NBER WORKING PAPER SERIES

FLIGHT TO QUALITY AND COLLECTIVE RISK MANAGEMENT

Ricardo J. Caballero
Arvind Krishnamurthy

Working Paper 12136
http://www.nber.org/papers/w12136

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2006

Flight to Quality and Collective Risk Management
Ricardo J. Caballero and Arvind Krishnamurthy
NBER Working Paper No. 12136
March 2006
JEL No. E30, E44, E5, F34, G1, G22, G28

ABSTRACT

We present a model of flight to quality episodes that emphasizes systemic risk and the Knightian uncertainty surrounding these episodes. Agents make risk management decisions with incomplete knowledge. They understand their own shocks, but are uncertain of how correlated their shocks are with systemwide shocks. Aversion to this uncertainty leads them to question whether their private risk management decisions are robust to aggregate events, generating conservatism and excessive demand for safety. We show that agents' actions lock-up the capital of the financial system in a manner that is wasteful in the aggregate and can trigger and amplify a financial accelerator. The scenario that the collective of conservative agents are guarding against is impossible, and known to be so even given agents' incomplete knowledge. A lender of last resort, even if less knowledgeable than private agents about individual shocks, does not suffer from this collective bias and finds that pledging intervention in extreme events is valuable. The benefit of such intervention exceeds its direct value because it unlocks private capital markets.

Ricardo J. Caballero
MIT Department of Economics
Room E52-252a
Cambridge, MA 02142-1347
and NBER
caball@mit.edu

Arvind Krishnamurthy
Kellogg Graduate School of Management
Northwestern University
2001 Sheridan Road
Evanston, IL 60208
a-krishnamurthy@northwestern.edu

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 and macroeconomic 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] In each of them, as hinted by Greenspan, the Fed intervened early and stood ready to intervene "as much as needed" to prevent a meltdown.

Two questions immediately arise when considering these examples and Greenspan's comments on flight to quality episodes. First, how can events that appear small relative to US private wealth create so much havoc? Second, how can interventions by the Fed, that also appear small relative to US wealth, be effective in preventing a meltdown? In this paper we develop a model to address these questions. In short, our answers are that during flight to quality episodes private capital becomes locked up, and much of the power of intervention derives from its ability to unlock this capital.

Our model builds on two observations about flight to quality events. First, at a basic level these episodes are about an increase in perceived "risk." However, the risk does not circumscribe a purely fundamental shock, and instead centers around the nerve center of the economy, 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*.

---

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

Second, the "risk" in these episodes is of an extreme nature. Most episodes are triggered by an unexpected event that causes agents to re-evaluate their models of the world. In the Fall of 1998, the comovement of Russian government bond spreads, Brazilian spreads, and U.S. Treasury bond spreads was a surprise to even sophisticated market participants. These high correlations rendered standard risk management models obsolete, leaving financial market participants searching for new models. Agents responded by making decisions using "worst-case" scenarios and "stress-testing" models, in a manner suggestive of *Knightian uncertainty*. Indeed, the extreme disengagement from risky activities that Greenspan highlights suggests that agents respond to more than just risk.

Our model focuses on agents' perceptions regarding the ability of financial intermediaries to deliver on their financial contracts. Agents contract with financial intermediaries to cover liquidity shocks that may arise in their markets. These risk management decisions are made with incomplete knowledge. Agents understand their own shocks, but are uncertain about the probability model describing shocks in markets different from theirs and treat this uncertainty as Knightian. If financial intermediaries have limited resources, agents grow concerned that shocks may arise in other markets that will deplete the resources of intermediaries and compromise their ability to deliver on their financial contracts. Such riskiness rises either through an increase in Knightian uncertainty or a fall in intermediaries' assets. Thus our model combines the financial system risk as well as the Knightian uncertainty that underlies flight to quality episodes.

The model captures two important aspects of flight to quality episodes. First, the increase in perceived riskiness generates conservatism and demand for safety. Second, we show that Knightian agents respond to their uncertainty regarding other markets by requiring financial intermediaries to lock-up some capital to devote to their own markets' shocks, regardless of what happens in other markets. By forcing intermediaries to dedicate capital to cover their own shocks, agents ensure that their shocks are covered regardless of who else receives shocks. However, once locked-up, intermediaries' capital is not free to move across markets in response to shocks, resulting in bottlenecks and market segmentation.

The capital bottlenecks have macroeconomic consequences, and lend support to Greenspan's call for central bank action during flight to quality. While each Knightian agent covers himself against an extreme shock, collectively these actions prevent intermediaries from moving capital across markets to expediently offset shocks as they arrive. We show that this inflexibility leaves the economy overexposed to (moderate) aggregate shocks that are manageable by the private sector in the absence of flight to quality. Moreover, the scenario that the collective of conservative agents are guarding against is impossible, and known to be so even given agents' incomplete knowledge. Collectively, agents make poor risk management decisions that lead to avoidable losses.

In this context, a lender of last resort, even if less informed than private agents about each agent's own market, can unlock private capital and stabilize the economy during moderate shocks. It does so

by committing to intervene during extreme events where the financial intermediaries' 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 free flexible resources to deal with moderate shocks. In this sense, 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, which unlocked *private* collateral.

Literature review. The model we develop and our prescriptions over the lender of last resort are most closely related to the Diamond and Dybvig (1983) model of bank runs. An important difference relative to the Diamond and Dybvig model is that ours does not not involve a coordination failure. We will discuss the relation to Diamond and Dybvig in greater depth after we introduce the model.

Our paper fills 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. Instead, we emphasize 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.

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.

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 to various extreme scenarios 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 decision rules that overweight a worst-case 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.

Easley and O'Hara (2005) study a model where ambiguity averse traders focus on a worst case scenario when making an investment decision. Like us, Easley and O'Hara point out that government intervention in a worst-case scenario can have large effects.

Finally, in our model agents are only Knightian with respect to systemic events. Epstein (2001) explores the home bias in international portfolios in a related setup, where agents are more uncertain about foreign than local markets.[2] As Epstein points out, this modelling also highlights the difference between high risk aversion and aversion to Knightian uncertainty. Moreover, our modelling 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, while in Section 3 we describe decisions, equilibrium and flight to quality. Section 4 derives the value of a lender of last resort in our economy. Section 5 illustrates the interaction between aggregate collateral and robustness considerations. Section 6 discusses moral hazard problems and other critiques associated with the lender of last resort. Section 7 concludes.

2   The Model

We study a model conditional on entering a period of turmoil where Knightian uncertainty is high. Our model is silent on what triggers the episode. In practice, we think that the occurrence of an unexpected

---

[2]Epstein's model is closely related to our model in Caballero and Krishnamurthy (2005).

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?

The model is a variant of a bank-run model (e.g. Diamond and Dybvig, 1983). It has a set of competitive agents and intermediaries. The agents are subject to liquidity shocks. As in the canonical Diamond and Dybvig (1983) banking model, an agent may be "late" in which case he has no special liquidity needs, or he may be "early" in which case he needs some liquidity immediately. We view the liquidity shocks as a parable for a sudden need for capital by a financial market specialist (e.g. a trading desk, hedge fund, market maker). The agents sign financial contracts with the intermediaries to provide them with liquidity insurance if they are "early." An intermediary can be thought of as a bank which extends a credit line to the specialists, or as the top-level capital allocator in an investment bank.

We are interested in situations where the agents perceive that there is some chance that intermediaries may not deliver the contracted liquidity insurance. In a bank run model, a central concern of agents is that they will arrive too late to claim liquidity from an intermediary, and possibly not receive the liquidity. We introduce this type of concern by assuming that early agents are further divided into those who receive shocks first and those who receive shocks second. Each intermediary is assumed to have a limited amount of liquidity/collateral, Z, which it delivers to the agents. We normalize things so that in aggregate liquidity (or collateral) is also Z, and no longer distinguish between one intermediary and the aggregate.

SHOCKS AND PROBABILITIES

Formally, there is a continuum of agents indexed by $\omega \in \Omega \equiv [0, 1]$. Agent ω receives utility:

$$U^\omega = \begin{cases} \bar{u} & \text{if late} & \text{Probability: } 1 - \phi_\omega(1) - \phi_\omega(2) \\ u(c_1) & \text{if early and first} & \text{Probability: } \phi_\omega(1) \\ u(c_2) & \text{if early and second} & \text{Probability: } \phi_\omega(2) \end{cases}$$

$u : \mathcal{R}_+ \to \mathcal{R}$ is twice continuously differentiable, increasing, strictly concave and satisfies the condition $u'(0) = \infty$.

Unlike the Diamond and Dybvig model, we assume that the shocks are correlated across agents. If shocks occur, there is a first wave, and then possibly a second wave. The first wave of shocks occurs with probability $\phi(1)$, and the second wave takes place with conditional probability $\phi(2|1)$. We assume that $\phi(2|1) < 1$, so that

$$1 \geq \phi(1) > \phi(2) > 0 \tag{1}$$

with $\phi(2) = \phi(2|1)\phi(1)$.

This condition states that, in aggregate, a single-wave event is more likely than the two-wave event. We

6

will refer to the two-wave event as an extreme event, capturing an unlikely but severe liquidity crisis in which many agents are affected.

The shocks are aggregate in that they simultaneously affect a mass of agents. We assume that if the first wave occurs, then one-half of the agents is hit by the shock. For agent ω, the probability that he receives a shock first, conditional on the first wave of shocks occurring, is equal to $\frac{\phi_\omega(1)}{\phi(1)}$. While for any given ω, the conditional probability need not be one-half, the average conditional probability across all agents must be equal to one-half:

$$\int_\Omega \phi_\omega(1)d\omega = \frac{\phi(1)}{2}.$$

If the second wave occurs, the remaining one-half of the agents is hit by the shock:

$$\int_\Omega \phi_\omega(2)d\omega = \frac{\phi(2)}{2}.$$

SECURITIES

There is a contracting time, date 0, where intermediaries offer financial contracts to provide liquidity insurance to the agents. We assume all shocks are observable and contractible. There is no concern that an agent will pretend to have a shock and collect on an insurance claim. We define two types of financial contracts that intermediaries offer:

- $s(\omega)$ is a safe claim. This claim pays one to agent ω if he receives a shock regardless of whether the shock is during the first or second wave. The claim costs $p(s, \omega)$.

- $r(\omega)$ is a risky claim. This claim pays one to agent ω if he receives a shock during the first wave, but not if he receives its shock during the second wave. The claim costs $p(r, \omega)$.

We refer to the r claim as risky because it depends on the order of the agent in the sequence of shocks. In the example of a bank-run, it reflects the feature where an agent may need liquidity but be second in line to claim the liquidity.

These two financial securities span the uncertainty in our economy. In principle one could introduce securities that depend on the shock realizations over every subset of Ω. But given the securities we have defined, such securities will not be traded.[3]

INFORMATION

Agents purchase financial claims from the intermediaries to insure against their liquidity shocks. In making the insurance decisions, agents have a probability model of their liquidity shocks in mind. We assume that agent-ω knows his probability of receiving a shock:

$$\phi_\omega = \phi_\omega(1) + \phi_\omega(2).$$

---

[3]For example, we could imagine a security that pays one if both ω and ω' receive a first shock. However, agent ω has no use for the ω' indexation.

However, and centrally to our model, agent ω is uncertain about the correlation between his own shock and aggregate conditions. In particular he does not know the relative likelihood of being among the first wave or second wave. Agents treat the latter uncertainty as Knightian.

We define,

$$\theta_\omega^\omega \equiv \frac{\frac{\phi_\omega^\omega(2)}{\phi_\omega^\omega(1)}}{\frac{\phi(2)}{\phi(1)}}. \tag{2}$$

Agent ω does not know the true probabilities $\phi_\omega(1)$ and $\phi_\omega(2)$. We use the notation $\phi_\omega^\omega(1)$, etc., to denote agent-ω's perception of the relevant true probability. Since agents know ϕ_ω, θ_ω^ω is a sufficient statistic for agent-ω's unknown probabilities. Normalizing these probabilities by the aggregate shock probabilities is convenient in the analysis.

Agents consider a range of probability-models θ_ω^ω in the set Θ, with support $[1 - K, 1 + K]$, and design insurance portfolios that are robust to their model uncertainty. We follow Gilboa and Schmeidler's (1989) Maxmin Expected Utility representation of the problem and write:

$$\max_{r,s} \min_{\theta_\omega^\omega \in \Theta} E_0[U^\omega | \theta_\omega^\omega] \quad s.t. \quad p(s, \omega)s(\omega) + p(r, \omega)r(\omega) \leq w_0 \tag{3}$$

K captures the extent of agents' uncertainty, while w_0 is the initial endowment of agents.

In a flight to quality event, such as the Fall of 1998 or 9/11, agents are unsure of how aggregate conditions will impinge on their activities. They may have a good understanding of their own markets, but are unsure of how the behavior of agents in other markets may affect them. For example during 9/11, market participants feared gridlock in the payments system. Each participant knew how much money he owed to others but was unsure whether money owed to him would arrive. In our modeling, agents are certain about the probability of receiving a shock, but are uncertain about the probability that their shocks will occur early relative to others, or late relative to others.

We view agents' max-min preferences in (3) as descriptive of their decision rules. The widespread use of worst-case scenario analysis in decision making by financial firms is an example of the robustness preferences of such agents.[4]

SYMMETRY

To simplify our analysis we assume that the agents are symmetric at date 0. While each agent's true θ_ω may be different, the θ_ω for every agent is drawn from the same Θ.

The symmetry applies in other dimensions as well: ϕ_ω, K, w_0, and $u(c)$ are the same for all ω. Moreover, this information is common knowledge.

We assume that the aggregate shock probabilities, $\phi(1)$ and $\phi(2)$, are known by all agents. We may think that agents observe the past behavior of the economy and form precise estimates of these aggregate

---

[4]See also Routledge and Zin (2004) for a discussion of max-min behavior among financial specialists.

probabilities. However the same past data does not reveal whether a given ω is more likely to be in the first wave or the second wave.

INTERMEDIARIES

Financial intermediaries have a limited capacity to deliver liquidity insurance. Each intermediary is endowed with some goods to deliver to its claimholders. These goods can be thought of as the intermediary's liquid funds or as the collateral/capital of an intermediary. A better capitalized/collateralized intermediary may have more goods and can thereby credibly sell more liquidity insurance. In aggregate, there are Z units of these goods. For much of our analysis, we assume that Z is constant. In Section 5 we study the effect of varying aggregate collateral Z (as in Kiyotaki and Moore, 1997).

Agents recognize that intermediaries have limited resources. Denote by R and S the aggregate amount of risky and safe claims sold by intermediaries, respectively. Agents can observe total liabilities and collateral assets for each intermediary. The agents require that each intermediary has sufficient assets to cover the liabilities in all states of the world. If the first wave occurs, one-half of the agents will be affected and the intermediaries pay out $\frac{R+S}{2}$.[5] If the second wave occurs, the remaining one-half is affected and the intermediaries pay out $\frac{S}{2}$. Then, the collateral constraint for the intermediaries is:

$$\frac{R}{2} + S \le Z \tag{4}$$

A representative intermediary's objective is to maximize the date 0 revenue from the sale of the insurance less a cost of the insurance resources disbursed to both first and second shock agents, discounted at the rate β:

$$U^I = c_0^I + E_0[\beta(c_1^I + c_2^I)], \qquad \beta \ge 0 \tag{5}$$

and subject to the collateral constraint in (4). The intermediary is a standard expected utility maximizer, with linear preferences. It does not need to know the true θ_ω of each agent in selling insurance. Knowledge of $\phi(1)$, $\phi(2)$, and ex-ante symmetry of agents are sufficient to compute the expectation in equation (5).

RELATION TO DIAMOND AND DYBVIG (1983)

An important distinction between our model and the Diamond and Dybvig (1983) model of bank runs is our assumption of complete state-contingent markets. In Diamond and Dybvig the sequential service constraint creates a possible coordination failure. The coordination failure informs the discussion of bank runs and policy intervention. Our model does not impose a sequential service constraint. Agents are able to write complete state-contingent contracts on the order of the shocks. The amount of early liquidation by agents is predetermined by date 0 contracts rather than by coordination failure among agents.

---

[5] Since it is common knowledge that agents are symmetric, agents know that $r(\omega) = R$. Since one-half of the agents receive the shock in the first wave, $R/2$ of risky claims are settled by the intermediaries. Also note that since each intermediary has a continuum of clients, it faces no uncertainty on the quantities to be delivered in each aggregate state.

In Diamond and Dybvig, the bank run triggers unanticipated withdrawals and default by the bank. In our model, agents recognize that intermediaries may default because they have limited Z and impose the collateral constraint of equation (4). However, in our complete markets model default never occurs in equilibrium. This is because, with complete markets, agents anticipate the state of the world where the intermediary's liabilities lead it to default, and rewrite contracts so as to reduce the liabilities exactly to the point where no default occurs. Thus default is central in the calculations of agents but, and as a result, does not occur in equilibrium.[6]

3 Knightian Uncertainty and Flight to Quality

In this section we describe agents' decisions and equilibrium, and connect robustness concerns with flight-to-quality episodes.

3.1 The cost of locking collateral

The problem of a financial intermediary is to sell $\widehat{r}(\omega)$ risky claims and $\widehat{s}(\omega)$ safe claims, at prices $p(r, \omega)$ and $q(r, \omega)$, in order to maximize revenue less the actuarial cost of providing liquidity insurance, subject to the collateral constraint. Following the objective in (5), intermediaries solve:

$$\max_{\widehat{r}(\omega), \widehat{s}(\omega)} \int_\Omega \left[(p(r, \omega) - \beta\phi_\omega(1)) \widehat{r}(\omega) + (p(s, \omega) - \beta(\phi_\omega(1) + \phi_\omega(2))) \widehat{s}(\omega) \right] d\omega.$$

Because agents are ex-ante symmetric, their decisions are identical and intermediaries offer insurance at the same prices to all agents. For brevity, we omit the argument ω in the decisions and price functions and let R and S denote the aggregated claims sold to agents.

Recall that it is common knowledge that if the first aggregate shock hits, one-half of the agents are affected. Likewise, if the second aggregate shock hits, the other one-half of the agents is affected. Since the intermediary is only concerned about the aggregate payout on the financial claims it has sold, we can rewrite its objective as:

$$\max_{R, S} \left(p(r) - \beta\frac{\phi(1)}{2} \right) R + \left(p(s) - \beta\frac{\phi(1) + \phi(2)}{2} \right) S.$$

---

[6] 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. While interesting, we deliberately removed these features from our framework through a complete Arrow-Debreu markets assumption which, as usual, allows us to state a dynamic problem as a static one.

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.

The intermediary must satisfy the collateral constraint in (4):

$$\frac{R}{2} + S \leq Z.$$

Because of the linearity of the objective function, if the collateral constraint does not bind the solution to the intermediary's problem yields that R and S are at an interior only if,

$$p(r) = \beta \frac{\phi(1)}{2} \quad \text{and} \quad p(s) = \beta \frac{\phi(1) + \phi(2)}{2}.$$

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

If the collateral constraint is binding, then:

$$p(r) = \beta \frac{\phi(1)}{2} + \frac{\mu}{2} \quad \text{and} \quad p(s) = \beta \frac{\phi(1) + \phi(2)}{2} + \mu. \tag{6}$$

where $\mu \geq 0$ is the Lagrange multiplier on the collateral constraint. A positive multiplier in the cost of insurance reflects that the intermediary has limited resources and hence the opportunity cost of selling one type of insurance claims is selling less of the other. Safe claims use up twice the amount of collateral as a risky claim, given our assumption that one-half of the agents are affected in a first shock. If the intermediary sells one less safe claim, then it can credibly sell two more risky claims. This logic leads the intermediary to factor twice the cost of the collateral constraint when selling safe claims over risky claims.[7]

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 first term on the right-hand-side of (6) to yield:

$$p(r) = \frac{p(s)}{2}. \tag{7}$$

This pricing expression is the main result from modelling the supply side of the economy. It captures the equilibrium cost of "locking" collateral. In Section 5 we revisit the case of $\beta > 0$.

3.2 The agent's decision problem

Facing prices $p(r)$ and $p(s)$ each agent solves the decision problem in (3):

$$\max_{r,s} \min_{\theta_\omega^\omega \in \Theta} \phi_\omega^\omega(1) u(c_1) + \phi_\omega^\omega(2) u(c_2) + (1 - \phi_\omega) \bar{u} \tag{8}$$

---

[7]Note that if the first and second waves were mutually exclusive events, then the intermediary could use the same collateral to back liquidity insurance in each event (i.e., it would separate the first and second wave collateral constraints). In this case the tradeoff of collateral would only be one for one between risky and safe claims, and the only binding constraint would be that for claims sold for agents hit on the first wave. In contrast, in our models waves take place sequentially and hence are not mutually exclusive.

where,

$$c_1 = r + s, \quad c_2 = s, \quad p(s)s + p(r)r \leq w_0,$$
$$\phi_\omega^\omega(1) = \frac{\phi(1)}{2}\frac{\phi(1) + \phi(2)}{\phi(1) + \theta_\omega^\omega \phi(2)}, \quad \phi_\omega^\omega(2) = \phi_\omega - \phi_\omega^\omega(1).$$

The expressions for c_1 and c_2 come from our earlier definitions of risky and safe claims: the risky claim pays off only in the first shock, while the safe claim pays off in both first and second shocks.

For the moment, let us consider the problem in (8) for a *given* value of θ_ω^ω (i.e we drop the min operator). The first order conditions are,

$$\phi_\omega^\omega(1)u'(c_1) = \psi p(r)$$
$$\phi_\omega^\omega(1)u'(c_1) + \phi_\omega^\omega(2)u'(c_2) = \psi p(s)$$

where ψ is the Lagrange multiplier on the agent's budget constraint.

Dividing the second equation by the first equation yields:

$$1 + \frac{\phi_\omega^\omega(2)}{\phi_\omega^\omega(1)}\frac{u'(c_2)}{u'(c_1)} = \frac{p(s)}{p(r)} \tag{9}$$

From the definition of θ_ω^ω in (2), we rewrite this expression as,

$$\theta_\omega^\omega \frac{\phi(2)}{\phi(1)}\frac{u'(c_2)}{u'(c_1)} = \frac{p(s)}{p(r)} - 1$$

The right hand side of this expression are equilibrium prices. From our analysis of the intermediary's problem, we know that the right hand side is equal to one. Thus

$$\frac{u'(c_2)}{u'(c_1)} = \frac{\phi(1)}{\phi(2)}\frac{1}{\theta_\omega^\omega} \tag{10}$$

That is, θ_ω^ω affects the agent's decisions by altering the perceived odds of being first or second.

An agent who considers a higher value of θ_ω^ω chooses higher c_2 and lower c_1. He can achieve this consumption pattern by holding relatively more safe claims and less risky claims. Since all claims have to be fully collateralized by the intermediary, in demanding more safe claims, the agent also forces the bank to set aside more collateral to cover shocks that will happen second. In terms of the bank-run example we have used, an agent who thinks that he is likely to be last in line at the bank, requires the bank to set aside resources for when he does arrive at the bank.

We now turn to the robustness step. The agent makes his insurance choices while being robust to alternative values of θ_ω^ω. In the game-theoretic language often used to describe max-min expected utility theory, the agent first chooses r and s, then "nature" chooses $\theta_\omega^\omega \in \Theta$ to minimize the utility of agent A, given his choices.

12

If $u(c_1) > u(c_2)$, from (8), we see that nature will set $\phi_\omega^\omega(2)$, and thus θ_ω^ω, as high as possible. That is, when the agent is more insured against being hit first than second, the worst-case for the agent is that his true θ_ω^ω is equal to $1 + K$. On the other hand, if $u(c_1) < u(c_2)$, nature will do the opposite and set $\theta_\omega^\omega = 1 - K$. At $u(c_1) = u(c_2)$, the expected utility from the agent's decisions are independent of θ_ω^ω.

The agent makes his decisions anticipating this behavior by nature. First, suppose that K is small so that θ_ω^ω is near one. In this case $c_1 > c_2$ so that nature will choose $\theta_\omega^\omega = 1 + K$. The agent chooses c_1 and c_2 to satisfy:

$$\frac{u'(c_2)}{u'(c_1)} = \frac{\phi(1)}{\phi(2)} \frac{1}{1+K} > 1$$

As K rises, the right-hand side of this expression falls and c_1/c_2 falls toward one. When,

$$K = \bar{K} \equiv \frac{\phi(1)}{\phi(2)} - 1,$$

the right-hand side is equal to one, and $c_1 = c_2$. We refer to this situation as the fully robust case since agents' decisions are robust to their uncertainty over $\phi_\omega(1)$ and $\phi_\omega(2)$.

Notice that if K rises past $\bar{K}$, the agent will not change his decision. This is because if the agent chooses $c_1 < c_2$, then nature will set $\theta_\omega^\omega = 1 - K$. Anticipating this action, the agent will prefer to choose $c_1 = c_2$.

Finally, note that equilibrium in this economy is unique. Given the supply determined relative price, $p(s) = 2p(r) = 2p$, a change in p only has income effects. Agents' demands are infinite at $p = 0$ and decrease monotonically toward zero as p rises. The following proposition summarizes these results:

Proposition 1 *The following insurance decisions constitute an equilibrium in the **robust economy***:

- *For $0 \leq K < \bar{K} \equiv \frac{\phi(1)}{\phi(2)} - 1$, agents' decisions are as if $\theta_\omega^\omega = 1 + K$:*

$$\frac{\phi(1)}{\phi(2)} \geq \frac{u'(c_2)}{u'(c_1)} = \frac{\phi(1)}{\phi(2)} \frac{1}{1+K} > 1$$

  We refer to this as the "partially robust" case.

- *For $K \geq \bar{K}$, agents' decisions are as if $K = \bar{K}$:*

$$\frac{u'(c_2)}{u'(c_1)} = \frac{\phi(1)}{\phi(2)} \frac{1}{1+\bar{K}} = 1$$

  We refer to this as the "fully robust" case.

3.3 Flight to quality

In the partially robust solution, every agent chooses $c_1 > c_2$. To achieve this consumption pattern, agents choose $r = c_1 - c_2$ and $s = c_2$. As K rises, so that the uncertainty for agents rises, c_1 and c_2 converge to

each other. Agents lower r and increase s. At the fully robust solution, agents set $c_1 = c_2 = s$ and $r = 0$. The pattern of shedding risky claims in favor of safe claims is central to flight to quality episodes.

Interpreting the financial claims of agents as credit lines from banks, the rise in K generates a "run" on banks' credit facilities.[8] The agents recognize that banks have limited resources to back up their credit lines. The concern of each agent is that his shock will occur after the average agents' shocks, at which point the bank's limited credit facility will have been depleted.[9] By setting $r = 0, s > 0$, agents insulate themselves from this concern.

In practice, the robustness action may be reflected by an agent preemptively drawing down a credit line at a date when the fear of other agents' shocks arises. As markets are complete in our model, the future action by each agent is prearranged at date 0. Alternatively, we can think of these actions as agents shifting their demands towards well capitalized intermediaries, where it is unlikely that the intermediary's resources will be depleted. This portfolio shift is common in flight to quality episodes.

We can also interpret the effect of a rise in K on equilibrium in terms of intermediaries' capital allocation decision. By choosing $r = 0, s > 0$, agents force intermediaries to reserve some resources to cover the agents' own shocks, regardless of when these shocks occur. s of the capital of an intermediary is locked-up ("committed capital"), and r of the capital is free ("trading capital") to allocate to the first agents who receive shocks. As θ_ω^ω rises, r falls and there is less trading capital to allocate flexibly. In the fully robust case, all of the intermediary's capital is locked-up and there is no capital mobility.

The latter interpretation captures a feature shared by most flight to quality episodes. Bottlenecks arise in the movement of capital and markets appear segmented. For example, in the Fall of 1998 episode the markets where hedge funds specialized were particularly affected by the crisis, as (abundant) capital did not flow into these markets.

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 5, but the intuition is clear enough. A fall in Z increases financial intermediary risk as agents grow concerned that intermediaries may not have enough resources to cover their shocks. Agents worry about being second and raise θ_ω. Some episodes, such as the Fall of 1998, can be thought of in terms of financial intermediary risk.

---

[8]Although the actions of agents in our model are most naturally interpreted as a run on banks' credit facility, at a deeper level, agents' actions are also related to the more commonly analyzed run on banks' deposits. As Diamond and Dybvig (1983) emphasize, a deposit contract implements an optimal shock-contingent allocation of liquidity. In the fully robust equilibrium of our model, agents choose an allocation that is non-contingent on each other's shocks. Agents each hoard liquidity to cover their own shocks, independent of other markets' shocks. In this sense, robust agents' preference for shock-independent liquidity allocations is related to the behavior of panicked depositors in a bank run.

[9]We can also interpret robustness in terms of the internal risk management of an investment bank. Each ω is a trading desk of an investment bank. Each ω lobbies for more risk capital, because of fear that other agents will suffer losses soon, reducing the bank's total risk capital.

4 Collective Bias and the Lender of Last Resort

In the fully robust equilibrium of Proposition 1 agents insure equally against first and second shocks. To arrive at the equal insurance solution, robust agents evaluate their first order conditions (equation 9) at *conservative* probabilities:

$$\phi_\omega^\omega(1) = \phi_\omega^\omega(2).$$

Suppose we compute the probability of one and two aggregate shocks using agents' conservative probabilities. Then,

$$\bar{\phi}(1) = 2 \int_\Omega \phi_\omega^\omega(1) d\omega.$$

where the two in this expression reflects the fact that only half of agents are affected by the first aggregate shock. Likewise,

$$\bar{\phi}(2) = 2 \int_\Omega \phi_\omega^\omega(2) d\omega.$$

Together, these computations imply that agents' conservative probabilities are such that,

$$\bar{\phi}(1) = \bar{\phi}(2).$$

Implicit in the conservative probabilities, the economy is perceived as equally likely to receive one or two aggregate shocks. But we know (and all agents know) that actually $\phi(1) > \phi(2)$, which implies that agents' conservative probabilities are *collectively biased*.

Since each agent is concerned about the scenario in which he receives a shock last and the intermediary's resources have been depleted, robustness considerations lead each agent to bias upwards the probability of receiving a shock later than the average agent. But collectively, every agent cannot be later than the "average."

These points carry over to the partially robust equilibrium. For any $K > 0$, robust agents set $\theta_\omega^\omega = 1 + \min[K, \bar{K}]$, implying that,

$$\phi_\omega^\omega(2) = (1 + \min[K, \bar{K}]) \frac{\phi(2)}{\phi(1)} \phi_\omega^\omega(1). \tag{11}$$

If we take the average of the conservative probabilities on the left and right hand side of equation (11), we find,

$$\frac{\bar{\phi}(2)}{\bar{\phi}(1)} = (1 + \min[K, \bar{K}]) \frac{\phi(2)}{\phi(1)} \geq \frac{\phi(2)}{\phi(1)}.$$

When $K > 0$ the last inequality is strict, showing that agents conservative probabilities are collectively biased.

Note that each agent's conservative probabilities are individually plausible. Given the range of uncertainty over θ_ω, it is possible that agent ω has a higher than average probability of being second. Only when viewed from the aggregate does it becomes apparent that agents' collective probabilities are implausible. These

observations motivate us to study how a central bank, which is interested in maximizing the collective, can improve on outcomes.

4.1 Central bank objective and information

The central bank maximizes the average utility that agents derive from their choices. In computing this average, we assume that the central bank uses its own probability assessments rather than the agents'. We also assume that the central bank equally weights the utility of every agent. The latter assumption follows because agents are observationally identical at date 0. The former assumption is more delicate.

When agents have non-Savage preferences, at least two approaches seem defensible: The central bank objective uses both the agents' preferences and the agents' conservative probability assessments; Or, the central bank takes a more paternalistic approach in which it is concerned with agents' ex-post average utility from consumption, but not evaluated at the agents' collectively biased conservative probability assessments. We follow the latter path but conclude the section by showing that both approaches have similar implications for the welfare gains associated to the presence of a lender of last resort.

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., worst-case scenario analysis). These preferences are useful because they reproduce the behavior of agents in flight to quality situations where they are faced with Knightian uncertainty. From this perspective, it is not obvious that a central bank should build biases into its objective function that may lead to obvious average losses, just because agents exhibit these biases.[10]

Thus, let us consider the following central bank objective function:[11]

$$V = \int_{\Omega} \left[\phi_{\omega}^{CB}(1)u(c_{1,\omega}) + \phi_{\omega}^{CB}(2)u(c_{2,\omega}) \right] d\omega.$$

$c_{1,\omega}, c_{2,\omega}$ are the consumption resulting from agents' insurance decisions, while $\phi_{\omega}^{CB}(1)$, etc., are the central bank's assessments of the probabilities of the relevant events. We note again that the central bank uses agents' utility function of $u(c)$ to evaluate their consumption decisions, but discounts $u(c)$ using its own probability assessments, rather than the agents' conservative probability assessments.

Importantly, we do *not* assume that the central bank knows the true θ_{ω}s of agents. Since the agents' biases stem from a lack of knowledge of how aggregates impinge on their activities, we also assume that the central bank lacks knowledge over θ_{ω}. In fact, for all of our results in this section we can assume that the

---

[10]Sims (2001) has made a related 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. Note, however, that the collective bias we identify would be solved by a central bank, even if it uses robust control for its decisions.

[11]We have omitted the constant term involving $\bar{u}$ in this objective.

central bank knows strictly *less* than the agents; i.e. that the K of the central bank is larger than the K of agents.

The central bank, like the agents, knows both aggregate probabilities and that one-half of the agents is affected by each aggregate shock:

$$\int_\Omega \phi_\omega^{CB}(1)d\omega = \frac{\phi(1)}{2} \quad \text{and,} \quad \int_\Omega \phi_\omega^{CB}(2)d\omega = \frac{\phi(2)}{2}.$$

We may imagine that the central bank and individual agents are able to form precise estimates over the aggregate behavior of shocks from observing past data.

4.2 Collective risk management and wasted collateral

Consider a central bank that alters agents' decisions by increasing $c_{1,\omega}$ by an infinitesimal amount, and decreasing $c_{2,\omega}$ by the same amount. This can be implemented, for a given amount of intermediaries' collateral, by requiring agents to hold *two* more units of r and *one* less unit of s. The value of the reallocation is:

$$V^{RE} = \int_\Omega \left[\phi_\omega^{CB}(1)u'(c_{1,\omega}) - \phi_\omega^{CB}(2)u'(c_{2,\omega}) \right] d\omega.$$

In the fully robust equilibrium, the first order condition for agents is,

$$\frac{u'(c_{2,\omega})}{u'(c_{1,\omega})} = 1.$$

The budget constraint for every agent is that,

$$p(r)r + p(s)s = w_0.$$

Since the central bank knows that all agents face the same prices and have the same wealth level, it can conclude that $c_{1,\omega} = c_1$ for all ω. Then, we can write,

$$
\begin{aligned}
V^{RE} &= u'(c_1) \int_\Omega \left[\phi_\omega^{CB}(1) - \phi_\omega^{CB}(2) \right] d\omega \\
&= \frac{1}{2}u'(c_1)(\phi(1) - \phi(2)) \\
&> 0.
\end{aligned}
$$

Similarly, in the partially robust case, the first order condition for agents is,

$$\frac{u'(c_{2,\omega})}{u'(c_{1,\omega})} = \frac{\phi(1)}{\phi(2)} \frac{1}{1+K} > 1.$$

Again, given that agents face the same prices and have the same wealth level, the central bank can conclude that $c_{1,\omega} = c_1$ for all ω. Then substituting the agents' first order condition into the central bank's objective,

we find,

$$
\begin{aligned}
V^{RE} &= \phi(1)u'(c_1)\int_{\Omega}\left[\frac{\phi_{\omega}^{CB}(1)}{\phi(1)} - \frac{\phi_{\omega}^{CB}(2)}{\phi(2)}\frac{1}{1+K}\right]d\omega \\
&= \phi(1)u'(c_1)\left(\frac{1}{2} - \frac{1}{2}\frac{1}{1+K}\right) \\
&> 0.
\end{aligned}
$$

We summarize these results in the following proposition:

Proposition 2 *For any $K > 0$, agent decisions are collectively biased. Agents choose too much insurance against receiving shocks second relative to receiving shocks first. A central bank that maximizes the expected (ex-post) utility of agents in the economy can improve outcomes by reallocating agent insurance away from safe claims and toward risky claims.*

The reallocation is valuable to the central bank because from the central bank's perspective agents are wasting aggregate collateral by using too many safe claims in a scenario where the more likely shock (the first one) is insurable with risky claims. The portfolio of insurance is important when aggregate collateral is limited because risky claims lock only half as much collateral as safe claims do (i.e. the collateral constraint is $R/2 + S \leq Z$). Thus by shifting agents' decisions toward more risky claims and less safe claims, intermediaries do not have to tie up as much of their resources to cover a second-shock that may never occur.

Finally, note that the central bank reaches these conclusions requiring only knowledge of aggregate probabilities. As we have remarked, the central bank may be much more confused than individual agents about individual θ_{ω}s. In this sense, the central bank may be the *least* informed agent of the economy. The important point is that the central bank does *not* suffer from collective bias.[12]

4.3 Lender of last resort

The somewhat abstract reallocation experiment discussed above highlights how central bank policy choices may differ from robust agents' policy choices. More concretely, we now focus on the value of a lender of last resort (LLR) in the robust equilibrium. We assume that the central bank obtains resources ex-post, at some cost, which it can credibly pledge to agents in the two-shock event.[13] In practice, this pledge may

---

[12] If the central bank were to be uncertain about the values of $\phi(1)$ and $\phi(2)$, then we could overturn our result. In particular, we may imagine a situation in which the central bank is uncertain about these probabilities, and its objective function overweights liquidity crises (i.e. the incidence of both shocks occurring). In this case, the biased central bank will also be subject to the "overinsurance" bias of agents. However, this "bias" is of a different nature than the one we emphasize as it would not be collectively inconsistent with conditional probabilities.

[13] As in Woodford (1990) and Holmstrom and Tirole (1998), the LLR has access to collateral that intermediaries do not (or at least, it has access at a lower cost). Woodford and Holmstrom and Tirole focus on the direct value of intervening using this collateral. Our analysis focuses on the gains beyond the direct value of the intervention.

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, the central bank credibly expands the collateral of the financial sector in the two-shock event by an amount Z^G. Thus, the collateral constraints on intermediaries are altered to

$$\frac{R}{2} + S \leq Z + Z^G.$$

Since we are interested in computing the marginal benefit of intervention, we study an infinitesimal intervention of Z^G.

If the central bank offers more insurance against the two-shock event, this insurance has a direct benefit in terms of reducing the disutility of an adverse outcome. The *direct* benefit of the LLR is,

$$V_{Z^G}^{direct} = 2 \int_\Omega \phi_\omega^{CB}(2) u'(c_{2,\omega}) \, d\omega = \phi(2) u'(c_2)$$

where the 2 in front reflects the fact that an extra unit of collateral for the second shock yields two units of consumption for the group (of measure one half) that is hit by the second wave.

However, the anticipation of the central bank's second-shock insurance leads agents to reoptimize their insurance decisions. Agents reduce insurance against the second-shock (reduce safe claims) and increase their first-shock claims (increase risky claims). The total benefit of the intervention includes both the direct benefit as well as any benefit from portfolio reoptimization:

$$V_{Z^G}^{total} = \int_\Omega \left[\phi_\omega^{CB}(1) u'(c_{1,\omega}) \frac{dc_{1,\omega}}{dZ^G} + \phi_\omega^{CB}(2) u'(c_{2,\omega}) \frac{dc_{2,\omega}}{dZ^G} \right] d\omega.$$

Since $c_{1,\omega} + c_{2,\omega} = (R+S) + S = 2(Z + Z^G)$, the reoptimization of each agents' portfolios must sum up to twice the increase in total resources:

$$\frac{dc_{1,\omega}}{dZ^G} + \frac{dc_{2,\omega}}{dZ^G} = 2.$$

As we have shown, a reoptimization of agents' portfolio away from safe claims and towards risky claims has a first-order benefit in the robust equilibrium. The following proposition confirms this additional value in the context of the lender of last resort:

Proposition 3 *For any $K > 0$, the total value of the lender of last resort policy exceeds its direct value:*

$$V_{Z^G}^{total} > V_{Z^G}^{direct}$$

Proof. We rewrite the total value of intervention as,

$$
\begin{aligned}
V_{Z^G}^{total} &= \int_\Omega \left[\phi_\omega^{CB}(1) u'(c_{1,\omega}) \frac{dc_{1,\omega}}{dZ^G} + \phi_\omega^{CB}(2) u'(c_{2,\omega}) \left(2 - \frac{dc_{1,\omega}}{dZ^G} \right) \right] d\omega \\
&= \phi(2) u'(c_2) + \int_\Omega \left[\phi_\omega^{CB}(1) u'(c_{1,\omega}) - \phi_\omega^{CB}(2) u'(c_{2,\omega}) \right] \frac{dc_{1,\omega}}{dZ^G} d\omega.
\end{aligned}
$$

We can now substitute using the first order condition for agents' decision to arrive at,

$$V_{Z^G}^{total} = V_{Z^G}^{direct} + \frac{\phi(1)}{2}u'(c_1)\left(1 - \frac{1}{1 + \min(K, \bar{K})}\right)\frac{dc_{1,\omega}}{dZ^G} > V_{Z^G}^{direct}$$

The last step follows since agents substitute some of the second-shock insurance of the LLR toward purchasing first-shock insurance. ∎

4.4 Welfare equivalence and paternalism

We conclude this section by revisiting the non-paternalistic case, where we compute welfare using agents' subjective and conservative probabilities. It turns out that there is still an important value to the LLR beyond its objective direct value. There is no indirect benefit from reallocation (by the envelope theorem), but the direct benefit rises since the agents collectively exaggerate the likelihood of the extreme two-shocks event

$$\begin{aligned}\tilde{V}_{Z^G}^{direct,\inf} &= 2\int_\Omega \phi_\omega^\omega(2)u'(c_{2,\omega})d\omega \\ &= \bar{\phi}(2)u'(c_2) \\ &> \phi(2)u'(c_2)\end{aligned}$$

That is, the LLR has an extra benefit which comes from reducing the "anxiety" of robust agents.

Moreover, in the fully robust case we note that

$$\bar{\phi}(2) = \frac{\phi(1) + \phi(2)}{2}$$

to imply:

$$\tilde{V}_{Z^G}^{direct,\inf} = \frac{\phi(1) + \phi(2)}{2}u'(c_2) = \tilde{V}_{Z^G}^{total}$$

This means that in the fully robust case the value of anxiety reduction (from the agents' perspective) is equal to the value of insurance reallocation (from the central bank's perspective). While the exact equivalence result carries over to the partially robust case only for the case of log-preferences, even in the partially robust case there is an extra benefit due to anxiety reduction.

Comparing the paternalistic and non-paternalistic perspectives, we note that agents value the intervention more because they exaggerate the probability of the event that is being subsidized, while the central bank values the intervention because the subsidy reduces private collateral waste. The latter effect means that if the economy were to be drawn infinite times, all agents would receive higher average ex-post utility beyond the direct ex-post benefit of intervention. The agent perceives this additional benefit ex-ante, although for the "wrong" reasons.

5 Price of Liquidity and Collateral Risk

Up to now we have analyzed a model where intermediaries factor zero actuarial cost when setting the prices of liquidity insurance (i.e., $\beta = 0$). For this case, the intermediaries' first order condition pins down the relative price of safe to riskless claims to be constant and independent of the agents' portfolio decisions. Although fixing relative prices allows for a particularly straightforward algebraic analysis, it hides some of the mechanisms underlying the model.

When $\beta > 0$, agents' robustness concerns have important price effects. In the first part of this section, we elaborate on these price effects to further clarify our model. In the second part we show that robustness considerations exacerbate the costs of collateral reductions. The latter point is the main substantive result in this section.

As in Kiyotaki and Moore (1997), reductions in collateral (or collateral values in their model) trigger a financial accelerator. Our model illustrates a further channel for amplification. Reductions in aggregate collateral lead to a flight-to-quality which locks already scarce collateral and amplifies the financial accelerator.

We set $\beta = 1$, so that intermediaries now factor in the actuarial cost of disbursing insurance resources in the first and second shocks.

5.1 Abundant collateral

Starting with a benchmark, suppose 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:

$$p(r) = \frac{\phi(1)}{2} \quad \text{and} \quad p(s) = \frac{\phi(1) + \phi(2)}{2}.$$

Facing these prices, robust agents make their insurance decisions to satisfy:

$$\theta_\omega^\omega \frac{\phi(2)}{\phi(1)} \frac{u'(c_2)}{u'(c_1)} = \frac{p(s)}{p(r)} - 1.$$

Given $p(r)$ and $p(s)$ we can rewrite the first order condition for agents as,

$$\theta_\omega^\omega \frac{u'(c_2)}{u'(c_1)} = 1$$

At $K = 0$, agents choose $c_1 = c_2$. It is easy to verify that for $K > 0$ agents continue to choose $c_1 = c_2$. This choice remains optimal because by choosing $c_1 = c_2$, agents' decisions are robust to the unknown probabilities, and the robust decision rule still does best.

When there is abundant intermediary collateral, the equilibrium is unaffected by agents' robustness concerns. Agents know that intermediaries will not default on their insurance claims. Thus, they are not

concerned about their shock being first or second relative to other agents. Their uncertainty about the shock ordering becomes irrelevant.[14]

Using the budget constraint for agents, we can calculate that facing the actuarially fair prices $p(r)$ and $p(s)$, agents will choose,

$$r = 0; \qquad s = c_1 = c_2 = \frac{2w_0}{\phi(1) + \phi(2)}.$$

In order for intermediaries to not be collateral constrained in equilibrium,[15]

$$Z \geq \underline{Z} \equiv \frac{2w_0}{\phi(1) + \phi(2)}.$$

5.2   Price of liquidity and aggregate shortage

Suppose next that Z falls below $\underline{Z}$. Then, at actuarially fair prices, agents' demand saturates the collateral constraint of the intermediaries and prices rise to reflect the collateral limitation:

$$p(r) = \frac{\phi(1)}{2} + \frac{\mu}{2} \quad \text{and} \quad p(s) = \frac{\phi(1) + \phi(2)}{2} + \mu$$

where $\mu > 0$ is the Lagrange multiplier on the intermediaries' collateral constraint.

We can compute these prices explicitly using the budget constraint for the agent $(p(r)r + p(s)s = w_0)$ and the intermediary collateral constraint $(\frac{r}{2} + s = Z)$:

Proposition 4 *When* $Z < \underline{Z}$, *the prices of risky and safe claims as a function of* c_1 *and* c_2 *are:*

$$p(r) = \frac{w_0}{2Z} + (\phi(1) - \phi(2))\frac{c_2}{4Z} \quad \text{and} \quad p(s) = \frac{w_0}{Z} - (\phi(1) - \phi(2))\frac{c_1 - c_2}{4Z}.$$

1. *A fall in* c_1 *and equal rise in* c_2 *causes the prices of both risky and safe claims to rise.*

2. *A fall in* c_1 *and equal rise in* c_2 *causes the relative price ratio,* $\frac{p(s)}{p(r)} - 1$, *to rise.*

3. *The economy has an **aggregate liquidity shortage**.*

As we have shown in the previous section, when agents' uncertainty rises they purchase more safe claims, thereby decreasing c_1 and increasing c_2. These actions lead to a rise in the equilibrium price of both financial claims (point 1), as well as to a rise in the relative price of safe claims over risky claims (point 2).

---

[14]The irrelevance result in the case of abundant collateral also relates to discussions in the asset pricing literature as to when the actions of "rational" agents will eliminate any effects stemming from the biases of "irrational" agents. In our model, the intermediaries are Bayesian and risk-neutral, but are subject to a constraint whereby they have to hold collateral to back-up any short-sales of financial securities. When there is sufficient intermediary collateral, the short-sale constraint does not bind and prices reflect the rational intermediaries' valuations. It is a feature of our model that agents' insurance decisions are not collectively biased at the actuarially fair prices.

[15]In the $\beta = 0$ case we have studied earlier $\bar{Z}$ is infinite because intermediaries are always at a corner solution in their insurance sales.

Point 3, regarding the aggregate liquidity shortage, is an important aspect of our model. To clarify this point, consider the following thought experiment. Suppose we introduce a small number of Savage (Bayesian) agents into the economy who also are subject to first and second shocks and wish to insure against these shocks. A rise in K increases the price of liquidity insurance for all agents in the economy, both robust and Savage agents. In this sense, flight to quality in our model is not just a problem affecting the decision rules of robust agents. It has a larger cost because the actions of robust agents raise insurance prices and thereby distort the insurance of all agents in the economy.

The mechanism behind the increase in prices is the collateral lock-up effect. This effect becomes particularly transparent if we examine the risky claim. The risky claim price rises as robustness rises; but, as we have shown earlier, robust agents actually *decrease* their demand for risky claims as robustness rises. Thus the increase in price of the risky claim is indirect. Agents increase their demand for safe claims, causing intermediaries to lock-up collateral to cover their shocks and thereby decreasing the amount of collateral that is free to back risky claims. That is, since safe claims lock up twice the amount of collateral as risky claims (i.e. $R/2 + S \leq Z$), as robustness concerns rise, agents force the intermediaries to lock up more collateral. The intermediaries' limited Z is used inefficiently, leading to a smaller effective Z of the economy and higher insurance prices.

We can think of claim prices as 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 agents. Our model shows that the actions of robust agents decreases the effective amount of liquidity provided to the entire economy.

Reductions in aggregate liquidity provision by financial intermediaries and a corresponding rise in financial liquidity premia are two central ingredients in most descriptions of flight to quality events. With a richer market structure – e.g. if each ω was a separate asset market – the rise in the cost of liquidity provision would also be reflected as higher liquidity premia in the asset markets where the agents are active.

5.3 Collateral reduction and flight to quality

We next study how a fall in Z, for fixed K, affects the equilibrium of the economy. Thus far we have described a flight to quality event in terms of increasing K. We now show that our model generates similar implications from a decrease in Z.

The first order condition for an agent's decision is:

$$\theta_\omega^\omega \frac{\phi(2)}{\phi(1)} \frac{u'(c_2)}{u'(c_1)} = \frac{p(s)}{p(r)} - 1 = \frac{2w_0 - (\phi(1) - \phi(2))c_1}{2w_0 + (\phi(1) - \phi(2))c_2} \tag{12}$$

where the second equality follows from Proposition 4. Additionally, from the collateral constraint of inter-

mediaries we know that

$$\frac{c_1}{2} + \frac{c_2}{2} = Z. \tag{13}$$

Equations (12) and (13) jointly define the equilibrium values of c_1 and c_2 as a function of Z. We are interested in examining how c_1 and c_2 change as Z falls below $\underline{Z}$. Let us write $c_2 = 2\alpha(Z)Z$ and $c_1 = 2(1 - \alpha(Z))Z$, and study the function $\alpha(Z)$. As we have remarked earlier, for $Z \geq \underline{Z}$, c_1 is equal to c_2 so that $\alpha(\underline{Z}) = \frac{1}{2}$.

Proposition 5 *Suppose that $K = 0$. Moreover, suppose that $-c\frac{u''(c)}{u'(c)}$ is weakly decreasing in c. Then comparing equilibria across economies with different levels of Z, the (unique) equilibrium value of α falls as Z falls. That is, as Z falls, the ratio of insurance devoted to the first shock over that to the second shock rises.*

Proof. When K equals zero, $\theta_\omega^\omega = 1$. From equation (12) we define,

$$F(\alpha, Z) \equiv \frac{\phi(2)}{\phi(1)} \frac{u'(2\alpha Z)}{u'(2(1-\alpha)Z)} - \frac{w_0 - (\phi(1) - \phi(2))(1-\alpha)Z}{w_0 + (\phi(1) - \phi(2))\alpha Z}.$$

The first order condition in (12) is satisfied when $F(\alpha(Z), Z) = 0$. Implicitly differentiating F, we find that $sign(\frac{d\alpha}{dZ})$ equals the negative of $sign\left(\frac{\frac{\partial F}{\partial Z}}{\frac{\partial F}{\partial \alpha}}\right)$.

$$\begin{aligned}
\frac{\partial F}{\partial Z} &= \frac{1}{Z} \frac{\phi(2)}{\phi(1)} \frac{u'(c_2)}{u'(c_1)} \left(c_2 \frac{u''(c_2)}{u'(c_2)} - c_1 \frac{u''(c_1)}{u'(c_1)} \right) \\
&+ \frac{w_0 - (\phi(1) - \phi(2))(1-\alpha)Z}{w_0 + (\phi(1) - \phi(2))\alpha Z} \left(\frac{(\phi(1) - \phi(2))(1-\alpha)}{w_0 - (\phi(1) - \phi(2))(1-\alpha)Z} + \frac{(\phi(1) - \phi(2))\alpha}{w_0 + (\phi(1) - \phi(2))\alpha Z} \right)
\end{aligned}$$

This term is positive given the regularity condition on $u(c)$.

$$\begin{aligned}
\frac{\partial F}{\partial \alpha} &= 2Z \frac{\phi(2)}{\phi(1)} \frac{u'(c_2)}{u'(c_1)} \left(\frac{u''(c_2)}{u'(c_2)} + \frac{u''(c_1)}{u'(c_1)} \right) \\
&- \frac{(\phi(1) - \phi(2))Z}{w_0 + (\phi(1) - \phi(2))\alpha Z} \left(1 - \frac{w_0 - (\phi(1) - \phi(2))(1-\alpha)Z}{w_0 + (\phi(1) - \phi(2))\alpha Z} \right)
\end{aligned}$$

This term is negative. Therefore $sign(\frac{d\alpha}{dZ})$ is positive. We also note that the equilibrium is unique since $\frac{\partial F}{\partial \alpha}$ is negative for all α. ∎

The $K = 0$ case corresponds to the central bank's solution we have derived in the previous sections. As Z falls, the allocation between first and second shock insurance reflects the limited availability of intermediary resources. Since the first shock is more likely then the second shock, the central bank allocates more resources to the first shock than to the second shock.

Consider next the robust economy where $K > 0$. We begin with a heuristic description of how equilibrium is affected by Z. For $Z \geq \underline{Z}$, agents set c_1 equal to c_2. Starting from this point, suppose that Z falls slightly below $\underline{Z}$. If agents act as if $\theta_\omega^\omega = 1$, they will decrease c_2 and increase c_1 slightly (see Proposition 5). However, this will leave them exposed to the possibility that θ_ω is high. Taking this into account, agents will consider

a larger θ_ω^ω when making decisions. For Z close to $\underline{Z}$, a small increase in θ_ω^ω is sufficient for agents to choose the fully robust insurance decisions. As Z falls sufficiently below $\underline{Z}$, θ_ω^ω reaches its maximum $1 + K$, with c_1 larger than c_2. In this case, agent decisions are partially robust.

Proposition 6 *Suppose that $-c\frac{u''(c)}{u'(c)}$ is weakly decreasing in c. Then comparing (the unique) equilibria across economies with different levels of Z, we have the following results. For a given value of Z, robust agents make equilibrium decisions as if,*

$$\theta_\omega^\omega = 1 + \min(K, \bar{K}(\underline{Z} - Z)),$$

where,

$$1 + \bar{K}(\underline{Z} - Z) \equiv \frac{\phi(1)}{\phi(2)} \frac{2w_0 - (\phi(1) - \phi(2))Z}{2w_0 + (\phi(1) - \phi(2))\bar{Z}}.$$

- *Since $\bar{K}(\underline{Z} - Z)$ increases as Z falls, θ_ω^ω increases (weakly) as Z falls.*

- *For Z such that $\theta_\omega^\omega = 1 + \bar{K}(\underline{Z} - Z)$, agent decisions are fully robust.*

- *For Z small enough that $\theta_\omega^\omega = 1 + K$, agent decisions are partially robust.*

Proof. We define,

$$F(\alpha, Z; \theta_\omega^\omega) \equiv \theta_\omega^\omega \frac{\phi(2)}{\phi(1)} \frac{u'(2\alpha Z)}{u'(2(1-\alpha)Z)} - \frac{w_0 - (\phi(1) - \phi(2))(1-\alpha)Z}{w_0 + (\phi(1) - \phi(2))\alpha Z}.$$

Nature will choose $\theta_\omega^\omega = 1 + K$ if $\alpha < \frac{1}{2}$; $\theta_\omega^\omega = 1 - K$ if $\alpha > \frac{1}{2}$; and is indifferent over $\theta_\omega^\omega \in [1 - K, 1 + K]$ if $\alpha = \frac{1}{2}$. An equilibrium in the robust economy is a value α that satisfies $F(\alpha, Z; \theta_\omega^\omega) = 0$, where θ_ω^ω is optimal for nature at that α.

We proceed as follows. First, we note that the proof of Proposition 5 directly generalizes so that $\frac{\partial F(\alpha, Z; \theta_\omega^\omega)}{\partial Z}$ is positive and $\frac{\partial F(\alpha, Z; \theta_\omega^\omega)}{\partial \alpha}$ is negative for any $\theta_\omega^\omega > 0$. Next, we define $\bar{\alpha}$ as the solution to $F(\bar{\alpha}, Z; \theta_\omega^\omega = 1 + K) = 0$. $\bar{\alpha}$ is uniquely determined since $\frac{\partial F}{\partial \alpha}$ is negative for all α. Similarly, we define $\underline{\alpha}$ as the solution to $F(\underline{\alpha}, Z; \theta_\omega^\omega = 1 - K) = 0$. Since $F(\alpha, Z; \theta_\omega^\omega)$ is increasing in θ_ω^ω, we have that $\bar{\alpha} > \underline{\alpha}$.

For a value of Z for which $\bar{\alpha} < \frac{1}{2}$, the equilibrium is at $\bar{\alpha}$. Nature sets $\theta_\omega^\omega = 1 + K$ if $\alpha < \frac{1}{2}$, and this value of α satisfies market clearing and agents' first order conditions (i.e. $F(\cdot) = 0$). Likewise, for a value of Z for which $\underline{\alpha} > \frac{1}{2}$, the equilibrium is at $\underline{\alpha}$. Nature sets $\theta_\omega^\omega = 1 - K$ if $\alpha > \frac{1}{2}$, and at this value of α satisfies market clearing and agents' first order conditions (i.e. $F(\cdot) = 0$). Finally, for a value of Z for which $\bar{\alpha} \geq \frac{1}{2} \geq \underline{\alpha}$, the equilibrium is at $\alpha = \frac{1}{2}$. Nature is indifferent over the choice of θ_ω^ω and sets θ_ω^ω such that $F(\alpha = \frac{1}{2}, Z; \theta_\omega^\omega) = 0$.

We know that at $Z = \underline{Z}$, $F(\alpha = \frac{1}{2}, Z = \underline{Z}; \theta_\omega^\omega = 1) = 0$. Since F is increasing in θ_ω^ω and F is decreasing in α, we must have that $\bar{\alpha} > \frac{1}{2}$ at $Z = \underline{Z}$. Moreover since decreases in Z decrease the function $F(\cdot)$, $\bar{\alpha}$ must

fall as Z falls. Two conclusions follow: It is not possible that $\underline{\alpha} > \frac{1}{2}$; and for Z sufficiently low $\bar{\alpha}$ must fall below $\frac{1}{2}$. The statements of the proposition follow immediately. ∎

The θ_ω^ω of an agent is a function of both K and Z. As Z falls, robust agents grow increasingly concerned about receiving a shock second and increase θ_ω^ω in making their own decisions. Likewise as K rises, the agents directly increaes θ_ω^ω. The agents' response to a fall in Z is similar to how agents respond to a rise in K. In this sense, our model generates flight to quality from either a rise in K or a fall in Z.

5.4   Risky collateral and policy

In practice, the collateral of financial intermediaries is risky and varies with the underlying state of the economy (as in Kiyotaki and Moore, 1997). Although it is beyond the scope of this paper to endogenize collateral risk, the comparative static with respect to Z provides some insight into the effects of collateral risk in our economy.

Starting from a state of the world where Z is plentiful and agents are confident that they hold safe financial claims, suppose there is a shock that reduces Z. As collateral falls, agents recognize that previously safe claims become risky as they lose their collateral backing. In response, agents act aggressively to shield themselves from risk by shedding risky claims and purchasing safe claims, in a manner consistent with a flight to quality. These actions lock-up the collateral of financial intermediaries, and thereby further reduce the effective collateral of the economy. Our model shows that the robust economy exacerbates the collateral shortage caused by reduced aggregate collateral, amplifying the standard financial accelerator highlighted by Bernanke and Gertler (1989) and Kiyotaki and Moore (1997), among others.

We can also think of the lender of last resort policy in these terms. The problem that agents respond to is increased collateral risk. To a large extent, private sector collateral will always be perceived as risky – i.e. intermediaries can never insulate themselves against default risk. A central bank promise, if credible, is "riskless." The lender of last resort policy derives benefits because agents are certain that the central bank will deliver on its promise. The value of such riskless promises rises when private sector collateral is risky.

6   Moral Hazard and Irrelevance Critiques

In this section we discuss our lender of last resort recommendation in light of two prominent critiques of central bank intervention. The moral hazard critique is that the anticipation of an explicit or implicit insurance policy leads the private sector to underinsure against shocks and may lead the central bank policy to backfire. The irrelevance critique is that the events warranting insurance, such as the Great Depression, are so rare in economies with deep and sound financial markets, that while in principle the central bank insurance is useful, in practice it has minimal ex-ante welfare implications.

In discussing these critiques, we also extend our model to the case of $N > 2$ waves of aggregate shocks.

6.1 Moral hazard critique

The moral hazard 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.[16] Public and private provision of insurance are complements in our model.

We 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.

More formally, let us return to the simpler $\beta = 0$ model and note that the *direct* benefit of intervention in the first shock is:

$$V_{Z^G}^{direct,first} = 2 \int_\Omega \phi_\omega^{CB}(1) u'(c_{1,\omega}) d\omega = \phi(1) u'(c_2).$$

The total benefit, $V_{Z^G}^{total}$, remains unchanged because the relevant constraint remains the constraint on two-shock insurance:

$$
\begin{aligned}
V_{Z^G}^{total} &= \int_\Omega \left[\phi_\omega^{CB}(1) u'(c_{1,\omega}) \frac{dc_{1,\omega}}{dZ^G} + \phi_\omega^{CB}(2) u'(c_{2,\omega}) \left(2 - \frac{dc_{1,\omega}}{dZ^G} \right) \right] d\omega \\
&= \int_\Omega \left[\phi_\omega^{CB}(1) u'(c_{1,\omega}) \left(2 - \frac{dc_{2,\omega}}{dZ^G} \right) + \phi_\omega^{CB}(2) u'(c_{2,\omega}) \frac{dc_{2,\omega}}{dZ^G} \right] d\omega \\
&= \phi(1) u'(c_1) - \int_\Omega \left[\phi_\omega^{CB}(1) u'(c_{1,\omega}) - \phi_\omega^{CB}(2) u'(c_{2,\omega}) \right] \frac{dc_{2,\omega}}{dZ^G} d\omega \\
&= V_{Z^G}^{direct,first} - \frac{\phi(1)}{2} u'(c_1) \left(1 - \frac{1}{1 + \min(K, \bar{K})} \right) \frac{dc_{2,\omega}}{dZ^G} < V_{Z^G}^{direct,first}
\end{aligned}
$$

The anticipation of the central bank's first-shock insurance leads agents to reoptimize their insurance decisions. Agents reduce insurance against the first-shock (reduce risky claims) and increase their second-shock claims (increase safe claims). This private insurance reallocation offsets some of the benefits of the LLR policy. That is, the lender of last resort facility, to be effective and improve private financial markets, has to be a *last* not an intermediate resort.

---

[16] Note 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.

6.2 Multiple shocks

It is clear that the LLR should not intervene during early shocks and instead should only pledge resources for late shocks; but if we move away from our two-shock model to a more realistic context with multiple potential waves of aggregate shocks, how late is late?

To answer this question we introduce multiple shocks into the model. We assume the economy may experience N waves of shocks, each affecting $1/N$th of the agents. The LLR policy takes the following form: The central bank injects $1/(N - T + 1)$ units of liquidity for all shocks after (and including) the Tth wave $(T \leq N)$. For simplicity we focus on the fully robust case.

The value of the intervention as a function of T is,

$$
\begin{aligned}
\tilde{V}_{Z^G}^{direct} &= \frac{1}{N - T + 1} \int_{\Omega} \sum_{i=T}^{N} \phi_{\omega}^{CB}(i) N u'(c_{i,\omega}) d\omega \\
&= \frac{1}{N - T + 1} \frac{1}{N} N \sum_{i=T}^{N} \phi(i) u'(c_i) \\
&= u'(c_1) \frac{1}{N - T + 1} \sum_{i=T}^{N} \phi(i),
\end{aligned}
$$

where we used the fact that one unit of additional collateral allows intermediaries to sell N units of additional insurance. As before, the anticipation of the central bank's insurance leads agents to reoptimize their private insurance decisions. Agents reduce insurance against the publicly insured shocks and increase their private insurance for the rest of the shocks. The total benefit of the intervention includes both the direct benefit as well as any benefit from portfolio reoptimization:

$$
\tilde{V}_{Z^G}^{total} = \int_{\Omega} \sum_{i=1}^{N} \phi_{\omega}^{CB}(i) u'(c_{1,\omega}) \frac{dc_{i,\omega}}{dZ^G} d\omega.
$$

with:

$$
\sum_{i=1}^{N} \frac{dc_{i,\omega}}{dZ^G} = N
$$

so that,

$$
\tilde{V}_{Z^G}^{total} = \frac{1}{N} \sum_{i=1}^{N} \phi(i) u'(c_i) \frac{dc_i}{dZ^G}.
$$

In the fully robust case, c_i and $\frac{dc_i}{dZ^G}$ are the same for all i. Then,

$$
\begin{aligned}
\tilde{V}_{Z^G}^{total} &= u'(c_1) \frac{dc_1}{dZ^G} \frac{1}{N} \sum_{i=1}^{N} \phi(i) \\
&= u'(c_1) \frac{1}{N} \sum_{i=1}^{N} \phi(i)
\end{aligned}
$$

Note that this expression is independent of the intervention rule T. In contrast, it is apparent that $\tilde{V}_{Z^G}^{direct}$ is decreasing with respect to T since the $\phi(i)$'s are monotonically decreasing. Then the ratio:

$$\frac{\tilde{V}_{Z^G}^{total}}{\tilde{V}_{Z^G}^{direct}} = \frac{\frac{1}{N}\sum_{i=1}^{N}\phi(i)}{\frac{1}{N-T+1}\sum_{i=T}^{N}\phi(i)}$$

is strictly greater than one for all $T > 1$ and is increasing with respect to T.

Of course, the above result does not suggest that intervention should occur only in the Nth shock. Instead it suggests that for any given amount of resources available for intervention, the LLR should first pledge resources to the Nth shock and continue to do so until it completely replaces private insurance, it should then move on to the $N - 1$st shock, and so on.

6.3 Irrelevance critique

The multiple shock model also clarifies another benefit of late intervention. As T rises, events that are being insured by the LLR become increasingly less likely. If we take the case where the shadow cost of the LLR resources for the central bank is constant, the expected cost of the LLR policy falls as T rises, while the expected benefit remains constant.

In other words, as T rises, it is the private sector that increasingly improves the allocation of scarce private resources to early and more likely aggregate shocks, thereby reducing the extent of the flight-to-quality phenomenon. In contrast, the central bank plays a decreasingly small role in terms of the expected value of resources actually disbursed, as T increases.

Thus, while a well designed LLR policy may indeed have a direct effect only in highly unlikely events (hence the irrelevance critique), its main benefits are felt during more likely and less extreme events. These benefits come from the impact of the policy in unlocking private collateral.

7 Final remarks

Flight to quality is a pervasive phenomenon that exacerbates the impact of recessionary shocks and financial accelerators. 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] Caballero, R.J. and A. Krishnamurthy, 2005, "Financial System Risk and Flight to Quality," NBER WP # 11996.

[4] Diamond, D.W. and P.H. Dybvig, 1983, "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy* 91, pp. 401-419.

[5] Dow, J. and S. Werlang, 1992, "Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio," *Econometrica 60*, p. 197-204.

[6] 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.

[7] Easley, David and Maureen O'Hara, 2005, "Regulation and Return: The Role of Ambiguity," working paper, Cornell University.

[8] Epstein, L., 2001, "Sharing Ambiguity," *American Economic Review*, 91(2), 45-50.

[9] Epstein, L. and M. Schneider, 2004, "Recursive Multiple Priors," *Journal of Economic Theory,* 113(1), 1-31.

[10] Epstein, L. and T. Wang, 1994, "Intertemporal Asset Pricing Under Knightian Uncertainty,' *Econometrica 62*, 283-322.

[11] Gatev, Evan, and Phil Strahan, 2005, "Banks' Advantage in Hedging Liquidity Risk: Theory and Evidence from the Commercial Paper Market," forthcoming, *Journal of Finance.*

[12] Gilboa, I. and D. Schmeidler, 1989, "Maxmin Expected Utility with Non-unique Priors," *Journal of Mathematical Economics,* 18, 141-153.

[13] Greenspan, A, 2004, "Risk and Uncertainty in Monetary Policy," Remarks for the AEA Meetings, January.

[14] Hansen, L. and T. Sargent, 1995, "Discounted Linear Exponential Quadratic Gaussian Control,' *IEEE Transactions on Automatic Control 40*, 968-971.

[15] Hansen, L.P. and T.J. Sargent, 2003, "Robust Control of Forward-looking Models," *Journal of Monetary Economics* 50(3), 581-604.

[16] Hansen, L.P., T.J. Sargent, G.A. Turmuhambetova, N. Williams, 2004, "Robust Control and Model Misspecification," April, Mimeo.

[17] Holmstrom, Bengt and Jean Tirole, "Private and Public Supply of Liquidity," *Journal of Political Economy*, 106(1), 2/98, pp. 1-40.

[18] 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.

[19] Kiyotaki, N. and J. Moore, 1997: "Credit Cycles," *Journal of Political Economy,* 105 (2), 211-248.

[20] Krishnamurthy, Arvind, 2003, "Collateral Constraints and the Amplification Mechanism," *Journal of Economic Theory*, 111(2), 277-292.

[21] Krishnamurthy, Arvind, 2002, "The Bond/Old-Bond Spread," *Journal of Financial Economics 66(2)*, 463-506, 2002.

[22] Kyle, A. and W. Xiong, 2001, "Contagion as a Wealth Effect", *Journal of Finance* 56, 1401-1440.

[23] Longstaff, Francis, 2004, "The Flight-to-Liquidity Premium in U.S. Treasury Bond Prices," *Journal of Business 77(3)*, 511-526.

[24] Routledge, Bryan and Stanley Zin, 2004, "Model Uncertainty and Liquidity," working paper, Carnegie Mellon University.

[25] Sims, Christopher A., 2001, "Pitfalls of a Minimax Approach to Model Uncertainty," working paper, Princeton University.

[26] Skiadas, Costis, 2003, "Robust Control and Recursive Utility," *Finance and Stochastics 7*, 475-489.

[27] Stock, J. and Watson, M., 1989. New indexes of coincident and leading economic indicators. NBER Macroeconomics Annual, 351-394.

[28] Woodford, M., "Public Debt as Private Liquidity," *American Economic Review, Papers and Proceedings 80*, May 1990, pp. 382-88.