Essays on financial crises, contagion and macro-prudential regulation[1]

Toni Ahnert[2]

April 2013

[1]A thesis submitted to the Department of Economics of the London School of Economics and Political Science for the degree of Doctor of Philosophy

[2]London School of Economics and Political Science, Department of Economics and Financial Markets Group, Houghton Street, London WC2A 2AE.
I am indebted to my advisers Margaret Bray and Dimitri Vayanos for their constant advice, guidance, and support. I am also indebted to Douglas Gale for his supervision and support, as well as for hosting me as a visiting PhD student at the Department of Economics at New York University.

Declaration

I certify that the thesis I have presented for examination for the Ph.D. degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of approximately 34,800 words.

Statement of conjoint work

I confirm that Chapter 3 on information contagion and systemic risk was jointly co-authored with Co-Pierre Georg and I contributed 60% of this work.

Statement of use of third party for editorial help

I can confirm that Chapter 1 of this thesis was copy edited for conventions of language, spelling and grammar by Nextgenediting (Nextgenology Ltd, 145-157 St. John Street, London, EC1V 4PW). Furthermore, I can confirm that Chapter 2 of this thesis was copy edited for conventions of language, spelling and grammar by Edit911 (Edit911, Inc., 2200 Winter Springs Blvd., Suite 106, Oviedo, FL 32765, USA).

Contents

List of Tables

List of Figures

Chapter 1

Crisis, Coordination, and Contagion

This chapter explores how private information acquisition affects the probability of a financial crisis. I study a global coordination game of regime change in which a crisis occurs if a sufficient number of creditors run on a bank or do not roll over debt to a sovereign. Creditors receive noisy private information about the solvency of the bank or sovereign and choose ex-ante whether to improve the quality of their information at a cost. Learning adverse public news about the solvency of the bank or sovereign increases strategic uncertainty among creditors. This induces a creditor to acquire private information to align his decision with that of other creditors. Since informed creditors are more likely to withdraw or refuse to roll over, information acquisition amplifies the probability of a financial crisis. Further extending the study to include news about another bank (Lehman) or sovereign (Greece), I demonstrate that the acquisition of private information makes bank runs systemic or sovereign debt crises contagious.

JEL Classifications: C7, D8, F3, G01, G21

Keywords: amplification, contagion, financial crisis, information acquisition

I am grateful to Christian Hellwig, John Moore, and Francesco Nava for comments and fruitful discussions. I thank seminar participants at Amsterdam (UvA), Bank of Canada, Bank of England, University of Bonn , Bowdoin, Bundesbank, Haskayne School of Business, Cleveland Fed, De Nederlandsche Bank, Federal Reserve Board of Governors, Frankfurt School of Finance, HEC Lausanne, HSE Moscow, LSE, Stockholm (Riksbank), Warwick Business School and Christoph Bertsch, Johannes Boehm, Amil Dasgupta, Jason Donaldson, Daniel Ferreira, Luca Fornaro, Co-Pierre Georg, Benjamin Nelson, and Kathy Yuan for feedback.

1.1 Introduction

Financial crises, such as systemic bank runs and sovereign debt crises, have substantial financial and economic costs.[1] A self-fulfilling financial crisis can be caused by the coordination failure of bank creditors or sovereign debt holders. In Diamond and Dybvig (1983), for example, the fear of premature withdrawal by other depositors induces a given depositor to withdraw prematurely, causing a bank run.[2] Pioneered by Carlsson and van Damme (1993), the global games literature establishes a unique equilibrium in coordination problems if the exogenous private information about the profitability of a bank, or the solvency of a debtor, is sufficiently precise.[3] While information is exogenous in this literature, a large part of advanced economies is devoted to the generation, collection, and processing of information (Veldkamp (2011)). This holds particularly for the financial industry and suggests a role for endogenous information.

This paper examines the consequences of private information acquisition on financial crises, such as bank runs or sovereign debt crises. I study a coordination game of regime change in which a crisis occurs if a sufficient number of creditors withdraw from a bank or do not roll over debt to a sovereign. Before making their decision, creditors receive noisy private information about the solvency of the bank or sovereign. This information can be made more precise ex-ante at a cost, for example by hiring analysts or by purchasing data and IT infrastructure. More precise information helps a creditor align his decision with that of other creditors. In short, information supports coordination.

This paper addresses three questions. First, I examine how changes in economic fundamentals affect incentives to acquire private information and consequently the probability of a financial crisis. For instance, how do disappointing Spanish unemployment figures alter the incentives of Spanish debt holders to gather information and, in turn, to roll over debt? Second, I study the strategic aspects of information acquisition. For example, in case of an investor in US municipal debt, how does the decision to find out more

[1] In a survey of banking crises in the member countries of the Basel Committee on Banking Supervision (BCBS) from 1985 to 2009, BCBS (2010a) find the median cost of banking crises to be 9% of pre-crisis GDP.

[2] This argument extends to the coordination problem associated with the roll-over decision of debt holders (Morris and Shin (2004)) and the attacking decision of currency speculators (Obstfeld (1986)).

[3] The unique equilibrium features a run on solvent but illiquid banks in Rochet and Vives (2004) or Goldstein and Pauzner (2005) and the refusal to roll over debt to solvent debtors in Morris and Shin (2004).

about the municipality's health depend on the amount of information gathering by other investors? Third, I analyze whether public information spillovers about another bank or sovereign cause contagion via their effect on information acquisition. For example, how does a debt restructuring in Greece affect Italian debt holders and their incentives to acquire information?

My main contribution is to demonstrate that the endogenous acquisition of private information after adverse public news increases the probability of a financial crisis. To illustrate this amplification effect, I focus on the case of a sovereign debt crisis, but the argument is also applicable to a bank run. Suppose that a sovereign's ability to repay its debt is high due to its taxation power or multinational support.[4] Adverse news has two effects on the probability of a crisis. First, it directly raises the probability of a crisis via a learning effect. If the sovereign is commonly perceived to be less solvent, a debt crisis can be triggered by fewer debt holders who do not roll over, and the probability of a crisis is increased. Second, adverse news changes the incentives for debt holders to acquire information. The effect information acquisition has on the probability of a crisis constitutes a novel amplification effect and is explained in detail below.

News about higher unemployment rates or higher borrowing costs increases uncertainty about whether there will be a sovereign debt crisis and, consequently, strategic uncertainty. A debt holder initially expects the sovereign to be in good fiscal condition such that a debt crisis is unlikely, and he expects few other debt holders not to roll over. After adverse news, however, a sovereign debt crisis is more likely and the actions of other debt holders more uncertain. Therefore, information about the sovereign's health becomes more valuable, as it helps a debt holder align his roll over decision with that of other debt holders. Since information supports coordination, a debt holder has a greater incentive to acquire information and more debt holders become informed.

Given the supposed high solvency of the sovereign, a larger proportion of informed debt holders leads to a higher probability of a crisis (lemma 1). Uninformed debt holders,

[4] As shown in the main text in section (1.4), the amplification effect after adverse news holds generally. Thus, the focus on strong fundamentals of the sovereign, arguably the empirically relevant case, is expositional. If fundamentals were weak, then adverse news *reduces* strategic uncertainty. Therefore, the incentives to acquire information is reduced, since there is less need for coordination with other investors. Fewer investors acquire information in equilibrium. This still increases the probability of a crisis since, conditional on adverse fundamentals, informed investors are *more* likely to roll over debt to the sovereign. Taken together, adverse information is amplified via the information acquisition choice of investors.

who have imprecise private information, only refrain from rolling over debt if their private information is very negative. By contrast, informed debt holders, who have precise private information, already refrain from rolling over if their information is mildly negative for two reasons: first, informed debt holders rely more on their private information as it is more precise; and second, they expect a larger proportion of informed debt holders to refrain from rolling over. Therefore, an informed debt holder rolls over less often than an uninformed debt holder for both fundamental and strategic reasons. Taken together, the acquisition of private information after adverse public news amplifies the probability of a sovereign debt crisis.

My second contribution is to show the existence of a novel contagion effect, based on the acquisition of private information, that makes bank runs systemic or sovereign debt crises contagious (proposition 3). I consider an extension that includes public information spillovers, such as news about another bank (the default of Lehman) or another sovereign (debt restructuring in Greece). This affects bank creditors or sovereign debt holders elsewhere, since the profitability of banks or the solvency of sovereigns is positively correlated across regions. Such positive correlation arises from asset commonality, since several banks invested in asset-backed securities or trade, financial, or political links between sovereigns. When the amplification effect is applied to a scenario with the spillover of public information, adverse news about another region changes the incentives to acquire information about the present region, increasing the probability of a financial crisis. Adverse news about Greek sovereign debt, for example, triggers sovereign debt holders of other peripheral European countries to acquire private information about their exposure, increasing the probability of a sovereign debt crisis on these countries. This novel contagion effect can lead to strong spillovers between loosely connected countries.

There is also a small technical contribution. The coordination motive between bank creditors or sovereign debt holders at the roll-over coordination stage translates into a coordination motive in information acquisition choices (lemma 2). Therefore, the incentive to acquire information increases in the proportion of investors who acquire information. Consequently, there are multiple equilibria for an interim value of the information cost (proposition 1). Section 1.1.1 places this result in the context of the closely related literature on multiple equilibria in other coordination games with endogenous information studied previously.

While the focus of the present paper is on bank runs and sovereign debt crises, the novel amplification and contagion effects apply more broadly. Coordination games of regime change have been used extensively to study a range of social and economic phenomena, including models of political regime change. Upon observing a revolution in a neighboring country, citizens acquire information about the strength of their political regime that enables them to coordinate their attempts at revolution (such as in the Arab spring).[5] Another application is in investment complementarity, for instance in foreign direct investment.[6] Adverse news about the macroeconomic fundamentals in the country to be invested in results in foreign direct investors to seek additional information sources to evaluate their investment prospects.

The remainder of this paper is presented as follows. Section 1.1.1 discusses related literature. Section 1.2 describes the model and section 1.3 analyzes its equilibrium. I present the main result of how the acquisition of private information amplifies the probability of a financial crisis in section 1.4. An extension that includes news about a crisis elsewhere is analyzed in section 1.5, and it demonstrates how the acquisition of private information makes a bank run systemic and a sovereign debt crisis contagious. An extension with dispersed costs of becoming informed is considered in section 1.6, which further strengthens the amplification and contagion results. Conclusions are presented in section 1.7, while derivations and proofs can be found in the Appendix (1.8).

1.1.1 Literature

There exists a large literature on contagion. Goldstein and Pauzner (2004) consider a global coordination game with risk-averse speculators invested in two regions. After a crisis in the first region, speculators become more averse to strategic uncertainty and thus

[5] Edmond (2012) studies a coordination game of political regime change with endogenous information manipulation (propaganda) and demonstrates the existence of a unique equilibrium. Circumstances are analyzed under which an information revolution, such as the availability of social media that make both private information more precise and more sources of information available, results in a higher probability of political regime change. In contrast, the present paper considers information acquisition and examines the effect on deteriorating fundamentals on incentives to become informed, which lies at the heart of the proposed amplification and contagion mechanisms.

[6] Dasgupta (2007) studies the effects of an option to delay an investment decision. Delay reduces the payoff from investment, since the best projects are already taken by competitors, but which allows agents to use more accurate information. By contrast, here I consider the effects on incentives to acquire information and the consequences in terms of the amplification of a financial crisis and contagion. See also Chamley (1999).

withdraw their funds from the second region, which is a wealth effect. In contrast, the contagion mechanism proposed in this paper arises from changes in information acquisition choices by risk-neutral agents after adverse news about the first region, providing a fully complementary contagion mechanism. Studying interbank linkages, Allen and Gale (2000) provide a model of financial contagion as an equilibrium outcome. Dasgupta (2004) demonstrates that financial contagion occurs with positive probability in the unique equilibrium of a global game extension of the model proposed by Allen and Gale (2000), focusing on the coordination failure initiated by adverse information. Ahnert and Georg (WP) analyze the effect of ex-post interbank contagion on a bank's ex-ante portfolio choice and study the consequences for systemic risk.[7] In contrast to these papers, I demonstrate that endogenous acquisition of private information after adverse news about another bank increases the probability of a bank run, constituting a novel contagion channel.[8]

Within the global games literature pioneered by Carlsson and van Damme (1993), and further developed by Morris and Shin (2003) and Frankel et al. (2003), Colombo et al. (WP) study the efficiency of endogenous acquisition of private information in a general setting. In an influential paper, Morris and Shin (2002) show that more transparency can be detrimental to welfare, initiating a debate on the social value of public information. Revisiting this result, Colombo et al. (WP) demonstrate that the social value of public information is higher when there is private information acquisition, since the provision of public information crowds out private information acquisition.[9] In contrast, I analyze how changes in the endogenous acquisition of private information amplifies the probability of a crisis.

[7]Systemic risk due common exposures is considered in Acharya and Yorulmazer (2008b), who show that banks can have an ex-ante incentive to correlate their investment decision to avoid information contagion. Allen et al. (2012) analyze systemic risk resulting from the interaction of common exposures and funding maturity through an information channel. Asset commonality can lead to fire sales, and there is a large literature on contagion through such a pecuniary externality which builds on Shleifer and Vishny (1992).

[8]A complementary contagion mechanism arises in Ahnert and Bertsch (WP), where agents observe a crisis elsewhere, but are uncertain about their region's exposure to it. Surprisingly, there can be contagion after good news: the probability of a crisis increases after agents learn that there is zero exposure to the crisis region. By contrast, the ex-post interdependence between regions is non-stochastic and non-zero in my model, and I study how the acquisition choice of private information about regional fundamentals amplifies and spreads banking and sovereign debt crises.

[9]A related paper is Szkup and Trevino (WP) who also analyze the social value of public information in a setting with endogenous private information acquisition similar to the present paper. However, their work is not concerned with the amplification and contagion mechanisms at the heart of this work.

Recent studies analyze the multiplicity of equilibria in coordination games with endogenous information. My model features strategic complementarity in private information acquisition (lemma 2) that leads to multiple equilibria in information acquisition (proposition 1). This complements the closely related result of Hellwig and Veldkamp (2009), who obtain multiple equilibria in a beauty contest coordination game with a binary information choice. Hellwig and Veldkamp (2009) discuss the similarity in the strategic motives between choosing an action and deciding on how much information to acquire in a beauty-contest model. In their words, investors "who want to do what others do, want to know what others know" (p. 223). They also demonstrate that endogenous acquisition of public information gives rise to multiplicity, while uniqueness is obtained with endogenous acquisition of private information. Angeletos and Werning (2006) show that the aggregation of dispersed private information into a publicly observed market price, similar to Grossman and Stiglitz (1980), reestablishes multiplicity of equilibrium - even in global coordination games.[10] Angeletos et al. (2006) and Angeletos and Pavan (WP) examine how the endogenous public information from a policy intervention generates multiple equilibria. By contrast, Zwart (2007) studies the signaling effect of an IMF intervention and obtains a unique equilibrium.

Some recent papers investigate the effect of information acquisition on bank runs. Introducing insolvent banks into a Diamond and Dybvig (1983) setup, Nikitin and Smith (2008) study the effects of costly verification of a bank's solvency. A partial bank run occurs for interim verification costs, where investors become informed and run on insolvent banks only. This is in contrast to the present paper, which studies how news about bank solvency affects the incentives of investors to become informed, and the subsequent repercussions on the probability of a bank run.[11] He and Manela (WP) show that information acquisition after adverse news about a bank's liquidity leads to a run on solvent but illiquid banks. The authors also study how a stress test – a public provision of information about bank solvency – curbs private incentives to acquire information and thus prevents bank runs. The baseline model in the present paper shares the feature of multiple equilibria in information acquisition, but differs in terms of the focus on the coordination aspect of bank runs. That is, a larger proportion of informed agents increases

[10]See also Hellwig et al. (2006) and Ozdenoren and Yuan (2008).

[11]Furthermore, there are a multiple equilibria without information acquisition in Nikitin and Smith (2008), while in my model the equilibrium is unique without information acquisition.

or decreases the probability of a bank run (lemma 1), depending on the strength of the bank's fundamental relative to the strategic effect from coordination.

1.2 Model

The economy is inhabited by a unit continuum of agents indexed by $i \in [0,1]$. Agents play a coordination game of regime change in which a status quo is either maintained ($R = 0$) or abandoned ($R = 1$). Agents simultaneously choose between attacking the status quo by taking actions in favor of regime change ($a_i = 1$) or not attacking ($a_i = 0$). The status quo is abandoned if and only if the aggregate attack size $A \equiv \int_0^1 a_i \, di$ exceeds the strength of the status quo that is parameterized by $\theta \in \mathbb{R}$.

This paper focuses on two specific regime change events or crises: bank failure after a run and sovereign debt restructuring. Agents correspond to bank creditors who withdraw funds from their bank and debt holders who do not roll over short-term government debt, respectively. The strength of the status quo measures a bank's investment profitability, such as the health of its loan book, and a sovereign's solvency, liquidity, or regional competitiveness.

The core of the setup is a coordination motive between agents. The relative payoff from attacking depends on whether the status quo is abandoned and on the relative cost of attacking, parameterized by $r \in (0,1)$, such as foregone interest payments:

$$u(a_i = 1, A, \theta) - u(a_i = 0, A, \theta) = \begin{cases} 1 - r & A \geq \theta \\ & \text{if} \\ -r & A < \theta \end{cases} \tag{1.1}$$

Since behavior is pinned down by the payoff difference, the payoff from not attacking is constant and normalized to zero.[12]

This payoff specification makes the agents' actions strategic complements: an agent's incentive to attack the status quo increases in the mass of attacking agents. Specifically, an individual agent finds it optimal to attack the regime if and only if the probability of

[12] This is common in the literature on global coordination games; see, for example, Angeletos et al. (2006, 2007) and Edmond (2012).

regime change is no less than the cost of attacking. The role of coordination is particularly highlighted if there is complete information about the strength of the status quo. Then, the game has two equilibria in pure strategies for $\theta \in (0, 1)$: no agent attacks and the status quo is maintained ($A^* = 0 < \theta$), or all agents attack and the status quo is abolished ($A^* = 1 > \theta$). There is a unique equilibrium in dominant strategies for extreme values of the strength of the status quo. All agents attack $A^* = 1$ if the regime is weak ($\theta \leq 0$), while all agents refrain from attacking $A^* = 0$ if the regime is strong ($\theta \geq 1$).

Agents have incomplete information about the strength of the status quo, sharing a common prior:

$$\theta \sim \mathcal{N}\left(\mu, \frac{1}{\alpha}\right) \tag{1.2}$$

For example, such a common prior is induced by a public signal about the strength of the status quo: $\mu \equiv \theta + \eta$, where the independent noise term η is normally distributed with mean zero and precision $\alpha \in (0, \infty)$, and agents initially had an improper uniform prior. The mean of the common prior μ is also referred to as the fundamental. Following the global game literature pioneered by Carlsson and van Damme (1993), agents have dispersed private information about the strength of the status quo upon receiving a noisy private signal:

$$x_i \equiv \theta + \epsilon_i \tag{1.3}$$

where the idiosyncratic noise ϵ_i is identically and independently normally distributed across agents with zero mean, uninformed precision $\gamma_U \in (0, \infty)$, and is also independent of the fundamental. All distributions are common knowledge.

The game has two stages. At the coordination stage, agents simultaneously decide whether to attack the status quo upon observing the private signal. The coordination stage is preceded by an information acquisition stage during which agents simultaneously decide whether to become informed at a cost $c > 0$.[13] The information cost captures the financial cost of acquiring information and the cost of resources to process information.[14] Specific examples are the cost of hiring analysts and investment in IT infrastructure to improve data analysis. Informed agents will receive a more precise private signal at the coordination stage, raising their precision to $\gamma_I \in (\gamma_U, \infty)$. Let $n_i \in \{I, U\}$ denote

[13] See section 1.6 for an extension with ex-ante heterogeneity in the cost of becoming informed.

[14] This cost may also capture the mental cost of information processing, as highlighted in the rational inattention literature.

the information acquisition choice of agent i and $n = \int_0^1 \mathbf{1}\{n_i = I\}di$ the aggregate proportion of informed agents. Agents who acquire information are also called informed agents and the subscript $z \in \{I, U\}$ is used to distinguish an agent's information choice. The following timeline summarizes the above:

Stage 1: Information acquisition stage

- Agents have a common prior about the strength of the status quo.
- Agents simultaneously decide whether to become informed at a cost.

Stage 2: Coordination stage

- Agents receive a private signal about the strength of the status quo.
  - Informed agents obtain a more precise private signal.
- Agents simultaneously decide whether to attack.
- The state of the regime is observed and payoffs are realized.

1.3 Equilibrium

A Perfect Bayesian Equilibrium consists of: (i) individual information acquisition choices $n_i^* \in \{I, U\}$ for each agent; (ii) the aggregate proportion of informed agents $n^* \in [0, 1]$; (iii) individual attack choices $a_i^* \in \{0, 1\}$ for each agent at the coordination stage; and (iv) the aggregate mass of attacking agents $A^* \in [0, 1]$. The consistency of individual choices with aggregate proportions requires:

$$n^* = \int_0^1 \mathbf{1}\{n_i^* = I\}di \tag{1.4}$$

$$A^* = \int_0^1 a_i^* di \tag{1.5}$$

Furthermore, the behavior of each agent is individually rational. At stage 1, each agent optimally decides whether to acquire information, which depends on the aggregate proportion of informed agents, the common prior, and the information cost: $n_i^* = n_i^*(n^*; \mu, c)$.

At stage 2, each agent optimally decides whether to attack, which depends on the information choice, the aggregate proportion of informed agents, and the private information received: $a_i^* = a_i^*(n_i^*, n^*; x_{i2})$. The specifics about the individually optimal behavior will be described when analyzing the decisions at each stage.

To construct the equilibrium, taking a sequential approach is useful. I start by deriving the optimal behavior at the coordination stage for any given proportion of informed agents n^*, which summarizes the behavior at the information acquisition stage. Next, I derive the optimal behavior at the information acquisition stage. My main interest is to explore the link between information acquisition and coordination, and to analyse the effect of changes in the common prior that give rise to the amplification and contagion mechanisms presented in sections 1.4 and 1.5, respectively.

1.3.1 Stage 2: Coordination stage

For a given set of information choices $\{n_i^*\}_{i\in[0,1]}$, the coordination stage is a standard game of imperfect information, once an agent's signal x_i is established as his type. A strategy s_i is a mapping from the signal into the binary action space: $s_i : \mathbb{R} \to \{0, 1\}$ for a given individual information choice n_i^* and the aggregate proportion of informed agents n^*. An agent's expected utility from attacking conditional on his private information x_i, his information choice n_i^*, and the aggregate proportion of informed agents n^* is $E[u(a_i = 1)|x_i, n_i^*, n^*] = -r + \Pr\{A(s_{-i}) \geq \theta | x_i, n_i^*, n^*\}$. As each agent is atomistic, the aggregate attack size is unaffected by the individual attack decision. Optimality for agent i at the coordination stage requires that strategy s_i maximizes his conditional expected utility, taking all other agents' strategies s_{-i}^* as given.

I focus on symmetric equilibria in monotone threshold strategies at the coordination stage throughout. This is without loss of generality, as shown with an argument based on iterated deletion of strictly dominated strategies by Frankel et al. (2003), Morris and Shin (2003), and Goldstein (2005). Thus, the equilibrium at stage 2 for a given proportion of informed agents is fully characterized by an **attacking threshold** for the informed and uninformed agent $(\overline{x}_I(n^*), \overline{x}_U(n^*))$ and a **aggregate threshold** $(\overline{\theta}(n^*))$. Each agent optimally follows a threshold strategy, whereby an agent i attacks the status quo if and only if his signal is below an attacking threshold. This threshold depends on whether the agent is informed or uninformed and the aggregate proportion of informed agents:

$$a_i^* = 1 \Leftrightarrow x_i \leq \overline{x}(n_i^*, n^*) \equiv \overline{x}_I(n^*)\mathbf{1}\{n_i^* = I\} + \overline{x}_U(n^*)\mathbf{1}\{n_i^* = U\} \tag{1.6}$$

Turning to the aggregate level, the status quo is abandoned if and only if it is weaker than the aggregate threshold:

$$R_i^* = 1 \Leftrightarrow \theta \leq \overline{\theta}(n^*) \tag{1.7}$$

These three thresholds are determined by a critical mass condition at the aggregate level and an indifference condition for both the informed and uninformed agent as derived below.

First, an agent i uses his private signal to form a posterior about the strength of the status quo:

$$\theta | x_i, n_i^* = z \sim \mathcal{N}\left(\frac{\alpha\mu + \gamma_z x_i}{\alpha + \gamma_z}, \frac{1}{\alpha + \gamma_z}\right)$$

where normality is preserved, the posterior mean is a weighted average of the common prior and the private signal, and the posterior precision is the sum of the precisions of the prior and of the signal (see DeGroot (1970), for example). An informed agent ($z = I$) receives a more precise private signal and relies on it more than an uninformed agent. Furthermore, the precision of the posterior is higher for an informed agent. Next, an agent with information choice n_i^* assigns the following probability to an abandonment of the status quo:

$$\Pr\{\theta \leq \overline{\theta}(n^*) | x_i, n_i^* = z\} = \Phi\left(\sqrt{\alpha + \gamma_z}\left[\overline{\theta}(n^*) - \frac{\alpha\mu + \gamma_z x_i}{\alpha + \gamma_z}\right]\right)$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. As the probability of regime change strictly decreases in an agent's signal, a threshold strategy is indeed optimal.

An agent who receives the threshold signal $x_i = \overline{x}_z(n^*)$ is indifferent between attacking and not attacking the status quo. Therefore, the **indifference conditions** state that the probability of regime change evaluated at the attacking threshold must equal the cost of attacking for both informed and uninformed agents, which yields the attacking thresholds:

$$\overline{x}_z(n^*) = \overline{\theta}(n^*) + \frac{\alpha}{\gamma_z}[\overline{\theta}(n^*) - \mu] - \frac{\sqrt{\alpha + \gamma_z}}{\gamma_z}\Phi^{-1}(r) \tag{1.8}$$

Since all agents play the threshold strategy with thresholds $\overline{x}_I(n^*)$ if informed, and $\overline{x}_U(n^*)$ if uninformed, the aggregate attack size for any realised strength of the status quo θ is:

$$\begin{aligned} A(n^*; \overline{x}_I(n^*), \overline{x}_U(n^*), \theta) &= \int_0^1 \mathbf{1}\{x_i \leq \overline{x}_z(n^*) | \theta, n_i^* = z\} di \\ &= n^*\Phi\left(\sqrt{\gamma_I}\left[\overline{x}_I(n^*) - \theta\right]\right) + (1 - n^*)\Phi\left(\sqrt{\gamma_U}\left[\overline{x}_U(n^*) - \theta\right]\right) \end{aligned}$$

by the law of large numbers. The **critical mass condition** states that the aggregate attack size is just sufficient to bring down the status quo when its strength equals the aggregate threshold:

$$\overline{\theta}(n^*) = A\left(n^*; \overline{\theta}(n^*), \overline{x}_I(n^*), \overline{x}_U(n^*)\right) \tag{1.9}$$

Combining indifference conditions with the critical mass condition, the equilibrium aggregate threshold at the coordination stage $\overline{\theta} = \overline{\theta}(n^*; \mu, r)$, for any given proportion of informed agents, is implicitly defined by:

$$\begin{aligned} \overline{\theta}(n^*) &= n^*\Phi\left(\frac{\alpha}{\sqrt{\gamma_I}}\left[\overline{\theta}(n^*) - \mu\right] - \sqrt{1 + \frac{\alpha}{\gamma_I}}\Phi^{-1}(r)\right)\cdots \\ &\quad \cdots + (1 - n^*)\Phi\left(\frac{\alpha}{\sqrt{\gamma_U}}\left[\overline{\theta}(n^*) - \mu\right] - \sqrt{1 + \frac{\alpha}{\gamma_U}}\Phi^{-1}(r)\right) \end{aligned} \tag{1.10}$$

Similar to Morris and Shin (2003), a unique solution to equation (1.10) for any given n^* is ensured by a sufficiently precise private signal of the uninformed agent. Under this condition, the slope of the left-hand side of equation (1.10) exceeds the slope of the right-hand side for any proportion of informed agents, ensuring at most one crossing:

$$\gamma_U > \underline{\gamma}_U \equiv \frac{\alpha^2}{2\pi} \Rightarrow 1 > A_\theta \equiv \frac{\partial A(n^*, \overline{\theta}(n^*))}{\partial \theta} \;\forall\, n^* \in [0, 1]$$

Uniqueness follows from $A \in [0, 1]$ and the fact that the realised strength of the status quo is unbounded. Once the unique aggregate threshold is determined, the attack thresholds $\overline{x}_z(n^*)$ are backed out from the indifference conditions.

The object of interest is the probability of a crisis, which strictly increases in the aggregate threshold:

$$\Pr\{\theta \leq \overline{\theta}(n^*)\} = \Phi\left(\sqrt{\alpha}[\overline{\theta}(n^*) - \mu]\right) \tag{1.11}$$

For a given proportion of informed agents, the outcome of the coordination stage is characterized by the fundamental captured by the common prior μ, and considerations about the actions of other players captured by the cost of attacking r. Understanding how the aggregate threshold at the coordination stage varies with the fundamental and the cost of attacking is useful for constructing the equilibrium at the information acquisition stage. It also highlights the role of endogenous information acquisition. As explained below, a financial crisis is less likely if the status quo is stronger or the cost of attacking higher.

First, the equilibrium aggregate threshold decreases in the cost of attacking ($\frac{\partial\overline{\theta}(n^*)}{\partial r} < 0$). This *strategic effect* relates to an agent's incentive to attack the status quo based on other agents' actions. If the attack cost r is high, few other agents attack, and the expected aggregate attack size is low. Therefore, an individual agent tends not to attack the status quo himself. In summary, the strategic effect suggests that an agent has a small incentive to attack the status quo when the attack cost is high. In other words, even a weak regime is maintained if agents do not expect other agents to attack.

Second, the equilibrium aggregate threshold decreases in the common prior ($\frac{\partial\overline{\theta}(n^*)}{\partial \mu} = -\frac{A_\theta(\overline{\theta})}{1-A_\theta(\overline{\theta})} < 0$). This *fundamental effect* relates to an agent's incentive to attack the status quo based on its commonly believed strength. If the common prior is strong, the required aggregate attack size to abandon the regime is high, inducing an individual agent not to attack the status quo himself. In summary, the fundamental effect suggests that an agent has a small incentive to attack the status quo when it is commonly believed to be strong. In other words, even a weak regime is maintained if agents expect it to be strong.

The following parameter constraint ensures that a change in the proportion of informed agents affects the probability of crisis. That is, the *fundamental effect* and the *strategic effect* never cancel as the proportion of informed agents changes.

Condition 1. *Parameter constraint A:*

$$\mu \neq \hat{\mu}(r) \equiv \Phi\left(\kappa_1 \Phi^{-1}(r)\right) - \kappa_0 \Phi^{-1}(r) \tag{1.12}$$

$$\kappa_0 \equiv \frac{\sqrt{\gamma_I(\alpha+\gamma_U)} - \sqrt{\gamma_U(\alpha+\gamma_I)}}{\alpha[\sqrt{\gamma_I} - \sqrt{\gamma_U}]} \tag{1.13}$$

$$\kappa_1 \equiv \frac{\alpha\kappa_0}{\sqrt{\gamma_I}} - \sqrt{1-\frac{\alpha}{\gamma_I}} = \frac{\alpha\kappa_0}{\sqrt{\gamma_U}} - \sqrt{1-\frac{\alpha}{\gamma_U}} \tag{1.14}$$

Note that $\frac{\partial\hat{\mu}}{\partial r} < 0$ and $\hat{\mu} \to \infty$ if $r \to 0$ as well as $\hat{\mu} \to -\infty$ if $r \to 1$. Parameter constraint A excludes a part of the parameter space with zero measure only.

The responsiveness of the aggregate threshold at the coordination stage to changes in the proportion of informed agents is summarized in lemma 1. This result, which links the information acquisition stage to the coordination stage, is useful for continuing the construction of equilibrium at the information acquisition stage. Furthermore, it states a condition under which a larger proportion of informed agents increases the probability of a financial crisis.

Lemma 1. *Suppose that the private signal of the uninformed agent is sufficiently precise ($\gamma_U > \overline{\gamma}_U$). If parameter constraint A holds ($\mu \neq \hat{\mu}(r)$), then the threshold fundamental at the coordination stage responds to changes in the proportion of informed agents:*

$$\frac{\partial\overline{\theta}}{\partial n^*} \neq 0 \tag{1.15}$$

Furthermore, the threshold fundamental increases (or decreases) in the proportion of informed agents if the fundamental lies above (or below) the line $\hat{\mu}(r)$ defined by parameter constraint A:

$$\frac{\partial\overline{\theta}}{\partial n^*}\left(\mu - \hat{\mu}(r)\right) > 0 \tag{1.16}$$

See Appendix (1.8.1) for a proof. Figure 1.1 illustrates parameter constraint A and the results of lemma 1.

A larger proportion of informed agents implies a reduced reliance on the common prior, since more agents have access to precise private information. To illustrate this point, consider the special case of strong fundamentals ($\mu > \hat{\mu}(r)$), where the crisis probability increases in the proportion of informed agents. Given the high common prior about the strength of the status quo relative to the attack cost, an uninformed agent only

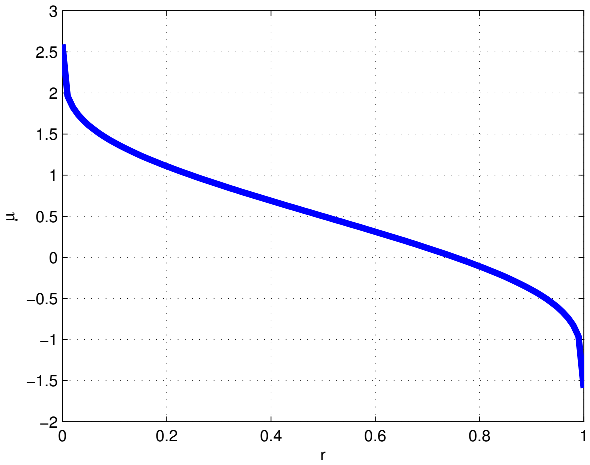

Figure 1.1: The chart depicts the combinations of the fundamental (μ) and cost of attacking (r) that *violate* parameter constraint A: $\mu = \hat{\mu}(r)$. If the regime is strong ($\mu > \hat{\mu}(r)$), then the aggregate threshold at the coordination stage increases in the proportion of informed agents ($\frac{\partial \overline{\theta}}{\partial n^*} > 0$). By contrast, the aggregate threshold decreases in the proportion of informed agents ($\frac{\partial \overline{\theta}}{\partial n^*} < 0$) if the regime is weak ($\mu < \hat{\mu}(r)$). Parameter values are $\alpha = 2$, $\gamma_U = 3$, and $\gamma_I = 4$.

attacks the status quo upon receiving a "particularly low" private signal. In contrast, an informed agent also attacks for "moderately low" private signals. This is for two reasons: (i) an informed agent relies more on the private signal relative to the prior, thus expecting the status quo to be weaker; and (ii) an informed agent expects more informed agents to join the attack as they must have received similar moderately low signals. Both the fundamental and the strategic reasons cause an informed agent to attack for a greater range of private signals than the uninformed agent. Therefore, a larger proportion of informed agents make a crisis more likely if the fundamental is strong.

1.3.2 Stage 1: Information acquisition stage

Equipped with the results from the coordination stage, I now turn to the information acquisition stage. To evaluate the incentives to become informed, the expected utilities of an informed ($n_i = I$) and uninformed ($n_i = U$) agent are compared. The expected utility has two terms. An agent receives the payoff $(1 - r)$ if he attacks ($x_i \leq \overline{x}_z$) when the status quo is abandoned ($\theta \leq \overline{\theta}$), while he incurs the cost of attacking r if he attacks when the status quo is maintained ($\theta > \overline{\theta}$):

$$\begin{aligned} EU^I(n^*) &= (1-r)\int_{-\infty}^{\overline{\theta}}\int_{-\infty}^{\overline{x}_I} f^I(x|\theta)dx\,dG(\theta) - r\int_{\overline{\theta}}^{\infty}\int_{-\infty}^{\overline{x}_I} f^I(x)dx\,dG(\theta) \\ EU^U(n^*) &= (1-r)\int_{-\infty}^{\overline{\theta}}\int_{-\infty}^{\overline{x}_U} f^U(x|\theta)dx\,dG(\theta) - r\int_{\overline{\theta}}^{\infty}\int_{-\infty}^{\overline{x}_U} f^U(x)dx\,dG(\theta) \end{aligned}$$

where $G(\theta)$ is the cumulative distribution function of the fundamental (that is distributed as $\mathcal{N}\left(\mu, \frac{1}{\alpha}\right)$), while $f^z(x)$ is the probability distribution function of private signals conditional on a realized fundamental θ and on the information choice z (that is distributed as $\mathcal{N}(\theta, \frac{1}{\gamma_z})$). The dependence of the three thresholds on the proportion of informed agents is suppressed for brevity.

Becoming informed, or receiving a more precise private signal, has two benefits. First, it allows an agent to form a more precise posterior about the fundamental. Second, it allows an agent to form a more precise posterior about the size of the aggregate attack. This follows from the fact that an informed agent has a more precise posterior about other agents' signals than an uninformed agent does. The benefit from becoming informed is measured in terms of the expected utility difference denoted by $D(n^*) \equiv EU^I(n^*) - EU^U(n^*)$, which depends on the aggregate proportion of informed agents:

$$D(n^*) = \underbrace{r\int_{\overline{\theta}}^{\infty}\Gamma(\theta)g(\theta)d\theta}_{\text{Type-I error}} - \underbrace{(1-r)\int_{-\infty}^{\overline{\theta}}\Gamma(\theta)g(\theta)d\theta}_{\text{Type-II error}} \tag{1.17}$$

$$\Gamma(\theta) \equiv \int_{-\infty}^{\overline{x}_U} f^U(x|\theta)dx - \int_{-\infty}^{\overline{x}_I} f^I(x|\theta)dx \tag{1.18}$$

where $\Gamma(\theta)$ captures the difference between the probability of an attacking uniformed agent and the probability of an attacking informed agent for a given realization of the fundamental.

Informed agents make fewer errors in their attacking decisions. To see this, observe that the benefit from being informed consists of two terms. The first term states that an informed agent attacks the status quo less often when there is no regime change (type I error). This term is proportional to the cost of attacking unsuccessfully. The second term states that an informed agent refrains from attacking the status quo less often when there is a regime change (type II error). This term is proportional to the foregone payoff from attacking successfully. To understand why informed agents make fewer errors in their attacking decisions, as reflected in the benefit from becoming informed, consider the special case of vanishing noise in the private signal if informed ($\gamma_I \to \infty$). Then, an informed agent never makes an error: he never attacks when the regime is maintained and always attacks when it is abandoned. Subsequently, the benefit from becoming informed in this special case reduces to the sum of the type I and type II errors of uninformed agents.

There is strategic complementarity in information acquisition choices. In other words, the benefit from becoming informed increases in the proportion of informed agents. Thus, the strategic complementarity in action present at the coordination stage translates into strategic complementarity at the information acquisition stage. In the words of Hellwig and Veldkamp (2009), who study a different coordination game: "Agents who want to do what others do, want to know what others know." (p. 223). This result carries over to a global game of regime change studied in the present paper and is summarized in lemma 2.

Lemma 2. *Suppose that the private signal of the uninformed agent is sufficiently precise ($\gamma_U > \overline{\gamma}_U$). Then, there is strategic complementarity in information acquisition:*

$$\frac{\partial D}{\partial n^*} \geq 0 \tag{1.19}$$

If parameter constraint A holds, the inequality is strict.

See Appendix (1.8.2) for a proof.

In order to understand the strategic complementarity in information acquisition, recall the two effects of becoming informed. First, informed agents forecast the fundamental θ more precisely, but this effect is independent of the proportion of informed agents. Second, informed agents forecast the behavior of other agents more precisely. In particular, an informed agent is better at forecasting other informed agents than other uninformed agents. Therefore, the second benefit of becoming informed increases in the proportion of informed agents, establishing the strategic complementarity in information acquisition. In short, information supports coordination.

Further insight can be obtained by comparing this paper to Grossman and Stiglitz

(1980). These authors study an economy with costly private acquisition of information about the value of a risky assets and the effect on subsequent trade in this asset. Uninformed traders only learn by observing the market price, which is more informative the larger the proportion of informed traders. Consequently, the private incentive to become informed decreases in the proportion of informed agents (strategic substitutes), since agents do not wish to incur the information cost if the publicly available market price reveals large amount of private information. In contrast, an agent's incentive to acquire information increases in the proportion of informed agents (strategic complements) in this paper, since information supports coordination. If many agents are informed, acquiring private information helps an agent align his actions with the action of others at the coordination stage.[15]

To construct equilibrium, it is useful to determine the boundaries of the benefit from becoming informed, summarized in lemma 3.

Lemma 3. *Suppose that the private signal of the uninformed agent is sufficiently precise ($\gamma_U > \overline{\gamma}_U$). The benefit from becoming informed is positive but smaller than unity:*

$$0 < D(n^*) < 1 \tag{1.20}$$

See Appendix (1.8.3) for a proof.

Consider an agent's optimal information acquisition choice. Given the binary action $n_i \in \{I, U\}$ at stage 1, and the fact that each agent is atomistic with no effect on the aggregate proportion of informed agents, an agent optimally acquires information if and only if the information cost is at most the benefit from becoming informed:

$$n_i^* = I \Leftrightarrow c \leq D(n^*) \tag{1.21}$$

where the benefit from becoming informed depends on the proportion of informed agents.[16] Thus, it is dominant for an agent to acquire information when the information cost is low ($c < D(0)$) and not to acquire information when the information cost is high ($c > D(1)$).

[15]Another difference relates to the existence of equilibrium. While Grossman and Stiglitz (1980) demonstrate the non-existence of a competitive equilibrium for a vanishing information cost, a unique equilibrium exists for that limiting case in the present paper.

[16]Heterogeneity in the information cost is analysed in section 1.6.

An agent's optimal information acquisition choice depends on the proportion of informed agents for an interim information cost ($D(0) \leq c \leq D(1)$), where an agent is indifferent between acquiring and not acquiring information if and only if $c = D(n^*)$.

The equilibrium at the information acquisition stage is constructed by combining these individual optimality conditions, the previous two lemmas, and the consistency between individually optimal information acquisition choices and the aggregate proportion of informed agents. It follows directly that there always exists an equilibrium and that the information cost determines the number of equilibria. Building on this, proposition 1 describes the equilibrium behavior.

Proposition 1. ***(Multiplicity)*** *Suppose that an uninformed agent's signal is sufficiently precise ($\gamma_U > \underline{\gamma}_U$) and that parameter constraint A holds ($\mu \neq \hat{\mu}(r)$). Then, the optimal behavior of agents at the coordination stage is uniquely pinned down for a given proportion of informed agents and is characterized by a threshold strategy. An informed (or uninformed) agent attacks if and only if his private signal falls short of the threshold $\overline{x}_I(n^*)$ (or $\overline{x}_U(n^*)$). There is a regime change if and only if the realized strength of the status quo falls short of the aggregate threshold ($\overline{\theta}(n^*)$). These thresholds are determined in equations (1.8) and (1.10).*

The number of equilibria in the overall game depends on the information cost:

- *If the information cost is low ($c < D(0)$), then there exists a unique symmetric equilibrium in which all agents acquire information ($n^* = 1$).*
- *If the information cost is high ($c > D(1)$), then there exists a unique symmetric equilibrium in which no agent acquires information ($n^* = 0$).*
- *If the information cost takes an interim value ($D(0) \leq c \leq D(1)$), then there exist three equilibria. The two symmetric equilibria described above prevail and there is also an asymmetric equilibrium. In the asymmetric equilibrium, agents are indifferent between becoming informed and remaining uninformed, and the aggregate proportion of informed agents is determined by the indifference of the marginal agent to become informed: $n^* = D^{-1}(c)$.*

Given the strict monotonicity of the benefit from becoming informed in the proportion of informed agents, the proportion of informed agents is uniquely determined in the asymmetric equilibrium. However, I will focus on the stable symmetric equilibria.

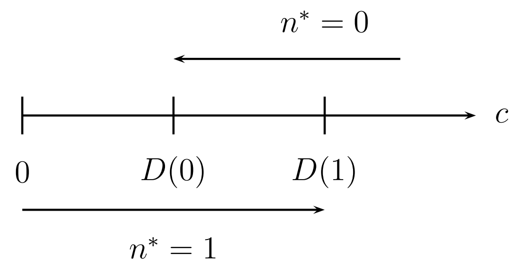

Figure 1.2: A visualization of the equilibrium with information acquisition. The information cost c and the benefit from becoming informed D are on the axis.

Similar to the case without information acquisition and with complete information, there is a unique equilibrium for extreme values of information cost, while there are multiple equilibria for interim values. Figure 1.2 illustrates the link between the information cost and the number of equilibria.

1.4 Amplification

This section shows how endogenous acquisition of private information after adverse news about the fundamental amplifies the probability of a crisis. Adverse news about the fundamental is modelled as a reduction in the common prior μ and has two consequences. First, there is the standard effect on the probability of a crisis given by the fundamental effect that increases the crisis probability ($\frac{\partial\overline{\theta}}{\partial\mu} < 0$). With endogenous information acquisition, there is an additional novel effect via the benefit from becoming informed.

Formally, the benefit from becoming informed changes with the mean of the fundamental, as shown in Appendix (1.8.4):

$$\frac{\partial D}{\partial \mu} = -g(\overline{\theta})\frac{\partial\overline{\theta}}{\partial n^*} \tag{1.22}$$

If the fundamental is strong ($\mu > \hat{\mu}(r)$) – above the line in figure 1.1 – a larger proportion of informed agents leads to a higher probability of crisis ($\frac{\partial\overline{\theta}}{\partial n^*} > 0$). Therefore, the benefit from becoming informed increases ($\frac{\partial D}{\partial \mu} < 0$). Since information supports coordination, becoming informed is particularly valuable when the common prior moves towards an interim value, which is the line $\hat{\mu}(r)$ specified by parameter constraint A. On this line, the behavior of other agents is least certain, that is strategic uncertainty is at its highest. Intuitively, information is more valuable the greater the uncertainty about other agents' behavior, as in the case of an interim common prior. By contrast, information is of little value if the behavior of other agents is quite certain, as in the case of a particularly high (or low) fundamental. Taken together, reducing the mean of the fundamental from a high level increases strategic uncertainty and therefore the benefit from becoming informed.

Changes to the benefit from becoming informed affect the equilibrium proportion of informed agents, as illustrated by figure 1.3. Specifically, there exists a range of information costs for which information acquisition ($n^* = 1$) becomes the unique equilibrium, where it was one of several equilibria previously. Similarly, there is a range of information costs for which information acquisition ($n^* = 1$) becomes one of several equilibria, where it was not an equilibrium previously. The equilibrium proportion of informed agents weakly increases for these information cost ranges.[17] If more agents become informed, such as a switch from $n^* = 0$ to $n^* = 1$, the probability of crisis increases. If parameter constraint A holds and the equilibrium proportion of informed agents changes, the probability of a crisis increases strictly.

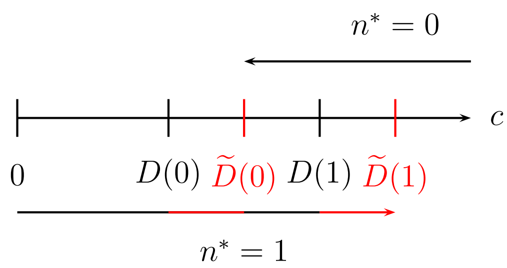

Figure 1.3: A visualization of the equilibrium with information acquisition equilibrium before and after a reduction in the common prior. Tildes are used to distinguish the latter case. The information cost c and various benefits from becoming informed D are on the axis.

Adverse news about the fundamental also increases the probability of a crisis if

[17]As shown in section 1.6, the equilibrium proportion of informed agents strictly increases if there is heterogeneity about the information cost.

the common prior is low ($\mu < \hat{\mu}(r)$) – below the line in figure 1.1. In this case, a smaller proportion of informed agents leads to a higher probability of crisis ($\frac{\partial \overline{\theta}}{\partial n^*} < 0$). Therefore, the benefit from becoming informed decreases ($\frac{\partial D}{\partial \mu} > 0$). Hence, there is a range of information costs for which no information acquisition ($n^* = 0$) becomes the unique equilibrium, where it was one of several equilibria previously. Similarly, there is a range of information costs for which no information acquisition ($n^* = 0$) becomes one of several equilibria, where it was not an equilibrium previously. The equilibrium proportion of informed agents weakly decreases for these information cost ranges. If fewer agents become informed, such as a switch from $n^* = 1$ to $n^* = 0$, the probability of crisis increases. If parameter constraint A holds and the equilibrium proportion of informed agents changes, the probability of a crisis increases strictly.

Proposition 2 summarizes the amplification effect after adverse news about the fundamental that reduces the common prior.

Proposition 2. ***(Amplification)*** *Suppose that an uninformed agent's signal is sufficiently precise ($\gamma_U > \underline{\gamma}_U$). Endogenous acquisition of private information amplifies the probability of a financial crisis:*

$$\frac{\Delta \overline{\theta}}{\Delta n^*} \frac{\Delta n^*}{\Delta D} \frac{\Delta D}{\Delta \mu} = -g(\overline{\theta}) \left(\frac{\Delta \overline{\theta}}{\Delta n^*} \right)^2 \frac{\Delta n^*}{\Delta D} \leq 0 \tag{1.23}$$

The inequality is strict if both parameter constraint A holds and there is a change in the equilibrium proportion of informed agents.

An important insight from this paper is that information supports coordination by allowing an agent to align his action with that of other agents at the coordination stage. But how is the incentive to acquire information affected by adverse news that reduces the common prior? The higher the degree of uncertainty, the more useful information is. Intuitively, there is not much need for information if the behavior of other agents is almost certain, as happens for extremely high or low common priors. In contrast, strategic uncertainty is highest for an interim common prior, that is when μ is close to $\hat{\mu}(r)$ specified by parameter constraint A. Therefore, the benefit from becoming informed increases when the common prior was high initially and is then reduced by adverse news. As strategic uncertainty is higher, the benefit from becoming informed increases, inducing information acquisition.

1.5 Contagion

This section considers an extension with a spillover of public information about another region, such as news about the default of Lehman or about a debt restructuring in Greece. I examine how public information spillovers affect incentives to acquire private information about the health of a related bank or sovereign. More specifically, I show that the probability of a crisis in one region increases after adverse news about another region due to changes in the incentives to acquire private information. In short, I provide a novel contagion mechanism based on endogenous acquisition of information.

Forbes (2012) suggests that there is a distinction between interdependence and contagion. Interdependence is defined as the correlation between regions in all states of the world, whereas contagion is the spillover of adverse shocks across (potentially interdependent) regions. My model features contagion in this broad definition, since adverse news about another region directly raises the probability of a crisis. However, the model also features contagion in a narrower sense that goes beyond interdependence. Adverse news induces the acquisition of private information that increases the probability of a crisis. Therefore, private information acquisition can lead to strong contagion between otherwise loosely connected regions.

Let us turn to the specification of the spillover of public information. There is another region, such as another bank or sovereign, with fundamental θ_2. Before making their private information acquisition choice, agents receive a noisy public signal about the other region y:[18]

$$
\begin{aligned}
y &\equiv \theta_2 + \nu && (1.24)\\
\nu &\sim \mathcal{N}\left(0, \frac{1}{\beta}\right) && (1.25)
\end{aligned}
$$

where the noise ν is independent of fundamental in second region θ_2, and the precision β captures the intensity of media coverage of the other region's event or the quality of a public announcement in the spirit of Morris and Shin (2002). Fundamentals are correlated between regions ($\rho \in (0, 1)$) to capture asset commonality between banks (such as joint investment in asset-backed securities) and trade, financial, or political links

[18]For example, agents play an otherwise identical coordination game of regime change in the other region. Once the aggregate attack size A_2 is observed, the strength of the status quo is inferred.

among sovereigns:[19]

$$\begin{aligned} \theta_2 &\equiv \rho\theta + (1-\rho)\mu + \xi \\ \xi &\sim \mathcal{N}\left(0, \frac{1}{\alpha_\xi}\right) \end{aligned} \tag{1.26}$$

where the fundamentals have the same mean and share the same precision if the noise term is scaled accordingly ($\alpha_\xi = \frac{\alpha}{1-\rho^2}$).

The public information about the other region changes the common prior among agents about the present region because of correlated fundamentals:

$$\begin{aligned} \theta|y &\sim \mathcal{N}\left(\tilde{\mu}, \frac{1}{\tilde{\alpha}}\right) \\ \tilde{\mu} &\equiv \mu + \frac{\rho}{1+\alpha/\beta}[y-\mu] \\ \tilde{\alpha} &\equiv \alpha\frac{1+\alpha/\beta}{1-\rho^2+\alpha/\beta} \end{aligned} \tag{1.27}$$

where receiving a public signal increases the precision of the common prior and shifts it towards the signal. The common prior is unchanged, however, if fundamentals are uncorrelated ($\rho \to 0$), or the public signal y is imprecise ($\beta \to 0$). Parameter constraint A generalizes accordingly, where (μ, α) are replaced by $(\tilde{\mu}, \tilde{\alpha})$. This restriction on the parameter space still has zero measure.

Adverse news about the other region (a lower y) has two effects. First, standard information spillover based on correlated fundamentals (interdependence or broad contagion according to Forbes (2012)) is a direct effect that is also present in models without information acquisition. It unambiguously raises the probability of a crisis, mirroring the previous fundamental effect:

$$\frac{\partial\overline{\theta}}{\partial y} < 0 \tag{1.28}$$

There is also a second effect if information is endogenous. Adverse news about another region, such as learning about a crisis elsewhere, is a wake-up call that induces agents to become informed. Indeed, adverse news increases an agent's incentive to become

[19]In case of the ongoing euro zone sovereign debt crisis, the correlation across countries also originates from the scarcity of rescue funds, such as the European Financial Stability Facility (EFSF) and the European Stability Mechanism (ESM).

informed for strong fundamentals ($\tilde{\mu} > \hat{\tilde{\mu}}$):

$$\frac{\partial D}{\partial y} = -\frac{\rho}{1+\frac{\alpha}{\beta}}\, g(\overline{\theta})\, \frac{\partial \overline{\theta}}{\partial n^*} \tag{1.29}$$

Intuitively, becoming informed particularly pays off when there is more uncertainty about other agents' behaviour, as in the run-up to a crisis. In turn, the probability of a crisis increases in the proportion of informed agents, since informed investors react more strongly to bad private signals than uninformed investors. Therefore, the endogenous acquisition of private information triggered by adverse news about a crisis elsewhere increases the probability of a crisis in the present region.

The novel contagion result based on endogenous information acquisition is summarized in proposition 3 and illustrated in figure 1.4.

Proposition 3. ***(Contagion)*** *Suppose that an uninformed agent's signal is sufficiently precise ($\gamma_U > \underline{\gamma}_U$). Endogenous acquisition of private information increases the probability of a financial crisis after observing adverse news about another region:*

$$\frac{\Delta \overline{\theta}}{\Delta n^*}\,\frac{\Delta n^*}{\Delta D}\,\frac{\Delta D}{\Delta y} = -\frac{\rho}{1+\frac{\alpha}{\beta}} g(\overline{\theta}) \left(\frac{\Delta \overline{\theta}}{\Delta n^*}\right)^2 \frac{\Delta n^*}{\Delta D} \leq 0 \tag{1.30}$$

The inequality is strict if both the adjusted parameter constraint A holds and there is a change in the equilibrium proportion of informed agents.

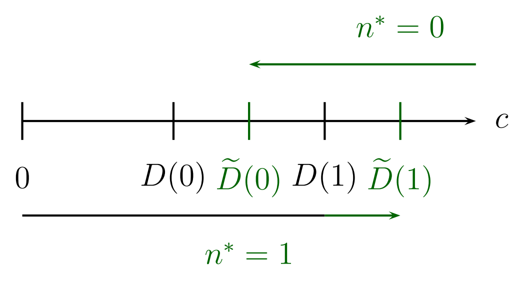

Figure 1.4: A visualization of the equilibrium with information acquisition after learning about a crisis elsewhere.

In effect, here I have generalized the amplification effect presented in the previous section to the case of public information spillovers, provided the additional public information is not uninformative ($\beta > 0$) and the regions are positively correlated ($\rho > 0$). The argument also extends to negatively correlated fundamentals.[20] Then, good news about another region is bad news for a given region, possibly due to a competition effect. The proposed contagion effect prevails after good news about another region. Furthermore, adverse news about another region is good news for agents in the present region that will affect their information acquisition in such a way that the probability of a crisis decreases. Again, information acquisition amplifies the initial response.

1.6 Ex-ante heterogeneity

This section considers an extension with ex-ante heterogeneity among agents that delivers a unique equilibrium and strengthens the amplification result. Agents are ex-ante identical in the baseline model and therefore make the same information acquisition choice in equilibrium, generating multiple equilibria. By contrast, this section explores the consequences of ex-ante heterogeneity, for example in the cost of becoming informed that reflects differences in the skill of generating and processing information. While either different or more general forms of ex-ante heterogeneity can be considered - and the result generalized accordingly, the following simple specification suffices to illustrate the effect of ex-ante heterogeneity on the information acquisition choice and on the amplification effect.

Let the individual cost of information acquisition c_i now be drawn independently from a uniform distribution over a unit interval, which is also the distribution of costs in the population by a law of large numbers:

$$c_i \sim U[0,1] \tag{1.31}$$

As before, an agent optimally acquires information if and only if the benefit weakly exceeds the cost:

$$n_i^* = I \Leftrightarrow c_i \leq D(n^*) \tag{1.32}$$

Optimal information acquisition is characterized by a threshold strategy. Let $\bar{c}$ be the cut-

[20]For instance, it is sometimes argued that a negative correlation exists between long-term German borrowing costs and adverse news about peripheral euro zone countries.

off value below which an agent acquires information. Then, the equilibrium proportion of informed agents is equal to the cut-off, given the distributional assumption ($n^*(\bar{c}) = \bar{c}$), which can be generalized to an increasing function in the cut-off level for other continuous distributions. Since the individual information cost equals the benefit from becoming informed for the marginal agent, the equilibrium cut-off value $\bar{c}$ is the fixed point of $D(n^*(\bar{c}))$.

Uniqueness requires that there is only one fixed point of $D(\cdot)$. Recall that $0 < D < 1$ (lemma 3) and $\frac{\partial D}{\partial n^*} \geq 0$, with strict inequality if and only if parameter constraint A holds (lemma 2). Appendix 1.8.2 shows that a strengthened lower bound on the precision of the uninformed agent suffices for $\frac{\partial D}{\partial n^*} < 1$, which completes the proof of the uniqueness of the cut-off level $\bar{c}$. Proposition 4 summarizes the new result on equilibrium uniqueness.

Proposition 4. ***(Ex-ante heterogeneity and uniqueness)*** *Let $\mu \neq \hat{\mu}(r)$ (parameter constraint A), $\gamma_U > \underline{\gamma}'_U \equiv \frac{\alpha^2}{2\pi+1-2\sqrt{2\pi}} > \underline{\gamma}_U$, and the cost of information acquisition be dispersed: $c_i \sim U[0,1]$. Then, there exists a unique equilibrium given by the unique cut-off value $\bar{c}$ implicitly defined by $\bar{c} = D(n^*(\bar{c}))$. Agents become informed at the information acquisition stage if and only if their individual information cost is no larger than the cut-off cost level:*

$$n_i^* = I \Leftrightarrow c_i \leq \bar{c} = n^* \tag{1.33}$$

Agents attack the status quo at the coordination stage if and only if their private signal falls short of the threshold $\overline{x}_U(\bar{c})$ if uninformed or $\overline{x}_I(\bar{c})$ if informed. The status quo is abandoned if and only if the fundamental falls short of the threshold $\overline{\theta}(\bar{c})$.

Ex-ante heterogeneity strengthens the amplification effect. In contrast to the baseline specification, the increase in the equilibrium proportion of informed agents is now strict whenever the benefit from becoming informed increases. To illustrate this, consider again the case of a strong common prior. Adverse news about the fundamental increase the benefit from becoming informed ($\frac{\partial D}{\partial \mu}$ <0), which *strictly* increases the equilibrium cut off value $\bar{c}$ and thus the equilibrium proportion of informed agents ($\frac{\partial n^*}{\partial \mu} < 0$). Since informed agents attack more often than uninformed agents if the common prior is strong, the probability of a financial crisis now increases *strictly.*

1.7 Conclusion

How does the acquisition of private information affect the probability of a bank run or sovereign debt crisis? Here I study a global coordination game of regime change in which a crisis occurs if a sufficient number of creditors withdraw from a bank or do not roll over debt to a sovereign. Creditors receive noisy private information about the solvency of the bank or sovereign and choose ex-ante whether to improve the quality of their information at a cost.

My main contribution is to demonstrate that endogenous information acquisition after adverse news increases the probability of a financial crisis. Adverse news about the solvency of the bank or sovereign increases the strategic uncertainty among creditors, making the behaviour of other creditors harder to predict. Since information supports coordination, an individual creditor wishes to acquire information to align his action with that of other creditors. If the initial fundamentals are high, informed creditors are more likely than uninformed creditors to withdraw funds from a bank or to not roll over debt to a sovereign. Thus, the acquisition of private information amplifies the probability of a bank run or a sovereign debt crisis.

My second contribution is the demonstration of a novel contagion effect based on endogenous information acquisition. Suppose that there is public information spillover, such as news about a crises elsewhere, which captures events like the failure of another bank (Lehman) or sovereign debt restructuring (Greece). Crises elsewhere induce creditors to become informed about the solvency of their bank or sovereign. Since the amplification effect established above increases the probability of a financial crisis in the initially unaffected region, the acquisition of private information is a powerful contagion mechanism for bank runs and sovereign debt crises.

1.8 Appendix

1.8.1 Responsiveness of aggregate threshold (proof of lemma 1)

This proof is in three steps. First, differentiating the equilibrium threshold fundamental with respect to the proportion of informed agents yields:

$$\frac{\partial \overline{\theta}}{\partial n^*} = \frac{\Phi\left(\frac{\alpha}{\sqrt{\gamma_I}}[\overline{\theta}-\mu] - \sqrt{1+\frac{\alpha}{\gamma_I}}\Phi^{-1}(r)\right) - \Phi\left(\frac{\alpha}{\sqrt{\gamma_U}}[\overline{\theta}-\mu] - \sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(r)\right)}{1-A_\theta(\overline{\theta})} \tag{1.34}$$

where the aggregate attack size A is equal to the right-hand side of equation (1.10). Second, note that $0 < A_\theta(\overline{\theta}) < 1$, where the second inequality follows from the sufficient condition for uniqueness at the coordination stage ($\gamma_U > \underline{\gamma}_U$). Third, parameter constraint A ensures that the numerator of the partial derivative is non-zero, which can be seen by contradiction. Suppose that the numerator is zero. Rewriting yields $\overline{\theta} = \hat{\theta} \equiv \mu + \frac{\kappa}{\alpha}\Phi^{-1}(r)$. Inserting this in the defining equation of $\overline{\theta}$, equation (1.10), yields $\mu = \hat{\mu}(r)$, as was to be shown. In summary, the overall partial derivative $\frac{\partial \overline{\theta}}{\partial n^*}$ is non-zero if parameter constraint A holds and the signal of the uninformed agent is sufficiently precise.

Using the same argument, the numerator of the partial derivative is positive if and only if the fundamental is strong relative to the strategic effect: $\mu > \hat{\mu}(r)$. Likewise, the numerator is negative if fundamentals are weak: $\mu < \hat{\mu}(r)$.

1.8.2 Strategic complementarity in becoming informed (proof of lemma 2)

I show that $\frac{\partial D}{\partial n^*} \geq 0$, establishing strategic complementarity in information acquisition. Using the Leibniz rule, the change in the benefit from becoming informed as more agents are informed is:

$$\frac{\partial D}{\partial n^*} = \left[1 - A_\theta(\overline{\theta})\right] g(\overline{\theta}) \left(\frac{\partial \overline{\theta}}{\partial n^*}\right)^2 \geq 0 \tag{1.35}$$

The sign arises since the probability distribution function g is always positive and $A_\theta(\overline{\theta}) \in (0,1)$ as implied by the sufficient condition for uniqueness at the coordination stage ($\gamma_U > \underline{\gamma}_U$). Furthermore, the inequality is strict if parameter constraint A holds, which ensures the responsiveness of the aggregate threshold with respect to the equilibrium proportion of informed agents (lemma 1).

A change in the proportion of informed agents has two effects on the benefit from becoming informed. First, there is a direct effect via a change in the threshold fundamental ($\frac{\partial D}{\partial \overline{\theta}} \frac{\partial \overline{\theta}}{\partial n^*}$). Second, there is an indirect effect via a change in the attacking thresholds ($\frac{\partial D}{\partial \overline{x}_z} \frac{\partial \overline{x}_z}{\partial \overline{\theta}} \frac{\partial \overline{\theta}}{\partial n^*}$). It can be shown that the indirect effect is always zero:

$$\frac{\partial D}{\partial \overline{x}_z} = r \int_{\overline{\theta}}^{\infty} f^z(\overline{x}_z) g(\theta) d\theta - (1-r) \int_{-\infty}^{\overline{\theta}} f^z(\overline{x}_z) g(\theta) d\theta = 0 \tag{1.36}$$

This is an envelope theorem result. The threshold $\overline{x}_z$ is chosen such that the marginal cost of attacking the status quo when it is maintained balances with the marginal benefit from attacking when the status quo is abandoned.

Using lemma 1, the change in the benefit from becoming informed as more agents are informed can be rewritten as:

$$\frac{\partial D}{\partial n^*} = \frac{g(\overline{\theta}) \Gamma(\overline{\theta})^2}{1 - A_\theta(\overline{\theta})} \tag{1.37}$$

where $g \leq \frac{1}{\sqrt{2\pi}}$ and $\Gamma < 1$ as shown in appendix 1.8.3. Therefore, $\frac{\partial D}{\partial n^*} < 1$ is ensured by $1 - A_\theta(\overline{\theta}) > \frac{1}{\sqrt{2\pi}}$ for all $\bar{c} \in [0,1]$. Since $\gamma_I > \gamma_U$, this inequality is hardest to satisfy for $n^* = 0$. Rewriting yields that $\gamma_U > \underline{\gamma}'_U \equiv \left(\frac{\alpha}{\sqrt{2\pi}-1}\right)^2$. In summary, $\gamma_U > \underline{\gamma}'_U$ suffices for $\frac{\partial D}{\partial n^*} < 1$.

1.8.3 Bounds on expected utility difference

I start by showing that $D < 1$. Note that Γ is the difference of two cumulative distribution functions such that $\Gamma \leq 1$ and $-\Gamma \leq 1$. Thus, $D \leq r \int_{\overline{\theta}}^{\infty} dG(\theta) - (1-r) \int_{-\infty}^{\overline{\theta}} dG(\theta) = r - G(\overline{\theta}) < r < 1$. Taken together, $D < 1$ without imposing parameter constraints.

It remains to be shown $D(n^*) > 0$ for all n^*, which is accomplished indirectly. First, note that $D \to 0$ as $\gamma_I \to \gamma_U$. Intuitively, there is no benefit from being informed if an informed agent receives a signal that is as precise as an uninformed agent's signal. Second, note that $D > 0$ as $\gamma_I \to \infty$. An informed agent no longer makes an error in his decision to attack, such that the benefit of becoming informed equals the strictly positive foregone expected utility of an informed agent due to his type I and type II errors. Third, the benefit from becoming informed strictly increases in the precision of the informed agent. Intuitively, being informed is more valuable when the improvement in the private signal is larger. Taking these three points together, it must be the case $0 < D(n^*) < 1$ for any n^* and any $\gamma_I \in (\gamma_U, \infty)$.

1.8.4 Effect of fundamental on benefit from becoming informed

A change in the fundamental μ affects the aggregate threshold $\overline{\theta}$, the attacking thresholds $\overline{x}_z$, and the distribution of fundamentals $g(\theta)$. As shown in Appendix 1.8.2, the effect on the attacking thresholds is zero by an envelope theorem argument. Using the Leibniz rule, the partial effect via the aggregate threshold is: $\frac{\partial D}{\partial \overline{\theta}} \frac{\partial \overline{\theta}}{\partial \mu} = -g(\overline{\theta}) A_\theta(\overline{\theta}) \frac{\partial \overline{\theta}}{\partial n^*}$. The partial effect via a change in the distribution is obtained by noting that $\frac{\partial g(\theta)}{\partial \mu} = g(\theta)\alpha(\theta - \mu)$. Partial integration of $\Gamma(\theta)\frac{\partial g(\theta)}{\partial \mu}$ and applying the envelope theorem once more yields: $\frac{\partial D}{\partial g(\theta)} \frac{\partial g(\theta)}{\partial \mu} = -g(\overline{\theta})[1 - A_\theta(\overline{\theta})] \frac{\partial \overline{\theta}}{\partial n^*}$. Therefore, the total effect on the benefit from becoming informed is:

$$\frac{\partial D}{\partial \mu} = -g(\overline{\theta}) \frac{\partial \overline{\theta}}{\partial n^*} \tag{1.38}$$

which is non-zero if parameter constraint A holds.

Chapter 2

Bank runs, liquidity, and macro-prudential regulation

This chapter examines the role of liquidity in an economy with many banks subject to runs and systemic liquidation costs. First, the presence of liquidity drives a wedge between the amount of withdrawals and liquidation. This restores multiple equilibria even when a global game refinement is used. Second, systemic liquidation costs imply that one bank's liquidity holding reduces the liquidation costs of other banks. The positive implication is the partial substitutability of private liquidity holdings as banks free-ride on other banks' liquidity. The normative implication is that banks hold insufficient liquidity relative to a constrained planner, interpreted as a macro-prudential authority that internalizes the system-wide effects of liquidity. Comparative statics analyses with respect to the expected investment return and the liquidation cost are performed.

JEL Classifications: G01, G11, G21, G28, G33

Keywords: Bank runs, liquidity, macro-prudential regulation, multiple equilibria

I wish to thank Fabio Castiglionesi, Co-Pierre Georg, Stefan Gissler, Todd Keister, and seminar participants at the London School of Economics and Universitat Pompeu Fabra for useful comments and discussion. I am also grateful to Benjamin Nelson for co-authored work that explores other aspects of macro-prudential regulation.

2.1 Introduction

A crucial concept in economics and finance is liquidity. An asset is liquid if it can be converted into cash quickly and at a low cost.[1] Holding enough liquid assets is important for financial intermediaries, as they may face sudden withdrawals from their investors. This applies to both the classical case of a run on a commercial bank by retail investors (Diamond and Dybvig (1983)) and modern-day runs, such as institutional investors who withdraw from money market mutual funds (e.g. Wermers (WP)). Liquidity also plays a major role in recent proposals for the regulation of financial intermediaries.[2]

This paper examines the role of liquidity in an economy with many banks subject to runs and systemic liquidation costs. Banks invest in a long-term project and hold some liquidity to prepare for early withdrawals from investors. The risky project has a higher expected return but is costly to liquidate before maturity. Liquidity drives a wedge between the amount of withdrawals from investors and the amount of liquidation, thereby trading off the opportunity cost of the higher expected investment return with the benefit from reducing costly liquidation. The profitability of a bank's project depends on aggregate economic conditions such as business cycle movements.[3] Investors have the option to withdraw before the maturity of the investment project. They receive noisy private news about the aggregate economic condition before deciding whether to withdraw. A bad economic condition results in a large number of investors with bad signals and therefore many withdrawals. This leads to run on a bank that has insufficient liquidity to serve all withdrawing investors and has to liquidate the project at a cost. Systemic liquidation costs, whereby one bank's liquidation cost increases in the other bank's liquidation volume, are also explored to analyze the system-wide dimension of liquidity.

The first contribution is to show that the presence of liquidity restores multiple

[1] An example is US sovereign debt, whose market is characterized by a large trading volume. Therefore, selling a given quantity can be realized fast and with a low price impact.

[2] This is in contrast to the financial regulation during the last three decades that focused mainly on capital. The current proposals include the introduction of rules governing the composition of banks' balance sheets envisaged under the Basel Committee on Banking Supervision's proposed Liquidity Coverage Ratio (LCR) or Net Stable Funding Ratio (BCBS (2010b,c)). Both regulatory tools seek to impose limits on the degree of liquidity mismatch on a bank's balance sheet by, for example, imposing a lower bound on banks' liquidity ratios.

[3] Business cycle movements affect the default rates of the loan portfolio of banks. Alternatively, there is a shock to an asset class in which banks are invested, such as asset-backed securities.

equilibria – even if a global game refinement is used (Proposition 5). When banks hold some liquidity to prepare for withdrawals from investors, there exists an equilibrium without liquidation if the economic condition is good. Some investors receive bad news about the economic condition, infer that their bank's profitability is low, and withdraw. But since the true state of the economy is good, this is a small number of investors and the available liquidity suffices to serve them. Most investors receive good news and do not withdraw, as such costly liquidation is avoided. Likewise, an equilibrium with liquidation exists if the economic condition is weak. Then many investors receive bad news and withdraw. As a result, liquidity is exhausted and costly liquidation occurs. In sum, I show that there exists an interim range of the economic condition that supports both equilibria.

The equilibrium with liquidation corresponds to the unique Bayesian equilibrium in other global game models of bank runs, such as Goldstein and Pauzner (2005) and Morris and Shin (2000). Why is the other equilibrium without liquidation absent in these papers? As investors receive noisy news, some investors will always receive a bad signal and withdraw – even if the economic condition is good. Without liquidity, there is always positive liquidation to serve withdrawing investors, ruling out the possibility of an equilibrium without liquidation. In fact, I show formally that the equilibrium without liquidation vanishes as the level of liquidity vanishes (Corollary 1). Therefore, liquidity is crucial for re-establishing multiple equilibria in bank run coordination games. The multiplicity result does not rely on endogenous information acquisition, which has also been shown to break uniqueness (e.g. Hellwig and Veldkamp (2009), Ahnert (WP)).

The second contribution is to demonstrate a role for a macro-prudential regulation of liquidity. To this end, I compare the privately optimal and socially constrained efficient levels of liquidity.[4] When liquidation costs are systemic, insufficient liquidity at one bank means more liquidation for a given amount of withdrawals and therefore a higher liquidation cost for other banks in the system. The positive implication is the partial substitutability of private liquidity holdings as banks free-ride on other banks' liquidity (Proposition 6). The normative implication is that the private banking system holds insufficient liquidity relative to a constrained planner (Proposition 7). As a planner

[4]In order to analyze the effects of ex-ante liquidity holdings, I need to select an equilibrium for economic conditions that support both equilibria. To focus on the macro-prudential implications of liquidity, I select the equilibrium with liquidation.

internalizes the system-wide effects of liquidity, this planner is naturally interpreted as a macro-prudential authority.

Proposition 8 summarizes comparative static results that illustrate the intuition of private and social liquidity choices. The level of liquidity held by a bank trades off the marginal cost in terms of foregone expected investment return with the marginal benefits in terms of avoiding costly liquidation, which reduces coordination failure among investors. A higher expected investment return (better economic conditions on average) increases the opportunity cost of liquidity and therefore reduces a bank's optimal liquidity level. By contrast, a higher liquidation cost increases the marginal benefit from avoiding liquidation and therefore increases a bank's optimal liquidity level. By extension, the comparative statics for the constrained efficient liquidity choice yield the same signs as the constrained planner faces the same trade-off, just with a higher marginal benefit from liquidity.

Systemic liquidation costs, which generate a positive externality from liquidity and are at the core of my normative result, are micro-founded by a body of literature. Limited participation in asset markets can lead to cash-in-the-market pricing and therefore under-pricing of assets (Allen and Gale (1994)). Similarly, liquidation values are depressed after an industry-specific shock since distress sales take place to unlevered industry outsiders who value industry-specific assets less (Shleifer and Vishny (1992)). Finally, financial arbitrageurs cannot pick up assets in fire sales since they are constrained by losses and outflows themselves (Gromb and Vayanos (2002)).

The paper closest in terms of methodology is Morris and Shin (2000), who build on the seminal work of Carlsson and van Damme (1993), using global games techniques to analyze a withdrawal game in the spirit of Diamond and Dybvig (1983).[5] The Bayesian equilibrium of Morris and Shin (2000), which features runs on illiquid but solvent banks, is unique. By contrast, I show in my first contribution how the presence of liquidity breaks equilibrium uniqueness by allowing for another equilibrium without runs. Furthermore, I extend the analysis to multiple banks to explore the effect of systemic liquidation costs

[5] Multiple equilibria in Diamond and Dybvig (1983) occur because of the self-fulfilling beliefs. If an investors fears withdrawals by other investors, then this will imply costly liquidation of the bank's assets that reduces a non-withdrawing investor's payoff. Therefore, each investor finds it optimal to withdraw, constituting a bank-run equilibrium. Likewise for the no-run equilibrium.

on ex-ante incentives to hold liquidity.[6] Vives (WP) and Morris and Shin (WP) also analyze investor withdrawal games and the effect of liquidity. However, they abstract from conditions that can induce the no-liquidation equilibrium and are not concerned with the ex-ante portfolio choice.

Allowing for multiple banks, my second contribution is to examine the consequences of systemic liquidation costs for ex-ante liquidity choices, both privately and socially. Other consequences of systemic liquidation costs have already been analyzed. Wagner (2011) studies the diversification-diversity trade-off in the ex-ante portfolio choice. Since joint liquidation is costly ex-post, investors have an incentive to hold diverse portfolios.[7] In contrast, I examine the consequences for ex-ante liquidity holdings in the presence of systemic liquidation costs and analyze the consequences for financial intermediaries that may be subject to runs. Uhlig (2010) analyzes endogenous liquidation costs in a model with outside investors and a two-tiered banking sector. The arising system-wide externality generates strategic complementarities in the depositors' withdrawal decisions also present in my model. His focus is on a positive analysis of the previous financial crisis and discusses some ex-post policy interventions. By contrast, my focus is on optimal (liquidity) regulation from an ex-ante perspective. Studying ex-ante policy has the advantage of precluding the issue of moral hazard arising from an ex-post policy intervention, a theme also stressed by Farhi and Tirole (2012).

The literature on macro-prudential regulation is growing fast. Korinek (2011) analyzes risk-taking in an economy in which systemic externalities take the form of pecuniary fire sales and provides a micro-founded rationale for macro-prudential policy, such as a Pigouvian tax on risk-taking or capital requirements. Korinek (WP) contrasts ex-ante macro-prudential regulation with ex-post policy interventions. In line with the present paper, Farhi and Tirole (2012) highlight the importance of a macro-prudential approach to contain a crisis ex-ante.

[6] Goldstein and Pauzner (2005) also use global games techniques to generate a unique equilibrium in a setup closer to the original model of Diamond and Dybvig (1983), for example preserving the sequential service constraint. The same comments apply.

[7] Wagner (2009) also stresses the role of endogenous liquidation costs, showing that they give rise to cross-bank externalities. The implications for optimal bank portfolios are ambiguous, however, as banks may be 'too correlated' (as in the standard case) or 'too diversified' under laissez faire, implying that regulatory treatment should be heterogeneous.

2.2 The Model

The economy extends over three dates labelled as initial ($t = 0$), interim ($t = 1$), and final ($t = 2$), and it is inhabited by a continuum of investors and two banks $n \in \{A, B\}$. The notion of financial intermediation provided by banks is not limited to the traditional case of retail investors and commercial banks but incorporates, for instance, money market mutual funds and investment banks.[8] There is a single physical good used for consumption and investment.

Investors There is a unit mass of initially identical investors $i \in [0, 1]$ with idiosyncratic uncertainty about their consumption needs (Diamond and Dybvig (1983)). All investors are uncertain at the initial date and privately learn their consumption preference $\theta_i \in \{0, 1\}$ at the interim date. Each investor is either early ($\theta_i = 1$) and wishes to consume at the interim date or late ($\theta_i = 0$) and wishes to consume at the final date. Investors can store between the interim and the final date. The ex-ante probability of being an early investor is $\lambda \equiv \Pr\{\theta_i = 1\} \in (0, 1)$, which is identical across investors and also the share of early investors by the law of large numbers. A investor's utility function is:

$$U_i(c_1, c_2) = \theta_i c_1 + (1 - \theta_i)c_2 \tag{2.1}$$

where c_t is consumption at date t, and θ_i represents an idiosyncratic liquidity shock. Investors are endowed with two units at the initial date and randomly deposit at either bank; as such each bank receives one unit of deposits.

Investment opportunities Two investment opportunities in the form of constant-return-to-scale technologies are available at the initial date (Table 2.1). First, storage is universally available and yields a unit safe return. Since it matures after one period, storage is referred to as *liquidity*. Second, an investment project, such as lending to a productive sector, is available to banks.[9] A project matures at the final date and

[8] Also, the investors and banks of this model can be interpreted as local and global banks in the spirit of Uhlig (2010). Then a prematurely withdrawing investor represents a run of a local bank on a global bank, an arguably reasonable feature of the recent financial crises.

[9] The motive for the existence of banks is different from that in Diamond and Dybvig (1983). While these authors demonstrate a role for a bank as provider of liquidity insurance for risk-averse investors, banking in this model arises from a bank's enhanced access to investment projects because of an advantage

Asset	$t = 0$	$t = 1$	$t = 2$
Storage $(0 \to 1)$	-1	1	
Project $(0 \to 2)$	-1	l_n	$(1 - l_n)(r - \chi(l_n, l_{-n}))$

Table 2.1: Investment technologies

yields a stochastic return with mean $\bar{r} > 1$. Premature liquidation of the project at the interim date is costly. Similar to Morris and Shin (2000), liquidation of an amount $l_n \in [0, 1]$ by bank n at par reduces the final-date return by $\chi(l_n, l_{-n})$, where $\chi(\cdot) \geq 0$ is the cost function of premature liquidation. The reduction to the final-date payoff is implied by a lender-of-last-resort policy, for instance. Liquidation costs are modelled to be proportional to the total amount of liquidation:

$$\chi(l_n, l_{-n}) \equiv \chi[l_n + d\, l_{-n}] \tag{2.2}$$

where $\chi \in (0, 1)$ measures the cost of liquidation, and $d \in \{0, 1\}$ is a dummy that is one when systemic liquidation costs are present. To avoid strict dominance of the project, $\bar{r} < 1 + 2\chi$ is assumed throughout.

Bank-specific liquidation costs are the source of strategic complementarity between late investors of a given bank. According to Ahnert and Nelson (WP), individual liquidation costs are discussed in James (1991) and Mullins and Pyle (1994). These costs comprise direct liquidation expenses and a reduction in the 'going concern' value of bank assets under distress. The empirical literature typically finds these liquidation costs to be large: of the order of 30% of bank assets on average.[10]

Systemic liquidation costs or fire-sales, if present, are the source of strategic complementarity between late investors across banks and create an externality in a bank's liquidity choice. Liquidation costs are systemic if there is limited liquidity in the market (Allen and Gale (1994)), a fire sale to industry outsiders (Shleifer and Vishny (1992)),

in monitoring, for example.

[10]Mullins and Pyle (1994) and Brown and Epstein (1992) present estimates of direct liquidation expenses of around 10%, varying between 17% for assets relating to owned real estate to 0% for liquid securities for assets in receivership at the FDIC. Adding to direct expenses losses associated with forced liquidation, James (1991) gives an average total cost of 30% of a failed bank's assets. Similar orders of magnitude are reported in Bennett and Unal (WP), whose sample runs for much longer, covering 1986-2007.

or financial constraints of arbitrageurs (Gromb and Vayanos (2002)). Since liquidation depresses not only a given bank's liquidation value, but also another bank's liquidation value, there is a negative externality from liquidation.

Information structure The investment return $\boldsymbol{r}$, a measure of the economic condition such as a key macroeconomic variable, is distributed normally with precision $\alpha \in (0, \infty)$:

$$\boldsymbol{r} \sim \mathcal{N}\left(\bar{r}, \frac{1}{\alpha}\right)$$

The investment return is realised at the interim date but not publicly observed. However, each investor receives a private signal x_i about the return:

$$\begin{aligned} \boldsymbol{x}_i &\equiv \boldsymbol{r} + \boldsymbol{\epsilon}_i \\ \boldsymbol{\epsilon}_i &\sim \mathcal{N}\left(0, \frac{1}{\gamma}\right) \end{aligned}$$

where the idiosyncratic noise $\boldsymbol{\epsilon}_i$ has zero mean, precision $\gamma \in (0, \infty)$, and is independently and identically distributed across investors and independent of the investment return. All distributions are common knowledge.

Banks At the initial date banks simultaneously choose their liquidity holdings (y_A, y_B) and invest the remainder in the project. The liquidity choice is publicly observed. Investors that withdraw at the interim date receive unity, while investors that wait for the final date receive a pro-rata payment of a bank's assets, which includes the proceeds from investment in the project.[11] Abstracting from a misalignment of incentives between the bank manager and investors, a bank's objective is to maximize the expected utility of investors.[12]

[11] See, for example, Dasgupta (2004), Goldstein (2005), and Shapiro and Skeie (WP) for a similar assumption on the interim-date withdrawal payment. Therefore, banks are viable at the interim date as the promised payment does not exceed the liquidation value. The focus of the present paper is on the effect of liquidity on equilibrium multiplicity as well as the consequences of a fire-sale externality from one bank's liquidation decision on the liquidity choice of banks. My main results hold for alternative assumptions about the interim payment.

[12] This objective arises as an equilibrium outcome in a generalized model with competition for deposits. The competition between banks for investors implies that banks offers the best possible liquidity holding to investors. In a related paper, Gale (2010) shows that a bank's optimal behaviour under free entry and subject to the investors' participation constraint can be expressed as the solution to a contracting problem in which the welfare of investors is maximised subject to the zero-profit constraint of the bank. If any given bank were not to choose this investment plan, it would fail to attract any deposits. Given the

Banks serve any withdrawals by using liquidity first. Let $w_n \in [0, 1-\lambda]$ denote the amount of withdrawals by late investors of bank n. If withdrawals from late investors are sufficiently high ($w_n > y_n - \lambda$), the bank partially liquidates its investment project, where the liquidation amount is given by $l_n \equiv \max\{0, w_n + \lambda - y_n\} \in [0, 1 - y_n]$. The liquidation amount decreases in a bank's liquidity holding y_n and increases in the amount of withdrawals by late investors w_n.

To prevent costly liquidation, the bank may hold excess liquidity ($y_n > \lambda$) – more liquidity than required to serve withdrawals from early investors. Holding excess liquidity drives a wedge between the proportion of prematurely withdrawing late investors w_n and the liquidation volume of the investment project l_n. As it is never optimal to face certain liquidation, the lower bound of a bank's liquidity level is the share of early investors ($y_n \geq \lambda$).

Payoffs Early investors always withdraw at the interim date. Late investors that withdraw at the interim date receive the same payoff as early investors since the liquidity preference of investors is unobserved by banks. To shed more light on a late investor's payoff, consider the cases of no liquidation and positive liquidation in turn.

No liquidation ($w_n \leq y_n - \lambda$) If few late investors withdraw at the interim date, excess liquidity holdings $y_n - \lambda$ suffice to serve them. There is no liquidation ($l_n = 0$), and some excess liquidity is carried over to the final date. The payoff to a late investor at the final date is:

$$c_{2n} = \frac{[y_n - \lambda - w_n] + (1 - y_n)r}{1 - \lambda - w_n} \tag{2.3}$$

where the asset payments available to investors at the final date consist of remaining liquidity ($y_n - \lambda - w_n$) and proceeds from investment in the project (numerator), all of which to be shared with the proportion of investors that wait for the final date (denominator). The realisation of the stochastic investment project r enters this expression, while the amount of liquidation by the other bank (l_{-n}) has no effect on the payoff of late investors of bank n in the absence of liquidation ($l_n = 0$).

---

alignment of interest between a bank and its investors as well as the bank's enhanced access to projects, all depositors deposit in full.

Positive amount of liquidation ($w_n > y_n - \lambda$) If many late investors withdraw at the interim date, the excess liquidity holding $y_n - \lambda$ is drawn down such that some amount of the project is liquidated ($l_n = w_n + \lambda - y_n$) to serve withdrawing investors. The payoff to a late investor at the final date is:

$$c_{2n} = \frac{(1 - y_n - l_n)[r - \chi(l_n, l_{-n})]}{1 - \lambda - w_n} = r - \chi(l_n, l_{-n}) \tag{2.4}$$

A fire-sale externality, which is a negative liquidation externality, is present if and only if there are systemic liquidation costs ($d = 1$) and the other bank liquidates a positive amount ($l_{-n} > 0$).

Remark 1. *Conditional on a positive liquidation ($l_n > 0$), there is a strategic complementarity between the withdrawal decisions of late investors of the same bank ($\frac{\partial c_{2n}}{\partial w_n} < 0$). If there is also positive liquidation by the other bank ($l_{-n} > 0$) and systemic liquidation costs are present ($d = 1$), there is also strategic complementarity between the withdrawal decisions of late investors across banks ($\frac{\partial c_{2n}}{\partial w_{-n}} < 0$).*

There are two dimensions to the strategic behaviour of a late investor. The first dimension is the strategic complementarity between the withdrawal decisions of late investors of a given bank. More withdrawals by other late investors have two effects. First, the bank draws down its excess liquidity and then liquidates a larger share of the project. This effect is detrimental to a late investor who keeps his funds for the final date. Second, there are fewer late investors to share the remaining resources with at the final date. This effect is beneficial for a late investor who keeps his funds. In the positive-liquidation case, the first effect unambiguously dominates and the incentives to withdraw increase in the proportion of withdrawing late investors ($\frac{\partial c_{2n}}{\partial w_n} = -\chi < 0$). By contrast, the incentives to withdraw decrease in the proportion of withdrawing late investors if no liquidation takes place and the project return is sufficiently high ($r \geq 1$).

The second dimension is a strategic complementarity between the withdrawal decisions of late investors across banks in the presence of systemic liquidation costs. The more late investors in the other bank $-n$ withdraw, the more of the investment project of bank $-n$ is liquidated, the lower the final-date payoff to investors at bank n due to the fire-sale externality. This increases the incentive for a late investor of bank n to withdraw as well, conditional on positive liquidation at bank n ($\frac{\partial c_{2n}}{\partial w_{-n}} = -\chi < 0$).

Timeline The following timeline summarizes the model:

Initial date $t = 0$

- Investors receive their endowment and deposit at banks.
- Each bank holds liquidity y_n and invests the remainder $1 - y_n$.

Interim date $t = 1$

- Each investor privately observes his consumption preferences θ_i (early or late).
- Each investor receives a private signal x_i about the investment return and updates his forecast about the return and the proportion of withdrawing investors.
- Investors may withdraw, and the mass of late investors that withdraw is w_n.
- Banks serve withdrawals using liquidity first. If necessary, a bank (partially) liquidates the investment project (liquidation amount l_n).
- Early investors consume and withdrawing late investors store their withdrawals.

Final date $t = 2$

- The investment project matures.
- Remaining late investors receive an equal share of the investment proceeds.
- Late investors consume.

2.3 Equilibrium

There are two stages: a perfect-information portfolio choice stage between banks at the initial date and an imperfect-information withdrawal stage between investors at the interim date. As the portfolio choices of banks are observed by investors, the equilibrium is best characterised by working backwards, starting with the equilibrium in the withdrawal subgame.

An investor's *strategy* is a plan of action for each private signal x_i. A profile of strategies is a Bayesian Nash equilibrium in the subgame at the interim date if the actions described by investor i's strategy maximize his expected utility conditional on x_i, taking as given the strategies followed by all other investors. *Threshold strategies* are considered by which a late investor withdraws if and only if his private signal falls short of a bank-specific threshold (to be determined): $x_i < x_n^*$. These thresholds depend on the liquidity choices of banks at the initial date: $x_n^* = x_n^*(y_n, y_{-n})$.

Next, consider the game between banks who choose a liquidity level y_n. Each bank takes the effect of its liquidity choice on the bank-specific withdrawal threshold $x_n^*(y_n, y_{-n})$ into account. I will determine the Nash equilibrium (y_A, y_B) in the game between banks at the interim date, where each y_n maximizes the bank's objective function subject to the effect on the withdrawal threshold, taking as given the level of liquidity held by the other bank.

Each investor uses his private information x_i to update his forecasts about the investment return and the proportion of withdrawing late investors at either bank. The posterior distributions are derived in Appendix 2.7.1. Let $\boldsymbol{R_i} \equiv r|x_i$ denote the posterior distribution of the investment return as formed by an investor who receives the private signal x_i. The posterior mean R_i is equivalent to the signal x_i because there is a bijective mapping between them. The equilibrium posterior mean R_n^*, which is computationally more convenient than the equilibrium signal x_n^*, is used to describe the equilibrium conditions.[13] Likewise, $\boldsymbol{W_{i,n}^n} \equiv w_n|x_i$ and $\boldsymbol{W_{i,-n}^n} \equiv w_{-n}|x_i$ denote the posterior distributions of the proportions of withdrawing late investors at the investor's bank and the other bank, respectively. Similarly, the expected amount of liquidation by bank n is given by $\boldsymbol{L}_{i,n}^n$.

Consider the equilibrium withdrawal behaviour of investors at the interim date. Early investors always withdraw, while late investors may withdraw. The bank-specific threshold R_n^* is defined as the mean of the posterior return that makes a late investor

[13] Note that both converge as the private noise vanishes ($x_n^* - R_n^* \to 0$ as $\gamma \to \infty$).

indifferent between withdrawing and not withdrawing his funds:

$$1 = c_{2n}(R_n^*, R_{-n}^*) \tag{2.5}$$

where the left-hand side is the payoff from withdrawing and the right-hand side is the expected payoff from not withdrawing conditional on the threshold signal x_n^*. Equation (2.5) implicitly defines the best response function $R_n^*(R_{-n}^*)$, where investors take the other bank's threshold R_{-n}^* as given. The withdrawal threshold of investors of one bank depends on the threshold of investors in the other bank in case of positive liquidation ($l_n > 0$) and systemic liquidation costs ($d = 1$).

The subsequent subsections construct a complete description of equilibrium in the subgame by analysing the role of liquidity for equilibrium multiplicity and the effect of systemic liquidation costs. In line with the global games literature (e.g. Morris and Shin (2003)), I shall assume vanishing private noise ($\gamma \to \infty$) throughout.

2.3.1 No systemic liquidation costs

First consider the case without systemic liquidation costs ($d = 0$).

No expected liquidation

Suppose the marginal investor expects no liquidation in equilibrium ($L_n(x_n^*) = 0$). Then, the indifference condition yields $(1 - y_n)(R_n^* - 1)$. If there is no intermediation ($y_n = 1$ or "narrow banking"), then the withdrawal decision is irrelevant since the bank's assets are always unity and the investor receives unity irrespective of his withdrawal decision. If there is intermediation ($y_n < 1$), the withdrawal threshold is unity ($R_n^* = 1$). To be consistent with zero liquidation *as expected by the marginal investor*, the bank's liquidity level must be sufficiently high. The marginal investor expects half of the late investors of mass $1 - \lambda$ to withdraw as private noise vanishes ($W_n^n \to \frac{1-\lambda}{2}$). Therefore, liquidity must be abundant to serve withdrawals from early and late investors:

$$y_n \geq \lambda + \frac{1-\lambda}{2} \tag{2.6}$$

Zero *actual* liquidation arises if and only if the number of withdrawals does not exceed excess liquidity ($w_n \leq y_n - \lambda$), requiring a sufficiently high realisation of the investment return ($r \geq \tilde{r}_0$). As the distribution of signals conditional on the economy-wide return is $\mathcal{N}\left(r, \frac{1}{\gamma}\right)$, the *lower* bound on the realised investment return is:

$$\tilde{r}_0 \equiv 1 - \frac{\alpha}{\gamma}(\bar{r} - 1) - \frac{1}{\sqrt{\gamma}}\Phi^{-1}\left(\frac{y_n - \lambda}{1 - \lambda}\right) \rightarrow 1$$

where $\Phi^{-1}(\cdot)$ is the inverse of the cumulative probability function of the standard normal distribution and $\tilde{r}_0 < 1$. More liquidity ensures that more withdrawals are consistent with zero actual liquidation. Hence, the lower bound on the investment return decreases in the liquidity holding ($\frac{\partial \tilde{r}_0}{\partial y_n} < 0$). As private noise vanishes, the lower bound converges to unity ($\tilde{r}_0 \rightarrow 1$). Lemma 4 summarizes:

Lemma 4. *Consider the withdrawal subgame without systemic liquidation costs* ($d = 0$), *vanishing private noise* ($\gamma \rightarrow \infty$), *and abundant liquidity* $y_n \in [\frac{(1+\lambda)}{2}, 1]$. *Then any threshold equilibrium has the following features:*

- $L_n^* = 0$: *the marginal investor expects no liquidation;*
- $x_n^* \rightarrow 1$: *a late investor withdraws if and only if his signal falls short of unity;*
- $l_n^* = 0 \Leftrightarrow r \geq 1$: *no actual liquidation occurs if and only if the economic condition is sufficiently good.*

In sum, the level of liquidity determines whether the marginal investor expects positive liquidation to occur in equilibrium, while the realised economic condition determines whether liquidation actually occurs.

Positive expected liquidation

For the marginal investor to expect a positive amount of liquidation in equilibrium ($L_n(x_n^*) > 0$), the bank's liquidity level must be scarce:

$$y_n < \frac{1 + \lambda}{2} \tag{2.7}$$

As I show in section 2.3.3, liquidity is scarce in equilibrium if it has a high opportunity cost in terms of a high expected investment return ($\bar{r} \geq \bar{r}_L$). With positive expected liquidation the indifference condition of the marginal investor reduces to:

$$R_n^* = 1 + \chi[(1-\lambda)\Phi(\sqrt{\delta}[R_n^* - \bar{r}] + \lambda - y_n)]$$

where $\delta \equiv \frac{\alpha^2(\alpha+\gamma)}{\gamma(\alpha+2\gamma)}$ collects precision parameters. As in Morris and Shin (2000), uniqueness requires the slope of the left-hand side to exceed the slope of the right-hand side and vanishing private noise is sufficient for this requirement.[14] A closed-form expression for the threshold is obtained for vanishing private noise:

$$R_n^* \to 1 + \chi[\frac{1-\lambda}{2} + \lambda - y_n] \tag{2.8}$$

Coordination failure between investors induces the threshold to exceed unity ($R_n^* > 1$), the efficient liquidation level. If there is no cost of premature liquidation ($\chi \to 0$), then the strategic complementarity between investors of the same bank is absent and the efficient allocation obtains in the withdrawal game. For a positive liquidation cost ($\chi > 0$), however, there is coordination failure between investors that pushes the threshold above the efficient level. Fearing that other late investors withdraw prematurely and thereby cause costly liquidation, another late investor has an incentive to withdraw prematurely – even if the nominal investment return exceeds the payoff from withdrawing prematurely. Furthermore, the threshold is below the expected return of the project ($R_n^* < \bar{r}$) if the (individual) liquidation cost is sufficiently low relative to the expected return ($\bar{r} > 1 + \frac{\chi(1-\lambda)}{2}$).

Actual liquidation occurs in equilibrium if the realised investment return is sufficiently low ($r < \tilde{r}_1$). Finding the equilibrium proportion of withdrawals for a given investment return as in the previous case, the *upper* bound on the investment return is:

$$\tilde{r}_1 \equiv R_n^* - \frac{\alpha}{\gamma}(\bar{r} - R_n^*) - \frac{1}{\sqrt{\gamma}}\Phi^{-1}\left(\frac{y_n - \lambda}{1-\lambda}\right) \to R_n^*$$

[14]The requirements is $1 > D_n^* \equiv \chi(1-\lambda)\sqrt{\delta}\phi\left(\sqrt{\delta}[R_n^* - \bar{r}]\right) > 0$, where $\phi(\cdot)$ is the probability distribution function of the standard normal distribution. This condition is satisfied as the private noise vanishes since $\delta \to 0$.

Holding more liquidity has two effects. First, it allows to serve a larger proportion of withdrawing investors without liquidating the project. As liquidation is costly, this reduces the amount of coordination failure between late investors for a given investment return and thus the withdrawal threshold:

$$\frac{\partial R_n^*}{\partial y_n} = -\frac{\chi}{1 - D_n^*} < 0 \tag{2.9}$$

Second, more liquidity reduces the upper bound on the investment return for which the equilibrium with positive liquidation exists ($\frac{\partial \tilde{r}_1}{\partial y_n} < 0$). More liquidity implies more available resources for withdrawing investors and therefore requires a worse economic condition to sustain positive liquidation as supposed. Lemma 5 summarizes:

Lemma 5. *Consider the withdrawal subgame without systemic liquidation costs* ($d = 0$)*, vanishing private noise* ($\gamma \to \infty$)*, and scarce liquidity* $y_n \in \left(\lambda, \frac{(1+\lambda)}{2}\right)$*. Then any threshold equilibrium has the following features:*

- $L_n^* > 0$*: the marginal investor expects liquidation;*
- $x_n^* \to 1 + \chi[\frac{1+\lambda}{2} - y_n] > 1$ *: a late investor withdraws if and only if his signal falls short of this threshold;*
- $l_n^* > 0$ *if and only if* $r < x_n^*$*: actual liquidation occurs if and only if the economic condition is sufficiently bad.*

Taking the previous lemmas together, the overall threshold equilibrium in the withdrawal subgame, which depends on the realised economic condition r and the amount of liquidity y_n held by the bank, is described in proposition 5:

Proposition 5. *Consider the withdrawal subgame without systemic liquidation costs* ($d = 0$) *and vanishing private noise* ($\gamma \to \infty$).

- *If liquidity is abundant* ($y_n \in [\frac{(1+\lambda)}{2}, 1]$)*, then there exists a unique threshold equilibrium in the subgame. The marginal investor expects no liquidation to take place* ($L_n(x_n^*) = 0$) *and the implied withdrawal threshold is* $x_n^* = 1$*. No liquidation occurs if the economic condition is good, while some liquidation occurs if it is bad* ($l_n^* = 0 \Leftrightarrow r \geq 1$).

- *If liquidity is scarce* ($y_n \in \left[\lambda, \frac{(1+\lambda)}{2}\right)$), *however, then there exist multiple equilibria in the subgame. The marginal investor expects liquidation to take place* ($L_n(x_n^*) > 0$) *and the implied withdrawal threshold is* $x_n^* \to 1+\chi[\frac{1+\lambda}{2}-y_n] > 1$. *The no-liquidation equilibrium occurs for a good economic condition* ($r \geq 1$), *while the equilibrium with liquidation occurs for a bad economic condition* ($r < x_n^*$). *Therefore, multiple equilibria exist for a range of economic conditions* $[1, x_n^*]$.

- *The range of economic conditions that support multiple equilibria shrinks as the bank's liquidity increases* ($\frac{\partial x_n^*}{\partial y_n} < 0$).

How does the multiplicity result relate to bank run models that obtain a unique equilibrium with positive liquidation (Goldstein and Pauzner (2005), Morris and Shin (2000))? If there is no excess liquidity ($y_n \to \lambda$), as in these papers, the no-liquidation equilibrium disappears. In fact, the lower bound on the economic condition consistent with no liquidation becomes arbitrarily high ($\tilde{r}_0 \to \infty$) for bounded private noise ($\underline{\gamma} < \gamma < \infty$). Thus, any equilibrium features a positive amount of liquidation in these papers. By contrast, liquidity drives a wedge between the amount of withdrawals and the liquidation volume in the present paper. If liquidity is scarce and the economic condition good, this supports an equilibrium without liquidation apart from the usual equilibrium with liquidation.

Corollary 1. *If there is no liquidity to serve late investors* ($y_n \to \lambda$), *the equilibrium without liquidation vanishes* ($\tilde{r}_0 \to \infty$). *Therefore, models without liquidity and unique equilibria, such as Morris and Shin (2000), are a special case of my model with vanishing liquidity for late investors.*

Finally, consider the marginal benefits of liquidity on the threshold equilibrium in the withdrawal subgame. There is no marginal benefit of liquidity in the no-liquidation equilibrium since the lower bound of the economic condition is unaffected by liquidity. By contrast, the marginal benefit from liquidity is positive in the equilibrium with liquidation. On the one hand, liquidity reduces the range for which the equilibrium with liquidation exists (also for bounded private noise). On the other hand, more liquidity reduces the amount of withdrawals and therefore costly liquidation for a given investment return. The marginal cost of liquidity is the reduction in the payoff to a late investors conditional on no-liquidation.

2.3.2 Systemic liquidation costs

I now consider the case with systemic liquidation costs ($d = 1$). Suppose that the marginal investor expects liquidation ($L_n^n(x_n^*) > 0$) and liquidation by the other bank ($L_{-n}^n(x_n^*) > 0$). If the marginal investor expects no liquidation by the other bank, then systemic liquidation costs have no impact and the equilibrium threshold is determined as in the previous case. The indifference condition of the marginal investor becomes:

$$\begin{aligned} R_n^*(R_{-n}^*; y_n, y_{-n}) \quad &= \quad 1 + \chi \left[(1-\lambda)\Phi\left(\sqrt{\delta}[R_n^* - \bar{r}]\right) + \lambda - y_n \right] + \cdots \\ &\cdots + \chi \left[(1-\lambda)\Phi\left(\sqrt{\delta}(1+\frac{\gamma}{\delta})[R_{-n}^* - \bar{r}] - \frac{\gamma}{\delta}\sqrt{\delta}[R_n^* - \bar{r}]\right) \right] \end{aligned} \tag{2.10}$$

As the marginal investor takes the withdrawal threshold of investors in the other bank R_{-n}^* as given, equation 2.10 specifies a best-response function since there exists a unique solution R_n^* for any given R_{-n}^* as shown below.

Following Morris and Shin (2003) and Goldstein (2005), the uniqueness proof is in two steps. First, a unique solution R_n^* must be obtained for any R_{-n}^*, requiring that the slope of the left-hand side of the best response function exceeds the slope of the right-hand side. Second, there is a unique intersection of best response functions, requiring that the best response function is bounded and that its slope is strictly within zero and one. Since the cumulative distribution function lies within zero and one, these conditions are all satisfied if the private noise is sufficiently small, yielding a unique solution $R_A(y_A, y_B), R_B(y_B, y_A)$.[15]

Coordination failure again induces an inefficiently large withdrawal threshold ($R_n^* > 1$). Coordination failure now takes place not only between investors of a given bank, but also between investors of different banks. Fearing that late investors of another bank withdraw, thereby increasing the liquidation volume of the other bank and therefore the liquidation costs of a given bank, a late investor of the given bank has a higher incentive to withdraw at the interim date as well.

Furthermore, the threshold is below the expected return of the project ($R_n^* < \bar{r}$) if

[15] The slope of the right-hand side is now:

$$D_n \equiv \chi(1-\lambda)\sqrt{\delta}\{\phi\left(\sqrt{\delta}[R_n^* - \bar{r}]\right) - \frac{\gamma}{\alpha}\phi\left(\sqrt{\delta}(\frac{\alpha+\gamma}{\gamma})[R_{-n}^* - \bar{r}] - \frac{\gamma}{\alpha}\sqrt{\delta}[R_n^* - \bar{r}]\right)\}$$

the (total) liquidation cost is sufficiently low relative to the expected return ($\bar{r} > \bar{r}_L \equiv 1+\chi(1-\lambda)$). The behaviour of investors is consistent with a liquidation in equilibrium if the realised investment return is sufficiently bad ($r < \tilde{r}_1$). As the private noise vanishes, the symmetric thresholds converge to:[16] [17]

$$R_A^* = R_B^* = R^* \equiv 1 + \chi(1 + \lambda - y_A - y_B) \in (1, \bar{r}) \tag{2.11}$$

The banks' liquidity choices affects the withdrawal thresholds. More liquidity allows to serve more withdrawing investors and thus reduces the coordination failure among investors of a given bank. Thus, more liquidity held by bank n reduces the withdrawal threshold R_n^*. Because of systemic liquidation costs, there is also coordination failure among investors of different banks. More liquidity held by a given bank reduces the degree of this coordination failure and therefore the other bank's threshold R_{-n}^*:

$$\frac{\partial R_n^*}{\partial y_n} = \frac{\partial R_n^*}{\partial y_{-n}} = -\chi < 0$$

Lemma 6 summarizes the new results in the case of systemic liquidation costs.

Lemma 6. *Consider the withdrawal subgame with systemic liquidation costs ($d = 1$), vanishing private noise ($\gamma \to \infty$), and scarce liquidity ($y_n \in (\lambda, \frac{(1+\lambda)}{2})$. Then, the marginal investor expects liquidation to occur ($L_n(x_n^*) > 0$), and the withdrawal threshold is $x_n^* \to 1+\chi[1+\lambda-y_A-y_B] \in (1, \bar{r})$ if $\bar{r} > \bar{r}_L = 1+\chi(1-\lambda)$. There is actual liquidation ($l_n^* > 0$) if and only if th economic condition is bad ($r \leq \tilde{r}_1 \to x_n^*$). The equilibrium threshold highlights the system-wide effects of liquidity because more liquidity held at either bank reduces the withdrawal threshold of a given bank.*

[16]This can be proved by contradiction. Suppose that $R_A^* > R_B^*$. Then, $W_B^A \to 0$ and $W_A^B \to (1-\lambda)$. The implied expressions for the thresholds can never satisfy the supposed inequality $R_A^* > R_B^*$. The argument applies for $R_A^* < R_B^*$ as well. Therefore, $R_A^* = R_B^*$ as claimed.

[17]The symmetry in the liquidation cost function implies an equal weight of liquidity choices in the withdrawal threshold expression. This can be relaxed, for example by putting a larger weight on the own liquidation volume or with a convex liquidation specification. Either specification implies a larger weight of a bank's withdrawal threshold on its own liquidity. For example, a liquidation cost function that is linear in both the own and the total liquidation volume $x(l_n, l_{-n}) = \chi l_n(l_n + dl_{-n})$ yields $R_n^* \to 1 + \chi(1 + \lambda - y_A - y_B)(\frac{1+\lambda}{2} - y_n)$ in the case of systemic liquidation costs.

2.3.3 Optimal portfolio choice

I complete the characterisation of the equilibrium by studying the banks' privately optimal liquidity choice at the initial date. To generate macro-prudential implications, I consider the setup with systemic liquidation costs and scarce liquidity. A lower bound on the expected investment return derived below suffices to generate scarce liquidity in equilibrium. As multiple equilibria occur for scarce liquidity (see proposition 5), some equilibrium selection is required. Since liquidity has a beneficial effect in the equilibrium with liquidation, I assume that this equilibrium in the subgame is selected whenever it exists.[18] A bank's objective function is the expected utility of its investors derived in Appendix 2.7.2 and given by:

$$EU_n(y_n, y_{-n}) = y_n + (1 - y_n)\left[F(R^*) \cdot 1 + (1 - F(R^*)) \cdot \left(\bar{r} + \frac{f(R^*)}{\alpha(1 - F(R^*))}\right)\right] \quad (2.12)$$

where $f(r) = \phi(\sqrt{\alpha}[r - \bar{r}])$ is the probability distribution function of the normally distributed investment return and $F(r)$ the associated cumulative distribution function. The expected utility has two terms. The first term is the amount of liquidity, and the second term is the payoff from the investment $(1 - y_n)$. If the investment return falls short of the threshold R^*, which occurs with probability $F(R^*)$, the project is liquidated. Otherwise, the project is continued, which occurs with probability $1 - F(R^*)$, and the expected investment return conditional on continuation is $\mathbb{E}[r|r > R^*] = \bar{r} + \frac{f(R^*)}{\alpha(1-F(R^*))}$.

A lower withdrawal threshold improves expected utility as it implies a smaller area of inefficient withdrawals by reducing the extent of coordination failure ($\partial EU_n/\partial R^* < 0$), as derived in Appendix 2.7.2. This highlights the beneficial role of liquidity in the equilibrium with liquidation: more liquidity reduces coordination failure and therefore the withdrawal threshold, thereby indirectly improving the expected utility of an investor:

$$\frac{\partial EU_n}{\partial R^*}\frac{\partial R^*}{\partial y_n} > 0$$

There is also a detrimental role of liquidity. As the ex-ante opportunity cost of liquidity is the foregone higher expected investment return, holding more liquidity is costly. This is further exacerbated by optimal liquidation, shielding the investor from particularly

[18] This which can be generalized to any fraction $p \in [0, 1]$.

adverse outcomes of the project. The direct effect of liquidity is:

$$\frac{\partial EU_n}{\partial y_n} = -(1 - F(R^*))\left(\mathbb{E}[r|r \geq R^*] - 1\right) < 0$$

In case of continuation, which occurs with probability $1 - F(R^*)$, the expected investment return conditional on continuation exceeds the unit return to liquidity. In case of liquidation, which occurs with probability $F(R^*)$, the project and liquidity both yield a unit return.

The bank balances the beneficial and detrimental effects of liquidity. It takes the response of investors at the interim date $R^*(y_n, y_{-n})$ into account and the other bank's choice of liquidity y_{-n} as given. The optimal liquidity choice of bank n solves the following problem:

$$y_n^*(y_{-n}) \equiv \arg\max_{y_n} EU_n(y_n, y_{-n}) \text{ s.t. } R^* = R^*(y_n, y_{-n}) \tag{2.13}$$

where the best response function $y_n^*(y_{-n})$ is determined by the first-order condition:

$$\begin{aligned} \frac{dEU_n}{dy_n} = \frac{\partial EU_n}{\partial y_n} + \frac{\partial EU_n}{\partial R_n^*}\frac{\partial R_n^*}{\partial y_n} &= 0 \\ \underbrace{\chi(1-y_n)(R^*-1)f(R^*)}_{\text{marginal benefits from liquidity}} &= \underbrace{[1-F(R^*)]\left(\frac{1}{\alpha}\frac{f(R^*)}{1-F(R^*)} + \bar{r} - 1\right)}_{\text{marginal (opportunity) costs of liquidity}} \end{aligned} \tag{2.14}$$

I derive conditions on the expected investment return to ensure the existence of a unique best response function in Appendix 2.7.3. First, an upper bound on the expected investment return $\bar{r}_H$ ensures that the first-order condition has a solution for any feasible liquidity choice of the other bank. Second, a lower bound on the investment return $\bar{r}_L$ ensures that liquidity is indeed scarce as supposed, again for any feasible liquidity choice of the other bank. Finally, I show that the objective function EU_n is globally concave in the level of liquidity y_n. Therefore, a unique solution $y_n^*(y_{-n})$ exists for any level of liquidity held by the other bank.

There is strategic substitutability in liquidity holdings. If the other bank holds more liquidity, the liquidation cost of a given bank is reduced for any given level of liquidity. As holding liquidity is costly, the bank optimally reduces its liquidity level, free-riding on the other bank's liquidity. The other bank's liquidity holding is only useful for partially deterring a run since a potential liquidation cost is reduced, but not for serving investors when they do withdraw. Thus, the reduction in liquidity is less than one-for-one:

$$\frac{dy_n^*}{dy_{-n}} = -\frac{(R^* - 1) + \chi(1 - y_n)[1 + \alpha(R^* - 1)(\bar{r} - R^*)]}{2(R^* - 1) + \chi(1 - y_n)[1 + \alpha(R^* - 1)(\bar{r} - R^*)]} \in (-1, 0) \tag{2.15}$$

Since the slope of the best-response function lies strictly within the unit circle, is bounded and symmetric, there exists a unique and symmetric level of liquidity held at each bank: $y_n^* \equiv y^*$. It is implicitly given by $\frac{dEU_n}{dy_n}(y^*, y^*) = 0$. Proposition 6 summarizes.

Proposition 6. *Consider the overall game with systemic liquidation costs ($d = 1$), vanishing private noise ($\gamma \to \infty$), and an expected investment return within the range $(\bar{r}_L, \bar{r}_H)$. Suppose that the equilibrium with liquidation is selected if multiple equilibria exist in the withdrawal subgame. Then, there exists a unique and symmetric equilibrium in threshold strategies. It is characterized by a bank's liquidity choice $y_A^* = y_B^* \equiv y^* \in \left(\lambda, \frac{1+\lambda}{2}\right)$ at the initial date and withdrawal threshold of investors in the subgame that are (implicitly) given by:*

$$R^* = 1 + \chi(1 + \lambda - 2y^*) \in (1, \bar{r}) \tag{2.16}$$

$$\chi(1 - y^*)(R^* - 1) = \frac{1}{\alpha} + (\bar{r} - 1)\frac{1 - F(R^*)}{f(R^*)} \tag{2.17}$$

The boundaries on the expected investment return are $\bar{r}_L \equiv 1 + \chi(1 - \lambda)$ and $\bar{r}_H \equiv 1 + \frac{f(1+0.5\chi(1-\lambda))}{1-F(1+0.5\chi(1-\lambda))}[0.5\chi^2(1 - \lambda)^2 - 1/\alpha]$.

The equilibrium is characterised by partial free-riding on the respective other bank's liquidity.

2.4 Welfare

This section derives the liquidity choice of a social planner and compares it to the bank's optimal portfolio choice. As in Lorenzoni (2008), I adopt the notion of constrained efficiency: the social planner chooses the levels of liquidity but takes the optimal withdrawal decision of investors at the interim date as given. A direct choice of the threshold would achieve the first-best allocation ($R_n^* = 1$). In contrast to a bank, the planner internalizes the beneficial effects of liquidity for another bank's investors (system-wide effects of liquidity). Therefore, the constrained planner can be thought of as a macro-prudential authority.

The constrained *socially efficient* levels of liquidity (y_A^{SP}, y_B^{SP}) solve the planner's portfolio choice problem at the initial date, taking investors' responses at the interim date $R_n^*(y_A, y_B)$ into account:

$$(y_A^{SP}, y_B^{SP}) \equiv \arg\max_{y_A, y_B} SWF \equiv EU_A + EU_B \text{ s.t. } R_A^*(y_A, y_B) = R_B^*(y_A, y_B) \tag{2.18}$$

The first-order condition for the social planner's problem is:

$$\begin{aligned} 0 = \frac{dSWF}{dy_n} &= \frac{\partial EU_n}{\partial y_n} + \frac{\partial EU_n}{\partial R_n^*}\frac{\partial R_n^*}{\partial y_n} + \frac{\partial EU_{-n}}{\partial R_{-n}^*}\frac{\partial R_{-n}^*}{\partial y_n} \\ \underbrace{\chi(1 - y_A + 1 - y_B)(R^* - 1)f(R^*)}_{\text{social marginal benefits from liquidity}} &= \underbrace{[1 - F(R_n^*)]\left(\frac{1}{\alpha}\frac{f(R_n^*)}{1 - F(R^*)} + \bar{r} - 1\right)}_{\text{social marginal costs of liquidity}} \end{aligned} \tag{2.19}$$

The planner balances the social marginal cost of liquidity in terms of foregone investment return conditional on continuation $(\partial EU_n/\partial y_n < 0)$ with the social marginal benefits from liquidity in terms of lower withdrawal thresholds. The private and social marginal costs of liquidity coincide, while the social marginal benefits from liquidity exceed the private marginal benefit. Apart from the beneficial effect of liquidity on the investors of one bank $(\partial R_n^*/\partial y_n < 0)$, which is identical to the private benefit from liquidity, the planner also considers the beneficial effect of liquidity on the other bank's investors $(\partial R_{-n}^*/\partial y_n < 0)$. Recall that more liquidity allows to serve more withdrawing investors and therefore avoids costly liquidation for a given number of withdrawals, thereby mitigating the coordination failure between investors.

The optimization problem is fully symmetric. There is full substitutability between liquidity held at one bank and that held at another to reduce the withdrawal threshold $R_n^{SP} = R^{SP} = 1 + \chi[1 + \lambda - y_A^{SP} - y_B^{SP}]$. Furthermore, both first-order conditions yield the same condition (equation 2.19). Therefore, only the total amount of liquidity is determined $y_{total}^{SP} \equiv y_A^{SP} + y_B^{SP}$.

I derive conditions on the expected investment return to ensure the existence of a unique constraint efficient liquidity level. First, the upper bound on the expected investment return changes relative to the bank's portfolio choice, and the following upper bound on the investment return is now required:

$$\bar{r}_H < \bar{r}_H^{SP} \equiv 1 + \left(2\chi^2(1-\lambda)^2 - \frac{1}{\alpha}\right) \frac{f(1+\chi(1-\lambda))}{1-F(1+\chi(1-\lambda))} \tag{2.20}$$

which is strictly below the upper bound of $1 + 2\chi$ implied by no-dominance. Second, the upper bound $y_{total} \to 2\bar{y}$ is never optimal. Finally, the global concavity of the social welfare function in the total amount of liquidity is established in Appendix (2.7.4) for which the no-dominance bound on the expected investment return suffices. Taking these points together, there exists a unique level of total liquidity y_{total}^{SP} that maximizes social welfare and is implicitly given by $\frac{dSWF}{dy_n}(y_{total}^{SP}) = 0$.

Proposition 7 summarizes and compares the total amount of liquidity held by a planner with the total amount of liquidity held by banks:

Proposition 7. *Consider the overall game with systemic liquidation costs* ($d = 1$), *vanishing private noise* ($\gamma \to \infty$), *and an expected investment return* $\bar{r} < \bar{r}_H^{SP} \equiv 1 + \left(2\chi^2(1-\lambda)^2 - \frac{1}{\alpha}\right) \frac{f(1+\chi(1-\lambda))}{1-F(1+\chi(1-\lambda))}$. *A macro-prudential authority, the constrained social planner, chooses the liquidity level and investors respond optimally as before. Then, there exists a unique level of total liquidity* y_{total}^{SP} *that maximizes social welfare and is implicitly given by:*

$$\chi(2 - y_{total}^{SP})(R^{SP} - 1) = \left(\frac{1}{\alpha} + [\bar{r} - 1]\frac{1 - F(R^{SP})}{f(R^{SP})}\right) \tag{2.21}$$

where $R^{SP} = 1 + \chi[1 + \lambda - y_{total}^{SP}]$ *is the withdrawal threshold of investors at either bank.*

A macro-prudential authority holds more liquidity than the private banking system:

$$y_{total}^{SP} > y_A^* + y_B^* \tag{2.22}$$

Proof. The higher level of liquidity held by a macro-prudential authority remains to be proven. Relative to the bank's first-order condition, the right-hand side of the social planner's first-order condition (2.19) has an additional positive term, the positive externality of liquidity in terms of reducing the other bank's withdrawal threshold. Thus, the social benefits from liquidity exceed the social cost of liquidity when evaluated at the optimal level $y_A^* = y_b^* = y^*$:

$$\left.\frac{dSWF}{dy_n}\right|_{y_n = y_{-n} = y^*} > 0 \tag{2.23}$$

Given the strict global concavity of the objective function in the total amount of liquidity y_{total}, the planner's total amount of liquidity must be higher ($y^{SP}_{total} > y^*_A + y^*_B$), thereby internalising the positive system-wide externality of liquidity. $\square$

The difference between the constraint efficient and the optimal level of liquidity is interpreted as a *macro-prudential liquidity buffer*. A constrained planner, such as a macro-prudential authority, takes all economy-wide effects into account by holding more liquidity, internalising the social costs of liquidation that arise in the presence of systemic liquidation cost.

2.5 Comparative Statics

This sections studies how the equilibrium allocation y^* and the planner's allocation y^{SP}_{total} vary with the exogenous parameters of the model. Parameters of interest are the liquidation cost parameter χ, the expected investment return $\bar{r}$, and the proportion of early investors λ. Proposition 8 summarizes the results.

Proposition 8. *The private and social levels of liquidity vary according to*

(a) $\frac{\partial y^*}{\partial \chi} > 0$ *and* $\frac{\partial y^{SP}_{total}}{\partial \chi} > 0$ *such that a higher liquidation cost raises the liquidity held privately and socially;*

(b) $\frac{\partial y^*}{\partial \bar{r}} < 0$ *and* $\frac{\partial y^{SP}_{total}}{\partial \bar{r}} < 0$ *such that a higher investment return lowers the private and social levels of liquidity;*

(a) $\frac{\partial y^*}{\partial \lambda} > 0$ *and* $\frac{\partial y^{SP}_{total}}{\partial \lambda} > 0$ *such that a larger proportion of early investors induces higher liquidity holdings.*

See Appendix 2.7.5 for a proof. The intuition underlying these results is as follows. First, a larger proportion of early investors increases the liquidity held privately and socially. Since early investors wish to consume at the interim date and always withdraw, more liquidity is held to serve them.

Second, the strength of the liquidation cost is captured by the liquidation cost parameter χ. It affects the benefits from holding liquidity in terms of avoiding costly liquidation in case of elevated withdrawals, thereby reducing the withdrawal threshold.

Thus, if liquidation is more costly, such as in times of financial distress, then liquidity is particularly valuable and more liquidity is held both privately and socially.

Third, a higher expected investment return $\bar{r}$ affects the ex-ante opportunity cost of holding liquidity. Therefore, both banks and the planner hold more liquidity when the project pays a better return on average. Note that there is no effect of the expected investment return on the withdrawal threshold R_n^* as private noise vanishes. However, if the private noise is bounded ($\underline{\gamma} < \gamma < \infty$), a higher investment return also reduces the run threshold. This second effect would further reduce the level of liquidity held.

2.6 Conclusion

This paper examined the role of liquidity in an economy with many banks subject to runs and systemic liquidation costs. I showed that the presence of liquidity, which drives a wedge between the amount of withdrawals and the liquidation volume, restores multiple equilibria – even if a global game refinement is used. Apart from the usual equilibrium with liquidation (Morris and Shin (2000); Goldstein and Pauzner (2005)), a no-liquidation equilibrium exists for a range of economic conditions. Furthermore, systemic liquidation costs imply that one bank's liquidity holding reduces the liquidation costs of other banks. The positive implication is the partial substitutability of private liquidity holdings as banks free-ride on the liquidity holdings of other banks. The normative implication is that banks hold insufficient liquidity relative to the average liquidity holding of a constrained planner. Since a planner internalizes the system-wide effects of liquidity, I interpret the planner as a macro-prudential authority.

This framework provides a natural laboratory for studying macro-prudential policies in a micro-founded setting more generally. I abstracted from capital requirements, diversification, and taxes on withdrawals in this paper, but analyze some of these in other work. There are other elements relevant to the conduct of macro-prudential regulation omitted in this framework, such as limited liability, 'too big to fail', and perverse incentives arising from incentive schemes. These are all exciting avenues for subsequent research.

2.7 Appendix

2.7.1 Posterior distributions

Investment return The posterior mean of the investment project return is a weighted average of the mean of the prior distribution and the private signal, in which the relative weights are given by the respective precisions. The precision of the posterior distribution is the sum of the precisions of the prior and the signal. Normality is preserved:

$$\boldsymbol{R}_i^n \sim \mathcal{N}\left(\frac{\alpha\bar{r}+\gamma x_i}{\alpha+\gamma}, \frac{1}{\alpha+\gamma}\right) \tag{2.24}$$

The ratio of the precision of the prior (public signal) relative to the private signal, $\frac{\alpha}{\gamma}$, determines the extent to which the posterior mean depends on the private signal. The more precise the private signal relatively to the prior, the more the posterior is determined by the private signal. In the limit of vanishing private noise ($\frac{\alpha}{\gamma} \to 0$ as $\gamma \to \infty$), the posterior mean converges to the private signal.

Proportion of prematurely withdrawing late investors at bank n Using the definition of the proportion of withdrawing investors, the posterior distribution of the mean, and a law of large numbers, the posterior proportion of withdrawing late investors at a given investor's bank $W_{i,n}^n$ can be written as:

$$W_{i,n}^n = (1-\lambda)\Phi\left(\sqrt{\delta}\,[R_n^* - \bar{r}] + \sqrt{\frac{\gamma(\alpha+\gamma)}{\alpha+2\gamma}}\,[R_n^* - R_i^n]\right) \tag{2.25}$$

$$\delta \equiv \frac{\alpha^2(\alpha+\gamma)}{\gamma(\alpha+2\gamma)} \tag{2.26}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution and δ summarizes precision parameters. A late investor that receives the threshold signal $x_i = x_n^*$ thus forms the following posterior mean of the proportion of withdrawing late investors at his bank:

$$(W_n^n)^* \equiv W_{i,n}^n\big|_{x_i=x_n^*} = (1-\lambda)\Phi(z_{1n}) \tag{2.27}$$

$$z_{1n} \equiv \sqrt{\delta}\,[R_n^* - \bar{r}] \tag{2.28}$$

If the private noise vanishes ($\gamma \to \infty$), then $\delta \to 0$ and $(W_n^n)^* \to \frac{1-\lambda}{2}$.

Proportion of prematurely withdrawing late investors at bank $-n$ Withdrawal thresholds may differ across banks. Depending on the other bank's threshold R^*_{-n}, an investors at bank n expects the following proportion of withdrawing late investors at bank $-n$:

$$
\begin{aligned}
W^n_{i,-n} &= (1-\lambda)\Phi\left(\sqrt{\delta}\left[R^*_{-n}-\bar{r}\right]+\sqrt{\frac{\gamma(\alpha+\gamma)}{\alpha+2\gamma}}\left[R^*_{-n}-R^n_i\right]\right) && (2.29)\\
\left(W^n_{-n}\right)^* &\equiv W^n_{i,-n}\big|_{x_i=x^*_n}=(1-\lambda)\Phi(z_{2n}) && (2.30)\\
z_{2n} &\equiv \sqrt{\delta}\left[R^*_{-n}-\bar{r}\right]+\sqrt{\delta}\frac{\gamma}{\alpha}\left[R^*_{-n}-R^*_n\right] && (2.31)
\end{aligned}
$$

2.7.2 Derivation of expected utility EU_n

When private noise vanishes ($\gamma \to \infty$), equilibrium withdrawals by late investors at the interim date are:

$$
w^*_n(r)=(1-\lambda)\Phi\left(\frac{\alpha}{\sqrt{\gamma}}[R^*_n-\bar{r}]+\sqrt{\gamma}[R^*_n-r]\right)\to\begin{cases}0 & r>R^*_n\\ \frac{1-\lambda}{2} \quad \text{if} & r=R^*_n\\ 1-\lambda & r<R^*_n\end{cases} \tag{2.32}
$$

Therefore, there is no liquidation if the project return is above the threshold R^*_n, while the investment project is completely liquidated if the investment return is below the threshold. Late investors receive the continuation payoff c_{2n} in the former case and unity in the latter. Early investors always receive unity as promised. Adding these components up, the expected utility is:

$$
\begin{aligned}
EU_n(y_n,y_{-n}) &= \int_{-\infty}^{R^*_n} 1\cdot 1dF(r)+\int_{R^*_n}^{\infty}\lambda\cdot 1+(1-\lambda)\cdot\frac{y_n-\lambda+(1-y_n)r}{1-\lambda}dF(r) && (2.33)\\
&= y_n+(1-y_n)\left[F(R^*_n)\cdot 1+(1-F(R^*_n))\cdot\left(\bar{r}+\frac{f(R^*_n)}{\alpha(1-F(R^*_n))}\right)\right] && (2.34)
\end{aligned}
$$

The partial derivatives are:

$$
\begin{aligned}
\frac{\partial EU_n}{\partial y_n} &= -(1-F(R_n^*))\left(\mathbb{E}[r|r \geq R_n^*]-1\right) < 0 && (2.35)\\
\frac{\partial EU_n}{\partial R_n^*} &= -(1-y_n)(R_n^*-1)f(R_n^*) < 0 && (2.36)\\
\frac{\partial^2 EU_n}{\partial y_n^2} &= \frac{\partial EU_n}{\partial y_{-n}} = 0 && (2.37)\\
\frac{\partial^2 EU_n}{\partial y_n \partial R_n^*} &= (R_n^*-1)f(R_n^*) > 0 && (2.38)\\
\frac{\partial^2 EU_n}{\partial (R_n^*)^2} &= -(1-y_n)f(R_n^*)[1+\alpha(R_n^*-1)(\bar{r}-R_n^*)] < 0 && (2.39)
\end{aligned}
$$

where the signs are implied by the ordering $1 < R_n^* < \bar{r}$ (Lemma 6).

2.7.3 Unique best response $y_n^*(y_{-n})$

Let $\Lambda(R^*) \equiv \frac{1-F(R^*)}{f(R^*)} > 0$ and therefore $\Lambda'(R^*) = -\sqrt{\alpha} - \alpha(\bar{r}-R^*)\Lambda(R^*) < 0$. The first-order condition becomes:

$$\chi^2(1-y_n^*)(1+\lambda-y_n^*-y_{-n}) = \frac{1}{\alpha} + (\bar{r}-1)\Lambda(R_n^*) \tag{2.40}$$

where $R_n^* = 1+\chi(1+\lambda-y_n^*-y_{-n})$. Note that the left-hand side (LHS) of equation (2.40) is decreasing in the liquidity level y_n^*, while the right-hand side (RHS) is increasing in it.

First, existence of equilibrium requires that the LHS exceeds the RHS when evaluated at the lower bound $y_n^* \to \lambda$ for any liquidity level y_{-n}. This inequality is hardest to satisfy for $y_{-n} \to \bar{y} \equiv \frac{1+\lambda}{2}$. Rewriting yields an upper bound on the expected investment return:

$$\bar{r} < \bar{r}_H \equiv 1 + \frac{\frac{\chi^2(1-\lambda)^2}{2} - \frac{1}{\alpha}}{\Lambda(1+0.5\chi(1-\lambda))} \tag{2.41}$$

This upper bound is strictly below the level of $1+2\chi$ as implied by no dominance and replaces this upper bound.

Second, the supposed scarcity of liquidity requires that the marginal cost of liquidity just exceeds its marginal benefit as the liquidity level converges its upper bound $\bar{y}$. Therefore: $LHS(y_n^* \to \bar{y}) < RHS(y_n^* \to \bar{y})$ for any liquidity level y_{-n}. This inequality is hardest to satisfy for $y_{-n} \to \lambda$. Rewriting yields a lower bound on the expected investment return:

$$\bar{r} > \bar{r}'_L \equiv 1 + \frac{\frac{\chi^2(1-\lambda)^2}{4} - \frac{1}{\alpha}}{\Lambda(1+0.5\chi(1-\lambda))} < \bar{r}_H \tag{2.42}$$

As a consequence of the no-dominance constraint, this lower bound is strictly below the level of $\bar{r}_L = 1 + \chi(1-\lambda)$, which ensures $R_n^* < \bar{r}$. Therefore, the lower bound of $\bar{r}_L$ is maintained. Note that $y_n^* \neq \bar{y}$ implies $y_n^* < \bar{y}$ by the global concavity of the objective function, which can be seen by the sign of the second-order derivative of the objective function:

$$\frac{d^2 EU_n}{dy_n^2} = \frac{\partial R_n^*}{\partial y_n}\left[2\frac{\partial^2 EU_n}{\partial y_n \partial R_n^*} + \frac{\partial^2 EU_n}{\partial (R_n^*)^2}\frac{\partial R_n^*}{\partial y_n}\right] + \frac{\partial EU_n}{\partial R_n^*}\frac{\partial^2 R_n^*}{\partial y_n^2} < 0$$

where $\frac{\partial^2 R_n^*}{\partial y_n^2} = 0$ and the sign follows directly from the previously established signs on the partial derivatives of the withdrawal threshold R_n^* and the expected utility EU_n.

2.7.4 Global concavity of SWF

Consider the second derivative of the social welfare function:

$$\frac{d^2 SWF}{d(y^{SP})^2} = -\chi f(R^*)[\sqrt{\alpha}(\bar{r}-1)+\chi(3+\lambda-y_{total})-(\bar{r}-R^*)(1-\alpha\chi^2(1+\lambda-y_{total}(2-y_{total})))] < 0$$

The highest possible values is reached when $\alpha \to 0$ and $y_{total} \to 2\bar{y}$. Then, the second-order derivative is still negative as $1 + 2\chi > \bar{r}$ by no-dominance. Therefore, the second-order derivative is always negative, establishing global concavity of the social welfare function.

2.7.5 Comparative statics

Privately optimal liquidity level y^*

Parameters of interest are $\chi, \bar{r}, \lambda$. The effect of parameters on the withdrawal threshold $R^* = 1 + \chi(1 + \lambda - 2y^*)$ is:

$$\frac{\partial R^*}{\partial \chi} = (1 + \lambda - 2y^*) > 0 \tag{2.43}$$

$$\frac{\partial R^*}{\partial \bar{r}} = 0 \tag{2.44}$$

$$\frac{\partial R^*}{\partial \lambda} = \chi > 0 \tag{2.45}$$

The first-order condition for the private level of liquidity y^* can be written as $G(a, R^*, y^*) = 0$, where $a \in \{\chi, \bar{r}, \lambda\}$ is a parameter:

$$G(a, R^*, y^*) = (r-1)\frac{1-F(R^*)}{f(R^*)} + \frac{1}{\alpha} - \chi^2(1-y^*)(1+\lambda-2y^*) \tag{2.46}$$

Then, the effect of a parameter on the equilibrium liquidity level is given by $\frac{dy^*}{da} = -\frac{\partial G/\partial a}{\partial G/\partial y}$. Note that:

$$\frac{\partial G}{\partial y^*} = -2\chi(\bar{r}-1)\Lambda'(R^*) + \chi^2(3+\lambda-4y^*) > 0 \tag{2.47}$$

$$\frac{\partial G}{\partial \bar{r}} = (\bar{r}-1)[\sqrt{\alpha} + (\bar{r}-R^*)\Lambda(R^*)] + \Lambda(R^*) > 0 \tag{2.48}$$

$$\frac{\partial G}{\partial \chi} = (\bar{r}-1)(1+\lambda-2y^*)\Lambda'(R^*) - 2\chi(1-y^*)(1+\lambda-2y^*) < 0 \tag{2.49}$$

$$\frac{\partial G}{\partial \lambda} = -\chi^2(1-y^*) + \chi(\bar{r}-1)\Lambda'(R^*) < 0 \tag{2.50}$$

Therefore, the partial derivatives of the privately held liquidity levels have the signs as claimed.

Socially efficient liquidity level y^{SP}_{total}

The effect of parameters on the withdrawal threshold $R^{SP} = 1 + \chi(1+\lambda-y^{SP}_{total})$ is:

$$\frac{\partial R^{SP}}{\partial \chi} = (1+\lambda-y^{SP}_{total}) > 0 \tag{2.51}$$

$$\frac{\partial R^{SP}}{\partial \bar{r}} = 0 \tag{2.52}$$

$$\frac{\partial R^{SP}}{\partial \lambda} = \chi > 0 \tag{2.53}$$

The first-order condition for the social level of liquidity y^{SP}_{total} can be written as $\tilde{G}(a, R^{SP}, y^{SP}) = 0$, where $a \in \{\chi, \bar{r}, \lambda\}$ is a parameter:

$$\tilde{G}(a, R^{SP}, y^{SP}_{total}) = (r-1)\frac{1-F(R^*)}{f(R^*)} + \frac{1}{\alpha} - \chi^2(2-y^{SP}_{total})(1+\lambda-y^{SP}_{total}) \tag{2.54}$$

Then, the effect of a parameter on the equilibrium liquidity level is given by $\frac{dy^{SP}_{total}}{da} = -\frac{\partial\tilde{G}/\partial a}{\partial\tilde{G}/\partial y}$. As above, partial differentiation of $\tilde{G}$ proves the signs on the comparative statics of the total level of liquidity held by the planner as claimed.

Chapter 3

Information contagion and systemic risk

Information contagion can reduce systemic risk defined as the joint default probability of banks. This paper examines the effects of ex-post information contagion on both the banks' ex-ante optimal portfolio choices and the implied welfare losses due to joint default. Because of counterparty risk and common exposures, bad news about one bank reveals valuable information about another bank, thereby triggering information contagion. We find that information contagion reduces (increases) the joint default probability when banks are subject to counterparty risk (common exposures). When applied to microfinance, our model also provides a novel explanation for higher repayment rates in group lending.

JEL Classifications: G01, G21, O16

Keywords: information contagion, counterparty risk, common exposure, systemic risk, microfinance, group lending

Co-authored with Co-Pierre Georg, Deutsche Bundesbank, Wilhelm-Epstein Strasse 14, D-60431 Frankfurt am Main, and Oxford University, Park End Road, Oxford OX1 1HP, United Kingdom. Part of this research was conducted when Co-Pierre Georg was a visiting scholar at the New York University Stern School of Business. E-Mail: `co-pierre.georg@keble.ox.ac.uk`

The authors wish to thank Viral Acharya, Jose Berrospide, Falko Fecht, Douglas Gale, Prasanna Gai (Discussant), Xavier Giné, Paul Glassermann (Discussant), Itay Goldstein, Todd Keister, Anton Korinek, Yaron Leitner, Ralf Meisenzahl, Jean-Charles Rochet, Cecilia Parlatore Siritto, Javier Suarez, Ernst-Ludwig von Thadden, Dimitri Vayanos, and Tanju Yorulmazer, as well as seminar participants at the IMF, LSE, NYU, the Philadelphia Fed and conference participants at the RBNZ-UA "Stability and efficiency of financial systems" conference, the 2012 FDIC Bank Research Conference, and EEA 2011 meeting for fruitful discussions and comments.

3.1 Introduction

Systemic risk is defined as the joint default of a substantial part of the financial system and is associated with large social costs.[1] One major source of systemic risk is information contagion: when investors are sensitive to news about the health of the financial system, bad news about one financial institution can adversely spill over to other financial institutions. For instance, the insolvency of one money market mutual fund with a large exposure to the investment bank Lehman Brothers spurred investor fears and led to a wide-spread run on all money market mutual funds in September 2008.[2] As information contagion affects various financial institutions including commercial banks, money market mutual funds, and shadow banks, we adopt a broad notion of financial intermediaries called banks for short.

There are at least two reasons for an investor of a bank to find information about another bank's profitability valuable. On the one hand, the first bank may have lent to the second bank in the past, for example to share liquidity risk as in Allen and Gale (2000). Learning about the debtor bank's profitability then helps the investor assess the counterparty risk of the creditor bank. On the other hand, both banks may have some common exposure to an asset class, such as risky sovereign debt or mortgage-backed securities. Learning about another bank's profitability then helps the investor assess the profitability of its bank.[3]

We develop a model of systemic risk with information contagion. Our model features two banks, where systemic risk refers to the ex-ante probability of joint default. Due to both counterparty risk and common exposures, bad news about one bank can trigger the default of another bank. Information contagion in this setup is the amount of a bank's additional financial fragility caused by such bad news. We examine the effects of ex-post information contagion on the ex-ante optimal portfolio choice of a bank and the implied level of systemic risk.

[1]BCBS (1997) compares the cost of systemic bank crises in various developing and industrialized countries and document the range from about 3% of GDP for the savings and loan crisis in the United States to about 30% of GDP for the 1981-87 crisis in Chile.

[2]Lehman Brothers failed on September 15, 2008 and the share price of the Reserve Primary Fund dropped below the critical value of 1$ on September 16, 2008.

[3]For example, the funding cost of one bank increases after adverse news about another bank because of correlated loan portfolio returns in Acharya and Yorulmazer (2008b).

Our first result refers to information contagion due to counterparty risk. When information spillover is unanticipated, that is it occurs with zero probability, the ex-ante optimal portfolio is unchanged and systemic risk increases (lemma 7). By contrast, anticipated information spillover makes the ex-ante portfolio choice more prudent to counteract ex-post information. Banks expose themselves less to counterparty risk by engaging in less liquidity co-insurance and hold more liquidity themselves. This *reduces* systemic risk (Result 1) and is labelled as a *resilience effect*. The direct detrimental effect of information contagion on systemic risk is more than fully compensated by an indirect beneficial effect via the ex-ante portfolio choice. Overall, systemic risk in the financial system is reduced once information spillover that give rise to information contagion is present.

We also analyze information contagion due to common exposures. When information spillover is unanticipated, systemic risk again increases (lemma 8) similar to lemma 7. When information spillover is anticipated, however, systemic risk increases (Result 2), which is labelled the *instability effect*. Taking these results together, the consequences of information contagion for the level of systemic risk (via changes of the ex-ante optimal portfolio choice) depend on the nature of the interbank linkage: financial fragility increases (decreases) when banks are linked via common exposure (counterparty risk).

Our main contribution is the analysis of information contagion due to counterparty risk and its effects on the ex-ante optimal portfolio choice and systemic risk. Counterparty risk as a source of information contagion and its consequences for the ex-ante portfolio choice have not been consistently studied before.[4] Our counterparty risk mechanism builds on the literature of financial contagion due to balance sheet linkages. Building on Diamond and Dybvig (1983), Allen and Gale (2000) describe financial contagion as an equilibrium result.[5] Interbank lending insures banks against a non-aggregate liquidity shock and potentially achieves the first-best outcome. However, a zero-probability aggregate liquidity shock may travel through the entire financial system. While counterparty risk in our model also arises from the potential default on interbank obligations, we obtain the ex-ante optimal portfolio choice given that contagion may occur with pos-

[4]Cooper and Ross (1998) and Ennis and Keister (2006) study the effect of ex-post individual bank runs on the ex-ante liquidity choice and the design of deposit contracts. By contrast, we analyze how information contagion due to counterparty risk affects the ex-ante portfolio choice and deposit contract design of banks and examine the consequences for the joint default probability of banks.

[5]Freixas et al. (2000) consider spatial instead of intertemporal uncertainty about liquidity needs.

itive probability.[6] Dasgupta (2004) also demonstrates the presence of financial contagion with positive probability in the unique equilibrium of a global game version of the model described by Allen and Gale (2000), focusing on the coordination failure initiated by adverse information. By contrast, we analyse the impact of information contagion from counterparty risk on the ex-ante portfolio choice of financial intermediaries, which is only partially addressed in Dasgupta (2004). Furthermore, our focus is on the consequences for systemic risk and we also analyse the role of common exposures.

Our results also relate to the literature on information contagion due to common exposures. Information about the solvency of one bank is an informative signal about the health of other banks with similar exposure in Acharya and Yorulmazer (2008b).[7] The anticipation of ex-post information contagion induces banks to correlate their ex-ante investment decisions, endogenously creating common exposures. By contrast, we consider counterparty risk as a principal source of information contagion. We also allow for a larger set of portfolio choice options.[8] Leitner (2005) analyzes the ex-ante beneficial insurance effects of ex-post financial contagion in the absence of an explicit ex-ante risk sharing mechanism due to limited commitment. By contrast, we focus on the ex-ante effects of ex-post *information* contagion in a model with commitment. Allen et al. (2012) analyze systemic risk stemming from the interaction of common exposures and funding maturity through an information channel.[9] However, our focus is on the novel analysis of counterparty risk as a source of information contagion and its repercussions for systemic risk.[10]

[6]Postlewaite and Vives (1987) show the uniqueness of equilibrium with positive probability of bank runs in a Diamond and Dybvig (1983) setup with demand deposit contracts and four periods. By contrast, we analyse the impact of information contagion from counterparty risk and common exposures on the ex-ante optimal portfolio choice and the implied level of systemic risk.

[7]Other models of common exposure include Acharya and Yorulmazer (2008a), who analyze the interplay between government bail-out policies and banks' incentives to correlate their investments. Chen (1999) shows that bank runs can be triggered by information about bank defaults when banks have a common exposure. Uninformed investors use the publicly available signal about the default of another bank to assess the default probability of their bank. An early model of information-based individual fragility is Jacklin and Bhattacharya (1988).

[8]While interconnectedness of banks only arises through the endogenous choice of correlated investments in Acharya and Yorulmazer (2008b), we maintain the exogenous correlation of the bank's investment returns as in Acharya and Yorulmazer (2008a) but endogenize liquidity holdings, interbank liquidity insurance (*co-insurance* as in Brusco and Castiglionesi (2007)), and insurance of impatient investors against idiosyncratic liquidity shocks.

[9]Banks swap risky investment projects to diversify, generating different types of portfolio overlaps. Investors receive a signal about the solvency of all banks at the final date. Upon the arrival of bad news about aggregate solvency, roll-over of short-term debt occurs less often when assets are clustered, leading to larger systemic risk.

[10]Furthermore, we consider an investment allocation between a safe and a risky asset and information

Our results on the interaction of information spillovers and counterparty risk are not limited to systemic risk in banking of advanced economies. Counterparty risk also arises from joint liability in group lending contracts commonly used by the Grameen bank and other microfinance institutions in developing economies (see e.g. Stiglitz (1990), Varian (1990), or Morduch (1999)). The idea behind group borrowing is to employ peer monitoring to overcome asymmetric information. Thus, borrowers in a group will know each other quite well (either neighbors from the same village, or even family members) and information spillover occurs frequently. In particular, our resilience effect predicts that (i) group loans have a higher repayment rate than individual loans and (ii) group borrowers hold more liquid assets. As described in section 3.5, these predictions are verified in the empirical microfinance literature.

The remainder of this paper is as follows. The model is described in section 3.2 and its equilibrium is analyzed in section 3.3, including a discussion of special limiting cases that provide further intuition to our model. We present our results in section 3.4, which also contains extensive robustness checks. Our model is applied to microfinance in section 3.5, providing a novel explanation for empirical findings in that literature. Finally, section 3.6 concludes. Derivations, proofs, and tables are found in appendices 3.7, 3.8, and 3.9.

3.2 Model

The economy extends over three dates labelled as initial ($t = 0$), interim ($t = 1$), and final ($t = 2$) and consists of two regions ($k = A, B$) interpreted as geographic regions or asset classes. Each region is inhabited by a bank and a unit continuum of investors. Our notion of financial intermediation is broad, capturing both the traditional case of retail investors at commercial banks and institutional investors at money market mutual funds.[11] There is a single physical good used for consumption and investment. The focus of this paper is on systemic risk measured by the probability of the joint failure of banks at the initial date.

spillover about bank-specific solvency.

[11]In the language of Uhlig (2010), our banks corresponds to core banks, while our investors correspond to local banks.

Inhabitants of each region have access to two investment opportunities at the initial date. First, storage produces one unit at the following date per unit invested. Second, a risky regional investment project matures at the final date and produces a stochastic output R_k that exceeds the output from storage in expectation ($\mathbb{E}[R_k] > 1$), where $\mathbb{E}$ is the expectation operator. Liquidation of the project at the interim is costly, producing an inferior output $\beta \in (0,1)$ only. Since the recovery rate is positive, liquidation is optimal if the realized output is known to be low. We adopt a bivariate specification of the project output:

$$R_k = \begin{cases} R & \text{w.p. } \theta_k \\ 0 & \text{w.p. } 1 - \theta_k \end{cases} \tag{3.1}$$

where $R > 2$ and the regional fundamental is uniformly distributed ($\theta_k \sim U[0,1]$) and interpreted as a regional solvency shock. Let $corr(\theta_A, \theta_B)$ denote the correlation between the regional fundamentals, where $corr(\theta_A, \theta_B) = 1$ refers to a common exposure. Despite common exposure, the realised regional project outputs can differ because of the individual randomness of each project. We abstract from portfolio diversification motivated by limits to monitoring, for instance.

As in Diamond and Dybvig (1983), investors learn their liquidity preference privately at the interim date. Early investors wish to consume at the interim date, while late investors wish to consume at the final date. The ex-ante probability of being an early investor $\lambda \in (0,1)$ is identical across investors and equals the regional proportion of early investors by a law of large numbers. The investor's period utility function $u(c)$ is twice continuously differentiable, strictly increasing, strictly concave and satisfies the Inada conditions. Thus, the expected utility of an investor is:

$$\mathbb{E}_\lambda[U(c_1, c_2)] = \lambda u(c_1) + (1 - \lambda) u(c_2) \tag{3.2}$$

where c_t is the investor's consumption at date t. Investors in each region are endowed with one unit at the initial date to be invested or deposited at their regional bank.

The role for banks in our model is the traditional provider of liquidity insurance (Diamond and Dybvig (1983)), which arises from the smaller volatility of regional liquidity demand than individual liquidity demand. Banks offer demand deposit contracts that specify withdrawals (d_1, d_2) if funds are withdrawn at the interim or final date, where we set $d_2 \equiv \infty$ without loss of generality.[12] Bank pay an equal amount to all withdrawing investors (pro-rata) in case of default. There is free entry to the banking sector, ensuring that each bank maximizes the expected utility of a representative investor.[13] Investors deposit in full since their interest is fully aligned with their bank's.

A bank is illiquid if a sufficiently large proportion of late investors withdraws and the project has to be partially liquidated. A bank is insolvent if a yet larger proportion of late investors withdraws and the full liquidation of the project does not suffice to serve them. An important insight of Diamond and Dybvig (1983) is that the strategic complementarity in investors' withdrawal decisions generates multiple equilibria, of which the inefficient one features a bank run. We focus on essential bank runs as in Allen and Gale (2007), however, whereby a run takes place only if it is unavoidable. That is, the no-run equilibrium is selected if multiple equilibria exist. Let a_k be the default probability of an individual bank and $A \equiv a_A a_B$ be the probability of joint default, which is our measure of systemic risk.

Counterparty risk is introduced via interbank insurance as in Allen and Gale (2000) because of negatively correlated regional liquidity demand. A region has low liquidity demand ($\lambda_L \equiv \lambda - \eta$) or high liquidity demand ($\lambda_H \equiv \lambda + \eta$) with equal probability, where $\eta > 0$ is the size of the regional liquidity shock.[14] To exclude bank runs merely driven by aggregate liquidity shortage, we study negatively correlated liquidity shocks of equal size:

probability	region A	region B
$\frac{1}{2}$	$\lambda_A = \lambda + \eta$	$\lambda_B = \lambda - \eta$
$\frac{1}{2}$	$\lambda_A = \lambda - \eta$	$\lambda_B = \lambda + \eta$

At the initial date banks agree on mutual liquidity insurance interpreted as mutual lines of credit or cross-holding of deposits.[15] The bank with high liquidity demand receives an amount $b \geq 0$ from the bank with low liquidity demand at the beginning of the interim date. Repayment with interest ($\phi \geq 1$) takes place at the final date if the debtor bank is

[12] Since the liquidity preference of an individual investor is private knowledge, the deposit contract between the bank and the investor cannot be contingent on it.

[13] See also Gale (2010).

[14] Freixas et al. (2000) motivate interbank insurance by allowing for interregional travel of investors who learn the location of their liquidity demand at the beginning of the first period.

[15] Since banks are symmetric at the initial date, they wish to exchange the same amount of deposits.

solvent. A solvent debtor bank repays the creditor bank even if the latter is insolvent.[16] We make the common assumption of seniority of interbank loans at the final date only.[17] Non-defaulted interbank claims can be liquidated at rate β.

There is strategic interaction between banks in their portfolio choices. At the initial date banks simultaneously choose the amount of investment in the project $1 - y_k$, the demand deposit contract, and agree on the volume of interbank insurance. A bank's portfolio choice affects its solvency threshold $\bar{\theta}_k$ below which an essential bank run occurs. Furthermore, it affects another bank's solvency threshold $\bar{\theta}_{-k}$ due to counterparty risk and information contagion at the interim date. The optimal portfolio choices of banks at the initial date are determined as a symmetric pure-strategy Nash equilibrium.

Turning to the information structure of the model, all prior distributions are common knowledge. Before making their withdrawal decision at the interim date, investors may receive independent signals about the success probabilities (θ_A, θ_B) with probability (q_A, q_B). Therefore, investors may receive no, one or two signals. If a signal is received, it perfectly reveals the regional success probability to the investor. If no signal is received, nothing is learnt.

Information spillover occurs if investors of one bank learn about the solvency of another bank. Such information is valuable to investors for two reasons. In case of common exposure investment returns are correlated and the knowledge about one bank's solvency helps to predict another bank's solvency. In case of counterparty risk, learning about the debtor bank's solvency helps investors predict the solvency of the creditor bank. Information contagion occurs if investors run on a bank upon learning about another bank's solvency but would not have done so without the information. Our interest is in analyzing the effect of such information contagion at the interim date on the optimal portfolio choice at the initial date and the implied systemic risk.

We close the description of the model by determining the investors' payoffs. Starting with the high liquidity demand or debtor region (H), the payoffs are independent of the behavior in the low liquidity demand region. If the bank is insolvent, all funds are liquidated and the interbank loan is defaulted upon. The impatient investor's payoff is

[16] We assume the existence of a liquidator of the creditor bank to which the solvent debtor bank repays its debt at the final date. This assumption is natural as the liquidation of banks does not destroy value because of claims on viable institutions.

[17] See, for example, Dasgupta (2004).

$d_H \equiv y+(1-y)\beta+b$. There are never partial runs since all bank runs are essential. If the bank is liquid, no liquidation takes place and the interbank loan is repaid. The patient investor's payoffs is $c_{2H}^G \equiv \frac{(1-y)R+y-\lambda_H d_1-(\phi-1)b}{1-\lambda_H}$ in the good state and $c_{2H}^B \equiv \frac{y-\lambda_H d_1-(\phi-1)b}{1-\lambda_H}$ in the bad state. Superscripts (G, B) denote success (good state) and failure (bad state) of the investment project and occur with probability $1-\theta_H$ and θ_H, respectively.

The bank in the low liquidity demand or debtor region (L) pays b to the bank in the high liquidity demand region at the interim date. In the case of a bank run in L, all assets including the interbank claim are liquidated, yielding a payoff $y+(1-y)\beta-b+\beta\phi\widetilde{b}$. The repayment of the interbank claim $\widetilde{b}$ is uncertain: it yields b if H repays and zero otherwise. The resulting payoffs are $d_L^N \equiv y+(1-y)\beta+(\beta\phi-1)b$ and $d_L^D \equiv y+(1-y)\beta-b$. Superscripts (N, D) denote survival and default of the bank in H. The liquidation value of the interbank claim is positive in case of repayment only. Hence, patient investors receive $c_{2L}^{GN} \equiv \left(\frac{(1-y)R+(y-\lambda_L d_1)+(\phi-1)b}{1-\lambda_L}\right)$ and $c_{2L}^{GD} \equiv \left(\frac{(1-y)R+(y-\lambda_L d_1)-b}{1-\lambda_L}\right)$ in the good state as well as $c_{2L}^{BN} \equiv \left(\frac{(y-\lambda_L d_1)+(\phi-1)b}{1-\lambda_L}\right)$ and $c_{2L}^{BD} \equiv \left(\frac{(y-\lambda_L d_1)-b}{1-\lambda_L}\right)$ in the bad state.

Table (3.1) provides a timeline of the model.

Date 0	**Date 1**	**Date 2**
1. Endowed investors invest or deposit at regional bank	1. Regional liquidity shocks are publicly observed	1. Investment projects mature
2. Banks choose portfolio and initiate interbank deposits	2. Banks settle date-1 interbank claims	2. Banks settle date-2 interbank claims
	3. Investors privately observe liquidity preference	3. Banks serve remaining withdrawals
	4. Investors observe regional solvency signals	
	5. Investors decide whether to withdraw	

Table 3.1: Timeline of the model.

3.3 Equilibrium

In this section we compute the solvency thresholds below which investors withdraw from their bank, causing an efficient bank run. We also obtain the expected utility of investors and the level of systemic risk for the cases of counterparty risk and common exposures. Furthermore, we describe the derivation of the equilibrium allocations and consider several limiting parameter values that yield a simple analytical solution to provide intuition for our model.

3.3.1 Counterparty risk

Consider the case with counterparty risk ($\eta > 0$) and without common exposure ($corr = 0$). Suppose first that no information spillovers occurs (investors receive no signal about the solvency of the other bank), which we will relax below.

Start with the debtor bank (H) since the solvency threshold there is unaffected by events at the creditor bank (L). With probability q_H investors at the debtor bank are informed and observe the realisation of the solvency shock θ_H. Because of essential bank runs, all investors withdraw if and only if the expected utility from the stochastic final-date consumption $\theta_H u(c_{2H}^G) + (1 - \theta_H)u(c_{2H}^B)$ falls short of the utility from their share of the liquidated bank portfolio $u(d_H)$. Therefore, the solvency threshold at the debtor bank $\overline{\theta}_{\mathrm{H}}$ is:

$$\overline{\theta}_{\mathrm{H}} \equiv \frac{u(d_H) - u(c_{2H}^B)}{u(c_{2H}^G) - u(c_{2H}^B)} \tag{3.3}$$

An essential bank run with full liquidation occurs if and only if $\theta_H < \overline{\theta}_{\mathrm{H}}$. Thus, the default probability of the debtor bank if informed is also $\overline{\theta}_{\mathrm{H}}$. With probability $1 - q_H$ investors are uninformed and base their efficient withdrawal behaviour on the prior distributions. We assume throughout that no bank runs occur without new information at the interim date. That is, the prior is sufficiently good as implied by a lower bound on the project output in the good state ($R \geq \underline{R}$). Thus, the overall default probability of the debtor bank is $a_{1\mathrm{H}} \equiv q_H \overline{\theta}_{\mathrm{H}}$. As shown in appendix 3.7, integrating the investors' respective payoffs over all possible signals yields the expected utility of investors at the debtor bank

EU_H:

$$\begin{aligned} EU_H &= (1-q_H)\Big\{\lambda_H u(d_1) + (1-\lambda_H)\frac{1}{2}(u(c_{2H}^G) + u(c_{2H}^B))\Big\} \\ &\quad + q_H\Big\{\overline{\theta}_{\mathrm{H}} u(d_H) + (1-\overline{\theta}_{\mathrm{H}})\left(\lambda_H u(d_1) + (1-\lambda_H)\frac{1}{2}\left[u(c_{2H}^G) + u(d_H)\right]\right)\Big\} \end{aligned} \tag{3.4}$$

which completes the description of the debtor region.

The creditor bank is affected by a default of the debtor bank, both in terms of the repayment at the final date and the liquidation value of the interbank claim at the interim date. If investors are informed, the solvency threshold $\overline{\theta}_{\mathrm{1L}}$ is:

$$\overline{\theta}_{\mathrm{1L}} \equiv \frac{a_{\mathrm{1H}}[u(d_L^D) - u(c_{2L}^{BD})] + (1-a_{\mathrm{1H}})[u(d_L^N) - u(c_{2L}^{BN})]}{a_{\mathrm{1H}}[u(c_{2L}^{GD}) - u(c_{2L}^{BD})] + (1-a_{\mathrm{1H}})[u(c_{2L}^{GN}) - u(c_{2L}^{BN})]} = \overline{\theta}_{\mathrm{1L}}(\overline{\theta}_{\mathrm{H}}) \tag{3.5}$$

and the overall default probability of the creditor bank is $a_{\mathrm{1L}} \equiv q_L\overline{\theta}_{\mathrm{1L}}$.

Counterparty risk, the dependence of the creditor bank on the debtor bank, is reflected by the solvency threshold $\overline{\theta}_{\mathrm{1L}}(\overline{\theta}_{\mathrm{H}})$. A higher solvency threshold at the debtor bank makes a default on the interbank claim more likely, thus raising the probability of default at the creditor bank ($\frac{\partial\overline{\theta}_{\mathrm{1L}}}{\partial\overline{\theta}_{\mathrm{H}}} > 0$).[18]

As shown in appendix 3.7, the expected utility of investors in the creditor bank is:

$$\begin{aligned} EU_{1L} &= (1-q_L)\Big\{\lambda_L u(d_1) + (1-\lambda_L)\frac{1}{2}\Big((1-a_H)\left(u(c_{2L}^{GN}) + u(c_{2L}^{BN})\right) \\ &\quad + a_H\left(u(c_{2L}^{GD}) + u(c_{2L}^{GN})\right)\Big)\Big\} \\ &\quad + q_L\Big\{\overline{\theta}_{\mathrm{1L}}\left((1-a_H)u(d_L^N) + a_H u(d_L^D)\right) + \lambda_L(1-\overline{\theta}_{\mathrm{1L}})u(d_1) \\ &\quad + (1-\lambda_L)\frac{1}{2}\Big((1-\overline{\theta}_{\mathrm{1L}}^2)\left((1-a_H)u(c_{2L}^{GN}) + a_H u(c_{2L}^{GD})\right) \\ &\quad + (1-\overline{\theta}_{\mathrm{1L}})^2\left((1-a_H)u(c_{2L}^{BN}) + a_H u(c_{2L}^{BD})\right)\Big)\Big\} \end{aligned} \tag{3.6}$$

There is one main difference to the expected utility of investors at the debtor bank. Since no information spillover takes place, the expectation over whether the debtor bank

[18] A failure of the debtor bank constitutes a negative externality on investors of the creditor bank. Early investors at the creditor bank receive their share of the liquidation value d_L instead of the higher promised payment d_1. Late investors are paid out fewer resources. Consequently, the solvency threshold at the creditor bank strictly increases in the solvency threshold of the debtor bank.

defaults, which occurs with probability a_H, is taken.

Finally, the overall expected utility EU_{CR} and the level of systemic risk A_{CR} in the case of counterparty risk (CR) are:

$$A_{\mathrm{CR}} \equiv a_{1\mathrm{L}}a_{1\mathrm{H}} = q_H q_L \overline{\theta}_{\mathrm{H}} \overline{\theta}_{1\mathrm{L}} \tag{3.7}$$

$$EU_{CR} \equiv \frac{1}{2}(EU_H + EU_{1L}) \tag{3.8}$$

Since the regional solvency shocks are uncorrelated, the overall expected utility is the average of the expected utility of an investor at the debtor and creditor bank, respectively. This will be generalized once we allow for correlated solvency shocks.

We now allow for information spillover, that is news about the solvency of the bank in the other region. The efficient withdrawal behaviour of investors at the debtor bank is unchanged and so is their expected utility EU_H. By contrast, with probability q_H investors at the creditor bank are informed about the debtor bank's solvency and infer whether repayment at the final date occurs and whether the liquidation of the interbank claim yields revenue at the interim date.[19] Consequently, there are two solvency thresholds at the creditor bank: one if the debtor bank defaults ($\overline{\theta}^D_{2\mathrm{L}}$) and one if the debtor bank repays ($\overline{\theta}^N_{2\mathrm{L}}$):

$$\overline{\theta}^N_{2\mathrm{L}} \equiv \frac{u(d^N_L) - u(c^{BN}_{2L})}{u(c^{GN}_{2L}) - u(c^{BN}_{2L})} \tag{3.9}$$

$$\overline{\theta}^D_{2\mathrm{L}} \equiv \frac{q_H[u(d^D_L) - u(c^{BD}_{2L})] + (1-q_H)[u(d^N_L) - u(c^{BN}_{2L})]}{q_H[u(c^{GD}_{2L}) - u(c^{BD}_{2L})] + (1-q_H)[u(c^{GN}_{2L}) - u(c^{BN}_{2L})]} \tag{3.10}$$

If the information spillover is unanticipated, the spillover of information at the interim date has no effect on the optimal portfolio choice at the initial date. Then, the solvency thresholds can be ranked:

$$\overline{\theta}^N_{2\mathrm{L}} < \overline{\theta}_{1\mathrm{L}} < \overline{\theta}^D_{2\mathrm{L}} \tag{3.11}$$

which captures both information contagion if the debtor bank defaults and stabilization if the debtor bank repays the creditor bank.

[19]The creditor bank is not repaid if and only if investors at the debtor bank are informed and the solvency of the debtor bank is low, which is a consequence of the seniority of interbank claims.

The expected utility of an investor at the creditor bank EU_{2L} changes to:

$$\begin{aligned} EU_{2L} &= (1-q_L)\Bigg\{\lambda_L u(d_1) + (1-\lambda_L)\frac{1}{2}\Big[(1-a_H)\left(u(c_{2L}^{GN}) + u(c_{2L}^{BN})\right) \\ &\quad +a_H\left(u(c_{2L}^{GD}) + u(c_{2L}^{BD})\right)\Big]\Bigg\} \\ &\quad +q_L\Bigg\{\left(\overline{\theta}_{\mathrm{2L}}^N(1-a_H)u(d_L^N) + \overline{\theta}_{\mathrm{2L}}^D a_H u(d_L^D)\right) \\ &\quad +\lambda_L\left(a_H(1-\overline{\theta}_{\mathrm{2L}}^D) + (1-a_H)(1-\overline{\theta}_{\mathrm{2L}}^N)\right)u(d_1) \\ &\quad +(1-\lambda_L)\frac{1}{2}\Big((1-a_H)[(1-(\overline{\theta}_{\mathrm{2L}}^N)^2)u(c_{2L}^{GN}) + (1-\overline{\theta}_{\mathrm{2L}}^N)^2 u(c_{2L}^{BN})] \\ &\quad +a_H[(1-(\overline{\theta}_{\mathrm{2L}}^D)^2)u(c_{2L}^{GD}) + (1-\overline{\theta}_{\mathrm{2L}}^D)^2 u(c_{2L}^{BD})]\Big)\Bigg\} \end{aligned} \tag{3.12}$$

where the solvency thresholds in the creditor region now depend on whether the debtor bank defaults. That is, the uninformed solvency threshold $\overline{\theta}_{\mathrm{1L}}$ is replaced by the conditional thresholds $(\overline{\theta}_{\mathrm{2L}}^N, \overline{\theta}_{\mathrm{2L}}^D)$.

The overall expected utility of an investor EU_{CR+IC} and the level of systemic risk $A_{\mathrm{CR+IC}}$ in case of counterparty risk and information contagion are:

$$EU_{CR+IC} \equiv \frac{1}{2}(EU_H + EU_{2L}) \tag{3.13}$$

$$A_{\mathrm{CR+IC}} = q_H q_L \overline{\theta}_{\mathrm{H}} \overline{\theta}_{\mathrm{2L}}^D \tag{3.14}$$

which yields the following result:

Lemma 7. *If information spillovers are unanticipated, then information contagion due to counterparty risk unambiguously increases systemic risk:*

$$A_{CR+IC} > A_{CR} \tag{3.15}$$

3.3.2 Common exposure

Consider the case with common exposures ($corr = 1$) and no counterparty risk ($\eta = 0$). Thus, the payoffs are symmetric across regions but investors are potentially asymmetrically informed about the common solvency shock. Final-date consumption simplifies to

$c_2^G \equiv \frac{y-\lambda d_1+(1-y)R}{1-\lambda}$ in the good state and $c_2^B \equiv \frac{y-\lambda d_1}{1-\lambda}$ in the bad state and the liquidation payoff to $d_\beta \equiv y + (1-y)\beta$. Again we start without information spillover. The solvency threshold in either region becomes:

$$\overline{\theta} = \frac{u(d_\beta) - u(c_2^B)}{u(c_2^G) - u(c_2^B)} \tag{3.16}$$

where an efficient bank run occurs if and only if the solvency level is below its threshold ($\theta < \overline{\theta}$). As derived in appendix 3.7, the expected utility in either region EU_{CE} and the level of systemic risk A_{CE} in case of pure common exposure are:

$$\begin{aligned} EU_{CE} &= \frac{q_A+q_B}{2}\left[\overline{\theta}u(d_\beta) + (1-\overline{\theta})\left(\lambda u(d_1) + (1-\lambda)\frac{1}{2}[u(c_2^G) + u(d_\beta)]\right)\right] \quad (3.17)\\ &+\frac{1-q_A+1-q_B}{2}\left[\lambda u(d_1) + (1-\lambda)\frac{1}{2}(u(c_2^G) + u(c_2^B))\right] \\ A_{\mathrm{CE}} &= q_A q_B \overline{\theta} \quad (3.18) \end{aligned}$$

We now allow for information spillover. While payoffs and the solvency threshold are unchanged, the probability of being informed changes to $q_A + (1-q_A)q_B > q_A$. Naturally, information spillovers increases the probability of being informed. Therefore, the expected utility in case of common exposure and information contagion (CE+IC) places more weight on the two terms in which liquidation may take place (those involving $\overline{\theta}$) and a smaller weight on the term without information and liquidation:

$$\begin{aligned} EU_{CE+IC} &\equiv (q_A+q_B-q_Aq_B)\left[\overline{\theta}u(d_\beta) + (1-\overline{\theta})\left(\lambda u(d_1) + (1-\lambda)\frac{1}{2}[u(c_2^G) + u(d_\beta)]\right)\right] \\ &+(1-q_A)(1-q_B)\left[\lambda u(d_1) + (1-\lambda)\frac{1}{2}(u(c_2^G) + u(c_2^B))\right] \quad (3.19)\\ A_{\mathrm{CE+IC}} &= (q_A + (1-q_A)q_B)\overline{\theta} \quad (3.20) \end{aligned}$$

which leads to the following result that mirrors lemma 7:

Lemma 8. *If information spillovers are unanticipated, then information contagion due to common exposure unambiguously increases systemic risk:*

$$A_{CE+IC} > A_{CE} \tag{3.21}$$

3.3.3 Optimal portfolio choice

We solve for the optimal portfolio choice and the optimal demand deposit payment of banks at the interim date. A bank faces the following constraints on its choice variables. The mutual insurance between banks trades off liquidity insurance with counterparty risk. It is never optimal to insure more than the maximum amount to compensate for the regional liquidity demand shock ($b^* \leq \eta d_1^*$), where stars denote equilibrium allocations. Furthermore, it is never optimal to face certain costly liquidation, which places a lower bound on the amount of storage: $y^* + b^* \geq \lambda_H d_1^*$ and $y^* - b^* \geq \lambda_L d_1^*$. Combined with the optimal amount of interbank insurance, we obatin a lower bound on storage:

$$y^* \geq \underline{y} \equiv \lambda_H d_1^* - b^* \geq \lambda \tag{3.22}$$

The interim payment d_1 is bounded from above by the available resources and achieves risk sharing between early and late investors only if it is positive:

$$0 < d_1^* \leq \min\{R, \frac{y^* + (1-y^*)\beta + b^*}{\lambda_H}, \frac{y^* + (1-y^*)\beta - b^*}{\lambda_L}\} \tag{3.23}$$

Our model does not admit a tractable analytical solution for two reasons. First, corner solutions of the form of no interbank insurance ($b^* = 0$) or no investment ($y^* = 1$) are optimal for some parameter values, invalidating interior solutions and calling for a global approach. Second, solvency thresholds are non-monotonic in several choice variables. For example, more liquidity is valued when the investment project fails, while less liquidity is valued when the investment project succeeds. Also , the change in the solvency thresholds with respect to interbank liquidity insurance is in general ambiguous.[20] In sum, both corner solutions and the non-monotonicity of the solvency thresholds in the choice variables confound an analytical solution. However, we determine analytical solutions for several limiting parameter values in section 3.3.4 to provide intuition for the mechanics of our model.

We solve the optimization problem numerically. We find the global optimum of

[20]By contrast, more insurance against the idiosyncratic liquidity risk of a investor (higher d_1) raises payments at the interim date at the expense of payments at the final date, thus unambiguously increasing the solvency threshold. Goldstein and Pauzner (2005) study this trade-off between insurance on a higher interim payment and higher idiosyncratic financial fragility in a global games setup that allows for non-essential bank runs.

the expected utility by discretizing the choice variables (d_1, y, b) on a three-dimensional grid, where the expected utility is evaluated at each grid point. The grid point where the expected utility takes its global maximum value yields the best response for a given portfolio choice of the other bank. The intersection of the (symmetric) best response functions yields the (symmetric) equilibrium allocations. Even though we will incur a numerical error from discretizing, this error will be small for a sufficiently fine grid. We verify the validity of our numerical solution method in section 3.3.4. We compare the optimal choice variables obtained numerically with the optimal choices in the cases of limiting parameter values, which admit simple analytical solutions, and obtain negligible discrepancies only.

We use the following baseline calibration. The period utility function is CRRA, where $\rho > 0$ parameterizes the coefficient of relative risk aversion. Baseline parameter values are $\beta = 0.7$, $R = 5.0$, $\phi = 1.0$, $\lambda = 0.5$, $\eta = 0.25$, $\rho = 1.0$, and $q_H = q_L = 0.7$. Alternative specifications are considered in appendix 3.8 and in section 3.4.3 that discusses the variation of each parameter within its feasible bounds. Our results hold across these various specifications.

3.3.4 Limiting parameter cases

Our model admits an analytical solution for several limiting parameter values that are discussed in this section. These cases help us build intuition for the model and serve as a benchmark for the accuracy of our numerical solution.

First, let the project output in the good state fall short of unity ($R \leq 1$). Then, the investment project is dominated by storage ($y^* = 1$). We verify this result across all benchmark calibrations listed in Appendix 3.8.1 and obtain the numerical solution of $y^*_{num} = 0.98$. Second, let investors be risk-neutral ($\rho = 0$). Then, the project dominates storage as the former has a higher expected return and investors, who do not mind the uncertainty about the idiosyncratic liquidity shock, prefer to invest fully in the project ($d_1^* = 0 = y^*$). This result is confirmed numerically ($d^*_{1,num} = 0 = y^*_{num}$). Likewise, if investors are very risk averse ($\rho \to \infty$), they are not willing to bear any of the investment risk associated with the project or any liquidity risk. Consequently, no investment takes place ($y^* = 1$) and there is full insurance ($d_1^* = 1$). In a numerically feasible and economically useful implementation we set $\rho = 200$ and obtain the affirmative results $y^*_{num} = 0.98$ and $d^*_{1,num} = 0.98$.

Third, no risk-averse investor ($\rho > 0$) seeks liquidity insurance in the absence of regional liquidity shocks ($\eta = 0$) for any value of repayment ($\phi \geq 0$). From an ex-ante perspective, liquidity insurance in this case is a mean-preserving spread to both interim-date and final-date payoffs and is rejected by any risk averse investor. We confirm this intuition numerically ($b^*_{num} = 0$). We also consider the related situation of a positive liquidity shock ($\eta > 0$) but no repayment ($\phi = 0$). A risk averse investor would then be partially insured against this risk $b^* > 0$, which is pure ex-ante liquidity insurance. Note that we require $\phi > 0$ in the baseline calibration and all other calibrations to maintain a counterparty risk mechanism. Intuitively, the amount of liquidity insurance decreases in the degree of risk aversion. As investors become more risk averse, they hold more liquidity as part of the optimal portfolio composition of late investors. The available liquidity serves as self-insurance against regional liquidity shocks at the interim date and is a substitute for interbank insurance. For example, a CRRA coefficient of risk aversion of $\rho = 1.0$ in the baseline calibration yields $b^*_{num} = 0.15$, while the same calibration with $\rho = 2.0$ yields $b^*_{num} = 0.1$.

Fourth, if there are no early investors ($\lambda = 0$), there is no need for insurance against idiosyncratic liquidity shocks. The amount of liquidity held fully reflects the optimal portfolio allocation of late investors ($0 < y^* < 1$) and increases with the level of risk aversion (ρ). These predictions are confirmed numerically in the specification of $\lambda = 0.01$, where the amount of liquidity ranges from $y^*_{num} = 0.42$ in a baseline calibration with $\rho = 1.0$ to $y^*_{num} = 0.74$ in the baseline calibration with $\rho = 2.0$. Likewise, if there are only early investors ($\lambda = 1$) it is optimal not to invest into an asset that only matures at the final date and is costly to liquidate ($y^* = 1$). There is no role for liquidity insurance in this specification ($b^* = 0$) as there cannot be any liquidity shocks. Since all resources are used to serve early investors, the optimal interim payment must also be one ($d^*_1 = 1$). This intuition is confirmed numerically ($d^*_{1,num} = 0.99$).

Finally, the prior distribution is assumed not to induce liquidation in case of being uninformed. Hence, no liquidation takes place ($\overline{\theta}_1 = \overline{\theta}^N_{2\mathrm{L}} = \ldots = 0$) if investors are never informed in either region ($q_A = q_B = 0$), which is again confirmed numerically.

3.4 Results

This section summarizes our two main findings. In subsection 3.4.1 we present a *resilience effect* that arises when information contagion occurs due to counterparty risk. In subsection 3.4.2 we show the existence of an *instability effect* that emerges when information contagion occurs due to common exposures. Subsection 3.4.3 provides a global parameter analysis, verifying the robustness of our two main results across feasible parameter values.

3.4.1 Resilience effect

How does information contagion affect systemic risk stemming from counterparty risk? We start by considering unanticipated information spillovers similar to the aggregate liquidity shock in Allen and Gale (2000). In this case the ex-ante optimal portfolio choice of banks is unaffected and systemic risk strictly increases (lemma 7). This result arises directly from the fact that a failure of the creditor bank becomes more likely after adverse news about the solvency of the debtor bank. Therefore, information contagion strengthens the effect of counterparty risk, which leads to a lower level of expected utility and higher systemic risk. This immediate result is also obtained numerically by comparing entries in the tables in appendix 3.8.2, notably entry (1,1) for the case of pure counterparty risk with entry (1,2) for the case of counterparty risk and information contagion, where both are evaluated at the optimal portfolio choice of the pure counterparty risk case.

The focus of our analysis is on anticipated information contagion. Taking information contagion at the interim date into account, banks alter their portfolio choice at the initial date. Specifically, a bank makes a more prudent portfolio choice to insure risk-averse investors against potential information contagion at the interim date. First, banks reduce the exposure to counterparty risk by engaging in less liquidity co-insurance (lower b). To cover the liquidity demand from early investors, a bank increases the amount of storage (larger y), which is akin to liquidity self-insurance. This reduces the investment in the risky project and funds a larger amount of insurance against idiosyncratic liquidity risk (larger d_1). The more prudent portfolio choices reduces the range of solvency shocks ($[\overline{\theta}_{2\mathrm{L}}^{N}, \overline{\theta}_{2\mathrm{L}}^{D}]$) for which counterparty risk after information contagion occurs. These results are obtained numerically by comparing the case of pure counterparty risk (entry (1,1)) with the case of counterparty risk and information contagion (entry (2,2)) in the tables in appendix 3.8.2.

The crucial insight is that the direct positive effect information contagion on systemic risk (lemma 7) is more than fully compensated by an indirect negative effect via the change in the ex-ante portfolio choice. Therefore, the overall effect is a reduction in systemic risk if information contagion is anticipated. We label this the resilience effect and show in section 3.4.3 that it holds across all feasible parameter values.

Result 1. *Consider the setup with counterparty risk ($\eta > 0$). Anticipating ex-post information contagion induces a more prudent portfolio choice ex-ante. More specifically, liquidity co-insurance, which exposes banks to counterparty risk, is substituted by direct holdings of liquidity (self-insurance) that reduces the investment in the risky project. The overall effect is a reduction in both expected utility and systemic risk.*

3.4.2 Instability effect

We now analyze how information contagion affects systemic risk in a setup with common exposures ($corr = 1$). If information spillover is unanticipated, the optimal portfolio choice is unaffected, implying more systemic risk (lemma 8). Since efficient withdrawals by investors become more likely after adverse news about the solvency of another bank, unanticipated information spillover always leads to greater systemic risk.

If information spillovers is anticipated, the bank adjusts its ex-ante optimal portfolio choice. Specifically, the optimal interim-date payment is unchanged, while the optimal liquidity level is slightly lower (within numerical accuracy) across all baseline cases and feasible parameter choices. However, the changes to the portfolio is small, implying a small indirect effect on systemic risk only. Therefore, the level of systemic risk increases overall once information contagion is present. These results are obtained by comparing the case of pure common exposure (entry (3,3)) with the case of common exposure and information contagion (entry (4,4)) in the tables in appendix 3.8.2. This effect is again numerically robust, as demonstrated in section 3.4.3.

Result 2. *Consider the setup with common exposures ($corr = 1$). Then, anticipating information contagion has a small effect on the portfolio choice at the initial date. As such, systemic risk and the expected utility increases.*

Additional information allows the late investors to decide on early withdrawals in more states of the world and has two consequences. First, liquidation is optimal for late investors after a bad solvency shock. Second, liquidation is detrimental to early investors who only receive their share of the liquidation value and not the (strictly larger) promised interim payment. Therefore, late investors impose a negative externality on early investors. Since the level of liquidity in case of common exposures is high to self-insure against investment risk, the second effect is quantitatively small such that additional liquidation increases overall expected utility.

3.4.3 Robustness checks

This section shows that the resilience effect and the instability effect are robust to exogenous parameter variations. In particular, we discuss a global variation of parameters by considering the entire range of feasible parameters and analyse the effect on systemic risk and expected utility. Details and further analyses, including the optimal portfolio choice and withdrawal thresholds, are contained in figures 3.3 - 3.9 in appendix 3.9. Consider the resilience effect (result 1) first. Figure 3.1 displays the expected utility (dotted line) and systemic risk (dashed line) in the case of counterparty risk and information contagion as a fraction of their respective levels in case of pure counterparty risk. Hence, the resilience effect is present if relative systemic risk is *below* unity. We consider parameter changes of the key variables of the model: the liquidation value (β), the final-date return to the investment project in the good state (R), the proportion of early investors (λ), and the level of transparency (q). In all cases, the resilience effect prevails.

Turning to the instability effect (result 2), figure 3.2 displays the expected utility (dotted line) and systemic risk (dashed line) in the case of common exposure and information contagion as a fraction of their respective levels in case of pure common exposure. Hence, the instability effect is present if the relative systemic risk is *above* unity. We consider the same parameter changes again. In all cases, the instability effect prevails.

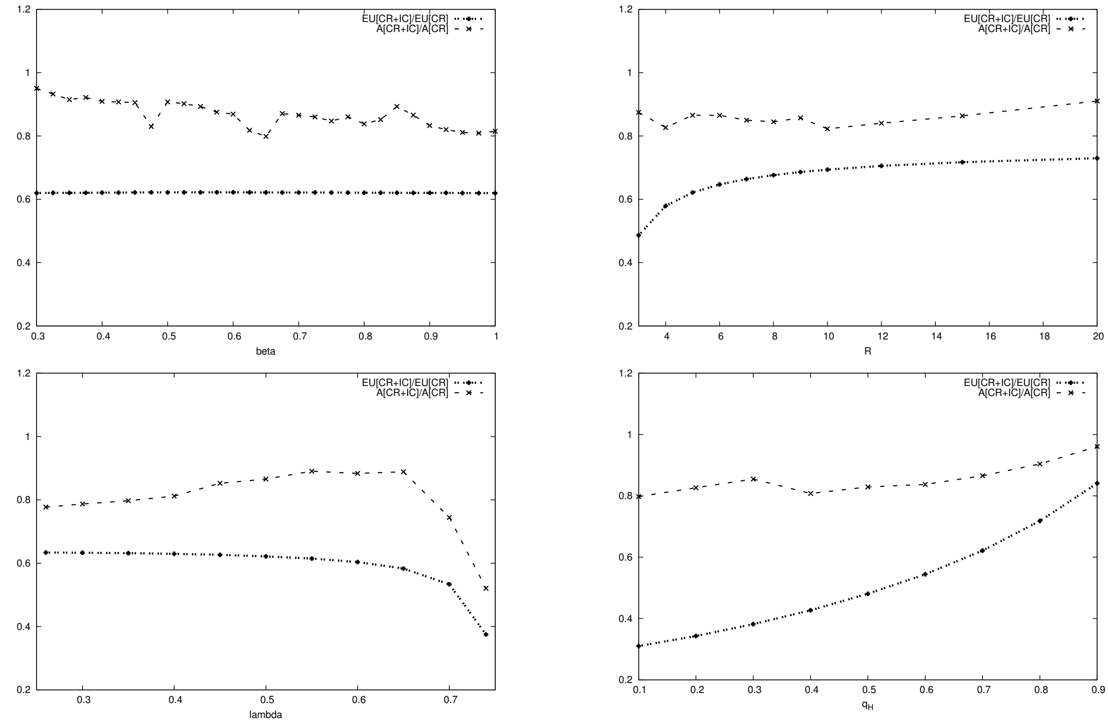

Figure 3.1: Robustness checks for the resilience effect (remma 1) consider a variation of β (top left), R (top right), λ (bottom left), and $q_H = q_L$ (bottom right). The figures display expected utility (dotted line) and the level of systemic risk (dashed line) in the case of counterparty risk and information contagion as a fraction of their respective levels in case of pure counterparty risk.

3.5 An application to microfinance

While our model focuses on the systemic risk in the financial system of advanced economies, it is also applicable to the microfinance industry prevalent in many emerging countries. Our model provides a novel theoretical explanation for several findings in the empirical microfinance literature. In particular, it predicts that (i) group loans have higher repayment rates than individual loans and (ii) group borrowers hold more liquid assets.

According to the Microcredit Summit Campaign (2012), microfinance institutions (MFIs) served over 205 million customers at the end of 2010, impacting the lives of an estimated 600 million household members. The growth of the microfinance industry is often attributed to group liability that is designed to overcome problems arising from asymmetric information (see e.g. Morduch (1999), or Armendáriz and Morduch (2010)) and beneficially transfers risks from the microlender to a group of borrowers (see e.g. Stiglitz (1990) and Varian (1990)). Group liability refers to an arrangement in which a lender grants a loan to a *group* of borrowers that monitor each other and jointly guarantee

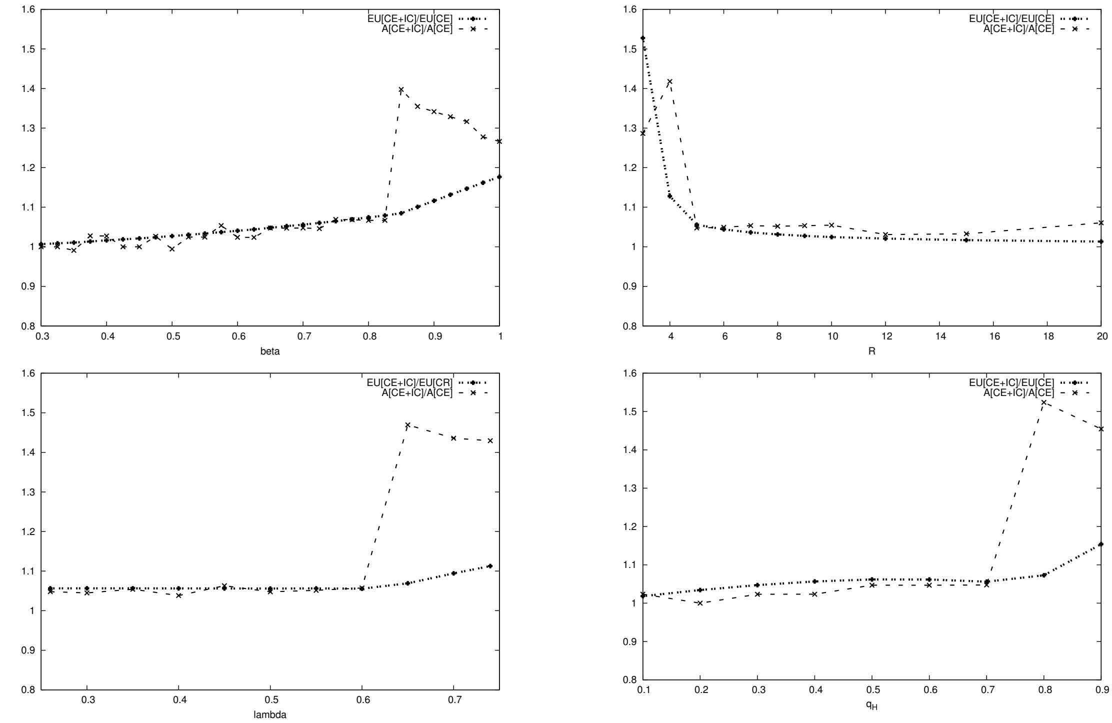

Figure 3.2: Robustness checks for the instability effect (remma 2) for a variation of β (top left), R (top right), λ (bottom left), and $q_H = q_L$ (bottom right). The figures display expected utility (dotted line) and the level of systemic risk (dashed line) in the case of common exposures and information contagion as a fraction of their respective levels in case of pure common exposures.

loan repayment. Borrowers are typically entrepreneurs from rural areas in developing countries that cannot pledge collateral.

The essential ingredients of microfinance are captured by our model. Due to joint liability, group lending is characterised by institutionalized counterparty risk. In particular, each group member guarantees the repayment of the entire loan even if another group member is unable (or unwilling) to repay, exposing an individual group member to a large amount of counterparty risk. Furthermore, group members often know each other well and are in close contact. This implies that news about one group member easily spreads to other group members, constituting information spillover.[21] Finally, the close proximity of group members gives rise to common exposures such as natural disasters (e.g. a flood or an earthquake).

The application to microfinance can be explicitly translated into our model setup.

[21]Since it is more costly for banks to acquire this kind of information about the borrowers, monitoring is delegated to the group and rewarded with lowered interest rates on group loans. See Stiglitz (1990) and Varian (1990) for a rationalisation of peer monitoring.

Consider two entrepreneurs $k = A, B$ that jointly wish to take out a group loan from a microfinance institution. Each entrepreneur has access to a safe storage technology (cash or durable goods) and is offered a risky investment opportunity R_k. This investment opportunity could be the start of a small local business (e.g. buying an ox to plow a field, or dwelling a well to sell the water) that has a probability to fail. In this interpretation, a region corresponds most naturally to a sector of the economy. The project pays R with a regional probability θ_k and zero with probability $(1-\theta_k)$. An alternative interpretation is that the investment project will always pay a safe return R but, with some probability $(1-\theta_k)$, the entrepreneur has to take this return to cover unexpected expenses such as an illness of a family member. Liquidation of investment projects is costly due to an alternative use argument similar to the banking case. In many cases, the MFI is unable to seize the investment project at all due to its remoteness from the borrower or due to social pressure (seizing assets from somebody who is already poor). The timeline of our model applied to microfinance is given in table 3.2.

Date 0	**Date 1**	**Date 2**
1. Microfinance institution (MFI) decides on group loan	1. Group loan institutionalizes counterparty risk	1. Investment projects mature
2. Entrepeneurs choose their portfolio	2. Entrepeneurs observe regional solvency signals	2. Group of entrepreneurs repays MFI
	3. Entrepeneurs decide whether to default	

Table 3.2: Timeline of the model applied to microfinance

The information structure is equivalent to the banking case. At the interim date, before the success or failure of the local business projects is determined, entrepreneurs receive a signal about the regional return of the other entrepreneur in the group.[22] Such a signal can be informative about the business prospects of the group partners or, in the alternative interpretation with safe investment projects, information about the health of the family of a group partner. In either case, this signal contains valuable information since both entrepreneurs are linked via joint liability. In the banking application, we focus

[22] We take the probability of receiving an informative signal q as being fixed exogenously. An extension could consider the extent of group member monitoring, modelled by a change in this probability.

on the impact of ex-post information contagion on ex-ante systemic risk when banks are subject to counterparty risk. Translated into the microfinance setting, we focus on the impact of ex-post information contagion on the ex-ante default probability of a group loan.

Strategic default by group members in the microfinance application is the equivalent of withdrawals by late investors in our banking model. Late investors make an efficient withdrawal decision. Likewise, entrepeneurs decide strategically whether to pay loan installments (interest and principal) to the MFI. The benefits of default (or diversion of funds) for an entrepreneur is not to repay his share of the group loan. Another benefit is not having to pay more upon default by other group members. In the alternative interpretation with safe investment projects, the benefits of default could be saving the life of a family member. The cost of default is exclusion from credit via group loans, foregoing future profits from investment projects. As default increases the burden on other group members, another cost of default is the possibility of facing hostile group loan cosigners.[23]

Similar to banks in our banking application, entrepreneurs decide about the portfolio shares of their funds ex-ante. When entrepeneurs decide between investment in their project and storage, they consider the possibility of a solvency shock, their business risk, and its effect on potential future exclusion from credit. The profits from future investment opportunities induce a precautionary motive for entrepeneurs. Hence, entrepreneurs try to avoid default by holding more of the safe asset (either cash or durable goods that have a high liquidation value). In our banking application, banks offer deposit contracts that may be accepted by investors. Likewise, in the microfinance application, entrepreneurs offer interest payments to a microfinance institution.

In the banking application, withdrawing late investors at the debtor bank exert an externality on late investors at the creditor bank. This correponds to the externality that one entrepreneur exerts on other members of the group loan when defaulting on its obligation. When making their ex-ante optimal portfolio choice, banks take this externality into account by holding more liquidity and this leads to reduced systemic

[23]There are news reports about large numbers of suicides that were caused by peer pressure after defaulting on a micro loan (see e.g. BBC News, "India's micro-finance suicide epidemic", 16 December 2010).

risk. Translating this *resilience effect (Result 1)* into the microfinance application, our model predicts that (i) group loans have a higher repayment rate than individual loans and (ii) group borrowers hold more liquid assets.

The empirical microfinance literature supports these predictions. For example, Giné et al. (2009) constructs a series of *"microfinance games"* conducted in an urban market in Peru. They show that loan repayment rates are higher in joint-liability games (0.88) than in individual-liability games (0.68). Wydick (1999) analyzes group lending in Guatemala and shows that group repayment rates are determined by the ability to monitor one another in the presence of asymmetric information. In particular, group loan repayment rates are higher when group members live in close geographic proximity or have knowledge about weekly sales of their peers. The resilience effect also implies that entrepreneurs will hold more liquidity (either in the form of cash or durable goods). This has been analyzed empirically by Banerjee et al. (2010) who show in a randomized experiment in India that households with an existing business at the time of the program invest more in durable goods.

The usefulness of our results for microfinance is highlighted by the empirical confirmation of our predictions. This relates to both the ex-ante portfolio choices of entrepeneurs and the repayment rates for group loans.

3.6 Conclusion

The aftermath of the Lehmann bankruptcy in September 2008 demonstrated that information contagion can be a major source of systemic risk, defined as the probability of joint bank default. One bank's investors find information about another bank's solvency valuable for two reasons. First, both banks might have invested into the same asset class like risky sovereign debt or mortgage backed securities. Learning about another bank's profitability then helps the investor assess the profitability of its bank. Second, one bank might have lent to the other, for instance as part of a risk-sharing agreement. Learning about the debtor bank's profitability then helps investors assess the counterparty risk of the creditor bank.

This paper presents a model of systemic risk with information contagion. Information about the health of one bank is valuable for the investors of other banks because of common exposures and counterparty risk. In each case, bad news about one bank adversely spills over to other banks and causes information contagion. We examine the effects of ex-post information contagion on the bank's ex-ante optimal portfolio choice and the implied level of systemic risk.

We demonstrate that information contagion can *reduce* systemic risk. When banks are subject to counterparty risk, investors of one bank may receive a negative signal about the health of another bank. Given the exposure of the creditor bank to the debtor bank, adverse information about the debtor bank can cause a run on the creditor bank. Such information contagion ex-post induces the bank to hold a more prudent portfolio ex-ante. Banks reduce their exposure to counterparty risk and rely more the self-insurance of liquidity instead of co-insurance. Overall, the level of systemic risk is reduced once information contagion is present.

Our model is also applicable to microfinance prevalent in many emerging countries. Group loans with joint liability agreements induce counterparty risk among the group members. Since group loan borrowers typically have a common bond (e.g. living in the same village), peer monitoring helps to overcome problems of asymmetric information. The common bond implies that group members receive information about their peers, constituting information spillover. We show that counterparty risk and information contagion after adverse news lead to reduced default rates of group loans and increased holdings of liquid assets by group borrowers because of the effect on the ex-ante portfolio choices. These predictions are verified in the empirical literature on microfinance, highlighting the applicability of our model to the microfinance setting.

We also show that the effects of information contagion on systemic risk depend on the source of the revealed information. In case of common exposures, ex-post information contagion increases systemic risk - similar to Acharya and Yorulmazer (2008a). This leads to the natural question about the overall effect of information contagion in a model that features both common exposures and counterparty risk. A *unified model of contagion* would be suited to identify the parameter regions characterized by higher (lower) levels of systemic risk and thus a less (more) stable financial system. Such a unified model of contagion would also contribute to our understanding of microfinance. While allowing for information spillover, the close geographic proximity between group lenders implies that they are subject to common exposures. Analysing joint liability agreements in the presence of informational spillovers and common exposure is an interesting research question. However, such a unified model of contagion is beyond the scope of the present paper and left for future research.

3.7 Derivations

3.7.1 Counterparty risk

Start with the debtor region (H). If no signal is received, early investors, of mass λ_H, receive the promised payment d_1 and late investors, of mass $1-\lambda_H$, receive high and low payoff, c_{2H}^G and c_{2H}^B, with equal probability. If a signal is received and is below the threshold $\overline{\theta}_{\mathrm{H}}$ is received, investors receive a share of the liquidation proceeds and obtain d_H. If a signal above the threshold $\overline{\theta}_{\mathrm{H}}$ is received, late houseshold obtain a weighted average of the high and low payoff, where the weights depend on the threshold and early investors again receive the promised payment.[24] Expected utility in the high liquidity demand region is given as:

$$
\begin{aligned}
EU_H \quad = \quad & (1-q_H)\left\{\lambda_H u(d_1) + (1-\lambda_H)\int_0^1 \left[\theta u(c_{2H}^G) + (1-\theta)u(c_{2H}^B)\right] \mathrm{d}\theta\right\} \qquad (3.24)\\
& + q_H\left\{\int_0^{\overline{\theta}_{\mathrm{H}}} u(d_H)\mathrm{d}\theta + \int_{\overline{\theta}_{\mathrm{H}}}^1 \lambda_H u(d_1) + (1-\lambda_H)\left[\theta u(c_{2H}^G) + (1-\theta)u(c_{2H}^B)\right] \mathrm{d}\theta\right\}
\end{aligned}
$$

which yields the expression in the text.

We proceed in the same way for the creditor region (L). The behaviour in the debtor region determines whether or not the creditor bank is repaid at the final date. This affects both the expected utility from liquidation and the expected utility from continuation. As the interbank loan is repaid with probability $a_{1\mathrm{H}}$, the expected utility from liquidation is $a_{1\mathrm{H}}u(d_L^D) + (1-a_{1\mathrm{H}})u(d_L^N)$. In the informed case, which happens with probability q_L, θ_L is known. Taking expectations over all possible fundamentals in the debtor region, the expected utility from continuation is the sum of two terms: (i) with probability $a_{1\mathrm{H}}$ the debtor bank defaults and patient investors at the creditor bank receive $\theta_L u(c_{2L}^{GD}) + (1-\theta_L)u(c_{2L}^{BD})]$; (ii) with probability $(1-a_{1\mathrm{H}})$ the debtor bank survives and patient investors at the creditor bank receive $\theta_L u(c_{2L}^{GN}) + (1-\theta_L)u(c_{2L}^{BN})$ at the final date. The withdrawal threshold is given in equation (3.5) and yields the expected utility of

[24] Note that in case of no bank run, the weights are equal because of the symmetry of the investent probabilities θ and $1-\theta$ when integrated between zero and unity. This symmetry vanishes once the lower integration bound is above zero.

investors at the creditor bank to be:

$$\begin{aligned} EU_{1L} &= (1-q_L)\Bigg\{\lambda_L u(d_1) + (1-\lambda_L)\int_0^1 \left[\theta\left(a_H u(c_{2L}^{GD}) + (1-a_H)u(c_{2L}^{GN})\right)\right. \quad (3.25) \\ &\quad \left. + (1-\theta)\left(a_H u(c_{2L}^{BD}) + (1-a_H)u(c_{2L}^{BN})\right)\right] \mathrm{d}\theta\Bigg\} \\ &\quad + q_L\Bigg\{\int_0^{\overline{\theta}_{1L}} \left(a_H u(d_L^D) + (1-a_H)u(d_L^N)\right)\mathrm{d}\theta \\ &\quad + \int_{\overline{\theta}_{1L}}^1 \lambda_L u(d_1) + (1-\lambda_L)\left[\theta\left(a_H u(c_{2L}^{GD}) + (1-a_H)u(c_{2L}^{GN})\right)\right. \\ &\quad \left. + (1-\theta)\left(a_H u(c_{2L}^{BD}) + (1-a_H)u(c_{2L}^{BN})\right)\right]\mathrm{d}\theta\Bigg\} \end{aligned}$$

which yields the expression in the text.

3.7.2 Common exposures

Turning to expected utility, using the short-hand notation for the continuation payoff: $\Gamma \equiv \lambda u(d_1) + (1-\lambda)[\theta u_2^G + (1-\theta)u_2^B]$, we find:

$$EU_C E \equiv \frac{1-q_A+1-q_B}{2}\int_0^1 \Gamma\mathrm{d}\theta + \frac{q_A+q_B}{2}\int_0^{\overline{\theta}} u(d_\beta)\mathrm{d}\theta + \frac{q_A+q_B}{2}\int_{\overline{\theta}}^1 \Gamma\mathrm{d}\theta \quad (3.26)$$

$$\begin{aligned} &\equiv \frac{q_A+q_B}{2}\left[\overline{\theta}u(d_\beta) + (1-\overline{\theta})\left(\lambda u(d_1) + (1-\lambda)\frac{1}{2}[u(c_2^G) + u(d_\beta)]\right)\right] \\ &\quad + \frac{1-q_A+1-q_B}{2}[\lambda u(d_1) + (1-\lambda)\frac{1}{2}(u(c_2^G) + u(c_2^B))] \quad (3.27) \end{aligned}$$

3.8 Tables

Subsection 3.8.1 contains the extreme parameter value benchmarks discussed in subsection 3.3.4 of the main text for additional baseline cases to show the robustness of our numerical implementation. Subsection 3.8.2 contains the results of section 3.4 of the main text.

3.8.1 Extreme parameter value benchmarks

	Baseline 1	Baseline 2	Baseline 3	Baseline 4
$R = 1.0$	$y^* = 0.98$	$y^* = 0.98$	$y^* = 0.98$	$y^* = 0.98$
$\rho = 0.0$	$d_1^* = 0.0$	$d_1^* = 0.0$	$d_1^* = 0.0$	$d_1^* = 0.0$
	$y^* = 0.0$	$y^* = 0.0$	$y^* = 0.0$	$y^* = 0.0$
$\rho = 200.0$	$d_1^* = 0.98$	$d_1^* = 0.98$	$d_1^* = 0.98$	$d_1^* = 0.98$
	$y^* = 0.98$	$y^* = 0.98$	$y^* = 0.98$	$y^* = 0.98$
$\eta = 0.0$	$b^* = 0.0$	$b^* = 0.0$	$b^* = 0.0$	$b^* = 0.0$
$\phi = 0.0$	$b^* = 0.15$	$b^* = 0.15$	$b^* = 0.15$	$b^* = 0.1$
$\lambda = 0.01$	$d_1^* = 1.06$	$d_1^* = 1.0$	$d_1^* = 1.1$	$d_1^* = 1.16$
	$y^* = 0.42$	$y^* = 0.36$	$y^* = 0.48$	$y^* = 0.74$
$\lambda = 0.99$	$d_1^* = 0.98$	$d_1^* = 0.98$	$d_1^* = 0.98$	$d_1^* = 0.98$
	$y^* = 0.98$	$y^* = 0.98$	$y^* = 0.98$	$y^* = 0.98$
$q_H = 0.0$	$A_1, \ldots, A_6 = 0.0$	$A_1, \ldots, A_6 = 0.0$	$A_1, \ldots, A_6 = 0.0$	$A_1, \ldots, A_6 = 0.0$

Table 3.3: Extreme parameter values for four baseline cases. Baseline 1: $\beta = 0.7$, $R = 5.0$ $\phi = 1.0$, $\lambda = 0.5$, $\eta = 0.25$, $\rho = 1.0$, $q_H = 0.7$. Baseline 2: $\beta = 0.7$, $R = 5.0$ $\phi = 1.0$, $\lambda = 0.5$, $\eta = 0.25$, $\rho = 0.9$, $q_H = 0.7$. Baseline 3: $\beta = 0.7$, $R = 5.0$ $\phi = 1.0$, $\lambda = 0.5$, $\eta = 0.25$, $\rho = 1.1$, $q_H = 0.7$. Baseline 4: $\beta = 0.3$, $R = 5.0$ $\phi = 1.0$, $\lambda = 0.5$, $\eta = 0.25$, $\rho = 1.1$, $q_H = 0.7$.

3.8.2 Results

	cr $(EU,\, d_1^*,\, y^*,\, b^*)$ $(\overline{\theta}_{\mathrm{H}},\, \overline{\theta}_{\mathrm{1L}}\,,\, A_{\mathrm{cr}})$	cr + ic $(EU,\, d_1^*,\, y^*,\, b^*)$ $(\overline{\theta}_{\mathrm{H}},\, \overline{\theta}_{\mathrm{2L}}^N\,,\, \overline{\theta}_{\mathrm{2L}}^D,\, A_{\mathrm{cr+ic}})$	ce $(EU,\, d_1^*,\, y^*,\, b^*)$ $(\overline{\theta},\, A_{\mathrm{ce}})$	ce + ic $(EU,\, d_1^*,\, y^*,\, b^*)$ $(\overline{\theta},\, A_{\mathrm{ce+ic}})$
cr	(0.172,0.88,0.73,0.08) (0.423,0.23,0.048)	(0.096,0.88,0.73,0.08) (0.423,0.212,0.252,0.052)		
cr + ic		(0.107,0.94,0.8,0.02) (0.379,0.211,0.222,0.041)		
ce			(0.13,1.0,0.77,0.0) (0.328,0.161)	(0.137,1.0,0.77,0.0) (0.328,0.161)
ce + ic				(0.137,1.01,0.76,0.0) (0.344,0.168)

Table 3.4: Equilibrium allocation for different forms of financial fragility for calibration β=0.7, R=5.0, ϕ=1.0, λ=0.5, η=0.25, ρ=1.0, q_H=0.7. Expected utility (EU), portfolio choice variables (d_1, y, b), withdrawal thresholds $(\overline{\theta}_H, \overline{\theta}_{1,L}, \overline{\theta}_{2,L}^N, \overline{\theta}_{2,L}^D, \overline{\theta})$, and systemic financial fragility $(A_{\mathrm{cr}}, A_{\mathrm{cr+ic}}, A_{\mathrm{ce}}, A_{\mathrm{ce+ic}})$ in the different model variants (cr: counterparty risk, ic: information contagion, ce: common exposure).

	cr $(EU,\ d_1^*,\ y^*,\ b^*)$ $(\overline{\theta}_{\mathrm{H}},\ \overline{\theta}_{\mathrm{1L}}\ ,\ A_{\mathrm{cr}})$	cr + ic $(EU,\ d_1^*,\ y^*,\ b^*)$ $(\overline{\theta}_{\mathrm{H}},\ \overline{\theta}_{\mathrm{2L}}^N\ ,\ \overline{\theta}_{\mathrm{2L}}^D,\ A_{\mathrm{cr+ic}})$	ce $(EU,\ d_1^*,\ y^*,\ b^*)$ $(\overline{\theta},\ A_{\mathrm{ce}})$	ce + ic $(EU,\ d_1^*,\ y^*,\ b^*)$ $(\overline{\theta},\ A_{\mathrm{ce+ic}})$
cr	(0.188,0.86,0.7,0.13) (0.482,0.304,0.072)	(0.105,0.86,0.7,0.13) (0.482,0.278,0.329,0.078)		
cr + ic		(0.117,0.93,0.78,0.06) (0.43,0.26,0.283,0.06)		
ce			(0.142,1.0,0.75,0.0) (0.373,0.183)	(0.154,1.0,0.75,0.0) (0.373,0.183)
ce + ic				(0.158,1.32,0.73,0.0) (0.5,0.245)

Table 3.5: Equilibrium allocation for different forms of financial fragility for calibration β=0.9, R=5.0, ϕ=1.0, λ=0.5, η=0.25, ρ=1.0, q_H=0.7. Expected utility (EU), portfolio choice variables (d_1, y, b), withdrawal thresholds $(\overline{\theta}_H, \overline{\theta}_{1,L}, \overline{\theta}_{2,L}^N, \overline{\theta}_{2,L}^D, \overline{\theta})$, and systemic financial fragility $(A_{\mathrm{cr}}, A_{\mathrm{cr+ic}}, A_{\mathrm{ce}}, A_{\mathrm{ce+ic}})$ in the different model variants (cr: counterparty risk, ic: information contagion, ce: common exposure).

	cr (EU, d_1^*, y^*, b^*) ($\overline{\theta}_{\mathrm{H}}$, $\overline{\theta}_{\mathrm{1L}}$, A_{cr})	cr + ic (EU, d_1^*, y^*, b^*) ($\overline{\theta}_{\mathrm{H}}$, $\overline{\theta}^N_{\mathrm{2L}}$, $\overline{\theta}^D_{\mathrm{2L}}$, $A_{\mathrm{cr+ic}}$)	ce (EU, d_1^*, y^*, b^*) ($\overline{\theta}$, A_{ce})	ce + ic (EU, d_1^*, y^*, b^*) ($\overline{\theta}$, $A_{\mathrm{ce+ic}}$)
cr	(0.343,0.84,0.69,0.14) (0.372,0.172,0.031)	(0.221,0.84,0.69,0.14) (0.372,0.15,0.206,0.038)		
cr + ic		(0.238,0.91,0.77,0.07) (0.318,0.139,0.166,0.026)		
ce			(0.274,1.0,0.75,0.0) (0.257,0.126)	(0.28,1.0,0.75,0.0) (0.257,0.126)
ce + ic				(0.28,1.01,0.74,0.0) (0.271,0.133)

Table 3.6: Equilibrium allocation for different forms of financial fragility for calibration β=0.7, R=10.0, ϕ=1.0, λ=0.5, η=0.25, ρ=1.0, q_H=0.7. Expected utility (EU), portfolio choice variables (d_1, y, b), withdrawal thresholds $(\overline{\theta}_H, \overline{\theta}_{1,L}, \overline{\theta}^N_{2,L}, \overline{\theta}^D_{2,L}, \overline{\theta})$, and systemic financial fragility $(A_{\mathrm{cr}}, A_{\mathrm{cr+ic}}, A_{\mathrm{ce}}, A_{\mathrm{ce+ic}})$ in the different model variants (cr: counterparty risk, ic: information contagion, ce: common exposure).

	cr $(EU,\ d_1^*,\ y^*,\ b^*)$ $(\overline{\theta}_{\mathrm{H}},\ \overline{\theta}_{\mathrm{1L}}\ ,\ A_{\mathrm{cr}})$	cr + ic $(EU,\ d_1^*,\ y^*,\ b^*)$ $(\overline{\theta}_{\mathrm{H}},\ \overline{\theta}^N_{\mathrm{2L}}\ ,\ \overline{\theta}^D_{\mathrm{2L}},\ A_{\mathrm{cr+ic}})$	ce $(EU,\ d_1^*,\ y^*,\ b^*)$ $(\overline{\theta},\ A_{\mathrm{ce}})$	ce + ic $(EU,\ d_1^*,\ y^*,\ b^*)$ $(\overline{\theta},\ A_{\mathrm{ce+ic}})$
cr	(0.262,0.83,0.6,0.07) (0.404,0.258,0.051)	(0.151,0.83,0.6,0.07) (0.404,0.249,0.271,0.054)		
cr + ic		(0.166,0.92,0.7,0.01) (0.35,0.231,0.234,0.04)		
ce			(0.182,1.01,0.68,0.0) (0.313,0.153)	(0.192,1.01,0.68,0.0) (0.313,0.153)
ce + ic				(0.192,1.02,0.66,0.0) (0.327,0.16)

Table 3.7: Equilibrium allocation for different forms of financial fragility for calibration β=0.7, R=5.0, ϕ=1.0, λ=0.3, η=0.25, ρ=1.0, q_H=0.7. Expected utility (EU), portfolio choice variables (d_1, y, b), withdrawal thresholds $(\overline{\theta}_H, \overline{\theta}_{1,L}, \overline{\theta}^N_{2,L}, \overline{\theta}^D_{2,L}, \overline{\theta})$, and systemic financial fragility $(A_{\mathrm{cr}}, A_{\mathrm{cr+ic}}, A_{\mathrm{ce}}, A_{\mathrm{ce+ic}})$ in the different model variants (cr: counterparty risk, ic: information contagion, ce: common exposure).

	cr $(EU,\, d_1^*,\, y^*,\, b^*)$ $(\overline{\theta}_{\mathrm{H}},\, \overline{\theta}_{\mathrm{1L}},\, A_{\mathrm{cr}})$	cr + ic $(EU,\, d_1^*,\, y^*,\, b^*)$ $(\overline{\theta}_{\mathrm{H}},\, \overline{\theta}_{\mathrm{2L}}^N,\, \overline{\theta}_{\mathrm{2L}}^D,\, A_{\mathrm{cr+ic}})$	ce $(EU,\, d_1^*,\, y^*,\, b^*)$ $(\overline{\theta},\, A_{\mathrm{ce}})$	ce + ic $(EU,\, d_1^*,\, y^*,\, b^*)$ $(\overline{\theta},\, A_{\mathrm{ce+ic}})$
cr	(0.232,0.82,0.69,0.0) (0.36,0.236,0.014)	(0.071,0.82,0.69,0.0) (0.36,0.236,0.236,0.014)		
cr + ic		(0.099,0.94,0.82,0.0) (0.331,0.207,0.207,0.011)		
ce			(0.121,1.0,0.79,0.0) (0.313,0.05)	(0.128,1.0,0.79,0.0) (0.313,0.05)
ce + ic				(0.128,1.0,0.78,0.0) (0.321,0.051)

Table 3.8: Equilibrium allocation for different forms of financial fragility for calibration β=0.7, R=5.0, ϕ=1.0, λ=0.5, η=0.25, ρ=1.0, q_H=0.4. Expected utility (EU), portfolio choice variables (d_1, y, b), withdrawal thresholds $(\overline{\theta}_H, \overline{\theta}_{1,L}, \overline{\theta}_{2,L}^N, \overline{\theta}_{2,L}^D, \overline{\theta})$, and systemic financial fragility $(A_{\mathrm{cr}}, A_{\mathrm{cr+ic}}, A_{\mathrm{ce}}, A_{\mathrm{ce+ic}})$ in the different model variants (cr: counterparty risk, ic: information contagion, ce: common exposure).

3.9 Details for robustness checks

This section provides further details about the robustness checks performed in subsection 3.4.3. In particular, we show the evolution of the portfolio choice variables and withdrawal thresholds when varying the exogenous parameters of the model.

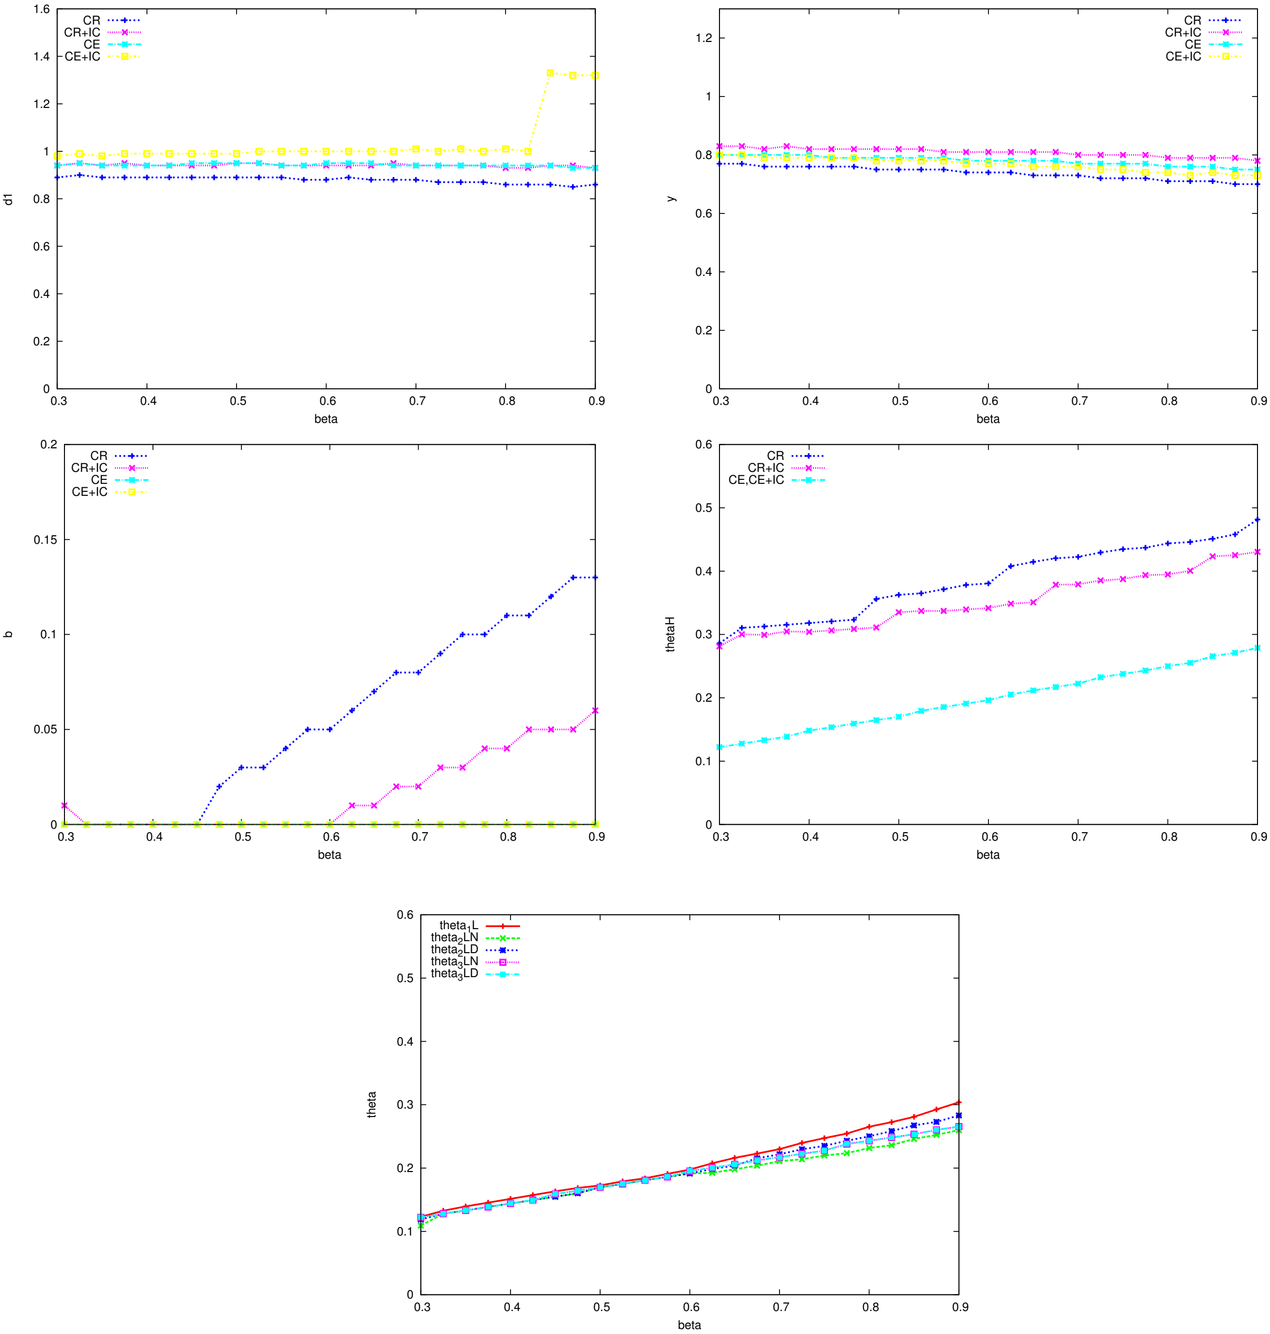

Figure 3.3: Details of portfolio choice: d_1, y (top), b and θ_H (middle), and various θ_L values for a variation of β. The baseline calibration is used for the non-varying parameters.

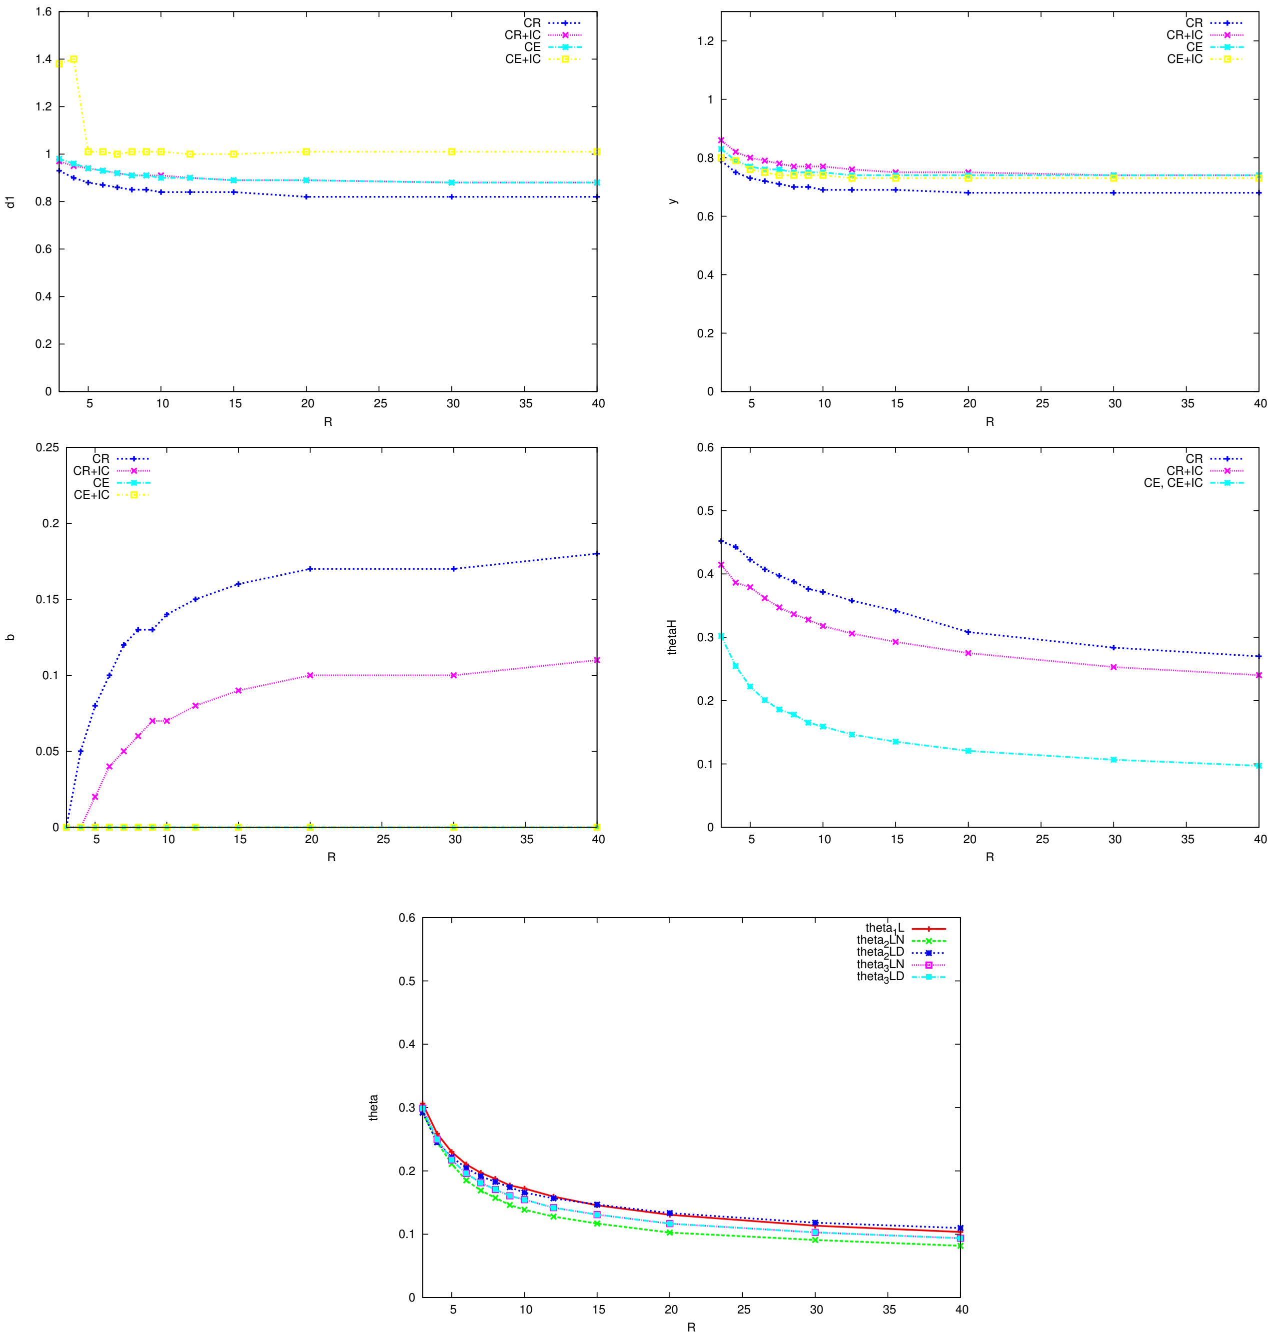

Figure 3.4: Details of portfolio choice: d_1, y (top), b and θ_H (middle), and various θ_L values for a variation of R. The baseline calibration is used for the non-varying parameters.

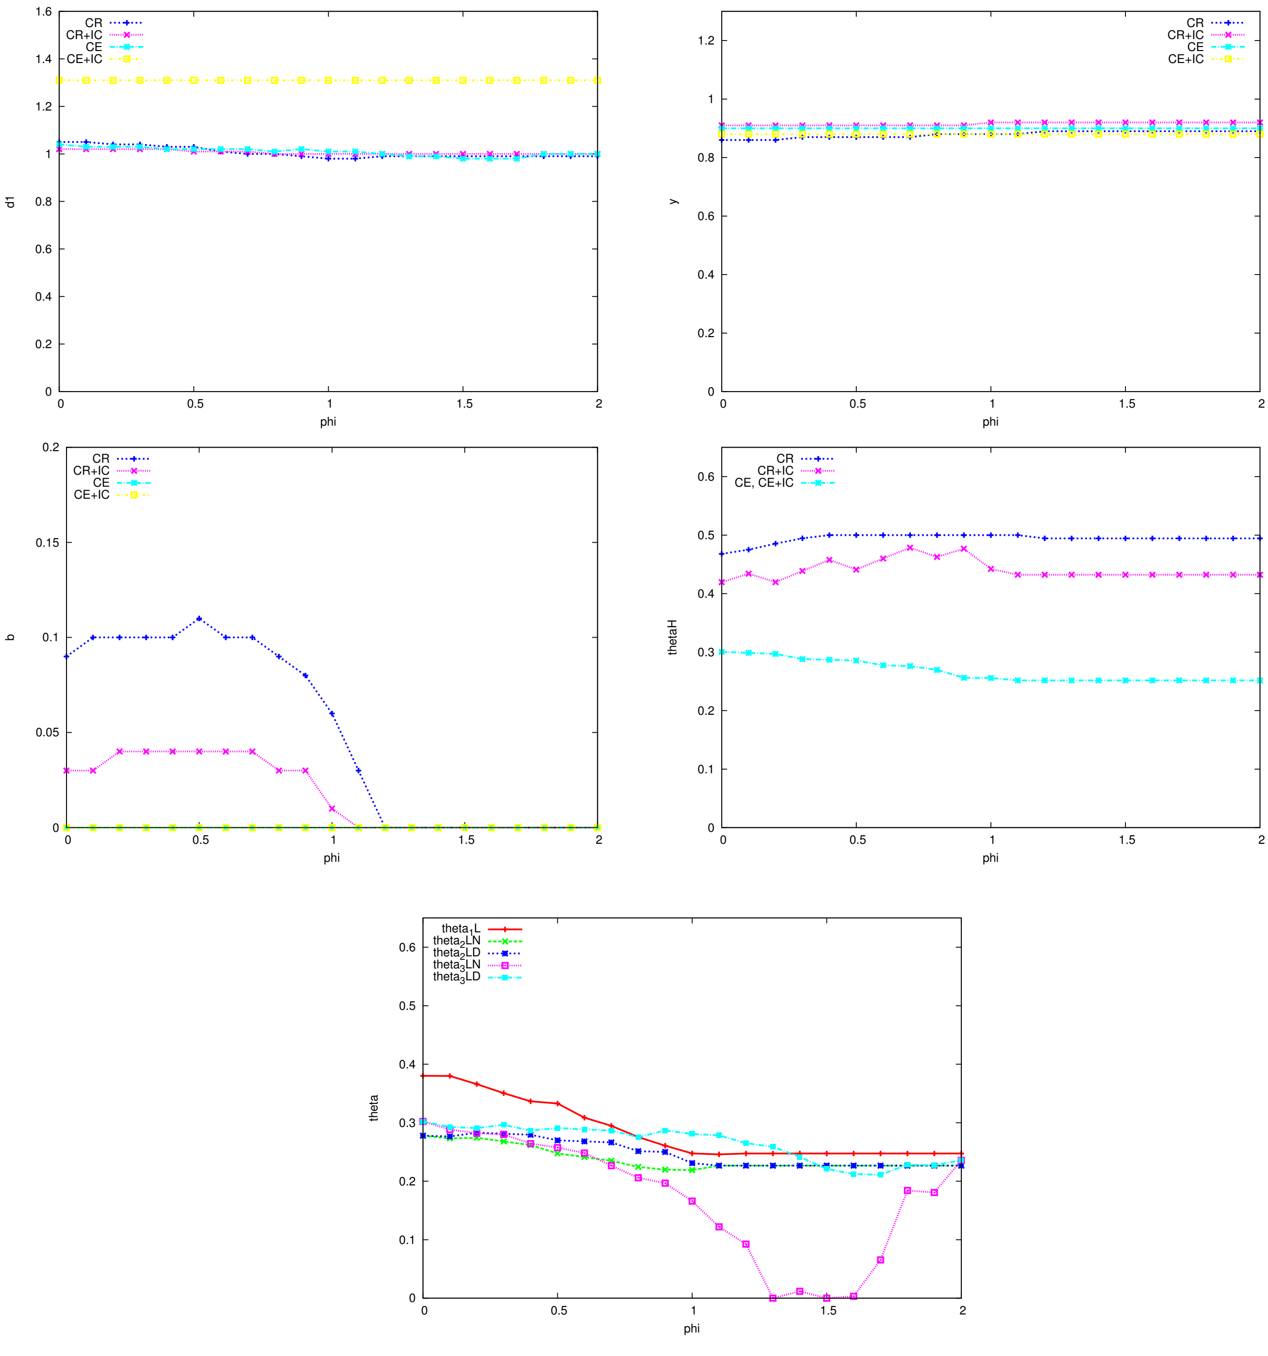

Figure 3.5: Details of portfolio choice: d_1, y (top), b and θ_H (middle), and various θ_L values for a variation of ϕ. The baseline calibration is used for the non-varying parameters.

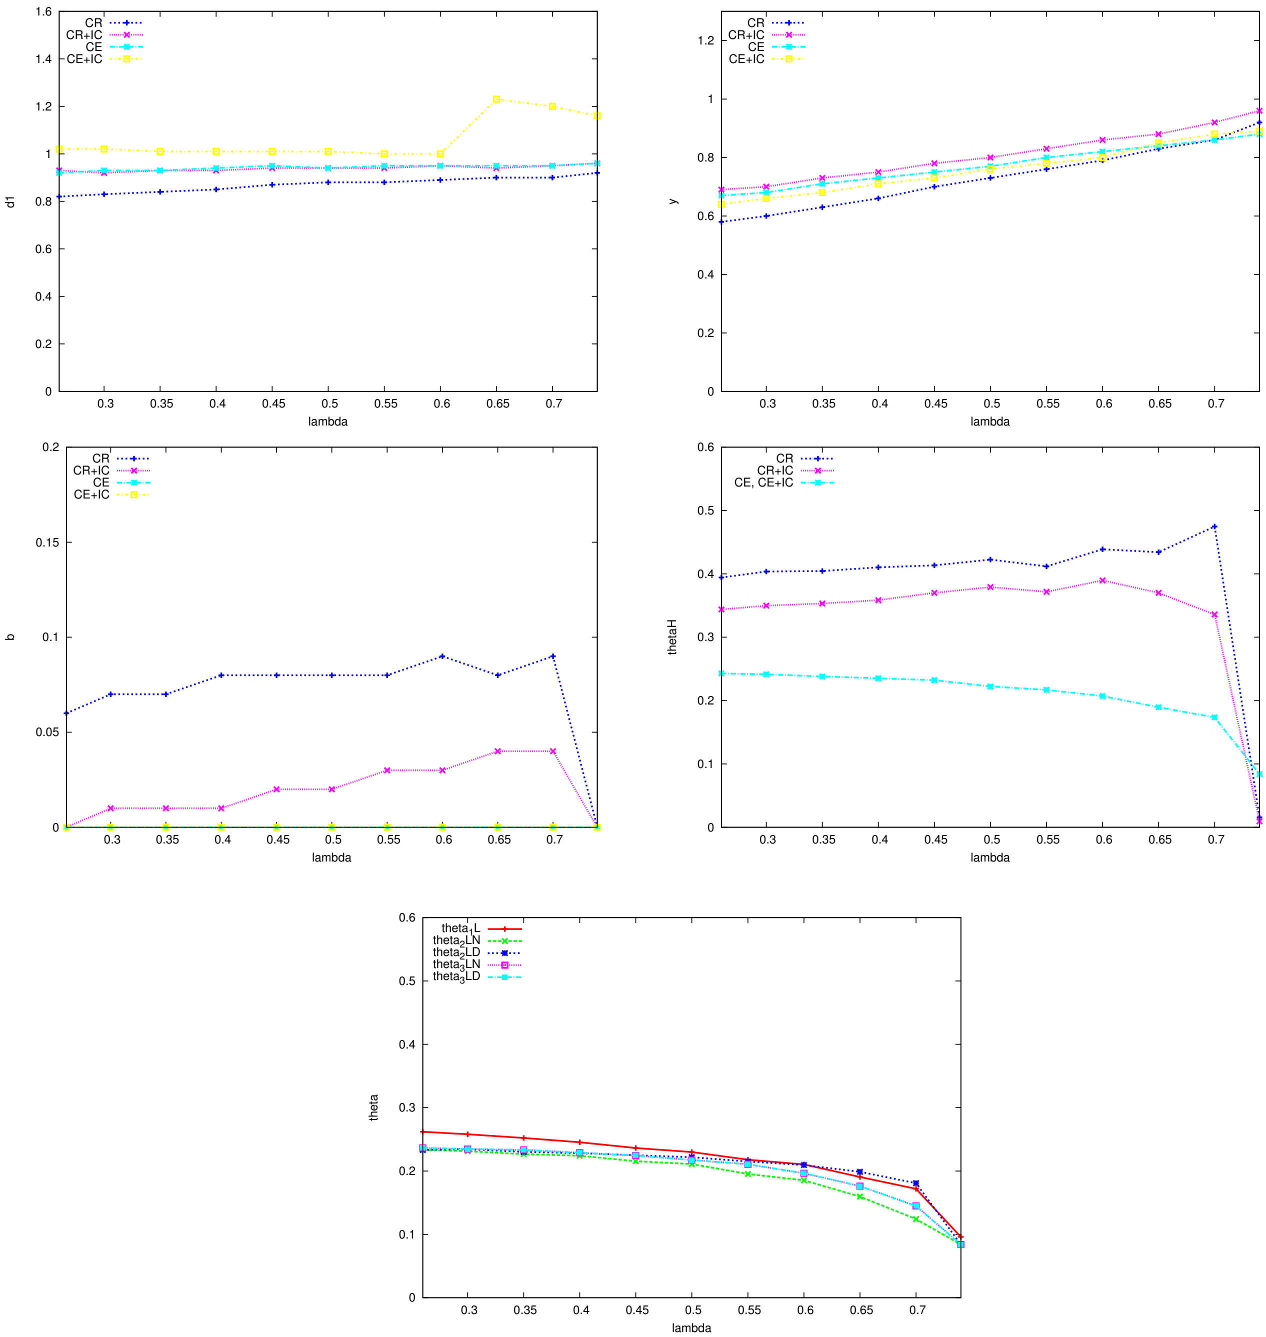

Figure 3.6: Details of portfolio choice: d_1, y (top), b and θ_H (middle), and various θ_L values for a variation of λ. The baseline calibration is used for the non-varying parameters.

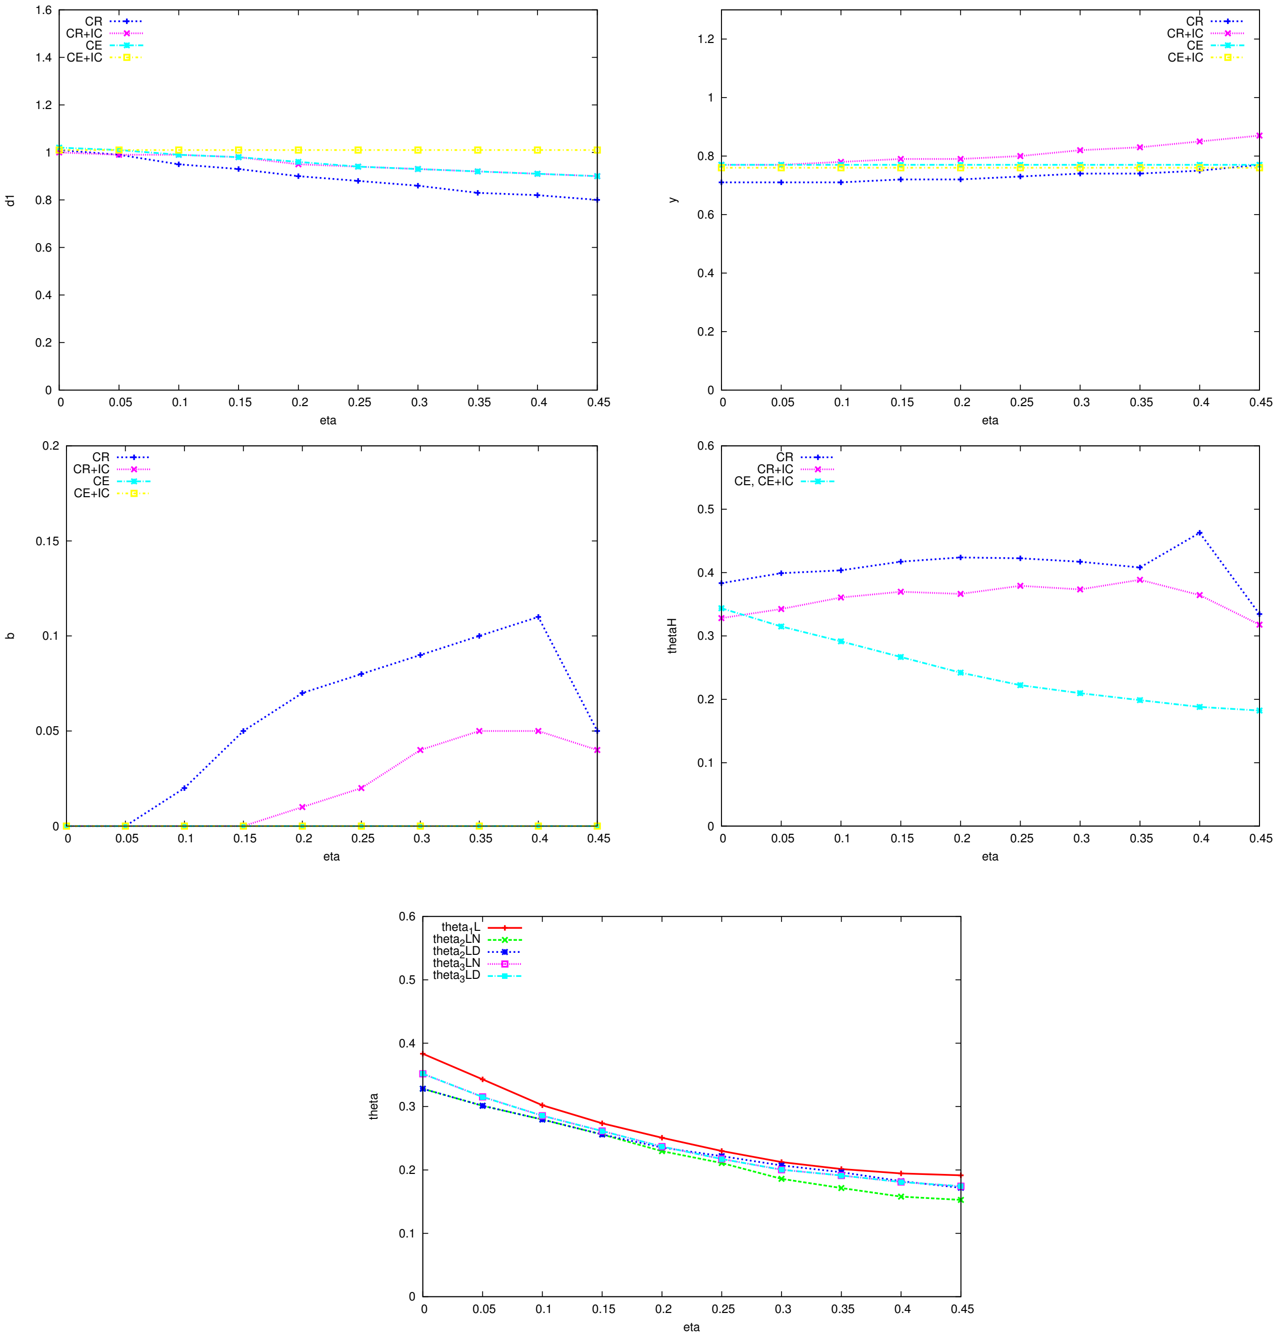

Figure 3.7: Details of portfolio choice: d_1, y (top), b and θ_H (middle), and various θ_L values for a variation of η. The baseline calibration is used for the non-varying parameters.

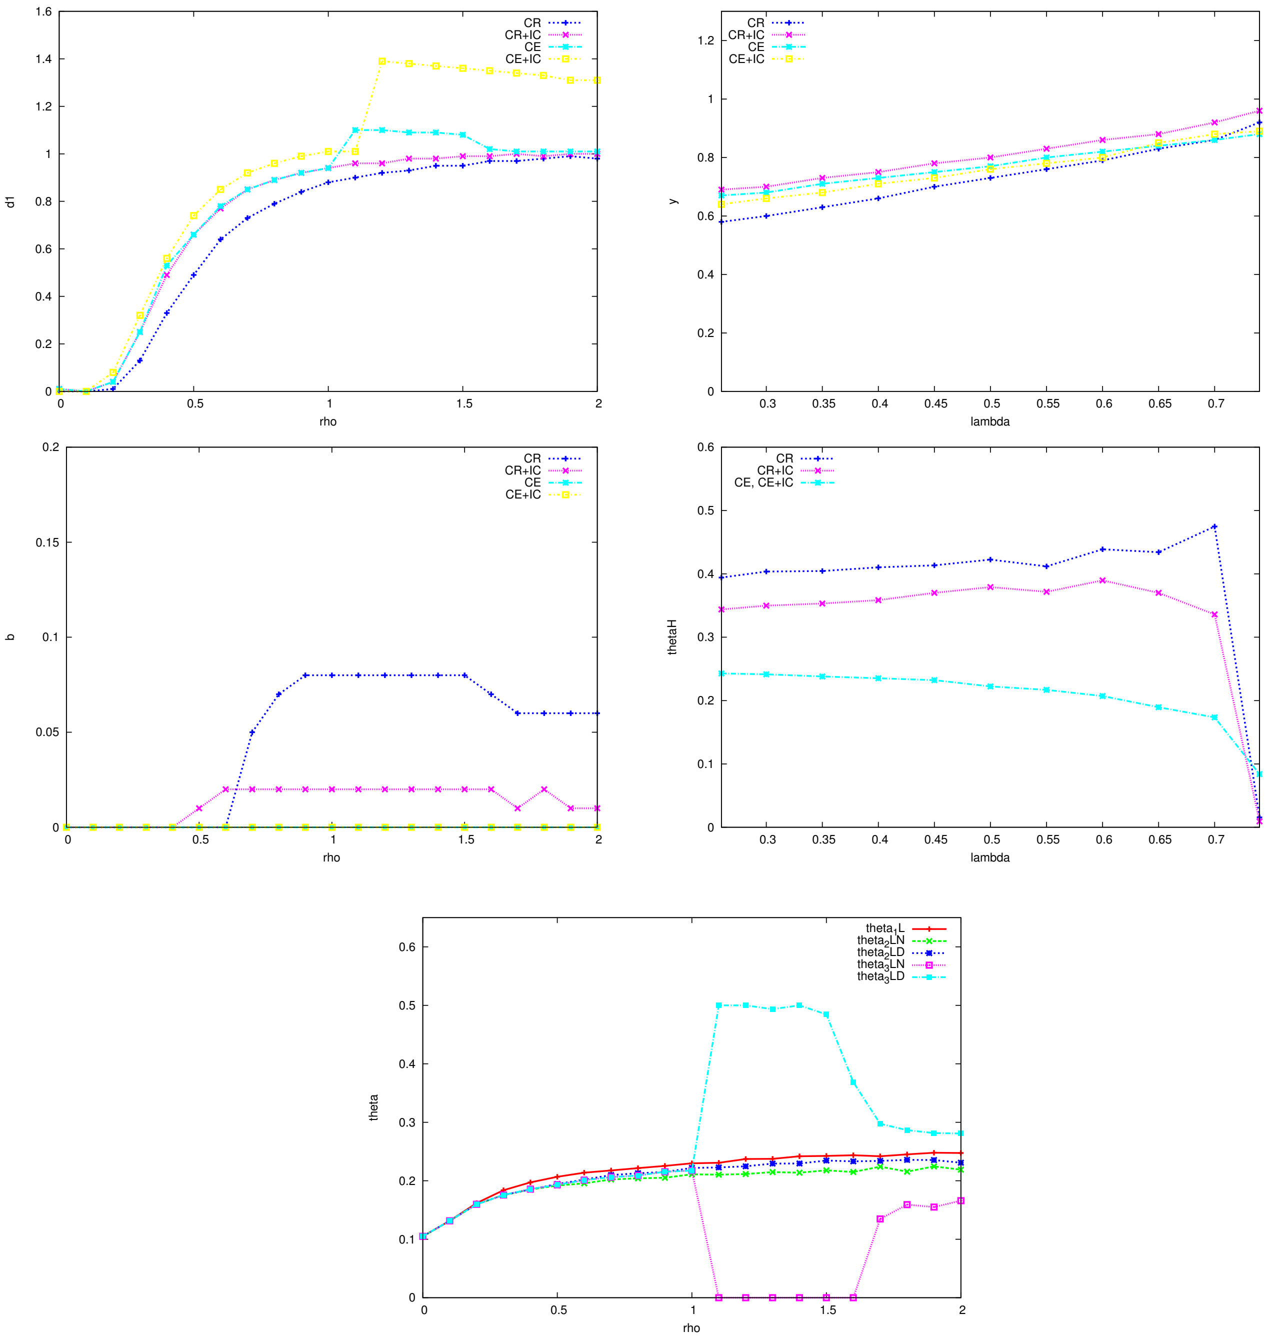

Figure 3.8: Details of portfolio choice: d_1, y (top), b and θ_H (middle), and various θ_L values for a variation of ρ. The baseline calibration is used for the non-varying parameters.

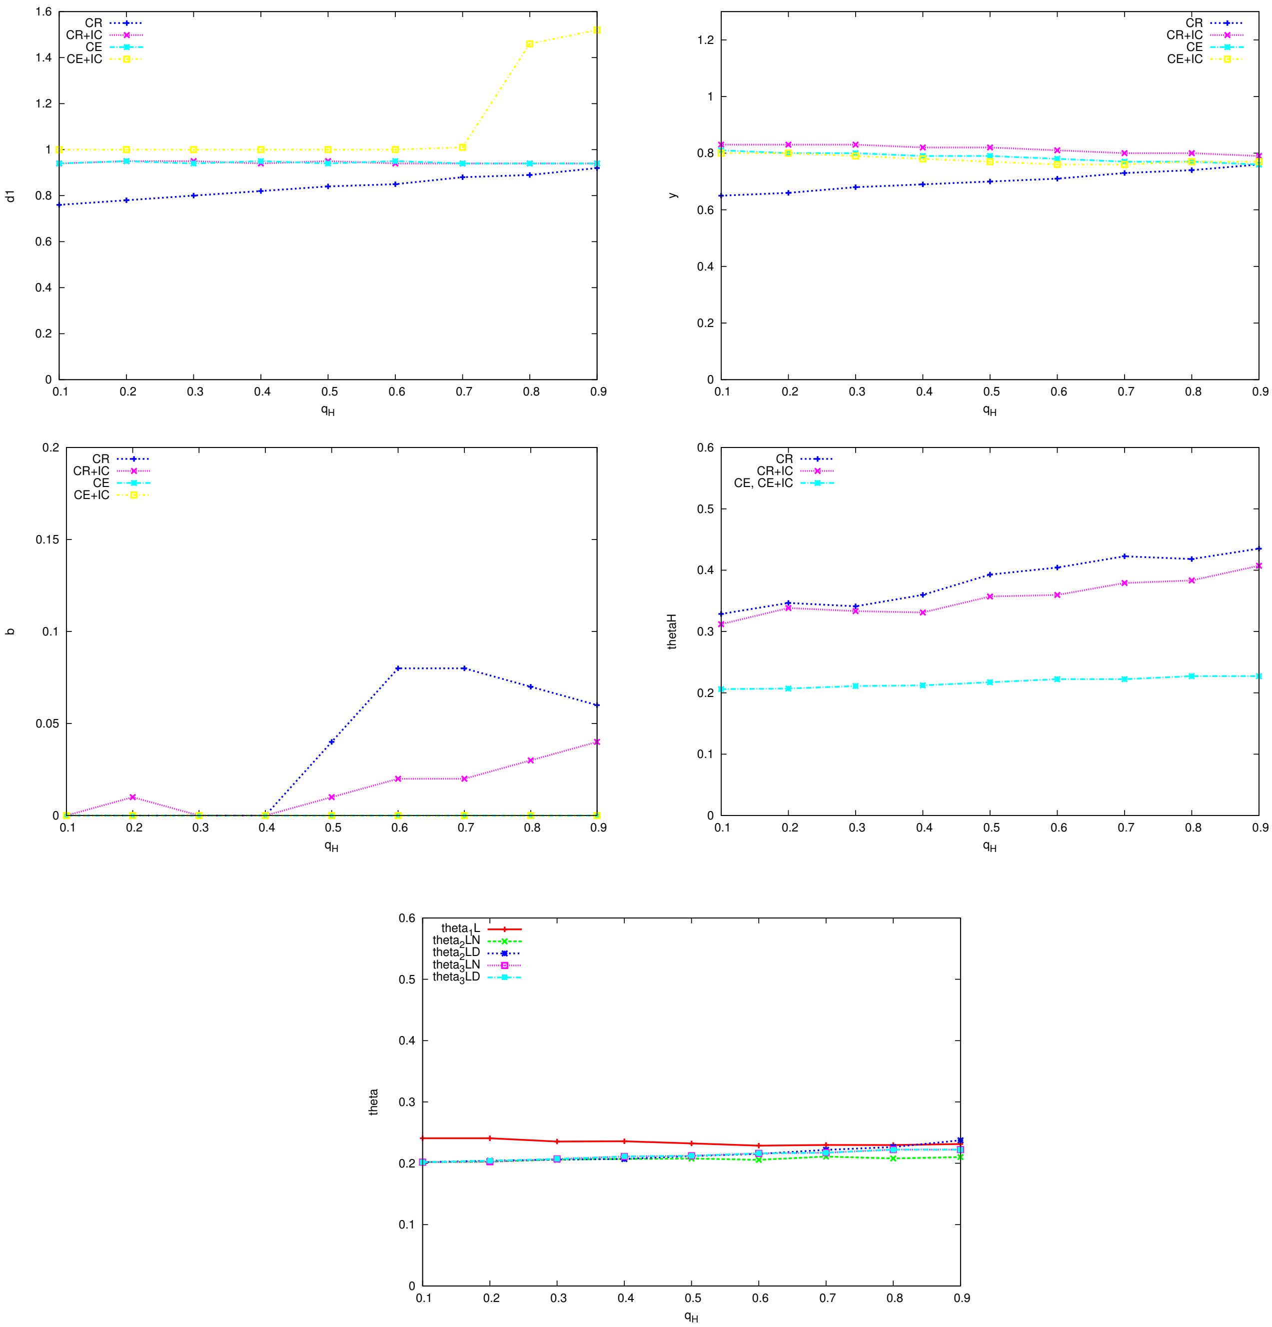

Figure 3.9: Details of portfolio choice: d_1, y (top), b and θ_H (middle), and various θ_L values for a variation of q_H. The baseline calibration is used for the non-varying parameters.

Bibliography

Viral V. Acharya and Tanju Yorulmazer. Cash-in-the-market pricing and optimal resolution of bank failures. *Review of Financial Studies*, 21(6):2705–2742, 2008a.

Viral V. Acharya and Tanju Yorulmazer. Information contagion and bank herding. *Journal of Money, Credit and Banking*, 40(1):215–231, 2008b.

Toni Ahnert. Crisis, coordination, and contagion. *Mimeo, London School of Economics and Political Science*, WP.

Toni Ahnert and Christoph Bertsch. A wake-up call: information contagion and speculative currency attacks. *Mimeo, London School of Economics and Political Science*, WP.

Toni Ahnert and Co-Pierre Georg. Information contagion and systemic risk. *Mimeo, London School of Economics and Political Science*, WP.

Toni Ahnert and Benjamin Nelson. Illiquidity, over-diversification, and macro-prudential regulation. *Mimeo, London School of Economics and Political Science*, WP.

Francis Allen and Douglas Gale. Financial contagion. *Journal of Political Economy*, 108: 1–33, 2000.

Franklin Allen and Douglas Gale. Limited market participation and volatility of asset prices. *American Economic Review*, 84(4):933–55, 1994.

Franklin Allen and Douglas Gale. Understanding financial crises. 2007.

Franklin Allen, Ana Babus, and Elena Carletti. Asset commonality, debt maturity and systemic risk. *Journal of Financial Economics*, 104:519–534, 2012.

George-Marios Angeletos and Alessandro Pavan. Predictions in global games with endogenous information and multiple equilibria. *Mimeo*, WP.

George-Marios Angeletos and Ivan Werning. Crises and prices: Information aggregation, multiplicity, and volatility. *The American Economic Review*, 96(5):1720–1736, 2006.

George-Marios Angeletos, Christian Hellwig, and Alessandro Pavan. Signaling in a global game: Coordination and policy traps. *Journal of Political Economy*, 114(3):452–484, 2006.

George-Marios Angeletos, Christian Hellwig, and Alessandro Pavan. Dynamic global games of regime change: Learning, multiplicity and timing of attacks. *Econometrica*, 75(3):711–756, 2007.

Beatriz Armendáriz and Jonathan Morduch. *The Economics of Microfinance*. The MIT Press, Cambridge, Massachusetts, 2010.

Abhijit Banerjee, Esther Duflo, Rachel Glennerster, and Cynthia Kinnan. The miracle of microfinance: Evidence from a randomized evaluation. 2010.

BCBS. Financial stability in emerging markets economies. 1997.

BCBS. An assessment of the long-term economic impact of stronger capital and liquidity requirements. 2010a.

BCBS. Basel III: A global regulatory framework for more resilient banks and banking systems. December 2010b.

BCBS. Basel III: International framework for liquidity risk measurement, standards and monitoring. December 2010c.

Rosalind L. Bennett and Haluk Unal. The cost effectiveness of the private-sector reorganization of failed banks. *FDIC Center for Financial Research Working Paper No 2009-11*, WP.

R.A. Brown and S. Epstein. Resolution costs of bank failures: An update of the FDIC historical loss model. *FDIC Banking Review*, 5(2):1–16, 1992.

Sandro Brusco and Fabio Castiglionesi. Liquidity coinsurance, moral hazard, and financial contagion. *The Journal of Finance*, 62(5):2275–2302, 2007.

Hans Carlsson and Eric van Damme. Global games and equilibrium selection. *Econometrica*, 61(5):989–1018, 1993.

Christophe Chamley. Coordinating regime switches. *Quarterly Journal of Economics*, 114(3):869–905, 1999.

Yehning Chen. Banking panics: The role of the first-come, first-served rule and information externalities. *Journal of Political Economy*, 107(5):946–968, 1999.

Luca Colombo, Gianluca Femminis, and Alessandro Pavan. Information acquisition and welfare. *Mimeo*, WP.

Russell Cooper and Thomas W. Ross. Bank runs: Liquidity costs and investment distortions. *Journal of Monetary Economics*, 41(1):27–38, 1998.

Amil Dasgupta. Financial contagion through capital connections: A model of the origin and spread of bank panics. *Journal of the European Economic Association*, 2(6):1049–1084, 2004.

Amil Dasgupta. Coordination and delay in global games. *Journal of Economic Theory*, 134:195–225, 2007.

Morris H DeGroot. *Optimal Statistical Decisions.* John Wiley and Sons (Wiley Classics Library Edition), 1970.

D.W. Diamond and P.H. Dybvig. Bank runs, deposit insureance and liquidity. *Journal of Political Economy*, 91:401–419, 1983.

Chris Edmond. Information manipulation, coordination, and regime change. *NBER Working Paper 17395*, 2012.

Huberto M. Ennis and Todd Keister. Bank runs and investment decisions revisited. *Journal of Monetary Economics*, 53(2):217–232, 2006.

Emmanuel Farhi and Jean Tirole. Collective moral hazard, maturity mismatch, and systemic bailouts. *American Economic Review*, 102(1):60–93, 2012.

Kristin J Forbes. The 'Big C': Identifying and Mitigating Contagion. *MIT Sloan Research Paper No. 4970-12*, 2012.

David Frankel, Stephen Morris, and Ady Pauzner. Equilibrium selection in global games with strategic complementarities. *Journal of Economic Theory*, 108(1):1–44, 2003.

Xavier Freixas, Bruno M Parigi, and Jean Charles Rochet. Systemic risk, interbank relations, and liquidity provision by the central bank. *Journal of Money, Credit and Banking*, 32(3):611–38, 2000.

Douglas Gale. Capital regulation and risk sharing. *International Journal of Central Banking*, 6(3):187–204, 2010.

Xavier Giné, Pamela Jakiela, Dean Karlan, and Jonathan Morduch. Microfinance games. 2009.

Itay Goldstein. Strategic complementarities and the twin crises. *Economic Journal*, 115: 368–390, 2005.

Itay Goldstein and Ady Pauzner. Contagion of self-fulfilling financial crises due to diversification of investment portfolios. *Journal of Economic Theory*, 119:151–183, 2004.

Itay Goldstein and Ady Pauzner. Demand deposit contracts and the probability of bank runs. *Journal of Finance*, 60(3):1293–1327, 2005.

Denis Gromb and Dimitri Vayanos. Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics*, 66(2):361–407, 2002.

Sanford J. Grossman and Joseph E. Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408, 1980.

Zhiguo He and Asaf Manela. Information acquisition in rumor-based bank runs. *Mimeo*, WP.

Christian Hellwig and Laura Veldkamp. Knowing what others know: Coordination motives in information acquisition. *Review of Economic Studies*, 76:223–251, 2009.

Christian Hellwig, Arijit Mukherji, and Aleh Tsyvinski. Self-fulfilling currency crises: The role of interest rates. *American Economic Review*, 96:1769–1787, 2006.

Charles J Jacklin and Sudipto Bhattacharya. Distinguishing panics and information-based bank runs: Welfare and policy implications. *Journal of Political Economy*, 96 (3):568–592, 1988.

Christopher James. The losses realized in bank failures. *Journal of Finance*, 46(4): 1223–1242, 1991.

Anton Korinek. Systemic risk-taking: Amplification effects, externalities, and regulatory responses. *ECB Working Paper 1345*, 2011.

Anton Korinek. Macroprudential regulation versus mopping up after the crash. *Mimeo, University of Maryland*, WP.

Yaron Leitner. Financial networks: Contagion, commitment, and private sector bailouts. 60:2925 – 2953, 2005.

Guido Lorenzoni. Inefficient credit booms. *Review of Economic Studies*, 75(3):809–833, 2008.

Microcredit Summit Campaign. The state of the microcredit summit campaign report. Technical report, 2012.

Jonathan Morduch. The microfinance promise. *Journal of Economic Literature*, 37:1569 – 1614, 1999.

S. Morris and H. S. Shin. Coordination risk and the price of debt. *European Economic Review*, 48:133–153, 2004.

Stephen Morris and Hyun Song Shin. Rethinking multiple equilibria in macroeconomic modelling. *Cowles Foundation Discussion Paper 1260*, 2000.

Stephen Morris and Hyun Song Shin. Social value of public information. *American Economic Review*, 92(5):1521–1534, 2002.

Stephen Morris and Hyun Song Shin. Global games: Theory and applications, 2003.

Stephen Morris and Hyun Song Shin. Illiquidity component of credit risk. *Mimeo, Princeton University*, WP.

Helena M. Mullins and David H. Pyle. Liquidation costs and risk-based bank capital. *Journal of Banking and Finance*, 18(1):113–138, 1994.

Maxim Nikitin and R. Todd Smith. Information acquisition, coordination, and fundamentals in a financial crisis. *Journal of Banking and Finance*, 32(6):907 – 914, 2008.

Maurice Obstfeld. Rational and self-fulfilling balance-of-payments crises. *American Economic Review*, 76(1):72–81, 1986.

Emre Ozdenoren and Kathy Zhichao Yuan. Feedback effects and asset prices. *The Journal of Finance*, 63(4):1939–1975, 2008.

Andrew Postlewaite and Xavier Vives. Bank runs as an equilibrium phenomenon. *Journal of Political Economy*, 95, 1987.

Jean-Charles Rochet and Xavier Vives. Coordination failures and the lender of last resort: was Bagehot right after all? *Journal of the European Economic Association*, 2 (6):1116–1147, 2004.

Joel Shapiro and David Skeie. Information management and banking crises. *Mimeo, Oxford University*, WP.

Andrei Shleifer and Robert W. Vishny. Liquidation values and debt capacity: A market equilibrium approach. *The Journal of Finance*, 47(4):1343–1366, 1992.

Joseph E. Stiglitz. Peer monitoring and credit markets. *The Worldbank Economic Review*, 4:351–366, 1990.

Michal Szkup and Isabel Trevino. Information acquisition and transparency in global games. *Mimeo*, WP.

Harald Uhlig. A model of a systemic bank run. *Journal of Monetary Economics*, 57: 78–96, 2010.

Hal R. Varian. Monitoring agents with other agents. *Journal of International and Theoretical Economics*, 146:153–174, 1990.

Laura Veldkamp. *Information choice in macroeconomics and finance*. Princeton University Press, 2011.

Xavier Vives. Strategic complementarity, fragility, and regulation. *Mimeo*, WP.

Wolf Wagner. Efficient asset allocations in the banking sector and financial regulation. *International Journal of Central Banking*, 5(1):75–95, 2009.

Wolf Wagner. Systemic liquidation risk and the diversity-diversification trade-off. *Journal of Finance*, 66:1141âĂŞ–1175, 2011.

Russ Wermers. Runs on money market mutual funds. *Mimeo*, WP.

Bruce Wydick. Can social cohesion be harnessed to repair market failures? Evidence from group lending in Guatemala. *Economic Journal*, 109:463 – 475, 1999.

Sanne Zwart. The mixed blessing of IMF intervention: Signalling versus liquidity support. *Journal of Financial Stability*, 3(2):149 – 174, 2007.