

Financial System Risk and Flight to Quality

Ricardo J. Caballero

Arvind Krishnamurthy*

This draft: October 18, 2005

Abstract

We present a model of flight to quality episodes that emphasizes financial system risk and the Knightian uncertainty surrounding these episodes. In the model, agents are uncertain about the probability distribution of shocks in markets different from theirs, treating such uncertainty as Knightian. Aversion to Knightian uncertainty generates demand for safe financial claims. It also leads agents to require financial intermediaries to lock-up capital to cover their own markets' shocks in a manner that is robust to uncertainty over other markets, but is wasteful in the aggregate. Locked collateral cannot move across markets to offset negative shocks and hence can trigger a financial accelerator. A lender of last resort can unlock *private* capital markets to stabilize the economy during these periods by committing to intervene should conditions worsen.

JEL Codes: E30, E44, E5, F34, G1, G21, G22, G28.

Keywords: Locked collateral, flight to quality, multiplier, excess volatility, value at risk, insurance, financial intermediaries, lender of last resort, collateral shocks, robust control.

*Respectively: MIT and NBER; Northwestern University. E-mails: caball@mit.edu, a-krishnamurthy@northwestern.edu. We are grateful to Marios Angeletos, Olivier Blanchard, Phil Bond, Xavier Gabaix, Jordi Gali, Michael Golosov, William Hawkins, Burton Hollifield, Bengt Holmstrom, David Lucca, Dimitri Vayanos, and seminar participants at Columbia, Northwestern, MIT, Wharton, NY Fed Liquidity Conference, UBC Summer Conference and the Central Bank of Chile for their comments. David Lucca provided excellent research assistance. Caballero thanks the NSF for financial support.

1 Introduction

“... Policy practitioners operating under a risk-management paradigm may, at times, be led to undertake actions intended to provide insurance against especially adverse outcomes..... When confronted with uncertainty, especially Knightian uncertainty, human beings invariably attempt to disengage from medium to long-term commitments in favor of safety and liquidity.... The immediate response on the part of the central bank to such financial implosions must be to inject large quantities of liquidity...” Alan Greenspan (2004).

Flight-to-quality episodes are an important source of financial instability. Modern examples of these episodes in the US include the Penn Central default of 1970; the stock market crash of 1987; the events of the Fall of 1998 beginning with the Russian default and ending with the bailout of LTCM; as well as the events that followed the attacks of 9/11. Behind each of these episodes lies the specter of a meltdown that may lead to a prolonged slowdown as in Japan during the 1990s, or even a catastrophe like the Great Depression.¹

From a broad-brush standpoint, flight-to-quality episodes involve an increase in perceived risk. However, the risk does not circumscribe a purely fundamental shock, and instead centers around the financial system. For example, the Russian default in the summer of 1998 eliminated a small fraction of the trillions of dollars of US wealth. Although small, the default created circumstances that severely strained the financial sector. As prices of illiquid assets fell, losses grew in commercial banks, investment banks, and hedge funds, leading investors to question the safety of the financial sector. Investors withdrew risk-capital from the affected markets and institutions and moved into short-term and liquid assets. Bottlenecks in the movement of capital emerged as sophisticated parts of the financial system were compromised while other sectors of the economy were relatively unaffected. The primary risk during this episode was *financial system risk*.

In this paper we develop a model of a flight-to-quality episode. In practice, the occurrence of an unexpected event triggers an increase in agents' perception of risk. We take this shock as exogenous, and focus on two important effects that arise during flight-to-quality episodes: agents demand safe and certain financial claims; and financial bottlenecks develop that hinder capital movement and segment financial markets. The increase in risk and demand for safety are a hallmark of flight-to-quality episodes. Bottlenecks transform flight to quality into a macroeconomic problem by triggering and amplifying collateral shortages and financial multipliers. It is this systemic dimension that motivates the Fed's reactions, as indicated by Greenspan's comments.

¹See Table 1 (part A) in Barro (2005) for a comprehensive list of extreme events in developed economies during the 20th century.

Our model centers on agent's perceptions regarding the ability of financial intermediaries to deliver on promises of credit. Agents contract with financial intermediaries to cover shocks that may arise in the agents' markets. However, agents are uncertain about the probability model describing shocks in markets different from theirs, treating such uncertainty as Knightian. If financial intermediaries have limited collateral, agents grow concerned that shocks may arise that will compromise intermediaries' promises of credit. Such riskiness rises either through an increase in Knightian uncertainty or a fall in intermediation collateral. The increase in perceived riskiness generates conservatism and a demand for safety. We study the case in which agents shield themselves from the increased risk by formulating plans that are robust to their uncertainty regarding other markets. In particular, we show that agents require financial intermediaries to lock-up some capital to devote to their own markets' shocks. Once locked-up, intermediaries' capital is not free to move across markets in response to shocks, resulting in bottlenecks and market segmentation.

The lock-up of intermediary capital is a best response of agents in their own market but is wasteful in the aggregate. Each agent covers himself against an extreme shock, but collectively these actions prevent intermediaries from moving capital across markets to efficiently offset shocks. We show that this inflexibility leaves the economy overexposed to moderate shocks (relative to the size of the economy) that trigger contractions. In this context, a lender of last resort, even if less informed than private agents about each agent's own market, can unlock private collateral and stabilize the economy during moderate shocks. It can do so by committing to intervene during extreme events where the financial intermediary's capital is depleted. Importantly, because these extreme events are highly unlikely, the expected cost of this intervention is minimal. If credible, the policy derives its power from a *private sector multiplier*: each pledged dollar of public intervention in the extreme event is matched by a comparable private sector reaction to insure against moderate shocks. Since the moderate shocks are more likely than the extreme ones, the benefit to cost ratio of the intervention is large.²

Our paper relates to several strands of literature. It helps to fill a gap in the literature on financial frictions in business cycle models. Papers such as Bernanke, Gertler, and Gilchrist (1998) and Kiyotaki and Moore (1997) highlight how financial frictions in firms amplify aggregate shocks. Our paper instead emphasizes how financial frictions lead to greater (Knightian) uncertainty in response to shocks, and how this rise in uncertainty feeds back into the financial accelerator.

Our paper also studies the macroeconomic consequences of frictions in financial intermediaries. In this sense, our paper is closer to Holmstrom and Tirole (1998) and Krishnamurthy (2003) that emphasize that with complete financial contracts, the *aggregate* collateral of intermediaries is ultimately behind financial amplification mechanisms.

²From our perspective, the LTCM bailout was important not for its direct effect, but because it served as a signal of the Fed's readiness to intervene should conditions worsen.

In terms of the policy implications, Holmstrom and Tirole (1998) study how a shortage of aggregate collateral limits private liquidity provision (see also Woodford, 1990). Their analysis suggests that a credible government can issue government bonds which can then be used by the private sector for liquidity provision. The key difference between our paper and those of Holmstrom and Tirole, and Woodford, is that we show how even a large amount of collateral in the aggregate may be inefficiently used, so that private sector liquidity provision is limited. In our model, the government intervention not only adds to the private sector's collateral but also, and more centrally, it improves the use of *private* collateral.

Our model illustrates how a lender of last resort commitment to insure a low probability extreme event stabilizes the economy. This result resembles Diamond and Dybvig (1983), where the benefit of a lender of last resort in bank runs is to rule out a "bad" equilibrium. However, our model has a unique equilibrium; the bad event is just a low probability extreme event. In Diamond and Dybvig, a run is triggered by depositors concern that they will be last in line to withdraw deposits. In our model, agents exaggerate the odds that intermediaries will have exhausted their capital, and fail to deliver on promises of credit. In Diamond and Dybvig, the bank run is a coordination failure. In our model, on the other hand, flight to quality results from agents overweighting the chance of an extreme event.

Flight-to-quality episodes have financial and real effects. On the financial side, papers such as Krishnamurthy (2002) and Longstaff (2004) document that spreads between illiquid/risky assets and liquid/safe assets widen during flight to quality. Gatev and Strahan (2005) document that investors shift their portfolios away from risky assets such as Commercial Paper and toward safe bank deposits during these episodes, while firms draw down credit lines from these banks. These actions are consistent with our model, as we predict that agents shift toward well capitalized financial intermediaries that are less sensitive to the robustness concerns of agents. A number of papers have also documented that the widening of the spread between Commercial Paper and Treasury Bills is a powerful leading indicator for cyclical downturns (see, e.g., Stock and Watson, 1989, Friedman and Kuttner, 1993, Kashyap, Stein, and Wilcox, 1993).

There is a growing economics literature that aims to formalize Knightian uncertainty (a partial list of contributions includes, Gilboa and Schmeidler (1989), Dow and Werlang (1992), Epstein and Wang (1994), Hansen and Sargent (1995, 2003), Skiadas (2003), Epstein and Schneider (2004), and Hansen, et al. (2004)). As in much of this literature, we use a max-min device to describe agents expected utility. Our treatment of Knightian uncertainty is most similar to Gilboa and Schmeidler, in that agents choose a worst case among a class of priors.

Our paper applies max-min expected utility theory to agents who we interpret as running financial firms. These firms typically stress-test their models before formulating investment policies. The widespread use of Value-at-Risk as a decision making criterion is an example of robust decision making in practice. Corporate liquidity management is also done with a worst case scenario for cash-flows in mind. We view the max-min

preferences of the agents we study as descriptive of their decision rules rather than as stemming from a deeper psychological foundation. In much of the paper we refer to agents as robust decision makers. This terminology most closely corresponds to the decision making process of the financial institutions that concern us in this paper.

Routledge and Zin (2004) also begin from the observation that financial institutions follow decision rules to protect against a worst case scenario. They develop a model of market liquidity in which an uncertainty averse market maker sets bids and asks to facilitate trade of an asset. Their model captures an important aspect of flight to quality: uncertainty aversion can lead to a sudden widening of the bid-ask spread, causing agents to halt trading and reducing market liquidity. Both our paper and Routledge and Zin share the emphasis on financial intermediation and uncertainty aversion as central ingredients in flight to quality episodes. But each paper captures different aspects of flight to quality.

In our model, agents are only Knightian with respect to other markets and not their own market. Epstein (2001) explores the home bias in international portfolios in a similar setup. As Epstein points out, this modeling also highlights the difference between high risk aversion and aversion to Knightian uncertainty. Moreover, our modeling shows that max-min preferences interact with macroeconomic conditions in ways that are not present in models with an invariant amount of risk aversion. We show that when aggregate intermediary collateral is plentiful, Knightian and standard agents behave identically. However when aggregate collateral falls, the actions of these agents differ, leading to flight to quality in the Knightian model.

In Section 2 we describe the environment and financial claims. Section 3 characterizes decisions and equilibrium. Section 4 explains the implications of our model for flight-to-quality episodes. Section 5 derives the value of a lender of last resort in our economy. Sections 6 and 7 illustrate the interaction between aggregate collateral and robustness. Section 8 concludes.

2 The Environment

We consider a continuous time economy with a single consumption good. At date 0, the economy enters a phase where liquidity shocks may take place, and it exits this phase at some random date τ , distributed exponentially with hazard rate δ . We study the economy conditional on entering this liquidity phase.

Liquidity demand

There are two groups of agents, each with unit measure, whose members we label as agent-type A and B. Each individual derives utility from consumption at date τ :

$$U^i = u(c_\tau^i) = \ln c_\tau^i \quad i \in \{A, B\}. \quad (1)$$

where the log-utility assumption is primarily made for expositional reasons (we retain the more general notation $u(\cdot)$ where convenient). There is no discounting.

Each agent is endowed at date 0 with w_0 in consumption goods. These goods are perishable, and are used to purchase financial claims – see below.

In the normal course of business, the agents can expect to earn a profit of Π when exiting the liquidity phase, so that $c_\tau = \Pi$. However, the agents may receive a “liquidity shock” that reduces profits to 0 unless they have some liquidity available to them immediately. With such liquid funds, agents can insure against the fall in profits. We normalize the return on the liquidity investment to one, so that investing l of liquid funds results in a consumption level of $c_\tau(l) = l$.

The liquidity shock is common to all agents of a given type. We will think of all agents of a type as belonging to a particular market, so that the shock has the interpretation of a market-specific shock. The shock occurs to each market at most once between 0 and τ . The arrival of the shock is Poisson, with intensity λ . The shocks are uncorrelated across markets.

Agents purchase financial claims that deliver some liquidity to them in the event of the liquidity shock and thereby insure against the fall in profits.

Liquidity supply and aggregate collateral constraint

The financial sector consists of intermediaries who sell liquidity insurance to A and B at date 0. We think of these intermediaries as having some capital available to allocate across A and B 's markets in the event of shocks.

Financial intermediaries are required to collateralize all sales of liquidity insurance. We assume that insurers have ready access to consumption goods at all dates, and are only restricted in their supply of liquidity insurance by their collateral. For now, we assume that collateral is riskless and worth Z , but later in the paper we relax this assumption and study situations where the collateral value varies over time and across states (see Sections 6 and 7).

The intermediaries' objective is to maximize the date 0 revenue from the sale of the insurance less a cost of the insurance resources disbursed at time s of β per unit:

$$U^I = c_0^I + \beta c_s^I, \quad \beta \geq 0 \tag{2}$$

and subject to the collateral constraint that, in any state, the insurer does not write insurance exceeding Z .

Financial claims

Each agent is uncertain about whether he will be hit by a shock or not and about the order in which shocks will take place. There is a complete set of Arrow-Debreu claims contingent on the shock realizations of both types of agents. We define $x(s)$ as a claim that pays one at date s , conditional on the shock at date

s being the first one. $y_s(t)$ is a claim paying one at date t , conditional on the first shock having occurred at date $s < t$.

Robustness

Agents purchase financial claims from the intermediaries to insure against their liquidity shocks. In making the insurance decisions, agents have a probability model of the liquidity shocks in mind. But, while each agent knows the intensity of the shock in his own market, he is unsure of the intensity of the shock in the other agent-types' market. We explore the case in which agents formulate decisions that are robust to the intensity of the other agents' shock.

We write all of our problems from the standpoint of agent(s) A , referring to the other agent-type as B , and denoting A 's assessment of B 's intensity as λ_B .

More formally, we follow Gilboa and Schmeidler (1989), and express agent A 's decision problem as:

$$\max_{\{x(s)\}, \{y_s(t)\}} \min_{\lambda_B \in \Lambda} E_0[u(c_T^A) \mid \lambda_B, \lambda_A]$$

where Λ is a class of alternative probability models that agent A deems as possible.³

Discussion of setup

Our model studies an economy conditional on entering a liquidity episode. It is silent on what triggers the episode. In practice, we think that the occurrence of an unexpected event, such as the LTCM or Enron crises, causes agents to re-evaluate their models and triggers robustness concerns. We focus on the mechanisms that play out during the liquidity episode: How do agents' robustness concerns affect prices and quantities? How do these robustness concerns interact with aggregate constraints? What is the role of an outside liquidity provider such as the central bank?

We think of the agents in our model as financial specialists who actively participate in a single market. For example, the agents may be firms in a single industry that is susceptible to a temporary drop in cash-flow. Alternatively, the agents may be banks that are susceptible to an outflow of deposits. Finally, the agents may be hedge-funds/market-makers who are susceptible to a shock to their trading capital. In all of these cases, we model the agents who are responsible for insuring against a temporary liquidity shortfall.

We view agents' max-min preferences as descriptive of their decision rules. The widespread use of value-at-risk as a decision making criteria is an example of the robustness preferences of such agents.⁴

³We study arrangements where the financial intermediary provides insurance, ruling out direct consumption insurance arrangements between A and B . In keeping with our collateral restrictions on the intermediary, we may think of A and B as having no collateral and therefore being unable to support such insurance arrangements. If we assigned $Z/2$ collateral to each of the agents to support mutual insurance, agents' robustness concerns would reduce such insurance. Intermediaries maximize the use of collateral.

⁴See, e.g., Routledge and Zin (2004) for a more extensive discussion of examples of max-min behavior among financial specialists.

Our main goal is to investigate why these robustness concerns are important in the aggregate and are affected by aggregate conditions. In principle, the demand for safety by a sector of hedge funds, say during the Russian default, can be easily accommodated by the rest of the capital market, with negligible effects on asset prices or quantities. Yet, in practice, capital does not move freely across markets to provide the necessary liquidity. There are bottlenecks, suggesting that there are constraints on the supply of capital. We model these constraints by assuming that financial intermediaries, who specialize in capital reallocation, are limited by a collateral constraint. The collateral of the intermediary measures the maximum amount of capital they can supply to a particular market. Since the intermediaries are risk-neutral, they will supply insurance to our agents. But, in aggregate, the collateral constraints place a limit on how much the supply-side can respond to the robustness-concerns triggered during a liquidity episode.

Our model assumes that there are complete liquidity-shock-contingent financial claims. Moreover, agents prearrange insurance against liquidity shocks at date 0. In practice, liquidity shocks trigger dynamic trading in asset markets – for example, by shedding risky assets and moving to Treasury Bills – so that agents’ responses are ex-post and not all ex-ante. We have chosen our strategy despite its interpretability costs because, as usual, with complete Arrow-Debreu markets a dynamic problem can be more easily stated as a static one. More substantively, our objective is to provide an *endogenous* explanation for why capital does not move across markets, as opposed to making an exogenous segmentation/incomplete markets assumption. The assumption of complete liquidity-shock-contingent claims isolates the mechanism we highlight.

As a matter of interpretation, the ex-ante insurance contracts can be thought of as collateralized contingent credit-lines.

3 Decisions and Equilibrium

In this section we describe agents’ decisions and equilibrium, and contrast them with a benchmark case where robustness is not a concern. Flight-to-quality, which we discuss more fully in the next section, is described in terms of the difference in agents’ equilibrium decisions across these two cases.

3.1 The cost of locking collateral

The problem of a financial intermediary is to sell $\hat{x}(s)$ first-shock claims and $\hat{y}_s(t)$ second-shock claims in order to maximize revenue less the cost of providing liquidity insurance, subject to the collateral constraint. Denote the price of the first-shock claim $p(s)$ and the price of the second-shock claim $q_s(t)$. Also, denote $\phi_1(s)$ as the (true) probability of the first event, and $\phi_s(t)$ as the (true) probability of the second event:

$$\phi_1(s) = \lambda e^{-(\delta+2\lambda)s} \quad , \quad \phi_s(t) = \lambda^2 e^{-\lambda s} e^{-(\delta+\lambda)t} .$$

We assume that the insurer has no robustness concerns and knows these true probabilities, but it turns out that for most of the conclusions we derive, this is without loss of generality. Then the intermediary solves,

$$\max_{\{\hat{x}(s)\}, \{\hat{y}(s)\}} \int_0^\infty \left((p(s) - \beta\phi_1(s))\hat{x}(s) + \int_s^\infty (q_s(t) - \beta\phi_s(t))\hat{y}_s(t)dt \right) ds \quad (3)$$

subject to,

$$\hat{x}(s) \leq Z \quad (4)$$

and

$$\hat{x}(s) + \hat{y}_s(t) \leq Z. \quad (5)$$

Constraints (4) and (5) are standard collateral constraints that agents impose on the intermediary. Consider an agent looking to purchase first-shock insurance. We assume that this agent can observe the balance sheet of the intermediary, i.e., the collateral levels of the intermediary and the claims the intermediary has sold to other agents. The agent only purchases first-shock insurance if the intermediary has sufficient collateral, Z , to cover all of his liabilities, $\hat{x}(s)$. This gives the constraint (4).

On the other hand, an agent looking to purchase second-shock insurance is concerned about both the first- and second-shock insurance that the insurer has sold. Since second-shock insurance is paid only after first-shock claims have been settled, the effective collateral that backs up second-shock insurance is $Z - \hat{x}(s)$. The second-shock claim buyers require that the effective collateral of the insurer covers all second-shock insurance sold. Thus, we arrive at (5).

Because of the linearity of the objective function, if the collateral constraints do not bind the solution to problem (3) yields that $\hat{x}(s)$ and $\hat{y}_s(t)$ are at an interior only if,

$$\begin{aligned} p(s) &= \beta\phi_1(s) \\ q_s(t) &= \beta\phi_s(t). \end{aligned}$$

For example, when $\beta = 1$, prices are actuarially fair at an interior solution.

We are interested in the case when insurance is limited by the insurer's collateral. In this case, assuming the second constraint binds in equilibrium for all (s, t) pairs, and denoting by $\mu_s(t) > 0$ the corresponding shadow values, the first order conditions yield:

$$\begin{aligned} q_s(t) &= \beta\phi_s(t) + \mu_s(t), \\ p(s) &= \beta\phi_1(s) + \mu(s); \quad \mu(s) \equiv \int_s^\infty \mu_s(t)dt. \end{aligned}$$

Prices now rise to reflect the scarcity of collateral ($\mu > 0$). Integrating the expression for $q_s(t)$ to obtain the price of a second-shock claims for all $t > s$, gives:

$$q(s) \equiv \int_s^\infty q_s(t)dt = \beta\phi_1(s)\frac{\lambda}{\lambda + \delta} + \mu(s).$$

The first term on the right hand side reflects the probability of receiving two shocks. This probability is equal to the probability of receiving a single shock, $\phi_1(s)$, multiplied by the conditional probability of receiving one more shock prior to the end of the liquidity episode, $\frac{\lambda}{\lambda+\delta}$. The second term on the right hand side is the multiplier on the collateral constraint, $\mu(s)$.

An important point to note about the expressions for $q(s)$ and $p(s)$ is that the multiplier is common to both and does not depend on the probability of each of these events. That is, a tight collateral constraint pushes up the price of contingent claims *regardless* of the probability of the insurance-event. Intuitively, this is because the collateral constraint requires that if $\hat{x}(s)$ is high, then insurance on the second shock is curtailed for all $t > s$. Thus the opportunity cost of writing more first-shock claims is that less second-shock claims can be offered. Both claim prices have to be in line because they both reflect a common cost of locking collateral. We can make this point more precise by solving out $\mu(s)$ and writing $p(s)$ in terms of $q(s)$:

$$p(s) = q(s) + \beta\phi_1(s) \left(1 - \frac{\lambda}{\lambda + \delta}\right). \quad (6)$$

The common cost of locking collateral arises because of the sequential shock structure of the economy. The second shock only occurs after the first shock has occurred. If the $\hat{x}(s)$ and $\hat{y}(s)$ claims corresponded to two mutually exclusive events, then the intermediary could use the same collateral to back liquidity insurance in each event. In such case, selling more $\hat{y}(s)$ claims does not impinge on the intermediary's ability to sell $\hat{x}(s)$ claims, so that the prices of claims need not be aligned.

We are most interested in the opportunity cost of using collateral, as opposed to the actuarial cost of writing liquidity insurance. To simplify some of our expressions, we assume for now that $\beta = 0$ and thereby drop the second term on the right-hand-side of (6) to yield:

$$p(s) = q(s). \quad (7)$$

This pricing expression is the main result from modelling the supply side of the economy. It captures the equilibrium cost of “locking” collateral.

3.2 A benchmark

Before turning to the agent's decision problem, we pause to establish a “benchmark” for the model when there are no robustness concerns. In this case, we can look directly at the planner's problem:

$$\max_{\{x(s)\}, \{y_s(t)\}} \frac{\delta}{\delta + \lambda} u(\Pi) + \int_0^\infty \phi_1(s) \left[u(x(s)) + \lambda \int_s^\infty e^{-(\lambda+\delta)(t-s)} u(y_s(t)) dt \right] ds$$

In this objective, $x(s)$ is the amount of liquidity insurance delivered to the first agent who receives a shock, while $y_s(t)$ is the amount of insurance delivered to the second agent who receives a shock. The first term in the objective is the expected utility in the event that no shocks arrive and agents consume Π . This term is

invariant to the $x(s)$ and $y_s(t)$ decisions, and where convenient we omit it.⁵ The second term is the expected utility from the first shock arriving at date s , and the second shock occurring at a date t prior to (exit) date τ . The expectation is taken at date 0 over the probability of the first shock occurring at s .

The constraints in the optimization are:

$$0 \leq x(s) \leq Z$$

and

$$0 \leq x(s) + y_s(t) \leq Z.$$

Since $\beta = 0$, the only limitation on liquidity insurance delivered to the agents is the supply of collateral of financial intermediaries. Thus the planner simply allocates the limited collateral across time and agents.

The second collateral constraint has to bind since it subsumes the first constraint, and there is no other cost of providing insurance than the collateral limitation.

$$y_s(t) \equiv y(s) = Z - x(s).$$

Thus we find that $y_s(t)$ allocations are not a function of the time interval between the first and second shock, t . The expression also reflects the cost of locking collateral that we referred to earlier. Z can either be used to back $x(s)$ insurance, or to back $y_s(t)$ insurance, but not both. In this sense, offering more of one kind of insurance restricts the amount of the other kind.

We can then rewrite the planner's problem as:

$$\max_{\{x(s)\}, \{y(s)\}} \frac{\delta}{\delta + \lambda} u(\Pi) + \int_0^\infty \phi_1(s) \left[u(x(s)) + \frac{\lambda}{\lambda + \delta} u(y(s)) \right] ds.$$

The term in front of $u(y(s))$ is the probability of the economy receiving a second shock before the end of the liquidity episode, conditional on having already received one shock. This probability is less than one since there is a chance that the liquidity episode ends ($\delta > 0$) before the second shock takes place.

We can solve the problem by maximizing pointwise for each s .

Proposition 1 *In the **benchmark economy**, the first order conditions characterizing insurance decisions are, for each time s ,*

$$u'(x(s)) = \frac{\lambda}{\lambda + \delta} u'(Z - x(s)) \quad (8)$$

For the log utility case, the quantity of insurance claims are:

$$x^{be} = \left(\frac{\lambda + \delta}{2\lambda + \delta} \right) Z \quad y^{be} = \left(\frac{\lambda}{2\lambda + \delta} \right) Z;$$

and hence satisfy:

$$\frac{x^{be}}{y^{be}} = \frac{\lambda + \delta}{\lambda}.$$

⁵We also omit a 2 that should multiply the objective function throughout to weigh the two groups.

In words, once an agent is hit by a liquidity shock, the planner has an incentive to allocate more than half of the resources to the agent in distress. This is because as long as $\delta > 0$ there is a chance that the second agent never gets hit and scarce liquidity goes wasted.

We note that $x(s)$ is constant and increasing with respect to δ . We are interested in liquidity crisis episodes where δ is large (temporary phenomena). In these cases, the benchmark reveals a strong incentive to inject funds aggressively to the first group of agents hit by a shock.

In the benchmark competitive equilibrium without robustness considerations, the incentive for agents to purchase more x than y claims follows directly from the pricing schedule of a collateral constrained intermediary that we derived in the first part of this section. As δ rises, the single-shock event becomes more probable than the two-shock event. However, because the prices of both first and second shock financial claims reflect a common cost of locking collateral, the prices of both claims remain equal, favoring the purchase of the first-shock claim over the second-shock claim.

3.3 The agents' decision problem

An agent type A (henceforth agent A , for short) solves a decision problem where he chooses the paths of $x(s)$ and $y_s(t)$, attempting to be robust to values of λ_B . In this subsection, we characterize the optimal choice of $x(s)$ and $y_s(t)$ given a particular value of λ_B .

We define the expected utility function for A given choices of $\{x(s)\}$ and $\{y_s(t)\}$ and some value of λ_B :

$$V(\{x(s)\}, \{y_s(t)\}, \lambda_B) = \int_0^\infty \left[\phi_1^A(s)u(x(s)) + \int_s^\infty \phi_s^A(t)u(y_s(t)) dt \right] ds$$

where, from the exponential distribution function, the probabilities (as perceived by A) are:

$$\begin{aligned} \phi_1^A(s) &\equiv \text{P(First shock is A's, at } s) = \lambda e^{-(\delta+\lambda+\lambda_B)s}, \\ \phi_s^A(t) &\equiv \text{P(Second shock is A's, at } t>s) = \lambda \lambda_B e^{-\lambda_B s} e^{-(\delta+\lambda)t} = \lambda_B e^{-(\delta+\lambda)(t-s)} \phi_1^A(s). \end{aligned}$$

The agent maximizes $V(\cdot)$ subject to his budget constraint:

$$\max_{\{x(s) \geq 0\}, \{y_s(t) \geq 0\}} V(\{x(s)\}, \{y_s(t)\}, \lambda_B) \quad \text{s.t.} \quad \int_0^\infty \left(p(s)x(s) + \int_s^\infty q_s(t)y_s(t) dt \right) ds \leq w_0 \quad (9)$$

The budget constraint reflects purchase of $x(s)$ and $y_s(t)$, given prices of $p(s)$ and $q_s(t)$ and the initial endowment of w_0 . The first order conditions for the optimization are:

$$\phi_1^A(s)u'(x(s)) = \psi p(s) \quad (10)$$

and

$$\phi_s^A(t)u'(y_s(t)) = \psi q_s(t), \quad (11)$$

where ψ is the Lagrange multiplier on the budget constraint.

We can simplify the characterization of the agent's decision using the result from our earlier analysis of the intermediary's problem. There we found that $p(s)$ and $q_s(t)$ are related to each other by a common cost of locking collateral. Integrating both sides of (11) with respect to t and using equation (7), we find that,

$$\lambda_B \phi_1^A(s) \int_s^\infty e^{-(\lambda+\delta)(t-s)} u'(y_s(t)) dt = \psi p(s)$$

Then, combining this relation with (10), we can eliminate $p(s)$ to find:

$$u'(x(s)) = \lambda_B \int_s^\infty e^{-(\lambda+\delta)(t-s)} u'(y_s(t)) dt \quad (12)$$

We can simplify this relation considerably by using the knowledge that in all equilibria, $y_s(t) = Z - x(s)$, so that $y_s(t)$ is only a function of s and not of t . Since all insurance has to be backed by the collateral of Z , for each s , the agent's insurance decisions sum to Z , in equilibrium.⁶

Since the probability of receiving a second shock prior to the end of the liquidity episode conditional on having received a first shock at date s is equal to $\frac{\lambda_B}{\delta+\lambda}$, we can simplify the right-hand side of (12), to find:

Proposition 2 *The solution to the agent's decision problem, (9), given a value of λ_B is characterized by the relation:*

$$u'(x(s)) = \frac{\lambda_B}{\delta + \lambda} u'(y(s)) \quad (13)$$

which for log utility implies:

$$\frac{x(s)}{y(s)} = \frac{\delta + \lambda}{\lambda_B}.$$

As one would expect, the solution to the agent's decision problem coincides with the solution in the benchmark, (8), when $\lambda_B = \lambda$.⁷ More generally, for $\lambda_B < \lambda + \delta$ agents insure more against the first-shock than the second-shock. Note, however, that the ratio of x to y claims is decreasing with respect to λ_B : The more likely the other agent is expected to be hit by a shock, the more agent A reallocates demand to

⁶Note that we are not imposing the requirement that $y_s(t) = Z - x(s)$ in agents' decisions. Instead, we are using the knowledge that since in equilibrium, $y_s(t) = Z - x(s)$, only certain price functionals are possible. For a given λ_B , all feasible equilibrium price functionals, $q_s(t)$, must take the form (with respect to t),

$$q_s(t) = c(s) e^{-(\lambda+\delta)(t-s)},$$

for some $c(s)$. From the first order condition, (11), and using the relation

$$\int_s^\infty \phi_s^A(t) ds = \frac{\lambda_B}{\delta + \lambda} \phi_1^A(s),$$

we conclude that, $y_s(t) = y(s)$, regardless of λ_B .

⁷Given equilibrium prices, $p(s)$ and $q(s)$, from (10) and (11), as well as the intertemporal budget constraint we find that:

$$x(s) = \frac{w_0}{p(s)} (\lambda + \delta) e^{-(\delta+\lambda_B+\lambda)s} \quad (14)$$

$$y(s) = \frac{w_0}{q(s)} \lambda_B e^{-(\delta+\lambda_B+\lambda)s}. \quad (15)$$

(Note that the multiplier on the budget constraint is $\psi = \frac{1}{w_0} \frac{\lambda}{\delta+\lambda}$ for the case of log utility.)

second-shock claims. Thereby ensuring that he still receives payments if he is hit by a liquidity shock as well.

3.4 Robustness and equilibrium

We now turn to the robustness step. In equation (13) we characterized the decision problem for agent A in choosing $x(s)$ relative to $y_s(t)$, given a particular value of λ_B . These choices defined an expected utility function $V(\{x(s)\}, \{y_s(t)\}, \lambda_B)$. The robustness step for agent A is to make these choices while being robust to alternative values of λ_B .

In the game-theoretic language often used to describe max-min expected utility theory, agent A first chooses $x(s)$ and $y_s(t)$, then “nature” chooses $\lambda_B \in \Lambda$ to minimize the utility of agent A, given his choices. The full problem for agent A is:

$$\max_{\{x(s) \geq 0\}, \{y_s(t) \geq 0\}} \min_{\lambda_B \in \Lambda} V(\{x(s)\}, \{y_s(t)\}, \lambda_B) \quad (16)$$

given equilibrium prices $p(s)$ and $q(s)$ and the budget constraint. We assume the set of alternative models over which agents would like to be robust is an interval defined by:

$$\max\{0, \lambda - K\} \leq \lambda_B \leq \lambda + K.$$

K indexes the robustness preferences of agents, with a larger K corresponding to a more extreme worst-case scenario.

The following proposition presents the main result of this section:

Proposition 3 *The following insurance decisions constitute an equilibrium in the **robust economy**:*

- For $K < \delta$, agents’ decisions are as if $\lambda_B = \lambda + K$:

$$\frac{x^{pr}}{y^{pr}} = \frac{\lambda + \delta}{\lambda + K}.$$

We refer to this case as the “partially robust” case.

- For $K \geq \delta$, agents’ decisions are as if $\lambda_B = \lambda + \delta$:

$$\frac{x^{fr}}{y^{fr}} = 1.$$

We refer to this equal claims case as the “fully robust” case.

- In both cases, agents’ decisions are time invariant.

When $K = 0$, the model collapses to the benchmark model without the robustness concern, and decisions are as if $\lambda_B = \lambda$. As K rises, agents become concerned that they will receive a shock second, in which case their larger purchases of first-shock insurance will be wasted. As a result, they reduce their purchases of first-shock insurance and increase their purchases of second-shock insurance. At the extreme, when K is large enough, they equate their purchases of first and second-shock claims, and thereby insulate themselves against their uncertainty over the likelihood of market B receiving a shock.

The formal proof reflects the intuition we have provided, but is complicated by two issues. First, agents make their insurance decisions anticipating that the worst case λ_B will depend on those decisions; and, second, equilibrium prices are a function of the anticipated worst case of agents (but in an individual agent's decision problem, prices are taken as given).

Proof. First, we derive the prices in the proposed equilibrium. Denote the λ_B on which agents base decisions in equilibrium as $\tilde{\lambda}$. The market clearing conditions imply,

$$\begin{aligned}x^A(s) + y_s^B(t) &= Z, \\x^B(s) + y_s^A(t) &= Z,\end{aligned}$$

where we have used the fact that the collateral availability of Z determines the supply of liquidity insurance.⁸ Since the relative demand satisfies

$$\frac{x(s)}{y(s)} = \frac{\lambda + \delta}{\tilde{\lambda}},$$

we can substitute to solve out for $x(s)$ and $y(s)$. To find prices, we substitute the derived quantities into the first order conditions for agent A, (10) and (11), to yield:

$$p(s) = \frac{w_0}{Z}(\lambda + \delta + \tilde{\lambda})e^{-(\lambda + \delta + \tilde{\lambda})s} \quad (17)$$

and,

$$q_s(t) = \frac{w_0}{Z}(\lambda + \delta + \tilde{\lambda})(\lambda + \delta)e^{-(\lambda + \delta + \tilde{\lambda})s - (\lambda + \delta)(t - s)} \quad (18)$$

Small K : Consider the robustness step next. For small K , we wish to show that, at equilibrium prices, the highest utility attainable when solving,

$$V = \max_{\{x(s) \geq 0\}, \{y_s(t) \geq 0\}} \min_{\lambda_B \in \Lambda} V(\{x(s)\}, \{y_s(t)\}, \lambda_B) \quad (19)$$

occurs at x^{pr} and y^{pr} , and given these choices by the agent, “nature” always chooses $\lambda_B = \lambda + K$. Moreover, no other choice of $\{x(s)\}$ and $\{y_s(t)\}$ can induce nature to choose a λ_B that results in a higher utility for the agent.

⁸If agent A receives a shock first, then his second-shock insurance claim disappears. The same applies to B. Thus, the insurer economizes his limited collateral by only backing the two possible shock sequences of A then B , and B then A .

Given an assumed worst case of $\lambda + K$, and equilibrium prices (17) and (18), the unique optimum to the agents max problem, (9), is:

$$x(s) = \frac{\lambda + \delta}{\lambda + \delta + \tilde{\lambda}} e^{-(\lambda_B - \tilde{\lambda})s} Z \quad \text{and} \quad y(s) \equiv y_s(t) = \frac{\lambda_B}{\lambda + \delta + \tilde{\lambda}} e^{-(\lambda_B - \tilde{\lambda})s} Z$$

At $\tilde{\lambda} = \lambda_B$, $x(s)$ and $y(s)$ are constant functions of time. Using this knowledge, we can derive the expected utility function for the agent as,

$$V(x^{pr}, y^{pr}, \lambda_B) = \frac{\lambda}{\delta + \lambda + \lambda_B} \left(u(x^{pr}) + \frac{\lambda_B}{\lambda + \delta} u(y^{pr}) \right), \quad (20)$$

and using the envelope theorem, we find that,

$$V_{\lambda_B} = \frac{\lambda}{(\delta + \lambda + \lambda_B)^2} [u(y^{pr}) - u(x^{pr})]$$

which is strictly negative around the benchmark ($\lambda_B = \lambda$) and as long as $x^{pr} > y^{pr}$. Thus, since for small values of K , $x^{pr} > y^{pr}$, it is apparent that the agent is particularly unprotected if nature chooses to increase λ_B , since such a choice makes first-shock claims less valuable to the agent, and the agent is over-weight in holding these claims. Moreover, since the constraint set in choosing $\lambda \in \Lambda$ is linear, for small enough K (i.e. the agent considers only small deviations in the model), the utility of the agent at choices (x^{pr}, y^{pr}) is minimized when $\lambda_B = \lambda + K$ (i.e. highest value possible).

Next, define $V^{pr} \equiv V(x^{pr}, y^{pr}, \lambda_B = \lambda + K)$. We now argue, by contradiction, that V^{pr} is the highest utility attainable in (19). Suppose there exist choices of $\{x(s)\}$ and $\{y_s(t)\}$ which result in $V > V^{pr}$. It should be clear that such choices can only increase utility if they induce nature to choose $\lambda_B < \lambda + K$. Note that since the optimization problem in (9) is strictly concave, the first order conditions define a unique optimum. If the agent makes any choice other than x^{pr} and y^{pr} and nature continues to choose a worst case of $\lambda_B = \lambda + K$, then the induced utility is strictly lower than V^{pr} . Moreover, since nature can always choose $\lambda_B = \lambda + K$ given the agent's deviation, nature's optimal choice of λ_B further lowers the utility of the agent. Thus, the agent can do no better than to choose x^{pr} and y^{pr} .

Large K : For large values of K , the gap between first-shock and second-shock claims narrows until it disappears once $K = \delta$. At this point, the agent becomes fully robust with respect to any λ_B , since when $x = y$,

$$V_{\lambda_B} = 0$$

for any λ_B . At $K = \delta$, the agent attains utility of $V^{fr} \equiv V(\frac{Z}{2}, \frac{Z}{2}, \lambda_B = \lambda + \delta)$. Further increases in K expand the constraint set for nature, and thereby weakly decrease the highest utility attainable in (19). However, the agent can guarantee himself V^{fr} by choosing $x = y = \frac{Z}{2}$. Thus, the fully robust choices continue to attain the highest possible utility in (19) for all $K \geq \delta$. ■

4 Flight-to-Quality Episodes

We now discuss the connection between our model and flight-to-quality episodes. As noted in the introduction, two central aspects of flight-to-quality episodes are the demand for certainty by agents and the capital bottlenecks that develop in the financial system. The latter transforms flight to quality into a macroeconomic problem. We discuss these points below.

4.1 Demand for certainty

In our model, agents effectively hold $\min\{x, y\}$ units of “safe financial claims.” These claims deliver resources to agents contingent on their own liquidity shock, but *not* on the other agents’ liquidity shocks. Moreover, since $x \geq y$ in all equilibria, the agent’s guaranteed insurance is:

$$y = \min\{x, y\}.$$

In addition to the minimum level of guaranteed insurance, agents hold $x - y$ of liquidity insurance that only delivers resources to the first market that receives a shock. These latter claims are thereby contingent on the intermediary having capital available to deliver to the agent.

As λ_B rises, y rises and $x - y$ falls. Agents demand more financial claims that are not contingent on the event they do not understand, and reduce risky claims.

The essence of a flight-to-quality episode is agents’ demand for certainty. The empirical counterpart of the demand for uncontingent insurance as λ_B rises is that agents may shift their demands toward extremely well collateralized financial institutions. With a less well collateralized institution, an agent fears that his financial claim is contingent on the other markets’ shocks. Thus our model is consistent with the increased demand for bank deposits and bank credit lines that accompanies flight to quality (see Gatev and Strahan, 2005).

In Section 7 we show that these actions by the agents trigger a reaction by insurers to ensure the safety of their own collateral. Insurers increase their valuation of riskless assets relative to risky assets, which is consistent with the widening of the spread between Commercial Paper and Treasury bills that accompanies flight to quality.

Thus a rise in λ_B triggers increased demand for safe assets. In practice, most accounts of flight to quality episodes begin with an unexpected event that leads agents to reassess their models. In 1970, the unexpected event was the Penn Central default on Commercial Paper. In 1987, it was the sharp drop in the stock market. In terms of our model, a reassessment of models can be thought of as a rise in K (for $K < \delta$), which then leads to a rise in λ_B .

For a fixed K , an alternative way to generate flight to quality in our model is to reduce Z . We explain this case in more detail in Section 6, but the intuition is clear enough. A fall in Z increases financial intermediary

risk as agents consider that intermediaries may not have enough collateral to cover their shocks. Agents become more concerned about being second and raise λ_B . The Fall of 1998 episode is a good example of a flight to quality through this sort of financial intermediary risk. As agents grew concerned about the capital levels of intermediaries, they withdrew their investments with complex intermediaries such as hedge funds or investment banks and retrenched to safe investments in Treasury bills or commercial banks.

4.2 Aggregate liquidity provision

We can substitute the equilibrium values of $x(s)$ and $y(s)$ into the first order conditions for agent A, (10) and (11), to find the prices of first and second-shock liquidity insurance:

$$p(s) = \frac{w_0}{Z}(\lambda + \delta + \lambda_B)e^{-(\lambda + \delta + \lambda_B)s} \quad (21)$$

and,

$$q_s(t) = \frac{w_0}{Z}(\lambda + \delta + \lambda_B)(\lambda + \delta)e^{-(\lambda + \delta + \lambda_B)s - (\lambda + \delta)(t-s)} \quad (22)$$

As λ_B rises, equilibrium prices of insurance change as well. For small s , the exponential term in (21) and (22) is dominated by the linear term so that the increase in K increases the prices of both claims:⁹

$$p(0) = q(0) = \frac{w_0}{Z}(\lambda + \delta + \lambda_B).$$

The increase in prices of the first-shock insurance contract is indirect, driven by increased demand for second-shock insurance. As we have shown earlier, when λ_B rises, robust agents decrease their demand for first-shock insurance. They increase demand for second-shock insurance, causing insurers to lock-up collateral for y and decrease the supply available to first-shock insurance, thereby indirectly driving up the price of first-shock insurance.

The prices, p and q , represent the marginal cost of liquidity provision in the economy. Since intermediaries have a limited amount of collateral to back liquidity provision, at the aggregate level, Z parameterizes the capacity of the economy to provide liquidity to markets. As robustness rises, the cost of liquidity provision rises because the actions of robust agents reduce the effective Z of the economy. With a richer asset structure, the rise in the cost of liquidity provision will also be reflected as higher liquidity premia in the asset markets where the intermediaries are active.

Suppose we introduce a small number of rational Savage agents into the economy. These agents, as in the benchmark economy, will calculate that it is optimal to insure more against the first-shock as opposed

⁹Because the model has only two shocks, for s sufficiently large, the increase in λ_B lessens the odds that the agent places on surviving until date s without a shock. In these cases, the exponential term in (21) and (22) dominates and the prices of financial claims fall. The boundary between the small and large s occurs at $s = \frac{1}{3\lambda + \delta}$; for δ large relative to λ , the boundary is at $\frac{1}{\delta}$ which is the half-life of the recessionary shock.

to the second-shock. However, notice that these agents will have to trade at the same inflated (for low s) liquidity insurance prices that the robust agents face. At these high prices, the non-robust agents will also reduce the amount of insurance they purchase, and leave themselves exposed to liquidity shocks.¹⁰

In any flight to quality episode, the capacity of financial intermediaries – commercial and investment banks, hedge funds, trading desks, market makers, etc. – to provide liquidity to markets is constrained. As a result, liquidity falls across many markets during these episodes. Market participants, including both uncertainty averse and Savage agents, have to deal with the adverse liquidity conditions.

At the abstract level of our model, Z reflects the capacity of all liquidity providers, in aggregate, to provide liquidity to markets. Our model does not distinguish among the various modes of liquidity provision we observe in practice. But an important aspect of flight to quality episodes is that the natural liquidity providers become liquidity demanders. In the crises of the Fall of 1998, hedge funds were retrenching from many markets and liquidating assets, thereby reducing the overall supply of liquidity (see Kyle and Xiong, 2001). We may think of the robustness concern in our model as corresponding in practice to hedge funds' (or their investors') fears that market liquidity will fall, causing the hedge funds to overprotect against adverse outcomes, further reducing market liquidity. Our model illustrates how Knightian uncertainty can lead agents to reduce the effective Z of the economy and compound an aggregate liquidity shortage.

4.3 Financial bottlenecks and collateral waste

In practice, the financial system is severely strained during flight-to-quality episodes. Bottlenecks arise in the movement of capital and markets appear segmented.

The fully robust equilibrium in our model, where $x = y = Z/2$, generates complete market segmentation. We can think of this equilibrium as follows: Agent A requires the insurer to keep $Z/2$ collateral dedicated to the shocks of market A , while agent B does likewise for market B . Insurers create two separately capitalized financial intermediaries, each dedicated to serving a particular market. Collateral is locked up in each market and is unable to move across markets in response to shocks, causing markets to become segmented.¹¹

Segmentation is the market's response to agents' robustness concerns. Since agents require financial claims that are uncontingent on shocks outside their own markets, the market provides such claims by committing to segment their capital.

But, segmentation has a macroeconomic cost. When agents lock-up resources in their desire for certainty, the financial system becomes inflexible in its response to shocks. In the benchmark economy, $x^{be} > y^{be}$, so

¹⁰Recall that the change in prices depends on s . For large s , prices fall since the rise in λ_B leads to a decline in the perceived probability of shocks taking place. Hence rational Savage agents lower early insurance (both of $x(s)$ and $y(s)$) as a result of the presence of robust agents (the point we have highlighted) but increase late insurance.

¹¹We may also interpret y as the intermediary's committed capital to one market, and $x - y$ as the trading capital available across both markets.

that $x^{be} - y^{be}$ of collateral is left to promise to the first market that receives a shock. The freedom is valuable since as long as $\delta > 0$, there is a good chance that the second shock never takes place.

An alternative way to see the collateral waste associated to a flight-to-quality episode is to re-write the constraint

$$x + y \leq Z$$

in terms of the risky and riskless claims Z can support:

$$(x - y) + 2y \leq Z.$$

The riskless claim consumes twice as much collateral as a risky claim. Having all agents holding only riskless claims, as in the fully robust equilibrium, locks up collateral that goes unused when $\delta > 0$.

5 Locked Collateral and Extreme Event Intervention

In this section we study the impact of a lender of last resort on the economy. We consider a central bank that obtains resources ex-post, at some cost, which it can credibly pledge to agents in the two-shock extreme event. In practice, this pledge may be supported by costly ex-post inflation or taxation and carried out by guaranteeing, against default, the liabilities of financial intermediaries who have sold financial claims to both markets. We analyze the impact and marginal benefit of such a guarantee.

Formally, we assume that the government credibly expands the collateral of the financial sector in the two-shock event by an amount G . Thus, the collateral constraints on intermediaries are altered to

$$x + y \leq Z + G.$$

Since we are interested in computing the marginal benefit of intervention, we study an infinitesimal intervention of G . The intervention has no effect on the constraint for first-shock insurance.

5.1 Welfare criterion

A delicate issue arises when agents have non-Savage preferences: In evaluating policy, should the central bank also adopt the agents' preferences? Or, should it take a more paternalistic approach in which it is concerned with agents' ex-post average utility from consumption?

As noted earlier in the paper, we think of the robustness preferences as a realistic depiction of the decision rules of financial specialists (e.g., value-at-risk constraints). From this perspective, it is not clear why a central bank should build biases into its objective function that may lead to obvious average losses, just because agents exhibit these biases. Sims (2001) has made this point in questioning the application

of robust control to central bank policymaking. He argues that max-min preferences are simply shortcuts to generate observed behavior of economic agents, but should not be seen as deeper preferences. In what follows, we begin by analyzing a welfare function based on agents' expected utility functions, and then turn to alternative, more paternalistic, welfare criteria.

5.2 Robust Expected utility

Consider first the case where the central bank uses the agents' expected utility functions to evaluate welfare. In this case, the envelope theorem tells us that the indirect effect due to the change of insurance decisions is second-order, and that the only benefit is the direct effect.

In our model, we have derived that the expected utility for an agent at the optimal choices of x and y (see equation (20)) is,

$$V(x, y, \lambda_B) = \frac{\lambda}{\delta + \lambda + \lambda_B} \left(u(x) + \frac{\lambda_B}{\lambda + \delta} u(y) \right)$$

Offering more insurance against the second-shock leads agents to increase both x and y . The effect on y is direct, whereas the effect on x is an indirect effect of reoptimization of agents' portfolios. However, using the envelope theorem, the utility gain from an extra unit of Z can be computed assuming all of the Z is spent increasing y :

$$V_G = \frac{\lambda}{\delta + \lambda + \lambda_B} \frac{\lambda_B}{\lambda + \delta} u'(y) = \frac{\lambda}{\delta + \lambda} \frac{1}{Z}, \quad (23)$$

since $y = \frac{\lambda_B}{\lambda + \delta + \lambda_B} Z$.

While there is a welfare benefit of providing a government guarantee, the value of the benefit is substantially the same in the benchmark and the robust economies.¹² In both cases, the economy lacks collateral to fully insure agents against shocks. The benefit of the extra unit of Z is limited to the direct effect of providing agents more insurance against an adverse outcome. It stems from our assumption that the government can create liquidity to insure a shock from sources that the private sector cannot, as in Woodford (1990) and Holmstrom and Tirole (1998). As we show next, this logic does not apply if we consider alternative welfare criteria.

¹²Recall that a liquidity shock eliminates all profits. We have explored a setup where the shock reduces profits to a level of $\pi > 0$. We find that in this case the benefit of the lender of last resort policy may be greater in the robust than in the benchmark economy, even if agents' expected utility is the welfare criterion. When π is sufficient large, agents in the benchmark economy may eschew insurance against the extreme two-shock event in favor of increasing insurance against the single-shock event. In contrast, agents in the robust economy always insure against the two-shock event. Thus robust agents further value insurance against an extreme event, while agents in the benchmark economy have limited (expected) valuation for such insurance. Notice that here the benefit of the policy comes from the agent's valuation of the extra insurance rather than from the reorganization of its portfolio in the presence of the policy.

5.3 Fully informed central bank

Consider a welfare criterion based on agent's decisions, but evaluated at the reference λ 's. In this case the envelope theorem argument breaks down for the robust economy and the reallocation of private insurance from second-shock to first-shock insurance has first order effects. Since in equilibrium the agent "exaggerates" the likelihood of a two-shock event relative to a single-shock event, the reallocation of resources toward first-shock insurance has a first order benefit.

The expected value of the *direct* effect of intervention for the central bank is:

$$\left(\frac{\lambda}{2\lambda + \delta} \frac{\lambda}{\lambda + \delta} \right) u'(y), \quad (24)$$

where the term in parentheses is the probability the central bank assigns to a two-shock event at the time of commitment (date 0).

The total (direct plus indirect) expected value of the government intervention is:

$$\tilde{V}_G = \frac{\lambda}{2\lambda + \delta} \left(u'(x) \frac{dx}{dG} + \frac{\lambda}{\lambda + \delta} u'(y) \frac{dy}{dG} \right). \quad (25)$$

where \tilde{V}_G denotes the expected value of intervention evaluated using the central banks's probability distribution.

The first order condition for agents' optimization is:

$$u'(x) = \frac{\lambda^B}{\lambda + \delta} u'(y)$$

Using this relation, as well as the constraint that $\frac{dx}{dG} + \frac{dy}{dG} = 1$, we find,

$$\tilde{V}_G = \frac{\lambda}{2\lambda + \delta} \frac{\lambda}{\lambda + \delta} u'(y) \left(1 + \left(\frac{\lambda^B}{\lambda} - 1 \right) \frac{dx}{dG} \right) \quad (26)$$

Finally, we divide this total effect of intervention by the direct effect of intervention to construct a *private sector multiplier* of the lender of last resort commitment:

$$\mathbf{M} = \left(1 + \left(\frac{\lambda^B}{\lambda} - 1 \right) \frac{dx}{dG} \right). \quad (27)$$

Proposition 4 *When $\lambda^B = \lambda$ (benchmark economy), the total benefit of intervention is limited to the direct benefit and*

$$\mathbf{M}^{be} = 1.$$

For $\lambda_B > \lambda$, there is an indirect benefit stemming from the reoptimization of agents' portfolios:

$$\mathbf{M}^r > 1.$$

Intuitively, the benefit of intervention in the robust economy stems from the reaction of the private sector to the central bank's guarantee. When agents make their insurance decisions, they are overly concerned about receiving the shock second, and use too high an assessment of the likelihood of the two-shock event. This breaks the envelope theorem argument under the central banks' objective function. By offering to insure the two-shock event directly, the central bank leads the agent to free up resources to insure the more likely one-shock event ($\frac{dx}{dG} > 0$).

5.4 Partially informed central bank

Our assumption that the central bank knows the true λ 's of the economy, despite the fact that agents may be unsure about these λ 's, is a strong requirement. In this section, we consider a weaker informational requirement on the central bank.

Agent A knows that $\lambda_B \in [\max\{\lambda - K, 0\}, \lambda + K]$ (and likewise for agent B). We suppose that the central bank is also uncertain about the values of the λ 's, and only knows that the λ 's are drawn from some non-degenerate symmetric joint distribution $F(\lambda_A^{cb}, \lambda_B^{cb})$. We study the case where $K > \delta$ (i.e. the fully robust economy).¹³

Because of symmetry in $F(\cdot)$, it is straightforward to show that the central bank restricts attention to symmetric interventions, i.e. interventions that are equal regardless of whether it is A or B that is hit first, or second.

In this case we define the multiplier as:

$$\mathbf{M} \equiv \frac{E_{\lambda_A \lambda_B}[\tilde{V}_G(\lambda_A, \lambda_B)]}{E_{\lambda_A \lambda_B}[DE(\lambda_A, \lambda_B)]},$$

where $\tilde{V}_G(\lambda_A, \lambda_B)$ and $DE(\lambda_A, \lambda_B)$ respectively stand for the total and direct expected effect of the intervention conditional on λ_A and λ_B .

Then, for the fully robust economy,

$$E_{\lambda_A \lambda_B}[\tilde{V}_G(\lambda_A, \lambda_B)] = \frac{1}{2}u'(Z/2)E_{\lambda_A \lambda_B}[\Pr(A; \lambda_A) + \Pr(B; \lambda_B)],$$

where $\Pr(i; \lambda_i)$ denotes the probability that agent i is hit by a shock given its intensity λ_i . In the fully robust economy, the promised resources in the two-shock event is split equally across the first and second shock event. Thus the marginal value of resources in the shock is $u'(Z/2)$ times one-half. This value is weighted by the probability that shocks may occur, to either agent.

$$E_{\lambda_A \lambda_B}[DE(\lambda_A, \lambda_B)] = u'(Z/2)E_{\lambda_A \lambda_B}[\Pr(A|B; \lambda_A, \lambda_B) + \Pr(B|A; \lambda_A, \lambda_B)],$$

¹³In the fully robust economy, the x and y decisions of agents are independent of λ , which simplifies our analysis because we avoid computing the central bank's expectations over agents insurance decisions.

where $\Pr(i|j; \lambda_A, \lambda_B)$ denotes the probability that agent i is hit second given intensities λ_A and λ_B . For later use, let $\Pr(i|Noj; \lambda_A, \lambda_B)$ denote the probability that agent i is hit by a shock while j is not, given intensities λ_A and λ_B . The direct effect reflects that resources are all spent in the event of the second shock.

Proposition 5 *If the central bank knowledge of the markets' probability models is limited to a non-degenerate symmetric joint distribution $F(\lambda_A^{cb}, \lambda_B^{cb})$ with some mass for finite and positive λ_A^{cb} and λ_B^{cb} , the multiplier in the fully robust economy is,*

$$\mathbf{M}^r > 1.$$

Proof. From the symmetry of $F(\lambda_A^{cb}, \lambda_B^{cb})$, we have that the multiplier can be written as

$$\mathbf{M}^r = \frac{E_{\lambda_A \lambda_B}[\Pr(A; \lambda_A)]}{2E_{\lambda_A \lambda_B}[\Pr(A|B; \lambda_A, \lambda_B)]}.$$

Thus, the proposition holds if:

$$E_{\lambda_A \lambda_B}[\Pr(A; \lambda_A)] > 2E_{\lambda_A \lambda_B}[\Pr(A|B; \lambda_A, \lambda_B)].$$

Expanding the left hand side of this expression, we have that:

$$\begin{aligned} E_{\lambda_A \lambda_B}[\Pr(A; \lambda_A)] &= E_{\lambda_A \lambda_B}[\Pr(A|B; \lambda_A, \lambda_B) + \Pr(A|NoB; \lambda_A, \lambda_B) + \Pr(B|A; \lambda_A, \lambda_B)] \\ &= E_{\lambda_A \lambda_B}[2\Pr(A|B; \lambda_A, \lambda_B) + \Pr(A|NoB; \lambda_A, \lambda_B)]. \end{aligned}$$

Replacing this expression back into the inequality, we have that the latter holds if and only if

$$E_{\lambda_A \lambda_B}[\Pr(A|NoB; \lambda_A, \lambda_B)] > 0,$$

which is satisfied as long as $F(\lambda_A^{cb}, \lambda_B^{cb})$ has some mass for finite λ_B^{cb} and positive λ_A^{cb} since $\delta > 0$. ■

Intuitively, the proof exploits the central bank's knowledge of symmetry (in the distribution). Agent A exaggerates the probability of being second. Agent B does the same. Thus the total benefit of the lender of last resort reflects exaggerated probabilities of agents being second. Despite not knowing the true λ 's, the central bank knows that they are drawn from similar distributions. With a symmetric distribution over λ , the probabilities the central bank uses reflects that the agents cannot both be second. This information affects the central bank's calculation of the direct benefit of intervention. The symmetry information is thereby sufficient to conclude that the multiplier on intervention is greater than one.

5.5 Moral hazard critique

Under the paternalistic view, the public provision of insurance leads to improved private sector insurance decisions. The complementarity between public and private provision of insurance in our model cuts against

the usual moral hazard critique of central bank interventions. This critique is predicated on agents responding to the provision of public insurance by cutting back on their own insurance activities. In our model, in keeping with the moral hazard critique, agents reallocate insurance away from the publicly insured shock. However, when flight to quality is the concern, the reallocation improves (ex-post) outcomes on average.¹⁴

Note that the central bank can achieve the same distribution of insurance if instead it commits to intervene in the first shock. However the expected cost of this policy is much larger than the extreme event intervention, since the central bank rather than the private sector bears the cost of insurance against the (likely) single-shock event. Agents would reallocate the expected resources from the central bank to the two-shock event, which is exactly the opposite of what the central bank wants to achieve. In this sense, interventions in intermediate events are subject to the moral hazard critique. The lender of last resort facility, to be effective and improve private financial markets, has to be a *last* not an intermediate resort.

6 Flight to Quality and Aggregate Collateral

Throughout the paper we have made assumptions so that the economy is always in a situation where aggregate collateral is limited. In this section we relax some of these assumptions and show that in this case a fall in intermediation collateral can trigger a flight to quality.

In this context, we derive a comparative static to contrast the response of the benchmark economy and the robust economy to a tightening of collateral. We find that the response to a collateral tightening is qualitatively similar to the increase in preference for robustness that we have discussed in earlier sections.

The only substantive modification we make in this section to our baseline model is to set $\beta > 0$ so that insurers now face a “consumption” cost of *using* insurance resources. We cannot derive a full analytical solution for this case but present some results characterizing the equilibrium. The section ends with numerical solutions illustrating the equilibria for both benchmark and robust economy.

6.1 Abundant collateral

Suppose first that Z is large enough so that, at equilibrium levels of demand, the financial sector is unconstrained in selling insurance. In this case, the prices of the two financial claims just reflect probabilities:

$$\begin{aligned} p(s) &= \beta\phi_1(s) = \beta\lambda e^{-(\delta+2\lambda)s} \\ q_s(t) &= \beta\phi_s(t) = \beta\lambda^2 e^{-\lambda s - (\lambda+\delta)t} \end{aligned}$$

¹⁴Note that if the direct effect of intervention is insufficient to justify intervention, then the lender of last resort policy is time inconsistent. This result is not surprising as the benefit of the policy comes precisely from the private sector reaction, not from the policy itself.

In the case of robust agents, using the first order conditions (10) and (11) we find that:

$$x(s) = \frac{1}{\beta\psi} e^{-(\lambda_B - \lambda)s} \quad y_s(t) = \frac{1}{\beta\psi} \frac{\lambda_B}{\lambda} e^{-(\lambda_B - \lambda)s}.$$

It is easy to verify that the solution for robust agents is to set $x(s) = y_s(t)$, and equal to a constant for all s .¹⁵ The argument resembles the proof of Proposition 3.

At equal levels of x and y , the agent is robust to any value of λ_B that nature chooses and V_{λ_B} in equation (20) is equal to zero. Picking the point $\lambda_B = \lambda$, we see that choosing the constant levels for x and y is the unique optimum for agents.

If agents are fully informed about λ_B , the same first order conditions, evaluated at $\lambda_B = \lambda$, determine the solution. Intuitively, since prices are actuarially fair, agents choose to insure equally against both the high likelihood one-shock event as well as the unlikely two-shock event.

The solution in the benchmark case is *identical* to the solution in the robust economy. Both economies respond equally to the typical (single shock) event, as well as to the extreme two-shock event.

6.2 Collateral tightening

Let

$$\underline{Z} \equiv 2w_0 \frac{\lambda + \delta}{\beta\lambda}.$$

That is, \underline{Z} represents the level of intermediation collateral which exactly satisfies agents insurance demand when prices are “fair” ($\mu = 0$).

Note at this point that when $\beta = 0$, as in the previous sections, $\underline{Z} \rightarrow \infty$ so that the economy is always in the collateral-constrained region. Henceforth let us specialize to the case with $\beta = 1$ so that when $Z > \underline{Z}$ prices correspond to the standard actuarially fair prices.

Suppose now that Z falls below \underline{Z} . Then, at actuarially fair prices, agents’ demand saturate the collateral constraint of the insurers and prices rise to reflect the collateral tightening:

$$\begin{aligned} p(s) &= \lambda e^{-(\delta+2\lambda)s} + \mu(s) \\ q_s(t) &= \lambda^2 e^{-\lambda s - (\lambda+\delta)t} + \mu_s(t), \end{aligned}$$

where $\mu(s) = \int \mu_s(t) > 0$.

The first order conditions for agents remain as in (10) and (11). For the x decision:

$$\left(\frac{1}{\psi} \lambda e^{-(\delta+\lambda+\lambda_B)s} \right) u'(x(s)) = \lambda e^{-(\delta+2\lambda)s} + \mu(s) \quad (28)$$

while for the $y_s(t)$ decision:

$$\frac{1}{\psi} \lambda e^{-\lambda_B s} \lambda_B e^{-(\lambda+\delta)t} u'(y_s(t)) = \lambda^2 e^{-\lambda s - (\lambda+\delta)t} + \mu_s(t).$$

¹⁵In addition, at this solution, the multiplier on the budget constraint is $\psi = \frac{\lambda}{(\lambda+\delta)w_0}$.

As before, we can use the knowledge that $y_s(t) = y(s)$ to integrate this equation and find:

$$\left(\frac{1}{\psi} \lambda e^{-(\delta+\lambda+\lambda_B)s} \right) \frac{\lambda_B}{\lambda+\delta} u'(y(s)) = \frac{\lambda^2}{\lambda+\delta} e^{-(\delta+2\lambda)s} + \mu(s). \quad (29)$$

We consider the effect of a collateral tightening by inspecting (28) and (29). The tightening increases the multiplier $\mu(s)$ and raises prices, leading to a fall in both $x(s)$ and $y(s)$.

More interestingly, for an equal change in $\mu(s)$, the relative fall in $x(s)$ and $y(s)$ depends on the ratio $\frac{\lambda_B}{\lambda+\delta}$. Beginning from the case where x and y are equal, from (29) we can see that y falls more than x in response to the tight collateral constraint since $\frac{\lambda_B}{\lambda+\delta} < 1$. Moreover, since in the robust economy $\lambda_B > \lambda$, the decrease in y is smaller in the robust economy than in the benchmark economy.

To prove these statements, for each of (28) and (29), we construct the difference between the case where $\mu(s) = 0$ and $\mu(s) > 0$ for the benchmark case. For each of (28) and (29), the only difference in the right-hand side across these two values of $\mu(s)$ is due to $\mu(s) > 0$. Thus, the difference in the left-hand sides across the two values of $\mu(s)$ are equal, for each of (28) and (29). Using this knowledge, we find that,

$$x^1(s) - x^0(s) = \frac{\lambda}{\lambda+\delta} (y^1(s) - y^0(s)) \frac{x^1(s)}{y^1(s)}.$$

Where $x^0(s)$ and $y^0(s)$ are the equilibrium values of x and y when the constraint does not bind, and $x^1(s)$ and $y^1(s)$ are equilibrium values when the constraint does bind. As before, when the constraint does not bind, $x^0(s) = y^0(s)$. Since $\frac{\lambda}{\lambda+\delta} < 1$, it follows that

$$x^1(s) - x^0(s) > y^1(s) - y^0(s).$$

The tightening of the collateral constraint leads to a smaller fall in x than in y .

Let us compare the robust economy next to the benchmark economy, where both economies are constrained. Subtracting (29) from (28) we find:

$$u'(x(s)) = \frac{\lambda_B}{\lambda+\delta} u'(y(s)) + \psi e^{(\lambda_B-\lambda)s} \frac{\delta}{\lambda+\delta}.$$

In the robust case, λ_B rises beyond λ . From the above expression we can see that increasing λ_B causes the right-hand side to rise and $x(s)$ to fall. Thus, $x(s)$ is lower in the robust economy than the benchmark economy.

To summarize:

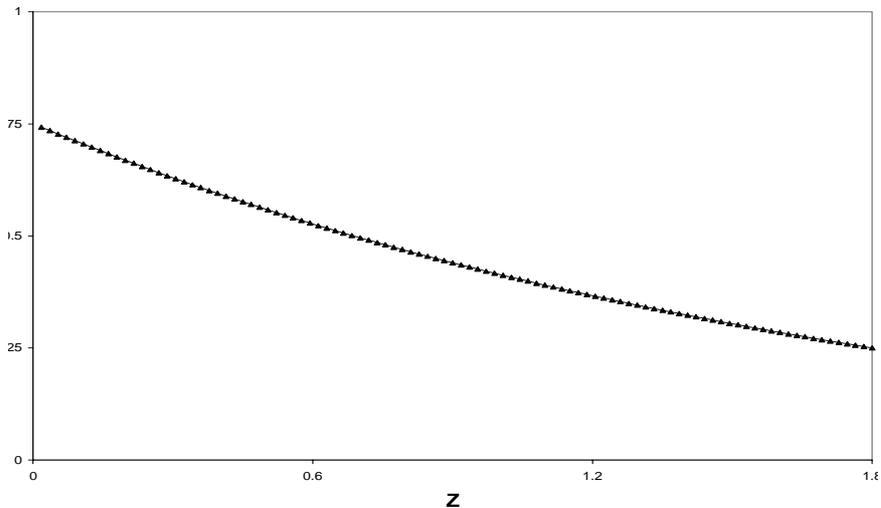
Proposition 6 *Comparing equilibria in the robust economy and the benchmark economy across two different levels of aggregate collateral, we find:*

- If $Z > \underline{Z}$, the equilibrium in the robust economy and the benchmark economy are identical.
- When $Z < \underline{Z}$, x is lower in the robust economy than in the benchmark economy.

The robustness concerns only arise when the aggregate constraint binds. Indeed, this is an important difference between “risk-aversion” and robustness in our setting. Comparing across two economies with different degrees of risk aversion yields differences independently of aggregate constraints.

6.3 Numerical example

Figure 1: λ_B and Collateral Squeeze



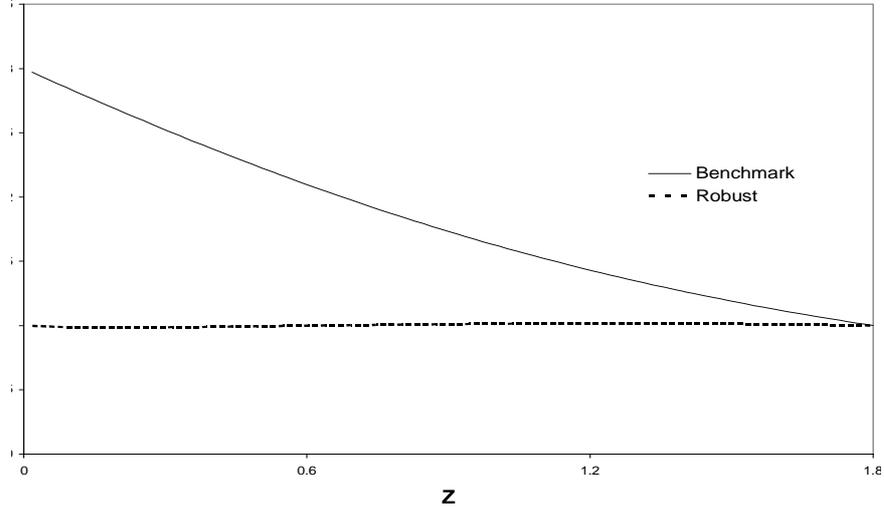
Notes: The figure illustrate the equilibrium selection of λ_B for $Z < \underline{Z}$. λ_B is on the vertical axis and always lies between λ (0.25) and $\lambda + \delta$ (0.75). Z is on the horizontal axis. As the financial sector’s collateral of Z falls, robust agents become increasingly focused on the chance that the other market will receive a shock first and therefore increase λ_B above λ .

In this subsection, we present a numerical example to quantitatively illustrate the workings of our model in the case where $\beta = 1$. We study the fully robust economy. We choose $\delta = 1/2$, to imply that recessions last on average 2 years. We choose $\lambda = 1/4$. If recessions occur once every 8 years, this choice of λ implies that the two-shock extreme event happens about every 48 years.

We normalize $\Pi = 1$. In order to prevent full insurance by the agents, we assume that $w_0 \frac{\lambda + \delta}{\lambda} < 1$, and choose $w_0 = 0.3$. With these parameters, $\underline{Z} = 1.8$.

Figure 1 illustrates the effect of squeezing the financial sector’s collateral on λ_B . As Z falls below \underline{Z} , λ_B begins to rise above λ until it reaches a maximum at $\lambda + \delta$. We note that this behavior is qualitatively similar to the effect of increasing K that we discussed earlier.

Figure 2: x/y and Collateral Squeeze



Notes: The figure contrasts the behavior of x/y between the robust economy and the benchmark economy, over the range $Z < \underline{Z}$. x/y is on the vertical axis. Z is on the horizontal axis. For the robust economy, x/y varies over s . We present the probability-weighted average value of x/y over all s . The average behavior of x/y is representative of the behavior of x/y for moderate values of s ($s < 3$ years). As the financial sector's collateral of Z falls, robust agents become increasingly focused on the chance that the other market will receive a shock first, and insure roughly equally against receiving the shock first or second. Agents in the benchmark economy account for the fact that collateral is scarce, and direct insurance purchases to the more likely single-shock event. Thus x/y rises as Z falls for the benchmark economy.

Figure 2 illustrates the effect of squeezing the financial sector's collateral on the agents' choice of x/y . For the robust economy, x/y varies over s . We present the probability-weighted average value of x/y over all s . This average behavior of x/y is representative of the behavior of x/y for moderate values of s ($s < 3$ years). As the financial sector's collateral of Z falls, robust agents become increasingly focused on the chance that the other market will receive a shock first, and insure roughly equally against receiving the shock first or second. Agents in the benchmark economy account for the fact that collateral is scarce, and direct insurance purchases to the more likely single-shock event. Thus x/y rises as Z falls for the benchmark economy.

These difference in insurance-portfolios translate into different output responses to a typical single shock recession. When there are no shocks during a liquidity episode, per-capita output in both the benchmark and robust economy is Π . When there is one shock during the liquidity episode, output falls to $\frac{\Pi+x}{2}$. As

noted earlier, x is lower in the robust economy than in the benchmark economy. As a result, output is lower in the robust economy. When there are two shocks, output is simply $\frac{x+y}{2}$. However, since the collateral constraint imposes that $x + y = Z$ in both benchmark and robust economy, the output upon two shocks is identical across both economies.

Thus, in terms of output, the benchmark economy does strictly better than the robust economy. The robust economy suffers deeper “typical” recessions, because robust agents become overly concerned with the two-shock extreme event. This “over-reaction” is more severe when aggregate collateral is more binding. For example, when Z is high and close to \underline{Z} , at 1.7, the robust economy sees output decline by only one percent of GDP more than the benchmark economy. Instead, when $Z = 1.2$, the decline is five percent of GDP more in the robust than in the benchmark economy.

7 An Extension to Risky Collateral

Kiyotaki and Moore (1997) observe that when collateral is both an input into production as well as security for financial claims, potentially large feedbacks emerge between the financial and real side. Negative shocks are amplified as financing conditions tighten, reducing demand for collateral as a productive input, reducing collateral values, and reinforcing the financial tightening.

We can consider how Kiyotaki and Moore’s insight applies within our model by assuming that the collateral of the financial sector is (exogenously) risky, but related to the shocks in market A and B . This exercise gives us some understanding of the effect of collateral risk, without working out a full-blown production model.

7.1 Modeling assumptions

The only uncertainty in the economy is in the number of shocks. We assume that if the number of shocks is zero or one, then the total collateral of the financial sector is Z , but if there are two shocks, then collateral falls to $z < Z$. This could happen, for example, if the second shock results in a fall in investment and demand for collateral assets (such as land or real estate).

The introduction of risky collateral requires us to rewrite the collateral constraint of insurers:

$$\hat{x} \leq Z \tag{30}$$

$$\hat{y} \leq \max\{z - \hat{x}, 0\}. \tag{31}$$

If only one shock hits the economy, insurers have enough collateral to cover up to x financial claims. The new constraint is that if the second shock affects the economy, not only has some of the collateral of insurers been pledged to cover the first shock, but also there is less collateral in total.

We also modify our model so that in the event of a liquidity shock, profits drop to $\pi > 0$ rather than $\pi = 0$. When $\pi > 0$, agents effectively are “endowed” with π units of first-shock and second-shock claims, so that their decisions are taken net of this π . In particular their decisions over $x(s)$ and $y_s(t)$ are as described thus far, less π (or zero, if the net is negative).

We assume that

$$\frac{\lambda}{\delta}Z < \pi < \frac{\lambda + \min\{K, \delta\}}{\delta - \min\{K, \delta\}}Z.$$

Then it is straightforward to show that for the benchmark case, π provides sufficient insurance against the (low probability) second-shock that agents use all of their wealth to purchase insurance against the (high probability) first shock:

$$\bar{x}^{be} = Z \quad \bar{y}^{be} = 0.$$

In contrast, agents in the robust economy continue to demand second-shock insurance, and in the particular case of the fully robust economy, $K \geq \delta$, we still have that,

$$\bar{x}^{fr} = \frac{Z}{2} \quad \bar{y}^{fr} = \frac{Z}{2}.$$

Finally, we revert to the case where $\beta = 0$.

7.2 Amplification

Since in the benchmark economy, $x^{be} = Z$ and $y^{be} = 0$, collateral risk has no effect on equilibrium because agents recognize that the risk in collateral is not relevant for their payoffs. The relevant amount of collateral constraining the financial sector continues to be Z . Of course, by not insuring y , the benchmark economy experiences large ex-post (distributional) utility costs if there is a two-shocks event. Nonetheless, these events are sufficiently rare that forgoing insurance is optimal.

In the robust economy, both x and y are positive. The risk in collateral values has a significant impact in the robust economy because when agents insure against two shocks, the relevant collateral constraint for the single-shock shift to (31). Offering more y claims requires intermediaries to lock-up collateral. But this reduces the amount of liquidity insurance available for x .

The effective collateral for the robust economy is the amount of collateral in the worst-case, two-shock event (z). Relative to the benchmark, insurance against the typical single-shock event falls by at least $Z - z$, amplifying the flight-to-quality problem we have identified.

7.3 Demand for Assets and Collateral Premia

Because in the robust economy agents evaluate insurers based on their collateral levels in the worst-case scenario of two shocks, insurers have an incentive to shore up their worst-case collateral so as to increase

their capacity. In a sense, the agents' robustness preferences induces "risk-aversion" in the insurers. In contrast, in the benchmark economy the credibility of the insurer is based on his collateral in the one-shock event. While still averse to losses, the insurer is more concerned about moderate losses.

We parameterize this effect by calculating the risk-premia that an insurer will assign to an infinitesimal amount of the following two Arrow-Debreu claims: (1) A claim that delivers an extra unit of collateral on the occurrence of the first shock; and (2) A claim that delivers an extra unit of collateral on the occurrence of the second shock. Denote by ϕ_1 and ϕ_2 the objective probabilities of each of these events.

We note that the claims are not typical Arrow-Debreu securities, since the events are not mutually exclusive. The second shock can occur only once the first shock has occurred, $\phi_1 > \phi_2$.

The valuations of the two claims are denoted as $\eta_1\phi_1$ and $\eta_2\phi_2$, where the η 's are the collateral premia intermediaries are willing to pay for first-and second-shock events (i.e. they are the premia on these claims after accounting for the probability of the events).

In the benchmark economy (recall that we have set $\beta = 0$ and that in this case the single shock constraint is binding), the extra collateral in the second-shock event is useless to the insurer so that

$$\eta_2^{be} = 0.$$

On the other hand, the extra collateral in the first shock is helpful since it loosens the collateral constraint on selling x claims. We note in the equilibrium we have constructed, agents pay w_0 in total to buy Z insurance claims. Thus the price of an x claim is w_0/Z . Then,

$$\eta_1^{be} = \frac{w_0}{Z} \frac{1}{\phi_1}.$$

In contrast, in the robust economy the extra collateral in the second-shock event is valuable. Following the same logic to calculate prices in the worst case scenario, the value of an extra unit of collateral is w_0/z . Then the collateral-premium assigned to this claim by the insurer is,

$$\eta_2^{fr} = \frac{w_0}{z} \frac{1}{\phi_2}.$$

Collateral in the first-shock event is also valuable in the robust economy, because the second-shock event is subsumed by the first-shock event (i.e. the extra collateral loosens the constraint in the worst-case scenario). However, the probability of the first shock is higher than the probability of the second-shock event, thus,

$$\eta_1^{fr} = \frac{w_0}{z} \frac{1}{\phi_1} < \eta_2^{fr}.$$

Proposition 7 *Collateral-premia are ordered as follows:*

$$\eta_2^{fr} > \eta_1^{fr} > \eta_1^{be} > \eta_2^{be} = 0.$$

Proof. Follows directly from $Z > z$ and $\phi_1 > \phi_2$. ■

The insurer in the robust economy is most concerned about the worst-case event and is willing to pay the highest collateral-premia for claims that help him shore up the worst-case collateral level. Moreover, $w_0/z > w_0/Z$, so that the robust insurer always pays a higher collateral-premium than the benchmark insurer – even against first-shock events. This pattern of prices is in keeping with an agent who is more “risk-averse” against losses.

An empirical regularity during flight-to-quality episodes is the increase in demand for safety and the rise in risk premia. Demand for riskless assets such as Treasury Bills or bank deposits increases. The spread between Commercial Paper and Treasury Bills typically widens during these episodes. Volatility sensitive financial measures, such as the prices of out-of-the money put options on the S&P500 index, or the VIX implied volatility index, rise.

Although our model is too stylized to replicate these observed patterns precisely, these phenomena are broadly consistent with increased aversion to the worst-case scenario. If the financial intermediaries we have identified are the marginal investors in these asset markets, then their risk-valuations affect prices, and plausibly match the observed phenomena.

Our model also provides a new perspective on the cyclical indicator property of the Commercial-Paper to Treasury-Bills spread (see e.g., Stock and Watson, 1989, Friedman and Kuttner, 1993, Kashyap, Stein, and Wilcox, 1993). In our model the rise in the spread is a symptom of agents’ robustness concerns, as opposed to a contributing factor to a recession. But the robustness concerns trigger actions that lead to underinsurance against the single-shock event, and may thereby lead to a recession.

8 Final remarks

Flight to quality is a pervasive phenomenon that exacerbates the impact of recessionary shocks. In this paper we present a model of this phenomenon based on robust decision making by financial specialists. We show that when aggregate intermediation collateral is plentiful, robustness considerations do not interfere with the functioning of private insurance markets (credit lines). However, when agents think that aggregate intermediation collateral is scarce, they take a set of protective actions to guarantee themselves safety, but which leave the aggregate economy overexposed to recessionary shocks.

In this context, a Lender of Last Resort policy is useful if used to support extreme rather than intermediate events. The main benefit of this policy comes not so much from the direct effect of the policy during extreme events, which are very rare, but from its ability to unlock private sector collateral during milder, and far more frequent, contractions.

The implications of the framework extend beyond the particular interpretation we have given to agents and policymakers. For example, in the international context one could think of our agents as countries and the policymaker as the IMF or other IFI's. Then, our model prescribes that the IMF not support the first country hit by a sudden stop, but to commit to intervene once contagion takes place. The benefit of this policy comes primarily from the additional availability of private resources to limit the impact of the initial pullback of capital flows.

References

- [1] Barro, R.J., 2005, "Rare Events and the Equity Premium," Mimeo, Harvard University.
- [2] Bernanke, B., M. Gertler, and S. Gilchrist, 1999, "The Financial Accelerator in a Quantitative Business Cycle Framework," *Handbook of Macroeconomics*, Taylor and Woodford, eds.
- [3] Diamond, D.W. and P.H. Dybvig, 1983, "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy* 91, pp. 401-419.
- [4] Dow, J. and S. Werlang, 1992, "Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio," *Econometrica* 60, p. 197-204.
- [5] Friedman, B., and Kuttner, K., 1993, "Why does the paper-bill spread predict real economy activity." In: Stock, J., Watson, M. (Eds.), *New research in Business Cycles: Indicators and Forecasting*. University of Chicago Press, Chicago.
- [6] Epstein, L., 2001, "Sharing Ambiguity," *American Economic Review*, 91(2), 45-50.
- [7] Epstein, L. and M. Schneider, 2004, "Recursive Multiple Priors," *Journal of Economic Theory*, 113(1), 1-31.
- [8] Epstein, L. and T. Wang, 1994, "Intertemporal Asset Pricing Under Knightian Uncertainty," *Econometrica* 62, 283-322.
- [9] Gatev, Evan, and Phil Strahan, 2005, "Banks' Advantage in Hedging Liquidity Risk: Theory and Evidence from the Commercial Paper Market," forthcoming, *Journal of Finance*.
- [10] Gilboa, I. and D. Schmeidler, 1989, "Maxmin Expected Utility with Non-unique Priors," *Journal of Mathematical Economics*, 18, 141-153.
- [11] Greenspan, A, 2004, "Risk and Uncertainty in Monetary Policy," Remarks for the AEA Meetings, January.
- [12] Hansen, L. and T. Sargent, 1995, "Discounted Linear Exponential Quadratic Gaussian Control," *IEEE Transactions on Automatic Control* 40, 968-971.
- [13] Hansen, L.P. and T.J. Sargent, 2003, "Robust Control of Forward-looking Models," *Journal of Monetary Economics* 50(3), 581-604.
- [14] Hansen, L.P., T.J. Sargent, G.A. Turmuhambetova, N. Williams, 2004, "Robust Control and Model Misspecification," April, Mimeo.

- [15] Holmstrom, Bengt and Jean Tirole, "Private and Public Supply of Liquidity," *Journal of Political Economy*, 106(1), 2/98, pp. 1-40.
- [16] Kashyap, Anil, Jeremy Stein and David Wilcox, "Monetary Policy and Credit Conditions: Evidence from the Composition of External Finance," *American Economic Review*, 83(1), 3/93, pp 78-98.
- [17] Kiyotaki, N. and J. Moore, 1997: "Credit Cycles," *Journal of Political Economy*, 105 (2), 211-248.
- [18] Krishnamurthy, Arvind, 2003, "Collateral Constraints and the Amplification Mechanism," *Journal of Economic Theory*, 111(2), 277-292.
- [19] Krishnamurthy, Arvind, 2002, "The Bond/Old-Bond Spread," *Journal of Financial Economics* 66(2), 463-506, 2002.
- [20] Kyle, A. and W. Xiong, 2001, "Contagion as a Wealth Effect", *Journal of Finance* 56, 1401-1440.
- [21] Longstaff, Francis, 2004, "The Flight-to-Liquidity Premium in U.S. Treasury Bond Prices," *Journal of Business* 77(3), 511-526.
- [22] Routledge, Bryan and Stanley Zin, 2004, "Model Uncertainty and Liquidity," working paper, Carnegie Mellon University.
- [23] Sims, Christopher A., 2001, "Pitfalls of a Minimax Approach to Model Uncertainty," working paper, Princeton University.
- [24] Skiadas, Costis, 2003, "Robust Control and Recursive Utility," *Finance and Stochastics* 7, 475-489.
- [25] Stock, J. and Watson, M., 1989. New indexes of coincident and leading economic indicators. NBER Macroeconomics Annual, 351-394.
- [26] Woodford, M., "Public Debt as Private Liquidity," *American Economic Review, Papers and Proceedings* 80, May 1990, pp. 382-88.