankers are strictly increasing concave functions of bank capital. Once bank capital reaches the threshold e^*, the economy achieves the first-best level of investment. Any bank capital beyond this point just reduces the amount of deposits that bankers need to raise, which increases their final payoff in period 2 but does not benefit workers.

[10] Technically, when financial intermediation is unconstrained at the aggregate level ($e > e^*$), there is a continuum of equilibrium allocations of k^i since the lending spread is zero and individual bankers are indifferent between intermediating or not. The equation gives the symmetric level of capital intermediation k^* for this case.

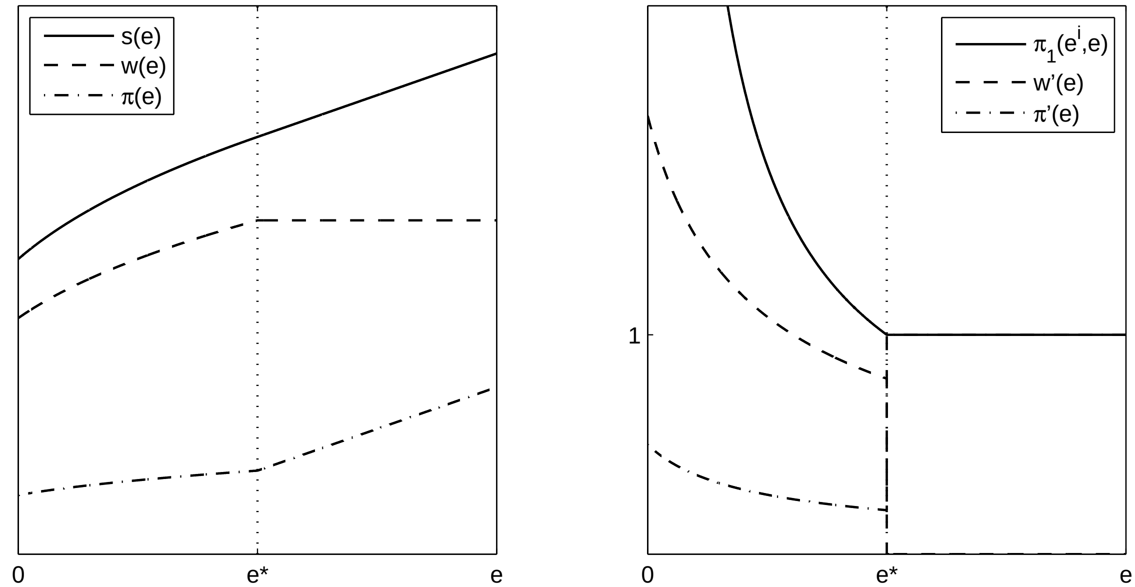

Figure 3.2: Welfare and marginal value of bank capital e: Social welfare $s(e)$, banker welfare $\pi(e)$, and worker welfare $w(e)$ are increasing and strictly concave in bank equity e below the critical level e^*. Above e^*, bankers benefit from additional equity, while workers do not. The right panel depicts marginal welfare, with π_1 depicting the marginal value to bankers of individual rather than aggregate bank equity. We generated this figure using parameters $A = 10$, $\alpha = 1/3$, and $\phi = 0.5$.

Beyond the threshold e^*, worker utility therefore remains constant and bank profits increase linearly in e. This generates a non-convexity in the function $\pi(e)$ at the threshold e^*. Our analytical findings on the value of bank capital are consistent with the empirical regularities of financial crises documented in e.g. Reinhart and Rogoff (2009).

3.3.2 Marginal Value of Bank Capital

How do changes in bank capital affect output and the distribution of surplus in the economy? If bankers are financially constrained in aggregate, i.e. if $e < e^*$, then a marginal increase in bank capital e allows bankers to raise more deposits and leads to a greater than one-for-one increase in capital investment k. Applying the implicit function theorem to (3.4) in the constrained region yields

$$k'(e) = \frac{1}{1 - \phi\alpha F_k} > 1 \quad \text{for} \quad e < e^* \tag{3.10}$$

If bankers are unconstrained, $e \geq e^*$, then additional bank capital e leaves physical capital investment unaffected at the first-best level k^*; therefore $k'(e) = 0$.

The effects of changes in bank capital for the two sectors differ dramatically depending on whether the financial constraint is loose or binding. In the unconstrained region $e \geq e^*$, the consumption value for bankers $\pi'(e) = 1$ is the only benefit of additional bank capital since $k'(e) = w'(e) = 0$. Bank capital is irrelevant for workers and the benefits of additional capital accrue entirely to bankers.

By contrast, in the constrained region $e < e^*$, additional bank capital increases physical capital intermediation k and output $F(k, 1)$. A fraction $(1 - \alpha)$ of the additional output F_k accrues to workers via increased wages, and a fraction α of the output net of the additional physical capital input accrues to bankers.[11] These

[11] Technically, these effects of bank capital on wages $w(e)$ and the return on capital $R(e)$ constitute *pecuniary externalities*. When atomistic bankers choose their optimal equity allocations, they take all prices as given and do not internalize that their collective actions will have general equilibrium effects that move wages and the lending rate.

effects are illustrated in Panel 2 of Figure 3.2.

When $e < e^*$, wages decline because labor is a production factor that is complementary to capital in the economy's production technology. Lending rates rise because the financial constraint creates scarcity, which drives up the return to capital investment. The difference between the lending rate and the deposit rate $r = 1$ allows bankers to earn a spread $R(e) - 1 > 0$. Observe that this spread plays a useful social role in allocating risk because it signals scarcity to bankers: there are extra returns available for carrying capital into constrained states of nature. However, the scarcity rents also redistribute from workers to bankers.

Equity Shortages and Redistribution. It is instructive to observe that small shortages of financial sector capital have first order redistributive effects but only second order efficiency effects. In particular, consider an economy in which bank capital is e^* so that the unconstrained equilibrium can just be implemented. Consider a wealth-neutral reallocation of the wealth of bankers across periods 1 and 2: bankers lose an infinitesimal amount ε of bank capital in period 1 so as to tighten their financial constraint, and regain it in period 2. The resulting payoffs for bankers and workers are $\pi(e^* - \varepsilon) + \varepsilon$ and $w(e^* - \varepsilon)$.

Lemma 3 (Redistributive Effects of Equity Shortages)**.** A marginal tightening of the financial constraint around the threshold e^* has first-order redistributive effects but only second-order efficiency costs.

Proof. See Appendix C for a proof of all lemmas and propositions. □

Intuitively, a marginal tightening of the constraint imposes losses on workers from lower wages that precisely equal the gains to bankers from higher lending spreads, i.e. the redistribution between workers and bankers occurs at a rate of one-to-one. Conceptually, this is because pecuniary externalities are by their nature redistributions driven by changes in prices. In our model, when financial constraints reduce the amount of capital intermediated and push down wages, the losses of workers equal the gains to firms. Similarly, when the lending rate rises, the losses to firms equal the gains to bankers. Since firms make zero profits, the losses to workers have to equal the gains to bankers. Put differently, since bankers are the bottleneck in the economy when the financial constraint binds, they extract surplus from workers in the form of scarcity rents.

3.3.3 Determination of Period 0 Risk Allocation

An individual banker i takes the lending rate R as given and perceives the constraint on deposits $d \leq \phi R k$ as a simple leverage limit. When a banker is constrained, she perceives the effect of a marginal increase in bank capital e^i as increasing her intermediation activity by $k_1\left(e^i, e\right) = \frac{1}{1-\phi R}$, which implies an increase in bank profits by

$$\pi_1\left(e^i, e\right) = 1 + \left[R\left(e\right) - 1\right] k_1\left(e^i, e\right) \tag{3.11}$$

In period 0, bankers decide what fraction x of their endowment to allocate to the risky project. In the laissez-faire equilibrium, banker i takes the aggregate levels of x and e as given and chooses $x^i \in [0,1]$ to maximize $\Pi^i\left(x^i; x\right) = E\left[\pi\left(e^i, e\right)\right]$. At an interior optimum, the optimality condition of bankers is

$$E\left[\pi_1\left(e^i, e\right)\left(\tilde{A}-1\right)\right] = 0, \tag{3.12}$$

i.e. the risk-adjusted return on the stochastic payoff $\tilde{A}$ equals the return of the safe storage technology.

The choice of x is determined by two opposing forces. Since $E[\tilde{A}] > 1$, the risky asset yields a higher expected return than the safe asset. Opposing this is a precautionary motive: following low realizations of $\tilde{A}$, aggregate bank capital is in short supply and bankers earn scarcity rents. As a result, bankers optimally trade off the opportunity to earn excess profits from the risky asset in period 0 versus excess profits from lending in period 1 when bank capital is scarce.

The stochastic discount factor π_1 in this expression is given by equation (3.11) and is strictly declining in e as long as $e < e^*$ and constant at 1 otherwise. Observe that each banker i perceives his stochastic discount factor as independent of his choices of e^i and x^i. However, in a symmetric equilibrium, $e^i = e$ as well as $x^i = x$ have to hold, and equilibrium is given by the level of x and the resulting realizations $e = \tilde{A}x + (1 - x)$ such that the optimality condition (3.12) is satisfied. As long as $E[\tilde{A}] > 1$, the optimal allocation to the risky project satisfies $x > 0$. If the expected return is sufficiently high, equilibrium is given by the corner solution $x = 1$. Otherwise it is uniquely pinned down by the optimality condition (3.12).

Denote by x^{LF} the fraction of their initial assets that bankers allocate to the risky project in the laissez-faire equilibrium. The resulting levels of welfare for entrepreneurs and workers are, respectively, $\Pi^{LF} = E\left[\pi\left(1 - x^{LF} + \tilde{A}x^{LF}\right)\right]$ and

$W^{LF} = E\left[w\left(1 - x^{LF} + \tilde{A}x^{LF}\right)\right]$. For a given risky portfolio allocation x, let $A^*(x)$ be the threshold of $\tilde{A}$ above which bank capital e is sufficiently high to support the first-best level of production. $A^*(x)$ satisfies

$$A^*(x) = 1 + \frac{e^* - 1}{x} \tag{3.13}$$

Well-Capitalized Banking System. If $e^* \leq 1$ (which can equivalently be read as $e_0 \geq e^*$ since $e_0 = 1$), then the safe return is sufficient to avoid the financial constraint and the first-best level of capital intermediation k^* would be reached for sure with a perfectly safe portfolio $x = 0$. This case corresponds to an economy in which the financial sector is sufficiently capitalized to intermediate the first-best amount of capital without any extra risk-taking. In this case, the risky project $\tilde{A}$ is a diversion from the main intermediation business of banks.[12]

For this case, bankers find it optimal to choose $x^{LF} > 1 - e^*$ (or, equivalently, $x^{LF} > e_0 - e^*$), i.e. they take on sufficient risk so that the financial constraint is binding for sufficiently low realizations of the risky return so that $A^*(x) > 0$. This is because the expected return on the risky project dominates the safe return, and bankers perceive the cost of being marginally constrained as second-order. Also observe that for $e^* < 1$, the function $A^*(x)$ is strictly increasing from $A^*(1 - e^*) = 0$ to $A^*(1) = e^*$, i.e. more risk-taking makes it more likely that the financial sector becomes constrained.

[12]Examples include a diversification from retail banking into investment banking, or loans by US banks to Latin American governments that offer extra returns at extra risk.

Under-Capitalized Banking System. If $e^* > 1$ (or, equivalently, $e_0 < e^*$), the economy would be constrained if bankers invested all their endowment in the safe return. This corresponds to an economy in which banks are systematically under-capitalized and risk-taking helps mitigate these constraints. In that case, the function $A^*(x)$ is strictly decreasing from $\lim_{x \to 0} A^*(x) = \infty$ to $A^*(1) = e^*$, i.e. more risk-taking makes it more likely that the financial sector becomes *un*constrained.

3.4 Pareto Frontier

We describe the redistributive effects of financial deregulation by characterizing the Pareto frontier of the economy, which maps different levels of financial risk-taking to different levels of welfare for the financial sector and the real economy. Financial regulation/deregulation moves the economy along this Pareto frontier.

Denote the period 0 allocation to the risky project that is collectively preferred by bankers by

$$x^B = \arg\max_{x \in [0,1]} E\left[\pi(\tilde{A}x + 1 - x)\right] \tag{3.14}$$

Similarly, let x^W be the level of risk-taking collectively preferred by workers, which maximizes $E[w(e)]$.

In a well-capitalized banking system, i.e. for $e^* \leq 1$ (equivalently, $e_0 \geq e^*$), workers prefer that risk-taking in the financial sector is limited to the point where financial constraints will be loose in all states of nature so that the first-best level of capital investment k^* can be implemented. This is guaranteed for any $x \in [0, 1 - e^*]$. Since workers are indifferent between all x within this interval but bankers benefit from risk-taking, the only point from this interval that is on the Pareto-frontier is $x^W = 1 - e^*$. In an under-capitalized banking system, i.e. for $e^* > 1$ (equivalently, $e_0 < e^*$), the optimal risk allocation for workers involves a positive level of risk-taking $x^W > 0$ – workers benefit from a little bit of risk because the safe return produces insufficient bank capital to intermediate the first-best amount of capital k^*, and risk-taking in period 0 increases the expected availability of finance in period 1.

Definition 4 (Pareto Frontier)**.** The Pareto frontier of the economy consists of all pairs of expected bank profits and worker wages $(\Pi(x), W(x))$ for $x \in \left[x^W, x^B\right]$.

To ensure that the Pareto frontier is non-degenerate, we assume that the optimal levels of risk-taking for workers and in the decentralized equilibrium are interior and satisfy $x^W < 1$ and $x^{LF} < 1$. This is a weak assumption that holds whenever the risk-reward trade-off associated with $\tilde{A}$ is sufficiently steep.

Proposition 13 (Characterization of Pareto Frontier)**.** (i) The risk allocations that are collectively preferred by workers and bankers, respectively, satisfy $x^W < x^B$.

(ii) Over the interval $\left[x^W, x^B\right]$, the expected utility of workers $W(x)$ is strictly decreasing in x, and the expected utility of bankers $\Pi(x)$ is strictly increasing in x.

(iii) The laissez-faire equilibrium satisfies $x^{LF} < x^B$. If $e^* \leq 1$ then $x^W < x^{LF} < x^B$.

The proposition characterizes the economy's Pareto frontier. Bankers prefer more risk-taking than workers because the risky technology offers higher returns and bankers capture all excess returns above e^*, whereas losses from a credit crunch are shared between bankers and workers. Any increase in risk-taking in $[x^W, x^B]$ benefits bankers and harms workers because it increases the incidence of binding financial constraints, which redistribute from workers towards bankers, as emphasized in Lemma 3. The laissez-faire equilibrium level of risk-taking x^{LF} will always be lower than x^B, and will often lie on the Pareto frontier between x^W and x^B. A sufficient condition for this is that $e^* \leq 1$ (or equivalently $e_0 \geq e^*$) so that the economy is initially well-capitalized.

3.4.1 Market Incompleteness and the Distributive Conflict

To pinpoint why there is a distributive conflict over the level of risk-taking, it is instructive to consider the consequences of removing one of the financial market imperfections. First, suppose there is no financial constraint on bankers in period 1. In that case, the profits or losses of bankers do not affect how much capital can be intermediated to the real economy and workers are indifferent about the level of risk-taking – bank capital does not generate any pecuniary externalities. In such an economy, Fisherian separation holds: financial risk-taking and financial intermediation are orthogonal activities and $x^W = x^B = x^{FB} = 1$, i.e. the distributive conflict disappears.

Alternatively, suppose that there is a complete insurance market in period 0 in which bankers and workers can share the risk associated with the technology $\tilde{A}$, but there is still a financial constraint in period 1. In that case, workers will insure bankers against any capital shortfalls so that bankers can invest in the risky technology without imposing negative externalities on the real economy. By implication all agents are happy to invest the first-best amount $x^W = x^B = x^{FB} = 1$ in the risky technology, and the distributive conflict again disappears. Introducing a risk market in period 0 puts a formal price on risk. If both sets of agents can participate in this market, this provides workers with a channel through which they can both share risk and transmit their risk preferences to bankers.

Even if both financial market imperfections are present, the distributive conflict also disappears if the constraint always binds.[13] In that case, bankers have the same risk exposure as workers and do not enjoy any asymmetric benefit on the upside since bank capital never exceeds the threshold where financial constraints are loose.[14] The distributive conflict is therefore generated by the combination of occasionally binding financial constraints and the lack of risk-sharing between bankers and workers. As argued in the introduction, both assumptions seem empirically highly relevant.

3.4.2 Financial Regulation

We interpret financial regulation in our framework as policy measures that affect risk-taking x. The unregulated equilibrium – in the absence of any other market distortions – is the laissez-faire equilibrium x^{LF}. If $x^{LF} \geq x^W$, then x^{LF} lies on the

[13]The constraint always binds if the initial endowment of bankers is insufficient to achieve the first-best level of capital k^*, even with $x = 1$ and the best realization of $\tilde{A}$, i.e. $\tilde{A}_{\max} e_0 \leq e^*$.

[14]This is the case in many macro models that are linearized around a steady state with binding constraints. Appendix D.1 provides an analytic exposition of this case.

Pareto frontier, and financial regulation moves the economy along the frontier.[15]

The two simplest forms of financial regulation of risk-taking are:

1. Regulators may impose a ceiling on the risk-taking of individual bankers such that $x^i \leq \bar{x}$ for some $\bar{x} < x^{LF}$. This type of regulation closely corresponds to capital adequacy regulations as it limits the amount of risk-taking per dollar of bank capital.

2. Regulators may impose a tax τ^x on risk-taking x^i so as to modify the optimality condition for the risk-return trade-off of bankers to $E[\pi_1 \cdot (\tilde{A} - \tau^x - 1)] = 0$. Such a tax can implement any level of risk-taking $x \in [0, 1]$. For simplicity, assume that the tax revenue is rebated to bankers in lump-sum fashion.

Financial regulators can implement any risk allocation $\bar{x} \leq x^{LF}$ by imposing $\bar{x}$ as a ceiling on risk-taking or by imposing an equivalent tax on risk-taking $\tau^x \geq 0$. As emphasized in the discussion of Proposition 13, $x^{LF} \geq x^W$ holds for a wide range of parameters, and always holds in the plausible case $e^* < 1$, i.e. when the economy is ex ante well-capitalized. For the remainder of this section, we assume that $x^W < x^{LF}$. In this case the distributive implications are straightforward:

Corollary 1 (Redistributive Effects of Financial Regulation)**.** Tightening regulation

[15] Observe that a financial regulator would not find it optimal to change the leverage parameter ϕ in period 1 of our setup. The parameter cannot be increased because it stems from an underlying moral hazard problem and banks would default or deviate from their optimal behavior. Similarly, it is not optimal to decrease ϕ because this would tighten the constraint on financial intermediation without any corresponding benefit.

by lowering $\bar{x}$ or raising τ^x increases worker welfare and reduces banker welfare for any $\bar{x} \in \left[x^W, x^{LF}\right]$.

Conversely, financial deregulation increases the ceiling $\bar{x}$ and redistributes from workers to bankers.

Scope for Pareto-Improving Deregulation. An interesting question is whether there exists a mechanism for Pareto-improving deregulation given additional instruments for policymakers other than the regulatory measures on x described in Corollary 1. Such a mechanism would need to use some of the gains from deregulation obtained by bankers to compensate workers for the losses they suffer during credit crunches.

First, consider a planner who provides an uncontingent lump-sum transfer from bankers to workers in period 1 to compensate workers for the losses from deregulation. The marginal benefit to workers is $1 - E\left[w'(e)\right]$, i.e. workers would obtain a direct marginal benefit of 1 in all states of nature, but in constrained states they would be hurt by a tightening of the financial constraint which reduces their wages by $w'(e)$. The uncontingent transfer thus entails efficiency losses from tightening the constraints on bankers. The planner needs to weigh the redistributive benefit of any transfer against the cost of the distortion introduced. This creates a constrained Pareto frontier along which the trade-off between the welfare of the two agents is less favorable than the original Pareto frontier. Compensating workers with an uncontingent payment without imposing these efficiency costs would require that the planner have superior enforcement capabilities to extract payments from bankers in excess of the financial constraint (3.3), which are not available to private markets.

Alternatively, consider a planner who provides compensatory transfers to workers contingent on states of nature in which bankers are unconstrained, i.e. in states in which they make high profits from the risky technology $\tilde{A}$. This would avoid efficiency costs but would again require that the planner can engage in state-contingent transactions that are not available via private markets. (It can be argued that this type of transfer corresponds to proportional or progressive profit taxation.)

In short, the planner can only achieve a Pareto improvement if she is either willing to provide transfers at the expense of reducing efficiency, or if she can get around one of the two market imperfections in our framework, i.e. mitigate either the financial constraint (3.3) or the incompleteness of risk markets.

3.5 Risk-Taking and Redistribution

We extend our baseline model to analyze the redistributive implications of four factors that are commonly viewed as reasons for risk-taking in the financial sector: agency problems, market power, financial innovation, and bailouts.

3.5.1 Asymmetric Compensation Schemes

It is frequently argued that asymmetric compensation schemes provide managers of financial institutions with excessive risk-taking incentives and that this may have played an important role in the build-up of risk before the financial crisis of 2008/09.

To illustrate this mechanism, we consider a stylized model of an incentive problem between bank owners and bank managers and analyze the distributive implications.

Assume that bank owners have to hire a new set of agents called bank managers to conduct their business. Bank managers choose an unobservable level of risk-taking x in period 0. Bank owners are able to observe the realization of profits and bank capital e in period 1, but do not observe $\tilde{A}$ and so cannot infer x. Bank owners instruct managers to allocate any bank capital up to e^* in financial intermediation, and to carry any excess capital $\max\{0, e - e^*\}$ in the storage technology. Financial intermediation versus storage can be viewed as representative of lending to real projects versus financial investments, or commercial banking versus investment banking.

Suppose that bank managers do not have the ability to commit to exert effort in period 1 and can threaten to withdraw their monitoring effort for both bank loans and storage in period 1. If they do not monitor, the returns on intermediation and storage (real projects and financial investments) are diminished by a fraction ε and $\delta\varepsilon$ respectively, where $\delta > 1$. In other words, the returns to financial investments are more sensitive to managerial effort than real investments. An alternative interpretation would be along the lines of Jensen (1986) that free cash provides managers with greater scope to abuse resources.

Assuming that managers have all the bargaining power, and given a symmetric equilibrium, the threat to withdraw their effort allows managers to negotiate an incentive payment from bank owners of

$$p\left(e^{i},e\right)=\varepsilon\min\left\{\pi\left(e^{i},e\right),\pi\left(e^{*},e\right)\right\}+\delta\varepsilon\max\left\{0,e^{i}-e^{*}\right\} \tag{3.15}$$

The marginal benefit of bank capital for an individual manager is $p_1(e,e)=\varepsilon\pi_1(e,e)$ for $e<e^*$ and $p_1(e,e)=\delta\varepsilon\pi_1(e,e)=\delta\varepsilon$ for $e\geq e^*$. Since financial investments deliver a greater incentive payment, the payoff of managers is more convex than the payoff of banks $\pi(e,e)$, and managers benefit disproportionately from high realizations of bank capital. Comparing this extension to our benchmark setup, $\Pi(x)$ is now the joint surplus of bank owners and managers, and the two functions $\Pi(x)$ and $W(x)$ remain unchanged compared to our earlier framework – the only thing that changes is the level of x that will be chosen by bank managers.

Managers internalize the asymmetric payoff profile when they choose the level of risk-taking in period 0. They maximize $E[p(e^i,e)]$ where $e^i=\tilde{A}x^i+1-x^i$. It is then straightforward to obtain the following result:

Proposition 14 (Agency Problems and Risk-Taking)**.** (i) The optimal choice of risk-taking of bank managers exceeds the optimal choice x^{LF} in our benchmark model if the payoff function of managers is asymmetric $\delta>1$.

(ii) If $x^W\leq x^{LF}$, the expected welfare of workers is a declining function of δ.

3.5.2 Financial Institutions with Market Power

Assume that there is a finite number n of identical bankers in the economy who each have mass $\frac{1}{n}$. Banker i internalizes that his risk-taking decision x^i in period 0 affects aggregate bank capital $e=\frac{1}{n}e^i+\frac{n-1}{n}e^{-i}$, where e^{-i} captures the capital of the other bankers in the economy. For a given e, assume that bankers charge the competitive market interest rate $R(e)$ in period 1.[16] Our results are summarized in Proposition 15.

Proposition 15. The optimal risk allocation x^n of bankers is a declining function of the number n of banks in the market, and $x^1 = x^B \geq x^\infty = x^{LF}$, with strict inequality excepting corner solutions.

Intuitively, bankers with market power internalize that additional equity when the economy is constrained reduces their lending spreads. This counteracts the precautionary motive to carry extra capital into constrained states of nature. Our example illustrates that socially excessive risk-taking is an important dimension of non-competitive behavior by banks.

3.5.3 Financial Innovation

An important manifestation of financial innovation is to allow financial market players to access new investment opportunities, frequently projects that are characterized by both higher risk and higher expected returns. For example, financial innovation may enable bankers to invest in new activities, as made possible by the 1999 repeal of the 1933 Glass-Steagall Act, or to lend in new areas, to new sectors or to

[16] By contrast, if bankers interacted in Cournot-style competition in the period 1 market for loans, they would restrict the quantity of loans provided for a given amount of bank equity e^i to $\min\left\{k\left(e^i\right), k^{*n}\right\}$ where $k^{*n} = k^*\left(\frac{n-(1-\alpha)}{n}\right)^{\frac{1}{1-\alpha}}$ to increase their scarcity rents. We do not consider this effect in order to focus our analysis on the period 0 risk-taking effects of market power.

new borrowers, as during the subprime boom of the 2000s.

Our setup can formally capture this type of financial innovation by expanding the set of risky assets to which bankers have access in period 0. For a simple example, assume an economy in which bankers can only access the safe investment project in period 0 before financial innovation takes place, and that financial innovation expands the set of investable projects to include the risky project with stochastic return $\tilde{A}$. Furthermore, assume that $e^* < 1$, i.e. the safe return in period 0 generates sufficient period 1 equity for bankers to intermediate the first-best level of capital. The pre-innovation equilibrium corresponds to $x = 0$ in our benchmark setup and this maximizes worker welfare.

Example 1 (Distributive Effects of Financial Innovation)**.** In the described economy, expanding the set of investment projects to include $\tilde{A}$ increases banker welfare but reduces worker welfare.

After financial innovation introduces the risky project, bankers allocate a strictly positive fraction of their endowment $x^{LF} > 1 - e^*$ to the risky project and incur binding financial constraints in low states of nature. This is their optimal choice because the expected return $E[\tilde{A}] > 1$ delivers a first-order benefit over the safe return, but bankers perceive the cost of being marginally constrained as second-order since $\pi_1(e^i, e)$ is continuous at e^*. Worker welfare, on the other hand, unambiguously declines as a result of the increased risk-taking.

This illustrates that financial innovation that increases the set of investable projects so as to include more high-risk/high-return options may redistribute from workers to bankers, akin to financial deregulation, even though total surplus may be increased. The problem in the described economy is that workers would be happy for bankers to increase risk-taking if they could participate in both the upside and the downside via complete insurance markets. Restrictions on the risk-taking activities of banks, e.g. via regulations such as the Volcker rule, may benefit workers by acting as a second-best device to complete financial markets. In the example described above this would be the case. Naturally, there are also some financial innovations that may increase worker welfare. In our framework, this may be the case for example for increases in ϕ, i.e. relaxations of the commitment problem of bankers.

3.5.4 Bailouts

Bailouts have perhaps raised more redistributive concerns than any other form of public financial intervention. This is presumably because they involve redistributions in the form of explicit transfers that are more transparent than other implicit forms of redistribution.

However, the redistributive effects of bailouts are both more subtle and potentially more pernicious than what is suggested by focusing on the direct fiscal cost. Ex post, i.e. once bankers have suffered large losses and the economy experiences a credit crunch, bailouts may actually lead to a Pareto improvement so that workers are better off by providing a transfer. However, ex-ante, bailout expectations increase risk-taking. This redistributes surplus from workers to the financial sector in a less explicit and therefore more subtle way, as emphasized throughout this paper.

Workers in our model find it ex-post collectively optimal to provide bailouts to bankers during episodes of severe capital shortages since this mitigates the credit crunch and its adverse effects on the real economy. Given an aggregate bank capital position e in period 1, the following policy maximizes ex-post worker welfare:[17]

Lemma 4 (Optimal Bailout Policy)**.** If aggregate bank capital in period 1 is below a threshold $0 < \hat{e} < e^*$, workers find it collectively optimal to provide lump-sum transfer $t = \hat{e} - e$ to bankers. The threshold $\hat{e}$ is determined by the expression $w'(\hat{e}) = 1$ or

$$\hat{e} = (1-\alpha)\left[1-(1-\phi)\alpha\right]^{\frac{\alpha}{1-\alpha}} e^* \tag{3.16}$$

The intuition stems from the pecuniary externalities of bank capital on wages: increasing bank capital via lump-sum transfers relaxes the financial constraint of bankers and enables them to intermediate more capital, which in turn expands output and increases wages. As long as $e < \hat{e}$, the cost of a transfer to workers is less than the collective benefit in the form of higher wages.

Bailouts constitute straight transfers from workers to bankers, but generate a Pareto improvement for $e < \hat{e}$ because they mitigate the market incompleteness that is created by the financial constraint (3.3) and that prevents bankers from raising deposit finance and intermediating capital to the productive sector. At the margin, each additional unit of bailout generates a surplus $F_k(e, 1) - 1$, of which $w'(e) - 1$

[17]This section focuses on bailouts in the form of lump-sum transfers. Appendix D.3 shows that our results apply equally if bailouts are provided via emergency lending or equity injections on subsidized terms.

accrues to workers and $\pi'(e)$ to bankers. For the last marginal unit of the bailout, the benefit to workers is $w'(\hat{e}) - 1 = 0$ – they are indifferent between providing the last unit or not. However, the marginal benefit to bankers for the last unit is strictly positive $\pi'(\hat{e}) = \frac{(1-\phi)\alpha}{1-\alpha}$.[18]

Period 0 Risk-Taking. Optimal discretionary bailouts impose a ceiling on the market interest rate $R^{BL}(e) \leq R(\hat{e}) = \frac{1}{1-(1-\phi)\alpha}$ since they ensure that aggregate capital investment is at least $k \geq k(\hat{e})$ at all times. This mitigates the precautionary incentives of bankers and increases their risk-taking, corresponding to an "income effect" of bailouts. This effect exists even if bailouts are provided in the form of lump-sum transfers and do not distort the optimality conditions of bankers. The adverse incentive effects of bailouts are aggravated if they are conditional on individual bank capital e^i, which distorts the risk-taking incentives of bankers, corresponding to a "substitution effect" of bailouts.[19] Denoting the amount of their endowment that bankers allocate to the risky project by x^{BL}:

Lemma 5 (Risk-Taking Effects of Bailouts)**.** Introducing bailouts increases period

[18] For the remainder of our analysis of bailouts, we assume that the parameters α, ϕ and A are such that $\hat{e} < 1$ (or, equivalently, $e_0 > \hat{e}$). This is a mild assumption that guarantees that the banking sector will not require a bailout if the period 0 endowment is invested in the safe project. It also implies that bailouts are not desirable in states of nature in which the risky project yields higher returns than the safe project. This is reasonable because typically bailouts occur only if risky investments have gone bad.

[19] This effect is well-understood in the literature on bailouts and operates in the same direction as the income effect. For details on this case, see Appendix D.4

0 risk-taking, $x^{BL} > x^{LF}$.

Intuitively, bailouts reduce the tightness of constraints and therefore the returns on capital π_1 in low states of nature. This lowers the precautionary incentives of bankers and induces them to take on more risk, even though the bailouts are provided in a lump-sum fashion. Observe that this effect is similar to the effects of any countercyclical policy or any improvement in risk-sharing via markets.[20]

Redistributive Effects. The welfare effects of introducing bailouts on bankers and workers can be decomposed into two parts, the change in expected welfare from introducing bailouts for a given level of risk-taking x^{BL}, corresponding to the market completion effect of bailouts, and the change in the level of risk-taking, corresponding to the incentive effects of bailouts:[21]

$$\Delta\Pi = \left[\Pi^{BL}(x^{BL}) - \Pi(x^{BL})\right] + \left[\Pi(x^{BL}) - \Pi(x^{LF})\right] \tag{3.17}$$

$$\Delta W = \underbrace{\left[W^{BL}(x^{BL}) - W(x^{BL})\right]}_{\text{market completion}} + \underbrace{\left[W(x^{BL}) - W(x^{LF})\right]}_{\text{incentive effect}} \tag{3.18}$$

Corollary 2 (Distributive Effects of Bailouts)**.** (i) Bankers always benefit from introducing bailouts.

(ii) Workers benefit from the market completion effect of bailouts, but are hurt by the incentive effects of bailouts if $x^W < x^{LF}$.

[20]Our framework does not explicitly account for bankruptcy because $\tilde{A}$ is bounded at 0; if period 0 investments could lead to bankruptcy, there may be an additional risk-taking incentive for banks.

[21]The first term for workers could be further separated into a negative term corresponding to the transfers that they make, and a larger positive term corresponding to the resulting increase in wages for given x.

Although the market completion effect is positive for both sets of agents, the increase in risk-taking benefits bankers at the expense of workers because $W'(x) < 0$. Bailouts increase banker welfare both directly because of the transfers received from workers and indirectly as a result of the higher risk-taking. Haldane (2010) emphasizes that the social cost of the 2008/09 credit crunch exceeded the fiscal cost of bailouts by an order of magnitude. This suggests that the effects of bailouts on risk-taking incentives may be far costlier to workers than the direct fiscal cost.

The introduction of bailouts is analogous to banker-biased technological change. The market completion effect is an outward shift of the Pareto frontier of the economy, whereas the higher risk-taking due to the incentive effect is a movement along the Pareto frontier towards bankers.

3.6 Conclusions

The central finding of our paper is that financial regulation has important redistributive implications. The majority of the literature on financial regulation focuses on the efficiency implications of financial regulation and disregards redistributive effects. Welfare is typically determined by a planner who picks the most efficient allocation under the assumption that the desired distribution of resources between different agents can be implemented independently.

We find that deregulation benefits the financial sector by allowing for greater risk-taking and higher expected profits. However, the downside is that greater risk-taking leads to a greater incidence of losses that are sufficiently large to trigger a credit crunch. If the financial sector is constrained in its intermediation activity, the real economy obtains less credit and invests less, lowering output and the marginal product of labor, which imposes negative externalities on wage earners. The degree of financial risk-taking and financial regulation therefore has first-order redistributive implications.

There are a number of issues that we leave for future analysis: First, since risk-taking is profitable, financial regulation generates large incentives for circumvention. If the regulatory framework of a country covers only one part of its financial system, the remaining parts will expand. In the US, for example, the shadow financial system grew to the point where it constituted an essential part of the financial sector, but it was largely unregulated and could engage in high levels of risk-taking. This made the sector vulnerable to the losses experienced during the 2008 financial crisis. And since the sector had become an essential part of the financial system, its losses generated strong adverse effects on the real economy.

Second, our results shed light on what types of financial innovation and financial regulation are most likely to increase the welfare of both the financial sector and the real economy in order to achieve a Pareto improvement. Our findings suggest two promising directions that correspond to alleviating the two main market imperfections in our framework: (i) innovations or regulatory interventions that increase risk-sharing between the two sectors on both the upside and the downside, such as taxes on financial profits during booms. These reduce the distributive conflict over risk-taking by allowing a more equitable sharing of the gains from financial risk-taking. (ii) innovations or regulatory interventions that reduce the likelihood of hitting binding constraints, for example better capitalized banks. These reduce the likelihood of credit crunches that have real implications, reducing the distributive conflict by alleviating the negative externalities from the financial sector on the real economy during such episodes.

Third, the paper mainly discusses the effects of financial risk-taking, but if the financial sector designs innovative ways of financing risky investment opportunities in the real economy and of sharing the associated risks so as to protect the economy from credit crunches, it is likely that both sectors benefit. An example would be innovations that increase the availability of venture capital. Thus it is important for regulators to distinguish between financial risk-taking and intermediating risk capital to the real economy.

Appendix A: Computational Algorithm (Chapter 1)

I begin by solving for the steady state equilibrium. For given parameters, this involves finding the job-finding rate p and borrowing constraint $\underline{a}$ such that the stationary distribution of assets resulting from the household decision rules satisfies $A = 0$.

A.1 Household problem

I solve the household problem for given $(p, \underline{a})$ using value function iteration.[1] For given values of (V, U), we can define consumption using (1.4) and (1.5). The consumption rules define saving rules by (1.1). If the implied saving rule violates the boorrowing constraint, saving at $\underline{a}$ is set to zero.

Letting (V_n, U_n) be the value function at step n of value function, we can define the next step of the iteration (V_{n+1}, U_{n+1}) by

$$\rho V_{n+1} = u(c_e) + \frac{dV_{n+1}}{da} \cdot \dot{a}_e + \frac{V_n - V_{n+1}}{dt} + s\left(U_{n+1} - V_{n+1}\right)$$
$$\rho U_{n+1} = u(c_u) + \frac{dU_{n+1}}{da} \cdot \dot{a}_u + \frac{U_n - U_{n+1}}{dt} + p\left(V_{n+1} - U_{n+1}\right)$$

where dt is the step size. Note that this is equivalent to iterating backward through time, with dt equal to the time step. The system of equations above can be quickly

[1] For a detailed description of this algorithm, see Achdou et al. (2013).

solved for the next step by inverting a single matrix. This process continues until the process converges, which generally takes on the order of 20 iterations.[2]

A.2 Calculating the stationary distribution

Once we have calculated the decision rules of households, we can quickly solve for the implied stationary distribution of assets. From (1.28) and (1.29), the stationary asset distribution satisfies

$$pm_u - sm_e - \frac{d}{da}(m_e \dot{a}_e) = 0$$
$$sm_e - pm_u - \frac{d}{da}(m_u \dot{a}_u) = 0$$

To calculate the stationary distribution numerically we discretize the state space to a set of points a_i, and let $\Delta = a_{i+1} - a_i$ be the distance between adjacent points. Then we approximate the derivative as

$$\frac{d}{da}(m\dot{a}) \approx \frac{1}{\Delta}\left[\max\left(m_{i-1}\dot{a}_{i-1}, 0\right) - \min\left(m_{i+1}\dot{a}_{i+1}, 0\right) - |m_i \dot{a}_i|\right]$$

This corresponds to taking the left difference when $\dot{a}_i > 0$, and the right difference when $\dot{a}_i < 0$. When $\dot{a}_i \approx 0$ so that $\dot{a}_{i-1} > 0$ and $\dot{a}_{i+1} < 0$, the expression corresponds to computing the first difference in the direction in which $m > 0$. This is important because $m_e \dot{a}_e$ may not be differentiable at a point where $\dot{a} = 0$, but the derivative in one direction will exist. Moreover, since we are in a two-state case, we know that $m = 0$ will hold in one direction when $\dot{a} = 0$. In particular, for employed households $m_{i+1} = 0$, and for unemployed households $m_{i-1} = 0$.

[2] Convergence of the algorithm requires an "upwind" approximation of the derivative of the value function dV/da, as described in Candler (2001).

In addition to the expression above, we also have the requirement that $\sum_i m_i^e = \frac{p}{p+s} = N$ and $\sum_i m_i^u = \frac{s}{p+s} = 1 - N$. Note that in this case m_i is the mass at point i, rather than the pdf, which would equal m_i/Δ.

Now we write these equations in matrix form as $Tm = v$, where $m = (m_e^1, ..., m_e^n, m_u^1, ..., m_u^n)$, and v a vector with N in the nth position, and $1 - N$ in the 2nth position, and otherwise zero. (Here n is the number of grid points). T is a matrix with entries in terms of p, s, and $\dot{a}_i$, that encodes the rate of transitions above, but with the nth and 2nth row all 0s and 1s to encode the summations $\sum_i m_i^e = \frac{p}{p+s} = N$ and $\sum_i m_i^u = \frac{s}{p+s} = 1 - N$.

Then we simply invert the matrix T to find the stationary distribution $m = T^{-1}v$. This matrix is invertible because $\dot{a}(a)$ is strictly decreasing, and so equals zero at just one point.

A.3 Finding a transition path

Once we have the steady state, we can calculate the transition path following a shock as follows. First we guess a path of hiring probabilities $p(t)$ and initial employment $n(0)$. Given this path of hiring probabiltiies, we can iterate the household problem backward from the steady state using one step of the value function iteration algorithm used to solve for the steady state. At every step we must also calculate the borrowing constraint implied by the previously computed value function for unemployed households U.

Once we have the sequence of decision rules, we can iterate the initial asset distribution and employment forward to compute the path of asset demand. We then search for the path of hiring probabilities that makes asset demand zero at every point in time.

Because changing the hiring probability at any point in time has nonlinear effects on asset demand before and after this time, it is difficult to define a simple rule for updating the path of $p(t)$. Directly searching for the path of p that satisfied $A(t) = 0$ at every t is prohibitively computationally expensive. I instead search for values of p at several points, and interpolate the intermediate values, with more points immediately following the shock to capture the more complicated dynamics in this region. I confirmed that this method produces a very close approximation to $A = 0$, and varying the number of gridpoints at the margin does not alter the resulting dynamics.

Appendix B: Proofs (Chapter 2)

B.1 Proofs from section 2.2

Proof of Proposition 1. We can solve (2.4) using the Lagrangian

$$L = \sum_{i=0}^{2} C_i + \mu_0 [A + H_0 - C_0 - I] + \mu_1 [H_1 - C_1 - p\lambda_L \rho_L I - (1-p)\lambda_H \rho_H I]$$

$$+ \mu_2 [H_2 + p\lambda_L \rho_1 I + (1-p)\lambda_H \rho_1 I + l - C_2] + \sum_{i=0}^{2} \nu_i C_i + \nu_3 I$$

The first-order condition with respect to C_i is

$$\mu_i = 1 + \nu_i$$

which implies $\mu_i > 0$, and so the period budget constraints hold with equality. Substituting them directly into the objective function, we obtain the new problem

$$\max_{\lambda_L, \lambda_H, I} \{A + H_0 + H_1 + H_2 + p\lambda_L (\rho_1 - \rho_L) I + (1-p)\lambda_H (\rho_1 - \rho_H) I - I\} \quad \text{(B.1)}$$

$$\text{s.t.} \quad H \geq I - A \quad \text{(B.2)}$$

$$H \geq p\lambda_L \rho_L I + (1-p)\lambda_H \rho_H I \quad \text{(B.3)}$$

$$I \geq 0$$

where $\lambda_s \in \{0,1\}$ for $s \in \{L,H\}$. The constraints (B.2) and (B.3) are the non-negativity constraints on C_0 and C_1 respectively. The Lagrangian of (B.1) is

$$L = A + H_0 + H_1 + H_2 + p\lambda_L(\rho_1 - \rho_L)I + (1-p)\lambda_H(\rho_1 - \rho_H)I - I$$
$$+ \mu_1[A + H - I] + \mu_2[H - p\lambda_L\rho_L I - (1-p)\lambda_H\rho_H I] + \mu_3 I \tag{B.4}$$

The first-order condition of (B.4) with respect to I is

$$p\lambda_L(\rho_1 - \rho_L) + (1-p)\lambda_H(\rho_1 - \rho_H) - 1 = \mu_2[p\lambda_L\rho_L + (1-p)\lambda_H\rho_H] + \mu_1 - \mu_3$$

By assumption, there will always be sufficient funds in period 1 to meet liquidity shocks, meaning $H_1 \geq p\lambda_L\rho_L I + (1-p)\lambda_H\rho_H I$, and so we have $\mu_2 = 0$. Thus the first-order condition reduces to

$$p\lambda_L(\rho_1 - \rho_L) + (1-p)\lambda_H(\rho_1 - \rho_H) - 1 = \mu_1 - \mu_3$$

The left-hand side $p\lambda_L(\rho_1 - \rho_L) + (1-p)\lambda_H(\rho_1 - \rho_H) - 1$ is the net return on investment. If the net return is positive, then $\mu_1 > 0$, the non-negativity constraint on C_0 binds, and the economy invests all available funds in period 0. If the net return is negative, then $\mu_3 > 0$, the non-negativity constraint on I binds, and the economy does not invest anything. We can express this investment rule as

$$I = \begin{cases} A + H_0 & p\lambda_L(\rho_1 - \rho_L) + (1-p)\lambda_H(\rho_1 - \rho_H) > 1 \\ [0, A + H_0] & p\lambda_L(\rho_1 - \rho_L) + (1-p)\lambda_H(\rho_1 - \rho_H) = 1 \\ 0 & p\lambda_L(\rho_1 - \rho_L) + (1-p)\lambda_H(\rho_1 - \rho_H) < 1 \end{cases}$$

Now we just need to determine the optimal choices of λ_L and λ_H. Increasing λ_L and λ_H from 0 to 1 results in increases in the objective function of $p\lambda_L(\rho_1 - \rho_L)I$

and $(1-p)\,\lambda_H\,(\rho_1-\rho_H)\,I$ respectively. Given our assumption that $\rho_1 > \rho_H > \rho_L$, both of these terms are positive for $I > 0$. Therefore it is optimal to choose $\lambda_L = 1$ and $\lambda_H = 1$. Since the return from a project is greater than the additional cost of bringing the project to completion after the liquidity shock is realized, the optimal continuation policy is to continue in all cases.

We can now define the unconstrained optimum. The solution is $\{C_0, C_1, C_2, I, \lambda\}$ with

$$C_0 = A + H_0 - I$$

$$C_1 = H_1 - p\rho_L I - (1-p)\rho_H I$$

$$C_2 = H_2 + p\rho_1 I + (1-p)\rho_1 I$$

$$\lambda_L = 1$$

$$\lambda_H = 1$$

$$I = \begin{cases} A + H_0 & R \geq 0 \\ 0 & R < 0 \end{cases}$$

where $R = p\,(\rho_1-\rho_L) + (1-p)\,(\rho_1-\rho_H) - 1$. □

Proof of Lemma 1. We need to show that it is always worthwhile to increase R^I_s in order to increase I as long as the constraint on pledgeable income in the s state does not bind. If this is true, then it follows that in any equilibrium with positive external financing, limited pledgeability binds, meaning that firms receive exactly the amount necessary for them to cooperate, and no more.

The first step is to argue that the period 1 investors will be paid exactly the funds necessary to finance meeting the liquidity shocks, and no more. This should not be controversial, since there is no other benefit to increasing payments to period 1 investors. To show this, we take the Lagrangian of the problem

$$
\begin{aligned}
L &= p\left(\lambda_L \rho_1 I - R_L^I - R_L^1 + \ell\right) + (1-p)\left(\lambda_H \rho_1 I - R_H^I - R_H^1 + \ell\right) \\
&\quad + \mu_1\left[R_L^1 - \lambda_L \rho_L I\right] + \mu_2\left[R_H^1 - \lambda_H \rho_H I\right] \\
&\quad + \mu_3\left[pR_L^I + (1-p)R_H^I - I - q\ell + A\right] \\
&\quad + \mu_4\left[\lambda_H \rho_0 I + \ell - R_H^I - R_H^1\right] + \mu_5\left[\lambda_L \rho_0 I + \ell - R_L^I - R_L^1\right]
\end{aligned}
$$

and differentiate with respect to R_L^1 and R_H^1. This yields conditions

$$
\begin{aligned}
\mu_1 &\leq \mu_4 + p \\
\mu_2 &\leq \mu_5 + 1 - p
\end{aligned}
$$

which hold with equality if the corresponding $R_s^1 > 0$. We also have $R_L^1 \geq \lambda_L \rho_L I$ and $R_H^1 \geq \lambda_H \rho_H I$, which indicates that if $I > 0$ and $\lambda_s > 0$, we have $R_s^1 > 0$. So we can conclude that in fact we have $R_s^1 = \lambda_s \rho_s I$, for $s \in \{L, H\}$. Substituting these terms directly, we can rewrite the Lagrangian as

$$
\begin{aligned}
L &= p\left(\lambda_L \rho_1 I - R_L^I - \lambda_L \rho_L I + l\right) + (1-p)\left(\lambda_H \rho_1 I - R_H^I - \lambda_H \rho_H I + l\right) \\
&\quad + \mu_3\left[pR_L^I + (1-p)R_H^I - I - ql + A\right] \\
&\quad + \mu_4\left[\lambda_H \rho_0 I + l - \lambda_H \rho_H I - R_H^1\right] + \mu_5\left[\lambda_L \rho_0 I + l - R_L^I - \lambda_L \rho_L I\right]
\end{aligned}
$$

Now we derive conditions for optimal λ_s. From the Lagrangian, we find that the marginal values of increasing λ_L and λ_H are respectively

$$p(\rho_1 - \rho_L)I + \mu_5(\rho_0 - \rho_L)I$$

$$(1-p)(\rho_1 - \rho_H)I + \mu_5(\rho_0 - \rho_H)I$$

The optimal choice of λ_s is 1 if the marginal value is positive, and 0 if the marginal value is negative. Since by assumption we have $\rho_1 > \rho_L$ and $\rho_0 > \rho_L$, the first condition is strictly positive, which means that we have $\lambda_L = 1$. The sign of the second is ambiguous since $\rho_0 < \rho_H$, and so the value of λ_H is not clear.

The next step is to observe that a pledgeability constraint will bind if (1) that shock is met, and (2) the multiplier on investment is greater than 1, i.e. $\mu_3 > 1$. This is logical because the direct cost of an increase in R_s^I is 1, so if the value of increasing investment is greater than 1, the optimal solution will be to increase R_s^I until some other constraint binds. This follows from the first-order conditions with respect to R_L^I and R_H^I, which are respectively

$$\mu_5 \geq p(\mu_3 - 1)$$

$$\mu_4 \geq (1-p)(\mu_3 - 1)$$

These show that the pledgeability constraints bind if and only if $\mu_3 > 1$.

Now we need only show that $\mu_3 > 1$ to establish the claim. To show this, we differentiate the Lagrangian with respect to I, which yield

$$\mu_3 \leq [p(\rho_1 - \rho_L) + \mu_5(\rho_0 - \rho_L)]\lambda_L + [(1-p)(\rho_1 - \rho_H) + \mu_4(\rho_0 - \rho_H)]\lambda_H$$

which holds with equality when $I > 0$. Since we have $\lambda_H = 0$ whenever $(1 - p)(\rho_1 - \rho_H) + \mu_4(\rho_0 - \rho_H) < 0$, the term $[(1-p)(\rho_1 - \rho_H) + \mu_4(\rho_0 - \rho_H)]\lambda_H$ is non-negative. Since $\lambda_L = 1$ and since by assumption $p(\rho_1 - \rho_L) > 1$, the term $[p(\rho_1 - \rho_L) + \mu_5(\rho_0 - \rho_L)]\lambda_L > 1$. Therefore we have $\mu_3 > 1$ as long as $I > 0$, which implies $\mu_4 > 0$ and $\mu_5 > 0$ and establishes the claim. □

Proof of Proposition 2. We first establish that the constraints

$$I \leq \frac{A - (q-1)\ell}{1 - p(\rho_0 - \rho_L) - (1-p)\lambda_H(\rho_0 - \rho_H)}$$

and $\lambda_H(\rho_H - \rho_0) I \leq \ell$ hold with equality. Consider the Lagrangian

$$L = p(\rho_1 - \rho_L) I + (1-p)\lambda_H(\rho_1 - \rho_H) I + \mu_1[\ell - \lambda_H(\rho_H - \rho_0) I]$$
$$+ \mu_2[A - (q-1)\ell - [1 - p(\rho_0 - \rho_L) - (1-p)\lambda_H(\rho_0 - \rho_H)] I]$$

The first-order condition with respect to ℓ yields

$$\mu_1 \leq (q-1)\mu_2$$

which holds with equality when $\ell > 0$. Since we have $\ell \geq \lambda_H(\rho_H - \rho_0) I$, as long as we have $\lambda_H = 1$ and $I > 0$, we will have $\ell > 0$. We proved in Lemma 1 that (2.11) binds at the solution, so we have $\mu_2 > 0$, which implies that $\mu_1 > 0$ as long as we have $q - 1 > 0$, $\lambda_H = 1$, and $I > 0$. If we have $q - 1 = 0$, then there is no cost of holding unnecessary ℓ, and so the optimal level of ℓ is undetermined. We assume without loss of generality that $\ell = \lambda_H(\rho_H - \rho_0) I$ holds with equality.

Now we substitute this expression for ℓ into the leverage constraint, and derive an expression for I

$$I = \frac{A}{1 - p(\rho_0 - \rho_L) - \lambda_H(q-p)(\rho_0 - \rho_H)}$$

which will hold for $\lambda_H = 1$. This same expression will also hold for $\lambda_H = 0$ and $\ell = 0$ from the leverage constraint with $\ell = 0$, so this will hold in either case.

All that is left is to determine when $\lambda_H = 1$ will be optimal. This will be true as long as

$$I_{\lambda_H=1} \geq pI_{\lambda_H=0}$$

which simplifies to

$$q - 1 \leq \frac{(1-p)\left[1 - p\left(\rho_H - \rho_L\right)\right]}{p\left(\rho_H - \rho_0\right)}$$

$\square$

Proof of Proposition 3. (i) follows directly from Proposition 2.

For (ii), we simply look at the demand for outside assets ℓ at $q = 1$. If $\ell < \bar{\ell}$, then we know from market clearing that equilibrium $q = 0$ because households must hold some of $\bar{\ell}$ in equilibrium. Substituting $q = 1$ into the expression for I in Proposition 2, we find that $I = \frac{A}{1-\rho_0+p\rho_L+(1-p)\rho_H}$. Since firms need to hold at least $\ell = (\rho_H - \rho_0)\, I$ in order to meet both shocks, we find that firm must hold at least $\ell = \frac{(\rho_H-\rho_0)A}{1-\rho_0+p\rho_L+(1-p)\rho_H}$, which is feasible if $\bar{\ell} \geq \frac{(\rho_H-\rho_0)A}{1-\rho_0+p\rho_L+(1-p)\rho_H}$, thus proving the statement.

For (iii), we want to find the range of $\bar{\ell}$ for which all firms meet both shocks and $q > 0$. Since $q > 0$, households do not hold any liquid assets, and so all $\bar{\ell}$ are held by firms. Since all firms meet both shocks and hold liquid assets, we have $\ell = \bar{\ell}$ by market clearing. Then we use $I = \frac{A}{1-p(\rho_0-\rho_L)-(q-p)(\rho_0-\rho_H)}$ and $I = (\rho_H - \rho_0)\,\ell$

from Proposition 2 to solve for implied q. This calculation yields

$$q-1=\frac{A}{\bar{\ell}}-\frac{1-\rho_0+p\rho_L+(1-p)\,\rho_H}{\rho_H-\rho_0}$$

This will be the equilibrium as long as firms are willing to meet both shocks at this level of q. By Proposition 2, firms are willing meet any shocks as long as

$$q-1\leq\frac{(1-p)\left[1-p\left(\rho_H-\rho_L\right)\right]}{p\left(\rho_H-\rho_0\right)}$$

Combined with the condition above, this yields the necessary level of $\bar{\ell}$ for this to be an equilibrium

$$\bar{\ell}\geq\frac{p\left(\rho_H-\rho_0\right)A}{1-p\left(\rho_0-\rho_L\right)}$$

If $\bar{\ell}$ is above this threshold, then $\ell=\bar{\ell}$, $I=\frac{\bar{\ell}}{\rho_H-\rho_0}$, and $\lambda_1=1$.

For (iv), if $\bar{\ell}$ is below this threshold then there is insufficient liquidity for all firms to meet all shocks. However, no firms meeting both shocks would not be an equilibrium, because in this case the available outside liquidity $\bar{\ell}$ would need to be held by households or firms that do not need it, which they would only do if we had $q=0$. But if we have $q=0$, then firms would prefer to buy outside liquidity at this price and meet both shocks. Therefore in equilibrium we must have a fraction of firms meeting both shocks, while the rest only meet the low shock. In order for this to be an equilibrium the firms must be indifferent between these two strategies. From Proposition 2, this will be true if

$$q-1=\frac{(1-p)\left[1-p\left(\rho_H-\rho_L\right)\right]}{p\left(\rho_H-\rho_0\right)}$$

Substituting this q into the expression for I from Proposition 2, we obtain

$$I=\frac{pA}{1-p\left(\rho_0-\rho_L\right)}$$

The corresponding amount of outside liquidity held by each firm must satisfy $\ell = (\rho_H - \rho_0) I$, so we have

$$\ell = \frac{pA(\rho_H - \rho_0)}{1 - p(\rho_0 - \rho_L)}$$

Let the fraction of firms that meet both shocks be ζ. Since $q > 0$, by market clearing the firms that meet both shocks must hold total liquidity $\bar{\ell}$. Therefore we have $\zeta \ell = \bar{\ell}$, and so $\zeta = \frac{\bar{\ell}}{\ell}$, or

$$\zeta = \frac{[1 - p(\rho_0 - \rho_L)]\,\bar{\ell}}{pA(\rho_H - \rho_0)}$$

The firms that meet only the low shock will follow the same strategy as when $p(\rho_H - \rho_L) > 1$, earning the same profits as the firms that meet both shocks. This proves the final statement. $\square$

B.2 Proofs from section 2.3

Proof of Lemma 2. Let the Lagrangian for the optimal contracting problem be

$$\begin{aligned}
L = {} & p\left(\lambda_L \rho_1 I - R_L^I - R_L^1 - \lambda_B R_L^B + \ell\right) + (1-p)\left(\lambda_H \rho_1 I - R_H^I - R_H^1 - \lambda_B R_H^B + \ell\right) \\
& + \mu_1 \left[R_L^1 - \lambda_L \rho_L I\right] + \mu_2 \left[R_H^1 + \lambda_B M - \lambda_H \rho_H I\right] \\
& + \mu_3 \left[p R_L^I + (1-p) R_H^I - I - q\ell + A\right] \\
& + \mu_4 \left[\lambda_L \rho_0 I + \ell - R_L^I - R_L^1 - \lambda_B R_L^B\right] + \mu_5 \left[\lambda_H \rho_0 I + \ell - R_H^I - R_H^1 - \lambda_B R_H^B\right] \\
& + \mu_6 \left[p R_L^B + (1-p) R_H^B - \lambda_B(\pi + (1-p)M)\right]
\end{aligned}$$

The first step is to show that (2.20) and (2.21) bind, meaning period 1 investors are paid just enough to finance the liquidity shock. Differentiating with respect to

R_L^1 and R_H^1 yields

$$\mu_1 \leq p + \mu_4$$

$$\mu_2 \leq (1-p) + \mu_5$$

which hold with equality if $R_L^1 > 0$ or $R_H^1 > 0$ respectively.

From (2.20) we have $R_L^1 \geq \lambda_L \rho_L I$, and so either $R_L^1 > 0$, which implies $\mu_1 = p + \mu_4 > 0$, or $R_L^1 = 0$ in which case $\lambda_L \rho_L I = 0$ (since $\lambda_L \rho_L I \geq 0$). In either case, the constraint (2.20) holds with equality.

From (2.21) we have $R_H^1 + \lambda_B M \geq \lambda_H \rho_H I$. Now we would like to argue as above, but first we must rule out the possibility that we have $R_H^1 = 0$ and $\lambda_B M > 0$, so that (2.21) does not hold with equality. To show that this is not the case, we differentiate the Lagrangian with respect to M, which yields

$$\mu_2 \lambda_B \leq \mu_6 (1-p) M$$

which holds with equality if $M > 0$. From this we can conclude that either (1) $M = 0$, or (2) $\mu_2 > 0$ and $\mu_6 > 0$, or (3) $\mu_2 = 0$ and $\mu_6 = 0$. If (1) or (2) then we are finished then (2.21) holds with equality and we are finished, so all that remains is to show that we cannot have $\mu_2 = \mu_6 = 0$ when $\lambda_B = 1$ and $M > 0$.

Differentiating the Lagrangian with respect to R_L^B and R_H^B, we obtain

$$\mu_6 \leq 1 + \frac{\mu_4}{p}$$

$$\mu_6 \leq 1 + \frac{\mu_5}{1-p}$$

which hold with equality if $R_L^B > 0$ or $R_H^B > 0$ respectively. From this we conclude that either $\mu_6 > 0$ or $R_L^B = R_H^B = 0$. If the latter, then (2.19) would be $\pi + (1-p)M =$

0, which is only possible if $\pi = M = 0$, contradicting our assumption that $M > 0$. Therefore we have $\mu_6 > 0$, and so $\mu_2 > 0$.

Note that the above argument also establishes that (2.19) holds with equality, since either $R_L^B = R_H^B = 0$ and so $\lambda_B(\pi + (1-p)M) = 0$, or else one of R_L^B or R_H^B is strictly positive, in which case $\mu_6 > 0$.

We have now established that (2.20) and (2.21) hold with equality, and so we can directly substitute them into the Lagrangian to obtain

$$
\begin{aligned}
L = {} & p\left(\lambda_L \rho_1 I - R_L^I - \lambda_L \rho_L I - \lambda_B R_L^B + \ell\right) + (1-p)\left(\lambda_H \rho_1 I - R_H^I - \lambda_H \rho_H I + \lambda_B (M - R_H^B) + \ell\right) \\
& + \mu_3 \left[p R_L^I + (1-p) R_H^I - I - q\ell + A\right] + \mu_4 \left[\lambda_L \rho_0 I + \ell - R_L^I - \lambda_L \rho_L I - \lambda_B R_L^B\right] \\
& + \mu_5 \left[\lambda_H \rho_0 I + \ell - R_H^I - \lambda_H \rho_H I + \lambda_B (M - R_H^B)\right] + \mu_6 \left[p R_L^B + (1-p) R_H^B - \lambda_B (\pi + (1-p)M)\right]
\end{aligned}
$$

We next establish that $\lambda_L = 1$. Since $\lambda_L, \lambda_H \in \{0, 1\}$, we can differentiate the Lagrangian with respect to each λ, and conclude that if this derivative is non-negative then that $\lambda = 1$. The derivatives of the Lagrangian with respect to λ_L and λ_H are respectively

$$[p(\rho_1 - \rho_L) + \mu_4(\rho_0 - \rho_L)]\, I$$

$$[(1-p)(\rho_1 - \rho_H) + \mu_5(\rho_0 - \rho_H)]\, I$$

Since by assumption $\rho_1 > \rho_L$ and $\rho_0 > \rho_L$, the first expression is positive for $I > 0$, and so we have $\lambda_L = 1$. The second expression may not be positive because $\rho_0 < \rho_H$, so we conclude that $\lambda_H = 1$ if and only if $[(1-p)(\rho_1 - \rho_H) + \mu_5(\rho_0 - \rho_H)]\, I \geq 0$.

Differentiating the Lagrangian with respect to R_L^I and R_H^I yields

$$\mu_4 \geq p\left(\mu_3 - 1\right)$$

$$\mu_5 \geq \left(1-p\right)\left(\mu_3 - 1\right)$$

which hold with equality when $R_L^I > 0$ or $R_H^I > 0$ respectively. Therefore (2.22) and (2.23) hold with equality as long as $\mu_3 > 1$.

Now to finish proving the proposition, it is enough to establish that $\mu_3 > 1$, since this would prove that (2.22), (2.23) and (2.18) hold with equality. Differentiating the Lagrangian with respect to I, we obtain

$$\mu_3 \geq p\left(\rho_1 - \rho_L\right) + \mu_4\left(\rho_0 - \rho_L\right) + \lambda_H\left[\left(1-p\right)\left(\rho_1 - \rho_H\right) + \mu_5\left(\rho_0 - \rho_H\right)\right]$$

By assumption, we have $p\left(\rho_1 - \rho_L\right) > 1$, so we have

$$\mu_3 > 1 + \mu_4\left(\rho_0 - \rho_L\right) + \lambda_H\left[\left(1-p\right)\left(\rho_1 - \rho_H\right) + \mu_5\left(\rho_0 - \rho_H\right)\right]$$

Since $\mu_4 \geq 0$ and $\rho_0 > \rho_L$, the term $\mu_4\left(\rho_0 - \rho_L\right) \geq 0$. As we found above, either $\lambda_H = 0$ or $\left(1-p\right)\left(\rho_1 - \rho_H\right) + \mu_5\left(\rho_0 - \rho_H\right) \geq 0$, so the term $\lambda_H\left[\left(1-p\right)\left(\rho_1 - \rho_H\right) + \mu_5\left(\rho_0 - \rho_H\right)\right] \geq 0$. Therefore we have $\mu_3 > 1$, and therefore $\mu_4 > 0$ and $\mu_5 > 0$. This proves that (2.22), (2.23) and (2.18) hold with equality. $\square$

Proof of Proposition 4. I solve this problem in two stages. In the first stage, I suppose that $\lambda_H = 1$ and solve for the optimal way to finance investment. In the second stage I check whether it is in fact optimal to meet both shocks.

Assuming $\lambda_H = 1$, the problem of an individual firm given (M, π, q) is

$$\max_{I,\ell,\lambda_B} \quad \{I\}$$
$$\text{s.t.} \quad I \leq \frac{A - (q-1)\,\ell - \lambda_B \pi}{\chi_1}$$
$$I \leq \frac{\ell + \lambda_B M}{\rho_H - \rho_0}$$

where $I, \ell \geq 0$ and $\lambda_B \in \{0, 1\}$. The lagrangian of this problem is

$$L = I + \lambda\left(\frac{A - (q-1)\,\ell - \lambda_B \pi}{\chi_1} - I\right) + \mu\left(\frac{\ell + \lambda_B M}{\rho_H - \rho_0} - I\right)$$

The first-order condition with respect to I is

$$\lambda + \mu \geq 1$$

which holds with equality if $I > 0$. This implies that one of λ or μ is greater than 0.

The first-order condition with respect to ℓ is

$$\frac{\mu}{\rho_H - \rho_0} \leq \frac{\lambda(q-1)}{\chi_1}$$

which holds with equality if $\ell > 0$. Since one of λ or μ is strictly positive, this implies that $\lambda > 0$. Moreover, $\mu = 0$ is only possible if either $\ell = 0$ or $q = 1$.

First suppose $q = 1$. Then since $\lambda > 0$, we have

$$I = \frac{A - \lambda_B \pi}{\chi_1}$$

and since $\pi \geq 0$, this term attains its maximum at $\lambda_B = 0$. This level of investment is feasible, and corresponds to $\ell = (\rho_H - \rho_0)I$.

Next suppose that $q > 1$. Then $\mu > 0$ unless $\ell = 0$, which only occurs when firms meet all liquidity shocks through bank financing. Differentiating the Lagrangian with respect to λ_B yields

$$\frac{\partial L}{\partial \lambda_B} = \frac{\mu M}{\rho_H - \rho_0} - \frac{\lambda \pi}{\chi_1}$$

Firms will only choose $\lambda_B = 1$ if $\partial L / \partial \lambda_B \geq 0$. But since $\lambda > 0$, this is only possible if either $\mu > 0$ or $\pi = 0$.

First suppose that $\pi = 0$. If $\ell = 0$ also, then investment satisfies $I = A/\chi_1$, and the liquidity constraint implies $M \geq (\rho_H - \rho_0)A/\chi_1$. Now we assumed that $M \leq (\rho_H - \rho_0)A/\chi_1$, so then either M is exactly at this level, in which case this is the solution, or else we have arrived at a contradiction and $\ell > 0$.

In the latter case, $\ell > 0$ implies $\mu > 0$, and we have

$$\frac{\mu}{\rho_H - \rho_0} = \frac{\lambda(q-1)}{\chi_1}$$

Substituting this into the FOC wrt λ_H yields

$$\frac{\partial L}{\partial \lambda_B} = \frac{\lambda}{\chi_1}\left[(q-1)M - \pi\right]$$

which is positive when $\pi/M \leq q - 1$. Moreover, $\mu > 0$ implies that

$$(\rho_H - \rho_0)I = \ell + \lambda_B M$$

Combining this with the leverage constraint arising from pledgeability, we obtain

$$I = \frac{A - \lambda_B\left[\pi - (q-1)M\right]}{\chi_1 + (q-1)(\rho_H - \rho_0)}$$

Further observe that this expression still holds for investment in the $q = 1$ and $\ell = 0$ cases considered above.

Now it only remains to determine whether it is optimal to meet the high shock. When the high shock is not met, firm expected profits are

$$\frac{pA}{1 - p(\rho_0 - \rho_L)}$$

and when both shocks are met expected profits are I as calculated above. The expression given in the proposition is a comparison of these values. $\square$

Proof of Proposition 5. First consider the case $\tilde{D}(q-1) \leq \bar{D}(q-1)$. We begin by arguing that $\tilde{D}(\vartheta)$ is increasing in ϑ. Recall that $\tilde{D}(\vartheta)$ is implicitly defined by

$$\vartheta\tilde{D} = C(\tilde{D}, K)$$

Implicitly differentiating this expression yields

$$\tilde{D}'(\vartheta) = \frac{\tilde{D}}{C_1(\tilde{D}, K) - C(\tilde{D}, K)/\tilde{D}}$$

which is positive by our assumptions $C(0, K) = 0$, $C_1 > 0$, and $C_{11} > 0$.

Now if any bank sets $\vartheta > q - 1$, it sells no credit lines, whereas any bank that sets $\vartheta \leq q - 1$ will sell some quantity of credit lines that cannot exceed $\tilde{D}(\vartheta)$, since this is the quantity of D at which its agency costs are binding. Since $\tilde{D}(\vartheta)$ is increasing in ϑ, total credit lines sold by all firms cannot exceed $\tilde{D}(q-1)$. Since this quantity is by assumption no greater than total demand for credit lines $\bar{D}(\vartheta)$, any bank that sets $\vartheta \leq q - 1$ will sell up to its agency costs $\tilde{D}(\vartheta)$. Thus profits are $\vartheta\tilde{D}(\vartheta)$ for $\vartheta \leq q - 1$, and 0 for $\vartheta > q - 1$, and since $\tilde{D}(\vartheta)$ is increasing in ϑ, the solution is $\vartheta = q - 1$.

Next consider the case $\tilde{D}(q-1) < \bar{D}(q-1)$. $\bar{D}(\vartheta) = \frac{(\rho_H - \rho_0)A}{\chi_1 + \vartheta(\rho_H - \rho_0)}$ is the maximum demand for credit lines from firms given a price of liquidity equal to ϑ. Clearly $\bar{D}(\vartheta)$ is decreasing in ϑ.

Given some distribution $\Phi(\cdot)$ of ϑ across banks, with $\Phi(q-1) = 1$, firms will purchase credit lines from the banks with the lowest ϑ until agency costs bind for those banks, and then move on to the next highest, and so on. Thus for a given ϑ, the total share of firms that have purchased credit lines from banks at a price less than or equal to this price will be

$$S(\vartheta) = \int_0^{\vartheta} \frac{\bar{D}(\vartheta')}{\tilde{D}(\vartheta')} d\Phi(\vartheta')$$

Firms purchase credit lines from banks up to a cutoff price $\bar{\vartheta}$. Since at ϑ the share of firms that have purchased credit lines is 1, this cutoff is defined by

$$\lim_{\vartheta \to \bar{\vartheta}^-} S(\vartheta) \leq 1 \leq \lim_{\vartheta \to \bar{\vartheta}^+} S(\vartheta)$$

where both inequalities hold with equality if there is not a point mass of firms at $\bar{\vartheta}$.

If there is a point mass $\bar{m}$ of banks at $\bar{\vartheta}$, then demand will be divided equally among them. Each will sell credit lines to a share $\phi = [1 - \lim_{\vartheta \to \bar{\vartheta}^-} S(\vartheta)]/\bar{m}$, of firms, and each firm purchases a credit line of size $\bar{D}(\bar{\vartheta})$.

Since $\bar{D}(\cdot)$ is decreasing in ϑ whereas $\tilde{D}(\cdot)$ is increasing, the term $\frac{\bar{D}(\vartheta)}{\tilde{D}(\vartheta)}$ is decreasing in ϑ. Moreover, because we are considering the case $\bar{D}(q-1) > \tilde{D}(q-1)$, this fraction is strictly greater than 1 for all $\vartheta \leq 1$. Therefore any bank that sets $\vartheta = q-1$ will receive a measure of customers that is strictly less than 1 (and may be 0).

Now consider the problem of a bank setting its price. Clearly the bank will never choose $\vartheta > \bar{\vartheta}$, since this yields no sales and no profits. If the bank chooses to just undersell the cutoff, it receives profits of just under $\bar{\vartheta}\tilde{D}(\bar{\vartheta})$. Finally, if the bank chooses $\vartheta = \bar{\vartheta}$, it receives profits $\phi\bar{\vartheta}\bar{D}(\bar{\vartheta})$.

Since all banks face the same demand curve which implies a unique optimal price, they will all choose the same ϑ. This point must satisfy $\tilde{D}(\bar{\vartheta}) = \phi\bar{D}(\bar{\vartheta})$. Since all banks choose the same price, $\phi = 1$ and this expression becomes $\tilde{D}(\bar{\vartheta}) = \bar{D}(\bar{\vartheta})$, which implicitly defines the price level set by banks. □

Proof of Proposition 6. Equilibrium is defined by firm behavior given by Proposition 4, bank pricing given by Proposition 5, and the market clearing condition $\left(\zeta\ell - \bar{\ell}\right)(q-1) = 0$, together with $\zeta\ell \leq \bar{\ell}$ and $q \geq 1$. Market clearing requires that when $q > 1$, firms must hold all available outside liquidity $\bar{\ell}$, since households will only hold ℓ if $q = 1$.

First suppose that $q = 1$. This immediately implies $\zeta = 1$. From bank behavior, we know that $C(M,K)/M \leq q - 1 = 0$, which is only possible when $M = 0$. Then from firm behavior we have $I = I_1(0) = A/\chi_1$. By assumption, $I_1(0) > pI_0$, so $\lambda_H = 1$ for all firms. Therefore firms meet all shocks by holding outside liquidity, and so total holdings must satisfy $\ell = (\rho_H - \rho_0)I_1(0)$. This will be the equilibrium as long as there is sufficient outside liquidity, i.e. $\bar{\ell} \geq (\rho_H - \rho_0)I_1(0)$.

Next suppose that $\bar{\ell} < (\rho_H - \rho_0)I_1(0)$, so that $q > 1$ in equilibrium. This implies that all outside liquidity is held by firms, and so $\zeta\ell = \bar{\ell}$. Now I claim that banks set $\vartheta = q - 1$. First suppose that $\bar{\ell} > 0$. From Proposition 5, either $\vartheta = q - 1$, or else $\vartheta < q-1$ and $\bar{M}(\vartheta) = \tilde{M}(\vartheta)$. In the latter case, all firms finance their desired liquidity holdings by holding credit lines, and so $\ell = 0$. Since $\bar{\ell} > 0$, we have $\zeta\ell < \bar{\ell}$, which violates market clearing given $q > 1$. Now if $\bar{\ell} = 0$, q may be set to any level as long as $\vartheta \leq q - 1$, so we can set them equal without loss of generality.

Since $\vartheta = q-1$, a firm that meets a high liquidity shock will choose investment

$$I = I_1(q-1) = \frac{A}{\chi_1 + (q-1)(\rho_H - \rho_0)}$$

and will be indifferent between purchasing a credit line and financing liquidity shocks entirely through holding outside liquidity. We assume without loss of generality that each firm that meets the high shock holds a credit line of size M. Then $\tilde{M} = \zeta M$ are total credit lines sold by banks. From Proposition 5, $\tilde{M}$ satisfies $\vartheta\tilde{M} = C(\tilde{M}, K)$. Combined with $\vartheta = q - 1$, this implies $q - 1 = C(\tilde{M}, K)/\tilde{M}$.

Now we consider two cases. First suppose that $\zeta = 1$, so that all firms meet the high shock. Then $\tilde{M} = M$ and $\ell = \bar{\ell}$. Then from firm behavior

$$(\rho_H - \rho_0)I_1(C(M, K)/M) = M + \bar{\ell}$$

This expression implicitly defines unique $M = M_1$. To prove this, frist note that by L'Hopital's rule, $\lim_{M\to 0} C(M, K)/M = 0$, since $C(0, K) = 0$ and $\lim_{M\to 0} C_1(M, K) = 0$. Then as $M \to 0$, the right-hand side converges to ℓ while the left-hand side converges to $(\rho_H - \rho_0)I_1(0)$. Since we are considering the case that $(\rho_H - \rho_0)I_1(0) > \bar{\ell}$, the left-hand side is greater than the right at $M \to 0$. Next we observe that since $I_1(q - 1)$ is strictly decreasing in $q - 1$, and since $C(M, K)/M$ is strictly increasing in M, the left-hand side is strictly decreasing in M whereas the right-hand side is strictly decreasing. Therefore if an M that makes the equation holds, it is unique. Finally, by assumption $\lim_{M\to\infty} C(M,K) = \infty$ and $\lim_{M\to\infty} C_1(M,K) = \infty$, so again by L'Hopital's rule $\lim_{M\to\infty} C(M,K)/M = \infty$. Since $\lim_{q\to\infty} I_1(q-1) = 0$, this implies that as $M \to \infty$ the left-hand side converges to 0 while the right-hand side converges to ∞. Thus M_1 exists.

The equilibrium will satisfy $M = M_1$ as long as the implied I is at least as large as pI_0. This condition can be written as

$$I_1(C(M_1,K)/M_1) \geq pI_0$$

Since $C(M,K)/M$ is strictly increasing in M, and $I_1(q-1)$ is strictly decreasing in $q-1$, this expression is equivalent to the condition $M_1 \leq M_2$, where M_2 satisfies $I_1(C(M_2,K)/M_2) = pI_0$. Note that M_2 is uniquely defined since by the arguments above $I_1(C(M,K)/M)$ is strictly decreasing in M, goes to 0 as $M \to \infty$, and goes to $I_1(0) > pI_0$ as $M \to 0$.

Suppose that instead $M_2 < M_1$. Then we have $\zeta < 1$. Moreover, since $\tilde{M}$ satisfies $C(\tilde{M},K)/\tilde{M} = q-1 > 0$, we have $\tilde{M} > 0$, and since $\tilde{M} = \zeta M$, we must have $\zeta > 0$. Thus $\zeta \in (0,1)$, and so firms must be indifferent between meeting the high shock or the low shock. Thus firms that meet the high shock choose $I = pI_0$. This implies that $q-1$ must satisfy

$$I_1(q-1) = pI_0$$

which uniquely defines a value of $q > 1$. Then $\tilde{M}$ is determined by $q - 1 = C(\tilde{M},K)/\tilde{M}$, which is what we're calling M_2. Now we know that total liquidity used satisfies $\tilde{L} = M_2 + \bar{\ell}$, and liquidity used by each firm that meets a high shock is $L = M + \ell = (\rho_H - \rho_0)I$. Therefore the fraction of firms that meet the high shock is

$$\zeta = \frac{M_2 + \bar{\ell}}{(\rho_H - \rho_0)I}$$

$\square$

B.3 Proofs from section 2.4

Proof of Proposition 7. At an interior equilibrium the supply of liquidity is

$$L^s(q-1) = M^s(q-1) + \bar{\ell}$$

where $M^s(q-1)$ is implicitly defined by $C(M,K)/M = q-1$, and demand for liquidity is

$$L^d(q-1) = \frac{(\rho_H - \rho_0)A}{\chi_1 + (q-1)(\rho_H - \rho_0)}$$

Equilibrium q is defined implicitly by $L^s(q-1) = L^d(q-1)$, and at the equilibrium we have $q - 1 = C(M,K)/M$ and $L^s = L^d = M + \bar{\ell}$.

The elasticity of liquidity supply at equilibrium is

$$\eta = \left(\frac{C/M}{C_M - C/M}\right)\left(\frac{M}{M + \bar{\ell}}\right)$$

and the elasticity of liquidity demand is

$$\epsilon = \left(M + \bar{\ell}\right)\frac{C/M}{A}$$

Now consider a marginal change in $\bar{\ell}$. Totally differentiating $L^s = L^d$ yields $L^s_q dq + d\bar{\ell} = L^d_q dq$, which implies

$$\frac{dq}{d\bar{\ell}} = -\left(\frac{C_M - C/M}{M}\right)\left(\frac{\eta}{\eta + \epsilon}\right)$$

The total change in $L = M + \bar{\ell}$ is $L_q^s dq$, since the demand curve is not affected directly by the change in $\bar{\ell}$. This implies $1 + \frac{dM}{d\bar{\ell}} = L_q^d \frac{dq}{d\bar{\ell}}$, or

$$\frac{dM}{d\bar{\ell}} = -\left(\frac{\eta}{\eta + \epsilon}\right)$$

Differentiating the expression for investment $I = (\rho_H - \rho_0)^{-1}\left(M + \bar{\ell}\right)$ yields $\frac{dI}{d\bar{\ell}} = (\rho_H - \rho_0)^{-1}\left(1 + \frac{dM}{d\bar{\ell}}\right)$, or

$$\frac{dI}{d\bar{\ell}} = (\rho_H - \rho_0)^{-1}\left(\frac{\epsilon}{\eta + \epsilon}\right)$$

□

Proof of Proposition 8. Equilibrium satisfies $L^s(q-1) = L^d(q-1)$, where L^s and L^d are as defined in the proof of Proposition 7. Then a marginal change in K produces a change in q equal to

$$\frac{dq}{dK} = \frac{L_K^d}{L_q^s - L_q^d}$$

We can calculate L_K^d by applying the implicit function theorem to $C(M, K) = (q-1)M$. From this we find $L_K^d = -\frac{C_K}{M} L_q^d$, which simplifies to

$$\frac{dq}{dK} = \left(\frac{\eta}{\eta + \epsilon}\right)\frac{C_K}{M}$$

Since K does not appear in the definition of L^d, we have $\frac{dM}{dK} = D_q \frac{dq}{dK}$, or

$$\frac{dM}{dK} = -\left(\frac{M}{C_M - C/M}\right)\left(\frac{\epsilon}{\epsilon + \eta}\right)\frac{C_K}{M}$$

and since $I = (\rho_H - \rho_0)^{-1}(M + \bar{\ell})$, we have

$$\frac{dI}{dK} = (\rho_H - \rho_0)^{-1}\frac{dM}{dK}$$

□

B.4 Proofs from section 2.5

Proof of Proposition 9. Since firms prefer the unconstrained optimum (UO) to any other allocation, it is sufficient to show that UO is feasible under the described set of transfers. The UO requires firms to invest $I = A + H_0$. Firms have an endowment of A, and so given the transfer H_0 from households, this level of investment is feasible. The UO also requires meeting both shocks. Since firms have no debt from period 0, they have sufficient pledgeable funds to borrow up to $\rho_0 I$ from households in period 1. Since $\rho_0 > \rho_L$, firms that experience the low shock are able to borrow enough to meet the shock. Firms that receive the high shock receive a transfer of $(\rho_H - \rho_0)I$ from households. Since these firms can borrow an additional $\rho_0 I$ from households, they are able to meet the high shock. Thus the unconstrained optimum is feasible under the described transfers. The Planner may then transfer $H_0 + (1-p)(\rho_H - \rho_0)(A + H_0)$ from firms to households in period 2 to effect a Pareto improvement. $\square$

Proof of Proposition 10. We solve for the constrained optimal allocation in two steps. First we fix the fraction ζ of firms that meet the high shock and define welfare $W(\zeta)$. Then we solve for ζ that maximizes $W(\zeta)$.

First consider the optimal choices of firms that choose $\lambda_H = 0$. This choice reduces to maximizing $\int_i p(\rho_1 - \rho_L)I_i$ s.t. $I_i \leq A/\left[1 - p(\rho_0 - \rho_L)\right]$. This yields optimal investment $I_i = A/\left[1 - p(\rho_0 - \rho_L)\right]$, and their welfare is

$$W_{\lambda=0} = \frac{R_0 A}{\chi_0}$$

where $R_0 = p(\rho_1 - \rho_L) - 1$ and $\chi_0 = 1 - p(\rho_0 - \rho_L)$.

Next consider the problem of firms that choose $\lambda_H = 1$. The Lagrangian of the maximization problem in Definition 3 with ζ given is

$$\begin{aligned} L = &\int_i R_1 I_i + \int_i \mu_1^i \left(\frac{A - \pi_i - (q-1)\ell_i}{\chi_1} - I_i \right) \\ &+ \int_i \mu_2^i \left(M_i + \ell_i - I_i(\rho_H - \rho_0) \right) + \mu_3 \left(\bar{\ell} - \int_i \ell_i \right) \\ &+ \mu_4 \left(\int_i \pi_i - C(\int_i M_i, K) \right) + \mu_5 (q-1) \end{aligned}$$

together with non-negativity constraints. The first-order conditions of the problem are:

$$\begin{aligned} \frac{\partial L}{\partial I_i} &= R_1 - \mu_1^i - \mu_2^i(\rho_H - \rho_0) \leq 0 \\ \frac{\partial L}{\partial M_i} &= \mu_2^i - C_1 \mu_4 \leq 0 \\ \frac{\partial L}{\partial \pi_i} &= \mu_4 - \frac{\mu_1^i}{\chi_1} \leq 0 \\ \frac{\partial L}{\partial \ell_i} &= \mu_2^i - \frac{\mu_1^i (q-1)}{\chi_1} - \mu_3 \leq 0 \\ \frac{\partial L}{\partial q} &= \mu_5 - \frac{\mu_1^i \ell_i}{\chi_1} \leq 0 \end{aligned}$$

Now I establish several facts about the solution.

First I claim that $\mu_1^i > 0$ for all i. I argue in two steps. First I show that at least one of $\mu_1^i, \mu_2^i > 0$. This follows from the expression for $\partial L / \partial I_i$ since $R_1 > 0$, $\rho_H > \rho_0$, and $\mu_1^i, \mu_2^i \geq 0$. Next I argue that $\mu_1^i > 0$. If this were not true, then by the previous claim $\mu_2^i > 0$. Then from the expression for $\partial L / \partial M_i$ it follows that $\mu_4 > 0$, and from the expression for $\partial L / \partial \pi_i$ it follows that $\mu_1^i > 0$, which is a contradiction.

Next I show that $q = 1$. From the expression for $\partial L / \partial q$ either $\ell_i = 0$ for all i, or else $\mu_5 > 0$. The latter implies $q = 1$, and if the former the value of q does not enter into the problem apart from the restriction that $q \geq 1$, and so we can choose $q = 1$ without loss of generality.

Next I argue that $\int_i \pi_i = C(\int_i M_i, K)$. This follows from the expression for $\partial L_i/\partial \pi_i$. If $\pi_i > 0$ for any i then $\mu_4 > 0$ and the claim follows. If $\pi_i = 0$ for all i, then the bank incentive compatability constraint becomes $C(\int_i M_i, K) \leq 0$, and since we have $C(\int_i M_i, K) \geq 0$ from non-negativity of M_i, the claim follows. Moreover, this argument together with the expression for $\partial L/\partial M_i$ implies that either $\mu_2^i > 0$, or $M_i, \pi_i = 0$ for all i.

Now observe that since total welfare satisfies $\int_i \left(\frac{A-\pi_i}{\chi_1}\right) R_1$, any feasible reallocation of π_i, M_i, ℓ_i across i will not affect total welfare. Thus we can WLOG restrict attention to the symmetric case where $M_i = M$, $\pi_i = \pi$, and $\ell_i = \ell$.

Combining all of the above, we can rewrite the maximization problem as

$$\begin{aligned} &\max \left\{ \frac{A - C(\zeta M, K)/\zeta}{\chi_1} \right\} \\ &\text{s.t.} \quad \frac{A - C(\zeta M, K)/\zeta}{\chi_1} \leq \frac{M + \ell}{\rho_H - \rho_0} \\ &\qquad \zeta \ell \leq \bar{\ell},\ M \geq 0 \end{aligned}$$

Clearly if $M = 0$ is feasible, the solution is $M = 0$ and $I = A/\chi_1$. This will be feasible only when $\frac{A}{\chi_1} \leq \frac{\bar{\ell}/\zeta}{\rho_H - \rho_0}$, which we can write as a threshold $\bar{\ell} \geq \zeta(\rho_H - \rho_0)\frac{A}{\chi_1}$.

Combining expressions, we obtain a single expression that implicitly defines $I_{\lambda=1}(\zeta)$:

$$\chi_1 I = A - \frac{C\left(\max(\zeta(\rho_H - \rho_0)I - \bar{\ell}, 0), K\right)}{\zeta}$$

and the welfare of firms that meet the high shock is $W_{\lambda=1}(\zeta) = R_1 I_{\lambda=1}(\zeta)$.

Now we turn to the determination of optimal ζ. We can now express total welfare as

$$W(\zeta) = \zeta W_{\lambda=1}(\zeta) + (1-\zeta) W_{\lambda=0}$$

Marginal welfare satisfies

$$W_\zeta = R_1 I_{\lambda=1} - \frac{R_0 A}{\chi_0} + \zeta R_1 \frac{dI_{\lambda=1}}{d\zeta}$$

First observe that if ζ is small enough that $\zeta(\rho_H - \rho_0)\frac{A}{\chi_1} \leq \bar{\ell}$, $I_{\lambda=1} = A/\chi_1$ and $W_\zeta = \frac{R_1 A}{\chi_1} - \frac{R_0 A}{\chi_0}$. Now we show that this expression is strictly positive under our assumptions. Recall that we assumed it was always profitable to meet the high shock in equilibrium when the liquidity premium was 0. This assumption can be written as $1/\chi_1 > p/\chi_0$, and is equivalent to $p(\rho_H - \rho_L) < 1$. Then

$$\begin{aligned}
\frac{R_1}{\chi_1} - \frac{R_0}{\chi_0} > \frac{pR_1}{\chi_0} - \frac{R_0}{\chi_0}& \\
&= \frac{1}{\chi_0}\left[p\left(\rho_1 - p\rho_L - (1-p)\rho_H - 1\right) - \left(p(\rho_1 - \rho_L) - 1\right)\right] \\
&= \frac{1-p}{\chi_0}\left[1 - p\left(\rho_H - \rho_L\right)\right] > 0
\end{aligned}$$

Thus for ζ sufficiently low, $W_\zeta > 0$.

Now we show that W_ζ is strictly decreasing in ζ for all ζ above this level. Applying the implicit function theorem to our expression for $I_{\lambda=1}(\zeta)$, we obtain

$$\frac{dI}{d\zeta} = \frac{C/\zeta - (\rho_H - \rho_0)C_1 I}{\zeta\left[(\rho_H - \rho_0)C_1 + \chi_1\right]}$$

which implies

$$W_\zeta = \frac{R_1 A}{(\rho_H - \rho_0)C_1 + \chi_1} - \frac{R_0 A}{\chi_0}$$

In the expression, ζ only appears in the argument of C_1 as ζI. As ζ increases, the term ζI increases:

$$\frac{d(\zeta I)}{d\zeta} = I + \zeta\frac{dI}{d\zeta} = \frac{\zeta\chi_1 I + C}{\zeta\left[(\rho_H - \rho_0)C_1 + \chi_1\right]} > 0$$

and as ζI increases, $C_1\left(\zeta(\rho_H - \rho_0)I - \bar{\ell}, K\right)$ decreases, since we assumed $C_{11} < 0$. Therefore $W_{\zeta\zeta} < 0$.

Since for some $\zeta > 0$ we have $W_\zeta > 0$, and since W_ζ is strictly decreasing in ζ, either $W_\zeta = 0$ for some $\zeta \leq 1$, or else $W_\zeta(1) > 0$. Thus optimal ζ is defined by

$$\frac{R_1}{(\rho_H - \rho_0)C_1\left(\zeta(\rho_H - \rho_0)I - \bar{\ell}, K\right) + \chi_1} \geq \frac{R_0}{\chi_0}$$

which will hold with equality unless $\zeta = 1$. By setting $\zeta = 1$ in this expression, we can solve for the cutoff level $\hat{\ell}$ below which $\zeta < 1$ will hold. Then this expression with equality implicitly defines optimal ζ for $\bar{\ell} < \hat{\ell}$. □

B.5 Proofs from section 2.6

Proof of Proposition 11. In turn,

(i) Suppose not. Then $RI - C(x)$ is maximized at $x = 0$. Let $f(x) = RI|_x$. Then since the function $f(x) - C(x)$ achieves a maximum at $x = 0$, f and C must satisfy

$$f(0) - C(0) \geq f(h) - C(h)$$

for every $h > 0$. Rearranging and dividing by h, we can write this as

$$\frac{C(h) - C(0)}{h} \geq \frac{f(h) - f(0)}{h}$$

this inequality will be preserved under taking limits

$$\lim_{h\to 0}\frac{C(h)-C(0)}{h}\geq \lim_{h\to 0}\frac{f(h)-f(0)}{h}$$

which is just the definition of the derivatives of C and f at $x=0$. Therefore we have

$$C'(0)\geq f'(0)$$

Since we have $C'(0)=0$, this implies that $f'(0)=R\left.\frac{dI}{dx}\right|_{x=0}\leq 0$. But from Proposition 7 we have $\frac{dI}{dx}>0$ for $\bar{\ell}>0$, so this is a contradiction.

(ii) Any interior solution satisfies the optimality condition

$$R\frac{dI}{dx}=D'(x)$$

which is a zero of the function

$$f(\cdot)=R\frac{dI}{dx}-D'(x)=0$$

The partial derivatives of f are

$$\begin{aligned}\frac{\partial f}{\partial x}&=R\frac{d^2I}{dx}-D''(x)\\ \frac{\partial f}{\partial \bar{\ell}}&=R\frac{d^2I}{dx}\\ \frac{\partial f}{\partial K}&=R\frac{d^2I}{dxdK}\end{aligned}$$

making use of the fact that $\frac{d^2I}{d\ell dx}=\frac{d^2I}{dx^2}$. Since this is a unique maximum of a continuously differentiable function the implicit function theorem is valid. Therefore we have

$$\begin{aligned}\frac{dx}{d\bar{\ell}}&=-\frac{R\frac{\partial^2 I}{\partial x^2}}{R\frac{\partial^2 I}{\partial x^2}-D''(x)}\\ \frac{dx}{dK}&=-\frac{R\frac{\partial^2 I}{\partial x\partial K}}{R\frac{\partial^2 I}{\partial x^2}-D''(x)}\end{aligned}$$

(iii) If the point x^* is an interior solution, then it must also be a local maximum. At a local maximum of a global function, the function is locally concave, meaning that the second derivative is negative. Here the function we are maximizing is $RI - D(x)$. Since this function is twice continuously differentiable, it will be concave at the point x^* iff the second derivative $R\frac{\partial^2 I}{\partial x^2} - D''(x) < 0$, which immediately implies the given condition.

(iv) Since we have $R\frac{\partial^2 I}{\partial x^2} - D''(x) < 0$, the denominators in the expressions for $dx/d\bar{\ell}$ and dx/dK are negative. Therefore each expression will be positive iff the numerator is positive. Therefore we obtain the given statement. □

Proof of Proposition 12. To prove (i), we can simply observe that x^* is implicitly defined by $R\phi L_x = D_x$, and that $L_x = (1 + \eta/\epsilon)^{-1}$. Therefore a higher value of η/ϵ at every value of x implies a higher L_x at every level of x. Thus at the old x^* we have $L_x > D_x$, and so it is optimal to increase x from that point, implying a higher value of x^*. Note that any lower value of x will not be the optimum, because at the old value of η/ϵ these points had lower cumulative welfare. This implies that integrating $RI_x - D_x$ between that earlier point and the old x^* yielded a positive number, and since we have now strictly increased I_x this integral has increased, and thus must still be positive. Thus the cumulative net benefit of raising x to at least x^* must be greater than before.

For (ii), we again use $L_z = (1 + \eta/\epsilon)^{-1}$. Let $\psi = \eta/\epsilon$ be the ratio of elasticities. We want to figure out what happens to ψ when we increase z. We can write this derivative as

$$\frac{\psi_z}{\psi} = \frac{\eta_z}{\eta} - \frac{\epsilon_z}{\epsilon}$$

Now we derive expressions for both terms. Using $\eta = \eta_M\left(\frac{L-z}{L}\right)$, we derive

$$\frac{\eta_z}{\eta} = \frac{\frac{d}{dz}(\eta_M)}{\eta_M} - \left(\frac{1}{L-z}\right)\left(\frac{(1+\psi)L-z}{(1+\psi)L}\right)$$

and using the expression $\epsilon = \frac{L\vartheta}{A}$, we derive

$$\frac{\epsilon_z}{\epsilon} = -\frac{1}{L}\left(\frac{1-\epsilon}{\eta+\epsilon}\right)$$

Combining these two expressions, we obtain

$$\frac{\psi_z}{\psi} = \frac{\frac{d}{dz}(\eta_M)}{\eta_M} - \left(\frac{1}{L-z}\right)\left(\frac{(1+\psi)L-z}{(1+\psi)L}\right) + \frac{1}{L}\left(\frac{1-\epsilon}{\eta+\epsilon}\right)$$

and after a bit of work, we obtain the given expression:

$$\frac{\psi_z}{\psi} = \frac{\frac{d}{dz}(\eta_M)}{\eta_M} + \frac{1-2\epsilon-\eta_M}{L(\epsilon+\eta)}$$

For (iii), first observe that if bank liquidity supply is isoelastic, then η_M is constant, and so $\frac{d}{dM}\eta_M = 0$. Then since $L(\epsilon+\eta) > 0$, for $L_{zz} < 0$ we need $\eta_M + 2\epsilon > 1$. Since $\epsilon = \vartheta L/A$, and at an interior solution we have $\vartheta > 0$, it follows that $\epsilon > 0$, and so $\eta_M \geq 1$ implies that $L_{zz} < 0$.

For (iv), observe that we can write the expression for L_K as

$$L_K = -\left(\frac{\vartheta_K}{\vartheta_M}\right)L_z$$

Taking the derivative yields

$$\frac{L_{Kz}}{L_z} = \left(-\frac{\vartheta_K}{\vartheta_M}\right)\left[\frac{L_{zz}}{L_z} - \left(\frac{\vartheta_{MK}}{\vartheta_K} - \frac{\vartheta_{MM}}{\vartheta_M}\right)(1-L_z)\right]$$

Then since $\vartheta_M > 0$ and $\vartheta_K < 0$, the term $-\frac{\vartheta_K}{\vartheta_M}$ is positive. Since $L_z > 0$, we find that $L_{Kz} > 0$ iff

$$\frac{L_{zz}}{L_z} - \left(\frac{\vartheta_{MK}}{\vartheta_K} - \frac{\vartheta_{MM}}{\vartheta_M}\right)(1 - L_z) > 0$$

which is equivalent to the given expression.

$\square$

Appendix C: Proofs (Chapter 3)

Proof of Lemma 3. We take the left-sided limit of the derivative of the payoff functions of bankers and workers $\pi\left(e^{*}-\varepsilon\right)+\varepsilon$ and $w\left(e^{*}-\varepsilon\right)$ as $\varepsilon\rightarrow 0$ to find

$$\begin{aligned}
\lim_{e\rightarrow e^{*-}} -\pi'(e)+1 &= \lim_{e\rightarrow e^{*-}}\left(1-\alpha\right)k'(e) &= \frac{1-\alpha}{1-\alpha\phi}\\
\lim_{e\rightarrow e^{*-}} -w'(e) &= \lim_{e\rightarrow e^{*-}} -\left(1-\alpha\right)k'(e) &= -\frac{1-\alpha}{1-\alpha\phi}
\end{aligned}$$

The marginal effect on total surplus is the sum of the two, $1-s'=1-\pi'-w'$, and is zero at a first-order approximation. □

Proof of Proposition 13. We first show that the marginal functions $\Pi'(x)$, $\Pi_1(x^i,x)$, and $W'(x)$ are strictly decreasing in x by differentiating each with respect to x,

$$\begin{aligned}
\Pi''(x) &= \int_0^{A^*}\left(\tilde{A}-1\right)^2\frac{(1-\phi)\alpha F_{kk}}{\left(1-\alpha\phi F_k\right)^3}dG(\tilde{A})<0\\
\frac{d}{dx}\Pi_1(x^i,x) &= \int_0^{A^*}\left(\tilde{A}-1\right)^2\frac{(1-\phi)F_{kk}}{\left(1-\phi F_k\right)^2\left(1-\alpha\phi F_k\right)}dG(\tilde{A})<0\\
W''(x) &= \left[\frac{(1-\alpha)}{(1-\phi)\alpha}\right]\Pi''(x)<0
\end{aligned}$$

Note that if it is indeed the case that $x^W<x^B$, then part (ii) of the proof follows immediately from this fact.

Next we show that $x^{LF}<x^B$ at an interior solution. At the point x^{LF} we

have $\Pi_1 = 0$. Then we find

$$\Pi'(x^{LF}) = \Pi'(x^{LF}) - \Pi_1(x^{LF}, x^{LF}) = -\int_0^{A^*} \frac{(1-\alpha)(1-\phi)(\tilde{A}-1)F_k}{(1-\phi\alpha F_k)(1-\phi F_k)} dG(\tilde{A})$$

Observe that the term $\frac{F_k}{(1-\phi\alpha F_k)(1-\phi F_k)}$ is strictly increasing in F_k. Now we define $\bar{R}$ as follows. If $A^* \leq 1$, so that the term $(\tilde{A}-1) < 0$ over the entire interval, we let $\bar{R}$ be the value of F_k when $\tilde{A} = A^*$. If instead we have $A^* > 1$, then let $\bar{R}$ be the value of F_k at $\tilde{A} = 1$. Then since F_k is decreasing in $\tilde{A}$, we have

$$-\int_0^{A^*} \frac{(1-\alpha)(1-\phi)(\tilde{A}-1)F_k}{(1-\phi\alpha F_k)(1-\phi F_k)} dG(\tilde{A}) > -\int_0^{A^*} \frac{(1-\alpha)(1-\phi)(\tilde{A}-1)F_k}{\left(1-\phi\alpha\bar{R}\right)(1-\phi F_k)} dG(\tilde{A})$$

Recall that at x^{LF} we have $\Pi_1 = 0$. We can write this as

$$\int_0^{A^*} (\tilde{A}-1)\frac{(1-\phi)F_k}{1-\phi F_k} dG(\tilde{A}) + \int_{A^*}^{\infty} (\tilde{A}-1) dG(\tilde{A}) = 0$$

Then since $\int_{A^*}^{\infty}(\tilde{A}-1)dG(\tilde{A}) > 0$, we must have $\int_0^{A^*}(\tilde{A}-1)\frac{(1-\phi)F_k}{1-\phi F_k}dG(\tilde{A}) < 0$. Thus we have

$$\Pi'(x^{LF}) > -\frac{(1-\alpha)}{\left(1-\phi\alpha\bar{R}\right)} \int_0^{A^*} \left(\tilde{A}-1\right) \frac{(1-\phi)F_k}{(1-\phi F_k)} dG(\tilde{A}) > 0$$

Thus we have $x^{LF} < x^B$. If $e^* \leq 1$ then $x^W = 1 - e^*$ because workers prefer avoiding any constraints whereas $x^{LF} > 1 - e^*$ because individual bankers would like to expose themselves to at least some constraints; therefore $x^W < x^{LF}$.

Finally, we show that $x^W < x^B$ for interior solutions to prove (i). Observe that

$$\Pi'(x) - \frac{(1-\phi)\alpha}{1-\alpha} W'(x) = \int_{A^*}^{\infty} \left(\tilde{A}-1\right) dG(\tilde{A}) > 0$$

Since at an interior solution we have $W'(x^W) = 0$, this implies $\Pi'(x^W) > 0$, and so $x^B > x^W$. □

Proof of Proposition 14. For (i), observe that we can write

$$p(e^i, e) = \epsilon\pi(e^i, e) + \epsilon(\delta - 1)\left(e^i - e\right)\mathbb{I}_{e^i \geq e^*}$$

where $\mathbb{I}_{e^i \geq e^*}$ is an indicator variable that is equal to 1 when $e^i \geq e^*$ and 0 otherwise. The preferred choice of x by managers, call it x^A, satisfies $P_1(x) = E\left[(\tilde{A} - 1)p_1(e^i, e)\right] \geq 0$. We can write this as

$$P_1(x) = \epsilon\Pi_1(x) + \epsilon(\delta - 1)E\left[(\tilde{A} - 1)\mathbb{I}_{e^i \geq e^*}\right]$$

where $\Pi_1(x) = E\left[(\tilde{A} - 1)\pi_1(e^i, e)\right]$ is the owner's first-order condition. The second term is strictly positive because

$$E\left[(\tilde{A} - 1)\mathbb{I}_{e^i \geq e^*}\right] = E\left[\tilde{A} - 1 | e^i \geq e^*\right] Pr(e^i \geq e^*)$$

Since e^i is strictly increasing in $\tilde{A}$, $E[\tilde{A} - 1 | e^i \geq e^*]$ is the expected value of the upper portion of a random variable, and so is strictly greater than $E[\tilde{A} - 1]$, which by assumption is strictly positive. Therefore we have $\Pi_1(x^{LF}) > 0$, and so $x^A > x^{LF}$.

To prove (ii), we begin by showing that x^A is strictly increasing in δ. Differentiating $P_1(x)$ with respect to δ yields $\epsilon E\left[(\tilde{A} - 1)\mathbb{I}_{e^i \geq e^*}\right]$, which is strictly positive. At the old preferred level of x, we now have $P_1(x) > 0$, and so x^A will increase. Now we observe that increasing x for $x > x^W$ will always make workers worse off. Then since $x^W < x^{LF} < x^A$, increasing δ will make workers worse off. □

Proof of Proposition 15. The marginal valuation of bank capital is now

$$\pi_1^{i,n}\left(e^i, e^{-i}\right) = \begin{cases} \frac{1}{n}\pi'(e) + \frac{n-1}{n}\pi_1^i\left(e^i, e\right) & \text{for } e < e^{*,n} \\ 1 & \text{for } e \geq e^{*,n} \end{cases}$$

This falls in between the marginal value of bank capital for the sector as a whole and for a competitive banker, i.e. $\pi' < \pi_1^{i,n} < \pi_1^i$.

Since we have $\pi_1^{i,n}(e^i, e) = \pi_1^i + \frac{1}{n}(\pi' - \pi_1^i)$, we can write the optimality condition for one of n large firms as

$$\Pi_1^{i,n} = \Pi_1(x) + \frac{1}{n}(\Pi' - \Pi_1) = 0$$

We immediately see that for $n = 1$, this reduces to $\Pi' = 0$, which has solution x^B, and for $n \to \infty$ this reduces to $\Pi_1 = 0$, which has solution $x^{LF} < x^B$.

Now suppose that for a given n, we have $x^n \in \left(x^{LF}, x^B\right)$. At x^n, we differentiate the optimality condition w.r.t. n and find

$$\frac{d}{dn}\Pi_1^{i,n} = -\frac{1}{n^2}(\Pi' - \Pi_1)$$

Since Π_1 and Π' are both strictly decreasing in x, and since they are zero at x^{LF} and $x^B > x^{LF}$ respectively, in the interval (x^{LF}, x^B) we have $\Pi_1 < \Pi'$. Therefore for higher n we have $\frac{d}{dn}\Pi_1^{i,n} < 0$, and so x^n is decreasing in n. □

Proof of Lemma 5. The welfare maximization problem of bankers under bailouts is

$$\max_{x^i \in [0,1]} \Pi^{BL}\left(x^i, x\right) = E\left[\pi^{BL}\left(e^i + t(e), e + t(e)\right)\right]$$

where $e^i = 1 - x^i + \tilde{A}x^i$ and $e = 1 - x + \tilde{A}x$ ($e^i = e$ in equilibrium). Let $\hat{A}$ be the level of $\tilde{A}$ that achieves the bailout threshold $\hat{e}$. The first partial derivative of the function Π^{BL} evaluated at x^{LF} satisfies

$$\begin{aligned}\Pi_1^{BL}\left(x^{LF},x^{LF}\right) = E\left[\left(\tilde{A}-1\right)\pi_1^{BL}\left(e^i,e\right)\right] = \\ = \pi_1(\hat{e},\hat{e})\int_0^{\hat{A}}(\tilde{A}-1)dG(\tilde{A}) + \int_{\hat{A}}^{\infty}(\tilde{A}-1)\pi_1 dG(\tilde{A}) > \\ > \int_0^{\hat{A}}(\tilde{A}-1)\pi_1 dG(\tilde{A}) + \int_{\hat{A}}^{\infty}(\tilde{A}-1)\pi_1 dG(\tilde{A}) = \Pi_1\left(x^{LF},x^{LF}\right) = 0\end{aligned}$$

Now we show why this inequality holds. First note that the second terms are identical and must be positive for $\Pi_1\left(x^{LF},x^{LF}\right)=0$ to hold. Thus if the first term in $\Pi_1^{BL}\left(x^{LF},x^{LF}\right)$ is positive, we are done. Suppose it is negative, which implies $E[\tilde{A}-1|A\leq\hat{A}]<0$. Then we need to show that

$$\int_0^{\hat{A}}(\tilde{A}-1)\pi_1\left(\hat{e},\hat{e}\right)dG(\tilde{A}) > \int_0^{\hat{A}}(\tilde{A}-1)\pi_1(e,e)dG(\tilde{A})$$

We can write this expression as $\int_0^{\hat{A}}(\tilde{A}-1)(\pi_1(\hat{e},\hat{e})-\pi_1)dG(\tilde{A})>0$, which is equivalent to $E[(\tilde{A}-1)(\pi_1(\hat{e},\hat{e})-\pi_1)|\tilde{A}\leq\hat{A}]>0$. This is the expectation of the product of two random variables, which equals

$$E\left[\tilde{A}-1|\tilde{A}\leq\hat{A}\right]\cdot E\left[\pi_1(\hat{e},\hat{e})-\pi_1|\tilde{A}\leq\hat{A}\right] + cov\left(\tilde{A}-1,\pi_1(\hat{e},\hat{e})-\pi_1|\tilde{A}\leq\hat{A}\right)$$

Since $\tilde{A}-1$ and $\pi_1(\hat{e},\hat{e})-\pi_1(e,e)$ are both strictly increasing in $\tilde{A}$ over the interval $[0,\hat{A})$, their covariance is strictly positive. Then since both $E[\pi_1(\hat{e},\hat{e})-\pi_1|\tilde{A}\leq\hat{A}]<0$ and $E[\tilde{A}-1|A\leq\hat{A}]<0$, this term is positive, and $\Pi_1^{BL}(x^{LF},x^{LF})>0$. Therefore individual bankers will choose to increase $x^{BL}>x^{LF}$ if there is a positive probability of bailouts. □

Proof of Lemma 4. The welfare of workers who collectively provide a transfer $t\geq 0$ to bankers is given by $w(e+t)-t$. An interior optimum satisfies $w'(e+t)=1$. We define the resulting equity level as $\hat{e} = e + t$, which satisfies equation (3.16). Observe that $w'(e)$ is strictly declining from $w'(0) = 1/\phi > 1$ to $w'(e^*) = \frac{1-\alpha}{1-\phi\alpha} < 1$ over the interval $[0, e^*]$ so that $\hat{e}$ is uniquely defined. If aggregate bank capital is below this threshold $e < \hat{e}$, workers find it collectively optimal to transfer the shortfall. If e is above this threshold, it does not pay off for workers to provide a transfer since $w' < 1$ and the optimal transfer is given by the minimum $t = 0$. $\square$

Appendix D: Variants of Baseline Model (Chapter 3)

D.1 Always Constrained Case

An interesting special case in which financial markets in period 0 are effectively complete is a two sector framework in which bankers own all the capital and workers own all the labor in the economy (i.e. there are no deposits $d = 0$ and no storage). By implication, bankers invest all their equity into real capital $k = e$. Given a Cobb-Douglas production technology, the two sectors earn constant fractions of aggregate output so that $\pi(e) = \alpha F(e, 1)$ and $w(e) = (1 - \alpha) F(e, 1)$ for $e = \tilde{A}x + (1 - x)$. As long as the two sectors have preferences with identical relative risk aversion (in our benchmark model both have zero risk-aversion), the optimal level of risk-taking for bankers simultaneously maximizes total surplus and worker welfare:

$$\arg\max_x E[\pi(e)] = \arg\max_x E[F(e, 1)] = \arg\max_x E[w(e)]$$

Bank capital still imposes pecuniary externalities on wages in this setting, but the pecuniary externalities under a Cobb-Douglas technology guarantee that both sets of agents obtain constant fractions of output, replicating the allocation under perfect risk-sharing. (Analytically, the constant capital and labor shares drop out of the optimization problem.) There is no distributive conflict.

D.2 Period 0 Production Function

This appendix generalizes our setup to a Cobb-Douglas production function that is symmetric across periods $t = 1$ and 2 of the form

$$\left[\tilde{A}_t x_t + 1 - x_t\right] F(k_t, \ell_t)$$

This allows us to account for the notion that the higher returns from risk-taking in the initial period are shared between workers and bankers.

We continue to assume that bankers choose the fraction x_t allocated to risky projects and firms choose the amount of capital invested k_t before the productivity shock $\tilde{A}_t$ is realized, i.e. in period $t-1$.

In period 0, bankers supply their initial equity e_0 to firms for physical capital investment so that $k_0 = e_0$. In period 1, the productivity shock $\tilde{A}_1$ is realized and firms hire $\ell = 1$ units of labor to produce output $\tilde{A}_1 F(e_0, 1)$. Bankers and workers share the productive output according to their factor shares,

$$e_1 = \alpha \left[\tilde{A}_1 x_1 + 1 - x_1\right] F\left(e_0, 1\right) \tag{D.1}$$

$$w_1 = (1-\alpha) \left[\tilde{A}_1 x_1 + 1 - x_1\right] F\left(e_0, 1\right)$$

where equation (D.1) represents the law-of-motion of bank capital from period 0 to period 1. Given the period 1 bank capital e_1, the economy behaves as we have analyzed in Section 3.3.1 in the main body of the paper, i.e. bankers and workers obtain profits and wages of $\pi(e_1)$ and $w(e_1)$. Observe that all agents are risk-averse with respect to period 2 consumption; therefore the optimal $x_2 \equiv 1$ and we can solve for all allocations as if the productivity parameter in period 2 was the constant $A_2 = E[\tilde{A}_2]$, as in our earlier analysis.

We express aggregate welfare of bankers and workers as a function of period 0 risk-taking x_1 as

$$\Pi(x_1) = E\left\{\pi(e_1)\right\}$$

$$W(x_1) = E\left\{w_1 + w(e_1)\right\}$$

where e_1 and w_1 are determined by risk-taking and the output shock, as given by equation (D.1).

Observe that in addition to the effects of risk-taking on period 2 wages $w(e_1)$ that we investigated earlier, period 1 wages now depend positively on risk-taking x_1 because wages are a constant fraction $(1-\alpha)$ of output and greater risk leads to higher period 1 output since $E[\tilde{A}_1] > 1$. Bankers do not internalize either of the two externalities on period 1 and period 2 wages.

Assuming an interior solution for x_1 and noting that $\pi'(e_1) - 1 = (\alpha F_k - 1)k'(e_1)$, the optimal level of risk-taking for the banking sector x_1^B satisfies

$$\Pi'(x_1^B) = E\left[\left(\tilde{A}_1 - 1\right)\pi'(e_1)\right] =$$
$$= E\left[\tilde{A}_1 - 1\right] + \int_0^{\hat{A}_1} \left(\tilde{A}_1 - 1\right)\left(\alpha F_k - 1\right) k'(e_1) dG(\tilde{A}_1) = 0$$

The banking sector prefers more risk than workers if $W'(x_1^B) < 0$:

$$W'\left(x^B\right) = E\left\{\left[(1-\alpha)F\left(e_0, 1\right) + w'\left(e\right)\right]\left(\tilde{A} - 1\right)\right\}$$
$$= \int_0^{\hat{A}} \left[w'\left(e\right) - (1-\alpha)F\left(e_0, 1\right)\left(\alpha F_k - 1\right) k'(e_1)\right]\left(\tilde{A} - 1\right) dG(\tilde{A}_1)$$

where we subtracted the expression $(1-\alpha)F(e_0,1)\,\Pi'(x_1^B)=0$ in the second line, which is zero by the optimality condition of bankers.

Let us impose two weak assumptions that allow us to sign this expression. First, assume $\phi > \alpha$, i.e. leverage is above a minimum level that is typically satisfied in all modern financial systems (1.5 for the standard value of $\alpha = 1/3$), and secondly, that $\hat{A} < 1$, i.e. only low realizations of the productivity shock lead to credit crunches. Note that these two assumptions are sufficient but not necessary conditions.

Now observe that the first term under the integral, $w'(e)$, is always positive. To sign the second term, notice that $F_k(k,1) \leq F_k(k(0),1) = 1/\phi \ \forall \ e \geq 0$ and so the assumption $\phi > \alpha$ implies that $\alpha F_k - 1 < 0$. Furthermore, by the second assumption, the term $(\tilde{A}-1)$ is negative since the integral is over the interval $[0,\hat{A}]$. As a result, the two conditions are sufficient to ensure that the expression is always negative and that workers continue to prefer less risk-taking than the banking sector.

Intuitively, our distributive results continue to hold when we account for production and wage earnings in both time periods because the distributive conflict stems from the asymmetric effects of credit crunches on bankers and workers, which are still present: workers are hurt by credit crunches but do not benefit from higher bank dividends in good times. Therefore workers prefer less risk-taking than bankers.

D.3 Different Forms of Bailouts

This appendix considers bailouts that come in the form of emergency lending and equity injections and shows that both matter only to the extent that they provide a subsidy (outright transfer in expected value) to constrained bankers that relaxes their financial constraint.[1]

Emergency Lending A loan d^{BL} that a policymaker provides to constrained bankers on behalf of workers at an interest rate r^{BL} that is frequently subsidized, i.e. below the market interest rate $r^{BL} \leq 1$. Such lending constitutes a transfer of $\left(r^{BL} - 1\right) d^{BL}$ in net present value terms.[2] Assuming that such interventions cannot relax the commitment problem of bankers that we described in Section 3.2.2, they are subject to the constraint

$$rd + r^{BL} d^{BL} \leq \phi R k \tag{D.2}$$

Equity Injections provide constrained bankers with additional bank capital/equity q in exchange for a dividend distribution D, which is frequently expressed as a fraction of bank earnings. The equity injection constitutes a transfer of $q - D$ from workers to bankers in net present value terms. Assuming that the dividend payment is subject to the commitment problem of bankers that we assumed earlier, it has to

[1] For a more comprehensive analysis of bank recapitalizations see e.g. Sandri and Valencia (2013). For a detailed analysis of the resulting incentives for rent extraction see Korinek (2013).

[2] In our framework, we assumed that default probabilities are zero in equilibrium. In practice, the interest rate subsidy typically involves not charging for expected default risk.

obey the constraint

$$rd + D \leq \phi Rk \tag{D.3}$$

Given our assumptions, both types of bailouts are isomorphic to a lump-sum transfer t from workers to bankers.[3]

In the following lemma, we will first focus on an optimal lump-sum transfer and then show that the resulting allocations can be implemented either directly or via an optimal package of emergency lending or equity injection.

Lemma 6 (Variants of Bailouts)**.** Both workers and bankers are indifferent between providing the bailout via subsidized emergency loans such that $\left(1 - r^{BL}\right) d^{BL} = t$ or via subsidized equity injections such that $q - D = t$. Conversely, emergency lending and/or equity injections that do not represent a transfer in net present value terms are ineffective in our model.

Proof. Let us first focus on an emergency loan package described by a pair (r^{BL}, d^{BL}) that is provided to bankers by a policymaker on behalf of workers. Since the opportunity cost of lending is the storage technology, the direct cost of such a loan to workers is $(1 - r^{BL})d^{BL}$. Bankers intermediate $k = e + d + d^{BL}$ where we substitute d from constraint (D.2) to obtain

$$k = \frac{e + \left(1 - r^{BL}\right) d^{BL}}{1 - \phi R\left(k\right)} = k\left(e + \left(1 - r^{BL}\right) d^{BL}\right)$$

Therefore the emergency loan is isomorphic to a lump sum transfer $t = \left(1 - r^{BL}\right) d^{BL}$ for bankers, workers and firms. For an equity injection that is described by a pair (q, D), an identical argument can be applied.

[3] Since labor supply is constant, a tax on labor would be isomorphic to a lump sum transfer.

These observations directly imply the second part of the lemma. More specifically, constraint (D.2) implies that an emergency loan of d^{BL} at an unsubsidized interest rate $r^{BL} = 1$ reduces private deposits by an identical amount $\Delta d = -d^{BL}$ and therefore does not affect real capital investment k. Similarly, constraint (D.3) implies that an equity injection which satisfies $q = D$ reduces private deposits by $\Delta d = -D$ and crowds out an identical amount of private deposits. □

This captures an equivalence result between the two categories of bailouts – what matters for constrained bankers is that they obtain a transfer in net present value terms, but it is irrelevant how this transfer is provided. From the perspective of bankers who are subject to constraint (3.3), a one dollar repayment on emergency loans or dividends is no different from a one dollar repayment to depositors, and all three forms of repayment tighten the financial constraint of bankers in the same manner. An emergency loan or an equity injection at preferential rates that amounts to a one dollar transfer allows bankers to raise an additional $\frac{\phi R}{1-\phi R}$ dollars of deposits and expand intermediation by $\frac{1}{1-\phi R}$ dollars in total.

Emergency loans or equity injections that are provided at 'fair' market rates, i.e. that do not constitute a transfer in net present value terms, will therefore not increase financial intermediation. We assumed that the commitment problem of bankers requires that they obtain at least a fraction $(1 - \phi)$ of their gross revenue. If government does not have a superior enforcement technology to relax this constraint, any repayments on emergency lending or dividend payments on public equity injections reduce the share obtained by bankers in precisely the same fashion as repaying bank depositors. Such repayment obligations therefore decrease the amount of deposits that bankers can obtain by an equal amount and do not expand capital intermediation.

Conversely, if government had superior enforcement capabilities to extract repayments or dividends, then those special capabilities would represent an additional reason for government intervention in the instrument(s) that relax the constraint most.

D.4 Bailouts Conditional on Individual Bank Capital

The adverse incentive effects of bailouts are aggravated if bailouts are conditional on individual bank capital e^i. Such bailouts provide bankers with an additional incentive to increase risk-taking in order to raise the expected bailout rents received.

To capture this notion, suppose that the bailout received by an individual banker i for a given level of individual and aggregate bank equity (e^i, e) is given by

$$t\left(e^i, e; \gamma\right) = \begin{cases} 0 & \text{if } e \geq \hat{e} \\ \hat{e} - (1-\gamma)\, e - \gamma e^i & \text{if } e < \hat{e} \end{cases}$$

where $\gamma \in [0, 1]$ captures the extent to which the bailout depends on individual bank equity. This specification nests our baseline model in which bailouts are entirely conditional on aggregate bank capital ($\gamma = 0$), but now also includes bailouts that are partially or wholly contingent on individual bank capital ($\gamma > 0$). Alternatively, if banks are non-atomistic and bailouts are conditional only on aggregate bank capital e, we can interpret the parameter γ as the market share of individual banks, since each bank will internalize that its bank equity makes up a fraction γ of aggregate bank equity.

We denote the amount of their endowment that bankers allocate to the risky project in period 0 by $x^{BL}(\gamma)$, and we find that bailouts have the following effects:

Proposition 16 (Risk-Taking Effects of Bailouts)**.** (i) Introducing bailout transfers increases period 0 risk-taking $x^{BL}(\gamma) > x^{LF}$ for any $\gamma \geq 0$.

(ii) Risk-taking $x^{BL}(\gamma)$ is an increasing function of γ.

Proof. Since we proved in Proposition 5 that $x^{BL}(\gamma) > x^{LF}$ holds for $\gamma = 0$, (ii) implies (i). To prove (ii), observe that the welfare maximization problem of bankers under bailouts for a given parameter γ is

$$\max_{x^i \in [0,1], e^i} \Pi^{BL}\left(x^i, x; \gamma\right) = E\left[\pi^{BL}\left(e^i + t\left(e^i, e; \gamma\right), e + t\left(e\right)\right)\right]$$

where $e^i = 1 - x^i + \tilde{A}x^i = e$ in equilibrium. Let us define $\hat{A}$ as the level of $\tilde{A}$ that achieves the bailout threshold $\hat{e}$. The optimal choice of $x^{BL}(\gamma)$ satisfies

$$\Pi_1^{BL}\left(x^{BL}, x^{BL}; \gamma\right) = (1-\gamma)\pi_1(\hat{e}, \hat{e}) \int_0^{\hat{A}} (\tilde{A} - 1) dG(\tilde{A}) + \int_{\hat{A}}^{\infty} (\tilde{A} - 1)\pi_1 dG(\tilde{A}) = 0$$

Differentiating the optimality condition at x^{BL} for a given γ yields

$$\frac{d\Pi_1^{BL}}{d\gamma} = -\pi_1(\hat{e}, \hat{e}) \int_0^{\hat{A}} (\tilde{A} - 1) dG(\tilde{A}) > 0$$

where the inequality holds since we assumed $\hat{A} < 1$. □

Point (ii) captures that the risk-taking incentives of bankers rise further because they internalize that one more dollar in losses will increase their bailout by γ dollars. This captures the standard notion of moral hazard, i.e. that bailouts targeted at individual losses increase risk-taking.

Redistributive Effects Corollary 2 showed that Bankers benefit by the introduction of bailouts, while workers benefit from the market-completion effect and are hurt by the incentive effect of bailouts. Since the market-completion effect does not depend on γ, higher γ acts as a pure incentive effect that raises risk-taking, and benefits bankers at the expense of workers. Therefore $\gamma > 0$ exacerbates the distributive effects of bailouts.

Bibliography

Acharya, V.V. and P. Schnabl. 2010. “Do Global Banks Spread Global Imbalances? Asset-Backed Commercial Paper during the Financial Crisis of 2007–09.” *IMF Economic Review* 58 (1):37–73.

Achdou, Yves, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2013. “Heterogeneous Agent Models in Continuous Time.” Unpublished Manuscript.

Admati, Anat R, Peter M DeMarzo, Martin F Hellwig, and Paul C Pfleiderer. 2013. “Fallacies, Irrelevant Facts, and Myths in the Discussion of Capital Regulation: Why Bank Equity is not Expensive.” *Stanford GSB Working Paper* (No. 2065).

Akerlof, George A. 1970. “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism.” *The quarterly journal of economics* :488–500.

Aruoba, S Borağan and Frank Schorfheide. 2013. “Macroeconomic Dynamics Near the ZLB: A Tale of Two Equilibria.” *NBER Working Paper* w19248.

Baker, Scott. 2014. “Debt and the Consumption Response to Household Income Shocks.” Unpublished Manuscript.

Beaudry, Paul, Dana Galizia, and Franck Portier. 2014. “Reconciling Hayek’s and Keynes Views of Recessions.” *NBER Working Paper* w20101.

Benhabib, Jess, Stephanie Schmitt-Grohé, and Martın Uribe. 2001. “The Perils of Taylor Rules.” *Journal of Economic Theory* 96 (1):40–69.

Benigno, Gianluca, Huigang Chen, Christopher Otrok, Alessandro Rebucci, and Eric Young. 2012. “Optimal Policy for Macro-Financial Stability.” .

Bernanke, B., C. Bertaut, L. DeMarco, and S. Kamin. 2011. “International Capital Flows and the Return to Safe Assets in the United States, 2003-2007.” *FRB International Finance Discussion Paper No. 1014* .

Bernanke, B. and M. Gertler. 1989. “Agency Costs, Net Worth, and Business Fluctuations.” *The American Economic Review* 79 (1):14–31.

Bernanke, Ben S and Michael Woodford. 1997. "Inflation Forecasts and Monetary Policy." *Journal of Money, Credit, and Banking* :653–684.

Bianchi, Javier and Enrique G Mendoza. 2010. "Overborrowing, Financial Crises and'Macro-prudential'Taxes." Tech. rep., National Bureau of Economic Research.

Blanchard, Olivier and Jordi Galí. 2010. "Labor Markets and Monetary Policy: a New Keynesian Model with Unemployment." *American Economic Journal: Macroeconomics* 2 (2):1–30.

Brunnermeier, Markus K and Lasse Heje Pedersen. 2009. "Market Liquidity and Funding Liquidity." *Review of Financial studies* 22 (6):2201–2238.

Brunnermeier, M.K. 2009. "Deciphering the Liquidity and Credit Crunch 2007-2008." *The Journal of economic perspectives* 23 (1):77–100.

Brunnermeier, M.K. and Y. Sannikov. 2010. "The I-Theory of Money." *work in progress* .

Caballero, Ricardo J and Guido Lorenzoni. 2014. "Persistent Appreciations and Overshooting: a Normative Analysis." *IMF Economic Review* 62 (1):1–47.

Caballero, R.J. 2010. "The 'Other' Imbalance and the Financial Crisis." .

Caggese, Andrea and Ander Perez. 2013. "Aggregate Implications of Financial and Labor Market Frictions." Unpublished Manuscript.

Calomiris, C.W. and C.M. Kahn. 1991. "The Role of Demandable Debt in Structuring Optimal Banking Arrangements." *The American Economic Review* :497–513.

Campbell, Jeffrey R and Zvi Hercowitz. 2009. "Welfare Implications of the Transition to High Household Debt." *Journal of Monetary Economics* 56 (1):1–16.

Candler, Graham V. 2001. "Finite–Difference Methods for Continuous–Time Dynamic Programming."

Carroll, Christopher D and Andrew A Samwick. 1997. "The Nature of Precautionary Wealth." *Journal of monetary Economics* 40 (1):41–71.

———. 1998. "How Important is Precautionary Saving?" *Review of Economics and Statistics* 80 (3):410–419.

Carroll, Christopher D, Jiri Slacalek, and Martin Sommer. 2012. "Dissecting Saving Dynamics: Measuring Wealth, Precautionary and Credit Effects." Unpublished Manuscript.

Challe, Edouard, Julien Matheron, Xavier Ragot, and Juan Rubio-Ramirez. 2014. "Precautionary Saving and Aggregate Demand." Unpublished Manuscript.

Chamley, Christophe. 2014. "When Demand Creates its own Supply: Saving Traps." *The Review of Economic Studies* 81 (2):651–680.

Clarida, Richard, Jordi Galí, and Mark Gertler. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *The Quarterly journal of economics* 115 (1):147–180.

Cochrane, John H. 2011. "Determinacy and Identification with Taylor Rules." *Journal of Political Economy* 119 (3):565–615.

DeMarzo, Peter and Darrell Duffie. 1999. "A Liquidity-Based Model of Security Design." *Econometrica* 67 (1):65–99.

Diamond, D.W. and P.H. Dybvig. 1983. "Bank Runs, Deposit Insurance, and Liquidity." *The Journal of Political Economy* 91 (3):401–419.

Diamond, D.W. and R.G. Rajan. 2001. "Liquidity Risk, Liquidity Creation, and Financial Fragility: A Theory of Banking." *Journal of Political Economy* 109 (2).

Dynan, Karen. 2012. "Is a Household Debt Overhang Holding Back Consumption?" *Brookings Papers on Economic Activity* 44 (1):299–362.

Dynan, Karen E, Jonathan Skinner, and Stephen P Zeldes. 2004. "Do the Rich Save More?" *Journal of Political Economy* 112 (2):397–444.

Eggertsson, Gauti and Neil Mehrotra. 2014. "A Model of Secular Stagnation." Unpublished Manuscript.

Eggertsson, Gauti B and Paul Krugman. 2012. "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach." *The Quarterly Journal of Economics* 127 (3):1469–1513.

Eggertsson, Gauti B and Michael Woodford. 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity* 34 (1):139–235.

Farhi, E. and J. Tirole. 2012. "Bubbly Liquidity." *Review of Economic Studies* 79 (2):678–706.

Farhi, Emmanuel and Iván Werning. 2013. "A Theory of Macroprudential Policies in the Presence of Nominal Rigidities." *NBER Working Paper* w19313.

Freixas, Xavier and Jean-Charles Rochet. 2008. "Microeconomics of Banking." *MIT Press Books* 1.

Galí, Jordi. 2011. "The Return of the Wage Phillips Curve." *Journal of the European Economic Association* 9 (3):436–461.

Gersbach, Hans and Jean-Charles Rochet. 2012. “Aggregate Investment Externalities and Macroprudential Regulation.” *Journal of Money, Credit and Banking* 44 (s2):73–109.

Gertler, Mark and Peter Karadi. 2011. “A Model of Unconventional Monetary Policy.” *Journal of monetary Economics* 58 (1):17–34.

Gertler, Mark, Nobuhiro Kiyotaki et al. 2010. “Financial Intermediation and Credit Policy in Business Cycle Analysis.” *Handbook of monetary economics* 3 (3):547–599.

Gorton, Gary and George Pennacchi. 1990. “Financial Intermediaries and Liquidity Creation.” *The Journal of Finance* 45 (1):49–71.

Gourinchas, Pierre-Olivier and Jonathan A Parker. 2002. “Consumption over the life cycle.” *Econometrica* 70 (1):47–89.

Greenwood, R., S.G. Hanson, and J.C. Stein. 2010. “A Comparative-Advantage Approach to Government Debt Maturity.” .

Guerrieri, Veronica and Guido Lorenzoni. 2011. “Credit Crises, Precautionary Savings, and the Liquidity Trap.” *NBER Working Paper* w17583.

Haldane, Andrew G. 2010. “The $100 Billion Question.” *Revista de Economía Institucional* 12 (22):83–110.

Hall, Robert E. 2005. “Employment Fluctuations with Equilibrium Wage Stickiness.” *American economic review* :50–65.

———. 2010. “Why Does the Economy Fall to Pieces after a Financial Crisis?” *The Journal of Economic Perspectives* :3–20.

———. 2011. “The Long Slump.” *American Economic Review* 101 (2):431–69.

———. 2014. “High Discounts and High Unemployment.” *NBER Working Paper* w19871.

Hansen, Alvin H. 1939. “Economic Progress and Declining Population Growth.” *The American Economic Review* :1–15.

Heathcote, Jonathan and Fabrizio Perri. 2014. “Wealth and Volatility.” Unpublished Manuscript.

Holmstrom, Bengt and Jean Tirole. 1998. “Private and Public Supply of Liquidity.” *Journal of Political Economy* 106 (1):1–40.

———. 2011. *Inside and Outside Liquidity*. MIT press.

Jeanne, Olivier and Anton Korinek. 2010a. “Excessive Volatility in Capital Flows: A Pigouvian Taxation Approach.” *American Economic Review* 100 (2):403–07.

———. 2010b. "Managing Credit Booms and Busts: A Pigouvian Taxation Approach." Tech. rep., National Bureau of Economic Research.

———. 2013. "Macroprudential Regulation versus Mopping Up After the Crash." Tech. rep., National Bureau of Economic Research.

Jensen, Michael C. 1986. "Agency Cost of Free Cash Flow, Corporate Finance, and Takeovers." *Corporate Finance, and Takeovers. American Economic Review* 76 (2).

Jordà, Òscar, Moritz Schularick, and Alan M Taylor. 2013. "When Credit Bites Back." *Journal of Money, Credit and Banking* 45 (s2):3–28.

Kaplan, Steven N and Joshua Rauh. 2010. "Wall Street and Main Street: What Contributes to the Rise in the Highest Incomes?" *Review of Financial Studies* 23 (3):1004–1050.

Karabarbounis, Loukas and Brent Neiman. 2014. "The Global Decline of the Labor Share." *The Quarterly Journal of Economics* 129 (1):61–103.

Kashyap, A.K. and J.C. Stein. 2000. "what Do a Million Observations on Banks Say about the Transmission of Monetary Policy?" *American Economic Review* :407–428.

Kennickell, Arthur B. 2009. "Ponds and Streams: Wealth and Income in the US, 1989 to 2007." *Finance and Economics Discussion Series, Federal Reserve Board, Washington, DC* .

Kimball, Miles S. 1990. "Precautionary Saving in the Small and in the Large." *Econometrica* 58 (1):53–73.

Kiyotaki, N. and J. Moore. 1997. "Credit Cycles." *The Journal of Political Economy* 105 (2):211–248.

———. 2008. "Liquidity, Business Cycles, and Monetary Policy." *Working Paper* .

Kocherlakota, Narayana. 2012. "Incomplete Labor Markets." Unpublished Manuscript.

Korinek, Anton. 2011. "Systemic Risk-Taking: Amplification Effects, Externalities, and Regulatory Responses." .

Korinek, Anton and Jonathan Kreamer. 2014. "The Redistributive Effects of Financial Deregulation." *Journal of Monetary Economics* 68:S55–S67.

Korinek, Anton and Enrique G Mendoza. 2014. "From Sudden Stops to Fisherian Deflation: Quantitative Theory and Policy." *Annu. Rev. Econ.* 6 (1):299–332.

Korinek, Anton and Alp Simsek. 2014. "Liquidity Trap and Excessive Leverage." *NBER Working Paper* w19970.

Kreamer, Jonathan. 2014. "Household Debt, Unemployment, and Slow Recoveries." *Manuscript, University of Maryland* .

Krishnamurthy, Arvind and Annette Vissing-Jorgensen. 2012. "The Aggregate Demand for Treasury Debt." *Journal of Political Economy* 120 (2):233 – 267.

Krugman, Paul R. 1998. "It's Baaack: Japan's Slump and the Return of the Liquidity Trap." *Brookings Papers on Economic Activity* 1998 (2):137–205.

Larrain, Mauricio. 2013. "Capital Account Liberalization and Wage Inequality." *Columbia Business School Research Paper* (12/48).

Lorenzoni, Guido. 2008. "Inefficient Credit Booms." *The Review of Economic Studies* 75 (3):809–833.

Mattesini, F., C. Monnet, and R. Wright. 2009. "Banking: a Mechanism Design Approach." *Working Papers* .

McKay, Alisdair, Emi Nakamura, and Jón Steinsson. 2015. "The Power of Forward Guidance Revisited." *NBER Working Paper* w20882.

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump*." *The Quarterly Journal of Economics* 128 (4):1687–1726.

Mian, Atif and Amir Sufi. 2010. "Household Leverage and the Recession of 2007–09." *IMF Economic Review* 58 (1):74–117.

———. 2014. "What Explains the 2007–2009 Drop in Employment?" *Econometrica* (forthcoming).

Michaillat, Pascal. 2012. "Do Matching Frictions Explain Unemployment? Not in Bad Times." *The American Economic Review* 102 (4):1721–1750.

Michaillat, Pascal and Emmanuel Saez. 2013. "A Model of Aggregate Demand and Unemployment." *NBER Working Paper* w18826.

Miles, David, Jing Yang, and Gilberto Marcheggiano. 2013. "Optimal Bank Capital*." *The Economic Journal* 123 (567):1–37.

Parker, Jonathan A and Bruce Preston. 2005. "Precautionary Saving and Consumption Fluctuations." *American Economic Review* 95 (4):1119–43.

Philippon, Thomas and Ariell Reshef. 2012. "Wages and Human Capital in the US Finance Industry: 1909–2006*." *Quarterly Journal of Economics* 127 (4).

———. 2013. "An International Look at the Growth of Modern Finance." *The Journal of Economic Perspectives* :73–96.

Pissarides, Christopher A. 2009. "The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?" *Econometrica* 77 (5):1339–1369.

Pozsar, Z. 2011. "Institutional Cash Pools and the Triffin Dilemma of the US Banking System." .

Pozsar, Z., T. Adrian, A. Ashcraft, and H. Boesky. 2010. "Shadow Banking." .

Ravn, Morten and Vincent Sterk. 2013. "Job Uncertainty and Deep Recessions." Unpublished Manuscript.

Reinhart, Carmen M and Kenneth Rogoff. 2009. *This Time is Different: Eight Centuries of Financial Folly*. princeton university press.

Rognlie, Matthew, Andrei Shleifer, and Alp Simsek. 2014. "Investment Hangover and the Great Recession." *NBER Working Paper* w20569.

Romer, Christina D and David H Romer. 2015. "New Evidence on the Impact of Financial Crises in Advanced Countries." *NBER Working Paper* w17344.

Samuelson, P.A. 1958. "An Exact Consumption-Loan Model of Interest With or Without the Social Contrivance of Money." *The Journal of Political Economy* 66 (6):467–482.

Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* :25–49.

Stein, J.C. 2012. "Monetary Policy as Financial Stability Regulation." *The Quarterly Journal of Economics* 127 (1):57–95.

Sufi, Amir. 2009. "Bank lines of credit in corporate finance: An empirical analysis." *Review of Financial Studies* 22 (3):1057–1088.

Tirole, Jean. 2011. "Illiquidity and All Its Friends." *Journal of Economic Literature* 49 (2):287–325.

Werning, Ivan. 2011. "Managing a Liquidity Trap: Monetary and Fiscal Policy." *NBER Working Paper* w17344.

Woodford, M. 1990. "Public Debt as Private Liquidity." *The American Economic Review* 80 (2):382–388.