Collective Risk Management in a Flight to Quality Episode

RICARDO J. CABALLERO and ARVIND KRISHNAMURTHY*

October 2, 2007

Abstract

Severe flight to quality episodes involve uncertainty about the environment and not only risk about asset payoffs. The uncertainty is triggered by unusual events and untested financial innovations that lead agents to question their world-view. We present a model of crises and central bank policy that incorporates *Knightian uncertainty.* The model can explain crisis regularities such as market-wide capital immobility, agents' disengagement from risk, and liquidity hoarding. We identify a social cost of these behaviors, and a benefit of a lender of last resort facility. The benefit is particularly high because public and private insurance are complements during uncertainty-driven crises.

*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, Jon Faust, Xavier Gabaix, Jordi Gali, Michael Golosov, Campbell Harvey, William Hawkins, Burton Hollifield, Bengt Holmstrom, Dimitri Vayanos, Ivan Werning , an associate editor and anonymous referee, as well as seminar participants at Columbia, DePaul, Imperial College, LBS, Northwestern, MIT, Wharton, NY Fed Liquidity Conference, NBER MIDM Conference, UBC Summer Conference, Bank of England, and the Central Bank of Chile for their comments. Vineet Bhagwat, David Lucca, and Alp Simsek provided excellent research assistance. Caballero thanks the NSF for financial support. This paper covers the same substantive issues as (and hence replaces) "Flight to Quality and Collective Risk Management," NBER WP # 12136.

"... 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, and 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 at by Greenspan (2004), the Fed intervened early and stood ready to intervene as much as needed to prevent a meltdown.

In this paper we present a model to study the benefits of central bank actions during flight to quality episodes. Our model has two key ingredients: capital/liquidity shortages and Knightian uncertainty (Knight (1921)). The capital shortage ingredient is a recurring theme in the empirical and theoretical literature on financial crises and requires little motivation. Knightian uncertainty is less commonly studied, but practitioners perceive it as a central ingredient to flight to quality episodes (see Greenspan's quote).

Most flight to quality episodes are triggered by unusual or unexpected events. In 1970, the Penn-Central Railroad's default on prime rated commercial paper caught the market by surprise and forced investors to re-evaluate their models of credit risk. The ensuing dynamics temporarily shut out a large segment of commercial paper borrowers from a vital source of funds. In October 1987, the speed of the stock market decline took investors and market makers by surprise, causing them to question their models. Investors pulled back from the market while key market-makers widened bid-ask spreads. 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-

[1] See Table 1 (part A) in Barro (2006) for a list of extreme events, measured in terms of fall in GDP, in developed economies during the 20th century.

testing" models. Finally, after 9/11, regulators were concerned that commercial banks would respond to the increased uncertainty over the status of other commercial banks by individually hoarding liquidity and that such actions would lead to gridlock in the payments system.[2]

The common aspects of investor behavior across these episodes – re-evaluation of models, conservatism, and disengagement from risky activities – indicate that these episodes involved *Knightian uncertainty* (i.e., immeasurable risk) and not merely an increase in risk exposure. The emphasis on tail outcomes and worst-case scenarios in agents' decision rules suggests uncertainty aversion. Finally, an important observation about these events is that when it comes to flight to quality episodes, history seldom repeats itself. Similar magnitudes of commercial paper default (Mercury Finance in 1997) or stock market pullbacks (mini-crash of 1989) did not lead to similar investor responses. Today, as opposed to in 1998, market participants understand that correlations should be expected to rise during periods of reduced liquidity. Creditors understand the risk involved in lending to hedge funds. While in 1998 hedge funds were still a novel financial vehicle, the large reported losses of the Amaranth hedge fund in 2006 barely caused a ripple in financial markets. The one-of-a-kind aspect of flight to quality episodes suggests that these events are fundamentally about uncertainty rather than risk.[3]

Section I of the paper lays out a model of financial crises based on liquidity shortages and Knightian uncertainty. We analyze the model's equilibrium and show that an increase in Knightian uncertainty or decrease in aggregate liquidity can reproduce flight to quality effects. In the model, when an agent is faced with Knightian uncertainty, he considers the worst-case among the scenarios over which he is uncertain. This modeling of agent decision making and Knightian uncertainty draws from the decision-theory literature, and in particular from Gilboa and Schmeidler (1989). When the aggregate quantity of liquidity is limited, the Knightian agent grows concerned that he will be caught in a situation where he needs liquidity, but there is not enough liquidity available to him. In this context, agents react by shedding risky financial claims in favor of safe and uncontingent claims. Financial intermediaries become self-protective and hoard liquidity. Investment banks and trading desks turn conservative in their allocation of risk capital. They lock up capital and become unwilling to

[2]See Calomiris (1994) on the Penn-Central default, Melamed (1988) on the 1987 market crash, Scholes (2000) on the events of 1998, and Stewart (2002) or McAndrews and Potter (2002) on 9/11.

[3]This observation suggests a possible way to empirically disentangle uncertainty aversion from risk aversion or extreme forms of risk aversion such as negative skew aversion. A risk averse agent behaves conservatively during times of high risk – it does not matter whether the risk involves something new or not. For an uncertainty averse agent, new forms of risk elicit the conservative reaction.

flexibly move it across markets.

The main results of our paper are in Sections II and III. As indicated by former Fed Chairman Greenspan's comments, the Fed has historically intervened during flight to quality episodes. We analyze the macroeconomic properties of the equilibrium and study the effects of central bank actions in our environment. First, we show that Knightian uncertainty leads to a *collective* bias in agents' actions: Each agent covers himself against his own worst-case scenario, but the scenario that the collective of agents are guarding against is impossible, and known to be so despite agents' uncertainty about the environment. We show that agents' conservative actions such as liquidity hoarding and locking-up of capital are macroeconomically costly because scarce liquidity goes wasted. Second, we show that central bank policy can be designed to alleviate the over-conservatism. A lender of last resort (LLR), even one facing the same incomplete knowledge that triggers agents' Knightian responses, finds that committing to add liquidity in the unlikely event that the private sector's liquidity is depleted is beneficial. Agents respond to the LLR by freeing-up capital and altering decisions in a manner that wastes less private liquidity. Public and private provision of insurance are complements in our model: each pledged dollar of public intervention in the extreme event is matched by a comparable private sector reaction to free-up capital. In this sense, the Fed's LTCM restructuring was important not for its direct effect, but because it served as a signal of the Fed's readiness to intervene should conditions worsen. We also show that the LLR must be a *last-resort* policy: If liquidity injections take place too often, the policy exacerbates the private sector's mistakes and reduces the value of intervention. This occurs for reasons akin to the moral hazard problem identified with the LLR.

Our model is most closely related to the literature on banking crises initiated by Diamond and Dybvig (1983).[4] While our environment is a variant of Diamond and Dybvig's, it does not include the sequential service constraint of Diamond and Dybvig, and instead emphasizes Knightian uncertainty. The change in modeling leads to different applications of our crisis model. The model applies to a wider set of financial intermediaries than commercial banks financed by demandable deposit contracts. More importantly, because our model of crises centers on Knightian uncertainty, its insights most directly apply to circumstances of market-wide uncertainty, such as the new financial innovations or events we have discussed above.

At a theoretical level, the sequential service constraint in the Diamond and Dybvig model creates a coordination failure. The bank "panic" occurs because each depositor runs con-

[4]The literature on banking crises is too large to discuss here. See Gorton and Winton (2003) for a survey of this literature.

jecturing that other depositors will run. The externality in the depositor's run decisions is central to their model of crises.[5] In our model, it is an increase in Knightian uncertainty that generates the panic behavior. Of course, crises in reality may reflect both the type of externalities that Diamond and Dybvig highlight and the uncertainties that we study. In the Diamond and Dybvig analysis, the LLR is always beneficial because it rules out the "bad" run equilibrium caused by the coordination failure.[6] As noted above, our model's prescriptions center on situations of market-wide uncertainty. In particular, our model prescribes that the benefit of the LLR is highest when there is *both* insufficient aggregate liquidity and Knightian uncertainty.

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 that aggregate collateral 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.

Routledge and Zin (2004) and Easley and O'Hara (2005) are two related analyses of Knightian uncertainty in financial markets.[7] Routledge and Zin 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

[5]More generally, other papers in the crisis literature also highlight how investment externalities can exacerbate crises. Some examples in this literature include Allen and Gale (1994), Gromb and Vayanos (2002), Caballero and Krishnamurthy (2003), or Rochet and Vives (2004). In many of these models, incomplete markets leads to an inefficiency that creates a role for central bank policy (see Rochet and Vives (2004) or Allen and Gale (2005))

[6]Rochet and Vives (2004) and Goldstein and Pauzner (2005) present variants of the Diamond and Dybvig (1983) which have a unique equilibrium. By applying global games techniques, they eliminate the multiplicity of equilibria in the Diamond and Dybvig model. Moreover, they show that an investment externality (coordination failure) remains among the bank's creditors in the unique equilibrium. Like us, Rochet and Vives show that the LLR is only necessary when the value of bank assets are below a threshold.

[7]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), Epstein (2001), 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.

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 uncertainty 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. Easley and O'Hara study how uncertainty aversion affects investor participation in stock markets, while the focus of our study is on uncertainty aversion and financial crises.

I. The Model

We study a model conditional on entering a turmoil period where liquidity risk and Knightian uncertainty coexist. Our model is silent on what triggers the episode. In practice, we think that the occurrence of an unusual event, such as the Penn Central default or 9/11, causes agents to re-evaluate their models and triggers robustness concerns. Our goal is to present a model to study the role of a centralized liquidity provider such as the central bank.

A. The Environment

PREFERENCES AND SHOCKS

The model has a continuum of competitive agents, which are indexed by $\omega \in \Omega \equiv [0,1]$. An agent may receive a liquidity shock in which he needs some liquidity immediately. We view these 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 shocks are correlated across agents. With probability $\phi(1)$, the economy is hit by a first-wave of liquidity shocks. In this wave, a randomly chosen group of one-half of the agents have liquidity needs. We denote the probability of agent ω receiving a shock in the first wave by $\phi_\omega(1)$, and note that,

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

Equation (1) states that on average, across all agents, the probability of an agent receiving a shock in the first wave is $\frac{\phi(1)}{2}$.

With probability $\phi(2|1)$, a second wave of liquidity shocks hits the economy. In the second wave of liquidity shocks, the other half of the agents need liquidity. Let $\phi(2) = \phi(1)\phi(2|1)$. The probability for agent ω of being in this second wave is $\phi_\omega(2)$, which satisfies,

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

With probability $1 - \phi(1) > 0$ the economy experiences no liquidity shocks.

We note that the sequential shock structure means that,

$$\phi(1) > \phi(2) > 0. \tag{3}$$

This condition states that, in aggregate, a single-wave event is more likely than the two-wave event. We refer to the two-wave event as an extreme event, capturing an unlikely but severe liquidity crisis in which many agents are affected. Relation (3), deriving from the sequential shock structure, plays an important role in our analysis.

We model the liquidity shock as a shock to preferences (e.g., as in Diamond and Dybvig (1983)). Agent ω receives utility:

$$U^\omega(c_1, c_2, c_T) = \alpha_1 u(c_1) + \alpha_2 u(c_2) + \beta c_T. \tag{4}$$

We define $\alpha_1 = 1, \alpha_2 = 0$ if the agent is in the early wave; $\alpha_1 = 0, \alpha_2 = 1$ if the agent is in the second wave; and, $\alpha_1 = 0, \alpha_2 = 0$ if the agent is not hit by a shock. We will refer to the first shock date as "date 1," the second shock date as "date 2," and the final date as "date T."

The function $u : \mathcal{R}_+ \to \mathcal{R}$ is twice continuously differentiable, increasing, strictly concave and satisfies the condition $\lim_{c\to 0} u'(c) = \infty$. Preferences are concave over c_1 and c_2 and linear over c_T. We view the preference over c_T as capturing a time, in the future, when market conditions are normalized and the trader is effectively risk neutral. The concave preferences over c_1 and c_2 reflect the potentially higher marginal value of liquidity during a time of market distress. The discount factor, β, can be thought of as an interest rate $(\frac{1}{\beta} - 1)$ facing the trader.

ENDOWMENT AND SECURITIES

Each agent is endowed with Z units of goods. These goods can be stored at gross return of one, and then liquidated if an agent receives a liquidity shock. If we interpret the agents of the model as financial traders, we may think of Z as the capital or liquidity of a trader.

Agents can also trade financial claims that are contingent on shock realizations. As we will show, these claims allow agents who do not receive a shock to insure agents who do receive a shock.

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. Markets are complete. There are claims on all histories of shock realizations. We will be more precise in specifying these contingent claims when we analyze the equilibrium.

PROBABILITIES AND UNCERTAINTY

Agents trade contingent claims to insure against their liquidity shocks. In making the insurance decisions, agents have a probability model of the liquidity shocks in mind.

We assume that agents know the aggregate shock probabilities, $\phi(1)$ and $\phi(2)$. We may think that agents observe the past behavior of the economy and form precise estimates of these aggregate probabilities. However, and centrally to our model, the same past data does not reveal whether a given ω is more likely to be in the first wave or the second wave. Agents treat the latter uncertainty as Knightian.

Formally, we use $\phi_\omega(1)$ to denote the true probability of agent ω receiving the first shock, and $\phi_\omega^\omega(1)$ to denote agent-ω's *perception* of the relevant true probability (similarly for $\phi_\omega(2)$ and $\phi_\omega^\omega(2)$). We assume that each agent ω knows his probability of receiving a shock either in the first or second wave, $\phi_\omega(1) + \phi_\omega(2)$, and thus the perceived probabilities satisfy:[8]

$$\phi_\omega^\omega(1) + \phi_\omega^\omega(2) = \phi_\omega(1) + \phi_\omega(2) = \frac{\phi(1) + \phi(2)}{2}. \tag{5}$$

We define,

$$\theta_\omega^\omega \equiv \phi_\omega^\omega(2) - \frac{\phi(2)}{2}. \tag{6}$$

That is, θ_ω^ω reflects how much agent ω's probability assessment of being second is higher than the average agent in the economy's true probability of being second. This relation also implies that,

$$-\theta_\omega^\omega = \phi_\omega^\omega(1) - \frac{\phi(1)}{2}.$$

Agents consider a range of probability-models θ_ω^ω in the set Θ, with support $[-K, +K]$ $(K < \phi(2)/2))$, and design insurance portfolios that are robust to their model uncertainty. We follow Gilboa and Schmeidler's (1989) Maximin Expected Utility representation of Knightian

[8]For further clarification of the structure of shocks and agents' uncertainty, see the event tree that is detailed in the Appendix.

uncertainty aversion and write:

$$\max_{(c_1,c_2,c_T)} \min_{\theta_\omega^\omega \in \Theta} E_0[U^\omega(c_1, c_2, c_T)|\theta_\omega^\omega] \tag{7}$$

where K captures the extent of agents' uncertainty.

In a flight to quality event, such as the Fall of 1998 or 9/11, agents are concerned about systemic risk and unsure of how this risk 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 he owed to others but was unsure whether resources owed to him would arrive (see, e.g., Stewart (2002), or McAndrews and Potter (2002)). 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 (7) 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.

It is also important to note that the objective function in (7) works through altering the probability distribution used by agents. That is, given an agent's uncertainty, the min operator in (7) has the agent making decisions using the worst-case probability distribution over this uncertainty. This objective is different from one which asymmetrically penalizes bad outcomes. That is, a loss aversion or negative skewness aversion objective function leads an agent to worry about worst-cases through the utility function U^ω directly. This asymmetric utility function model predicts that agents *always* worry about the downside. Our Knightian uncertainty objective predicts that agents worry about the downside in particular during times of model uncertainty. As discussed in the introduction, it appears that flight to quality episodes have a "newness/uncertainty" element, which our modeling can capture.

The distinction is also relevant because probabilities have to satisfy adding up constraints across all agents – i.e. $\int_\Omega \theta_\omega d\omega = 0$. Indeed, the term "collective" bias we use refers to a situation when agents' individual probability distributions from the min operator in (7) fails to satisfy an adding up constraint. As we will explain, the efficiency results we present later in the paper stem from this aspect of our model.

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, Z, and $u(c)$ are the same for all ω. Moreover, this information is common knowledge. As noted above, $\phi(1)$ and $\phi(2)$ are also common knowledge.

B. A Benchmark

We begin by analyzing the problem when $K = 0$. This case clarifies the nature of cross-insurance that is valuable in our economy as well. We derive the equilibrium as a solution to a planning problem, where the planner allocates the Z across agents as a function of shock realizations.

Figure 1 below describes the event tree of the economy. The economy may receive zero, one, or two waves of shocks. An agent ω may be affected in the first or second wave in the two wave case, or may be affected or not affected in the one wave event. We denote $s = (\ \#$ of waves, ω's shock) as the state for agent ω. Agent ω's allocation as a function of the state is denoted by C^s where in the event of agent ω being affected by a shock, the agent receives a consumption allocation upon incidence of the shock, as well as a consumption allocation at date T. For example, if the economy is hit by two waves of shocks in which agent ω is affected by the first wave, we denote the state as $s = (2, 1)$ and agent ω's allocation as (c_1, c_T^s). $\mathcal{C} = \{C^s\}$ is the consumption plan for agent ω (equal to that for every agent, by symmetry).

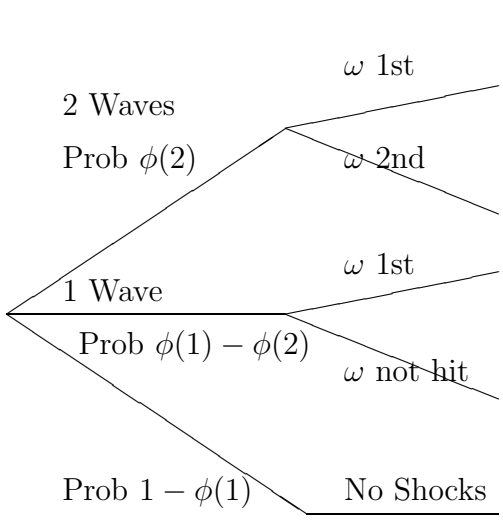
We note that c_1 is the same in both state $(2, 1)$ and state $(1, 1)$. This is because of the sequential shock structure in the economy. An agent who receives a shock first needs resources at that time, and the amount of resources delivered cannot be made contingent on whether the one or two wave event transpires.

Figure 1 also gives the probabilities of each state s. Since agents are ex-ante identical and $K = 0$, each agent has the same probability of arriving at state s. Thus we know that $\phi_\omega(2) = \phi(2)/2$, which implies that the probability of ω being hit by a shock in the second wave is one-half. Likewise, the probability of ω being hit by a shock in the first wave is one-half. These computations lead to the probabilities given in Figure 1.

The planner's problem is to solve,

$$\max_{\mathcal{C}} \sum p^s U^\omega(C^s)$$

subject to resource constraints that for every shock realization, the promised consumption



$s = (\#$ waves, ω's shock)	p^s	C^s
(2,1)	$\phi(2)/2$	$(c_1, c_T^{2,1})$
(2,2)	$\phi(2)/2$	$(c_2, c_T^{2,2})$
(1,1)	$(\phi(1) - \phi(2))/2$	$(c_1, c_T^{1,1})$
(1,no)	$(\phi(1) - \phi(2))/2$	$(c_T^{1,no})$
(0,no)	$1 - \phi(1)$	$(c_T^{0,no})$

Figure 1. Benchmark case. The tree on the left pictures the possible states realized for agent ω. The economy can go through zero (lower branch), one (middle branch), or two (upper branch) waves of shocks. In each of these cases, agent ω may or may not be affected. The first column lists the state, s, for agent ω corresponding to that branch of the tree. The second column lists the probability of state s occurring. The last column lists the consumption bundle given to the agent by the planner in state s.

amounts are not more than the total endowment of Z:

$$\begin{aligned} c^{0,no} &\leq Z \\ \frac{1}{2}\left(c_1 + c_T^{1,1} + c_T^{1,no}\right) &\leq Z \\ \frac{1}{2}\left(c_1 + c_T^{2,1} + c_2 + c_T^{2,2}\right) &\leq Z, \end{aligned}$$

as well as non-negativity constraints that for each s, every consumption amount in C^s is non-negative.

It is obvious that if shocks do not occur, then the planner will give Z to each of the agents for consumption at date T. Thus $c_T^{0,no} = Z$ and we can drop this constant from the objective. We rewrite the problem as:

$$\max_{c} \quad \frac{\phi(1) - \phi(2)}{2}\left(u(c_1) + \beta c_T^{1,1} + \beta c_T^{1,no}\right) + \frac{\phi(2)}{2}\left(u(c_1) + u(c_2) + \beta c_T^{2,1} + \beta c_T^{2,2}\right)$$

subject to resource and non-negativity constraints.

Observe that $c_T^{1,1}$ and $c_T^{1,no}$ enter as a sum in both objective and constraints. Without loss of generality we set $c_T^{1,1} = 0$. Likewise, $c_T^{2,1}$ and $c_T^{2,2}$ enter as a sum in both objective and

constraints. Without loss of generality we set $c_T^{2,1} = 0$. The reduced problem is:

$$\max_{(c_1, c_2, c_T^{1,no}, c_T^{2,2})} \quad \phi(1)u(c_1) + \phi(2)(u(c_2) + \beta c_T^{2,2}) + (\phi(1) - \phi(2))\,\beta c_T^{1,no}$$

subject to:

$$\begin{aligned} c_1 + c_T^{1,no} &= 2Z \\ c_1 + c_2 + c_T^{2,2} &= 2Z \\ c_1, c_2, c_T^{1,no}, c_T^{2,2} &\geq 0. \end{aligned}$$

Note that the resource constraints must bind. The solution hinges on whether the non-negativity constraints on consumption bind or not.

If the non-negativity constraints do not bind, then the first order condition for c_1 and c_2 yield:

$$c_1 = c_2 = u'^{-1}(\beta) \equiv c^*.$$

The solution implies that,

$$c_T^{2,2} = 2(Z - c^*), \quad c_T^{1,no} = 2Z - c^*.$$

Thus the non-negativity constraints do not bind if $Z \geq c^*$. We refer to this case as one of *sufficient aggregate liquidity*. When Z is large enough, agents are able to finance a consumption plan in which marginal utility is equalized across all states. At the optimum, agents equate the marginal utility of early consumption with that of date T consumption, which is β given the linear utility over c_T. A low value of β means that agents discount the future heavily and require more early consumption. Loosely speaking we can think of this case as one where an agent is "constrained" and places a high value on current liquidity. As a result, the economy needs more liquidity (Z) to satisfy agents' needs.

Now consider the case in which there is *insufficient* liquidity so that agents are not able to achieve full insurance. This is the case where $Z < c^*$. It is obvious that $c_T^{2,2} = 0$ in this case (i.e. use all of the limited liquidity towards shock states). Thus, for a given c_1 we have that $c_2 = c_T^{1,no} = 2Z - c_1$ and the problem is,

$$\max_{c_1} \quad \phi(1)u(c_1) + \phi(2)u(2Z - c_1) + (\phi(1) - \phi(2))\,\beta(2Z - c_1) \tag{8}$$

with first order condition,

$$u'(c_1) = \frac{\phi(2)}{\phi(1)} u'(2Z - c_1) + \beta\left(1 - \frac{\phi(2)}{\phi(1)}\right). \tag{9}$$

Since $u'(2Z - c_1) > \beta$ (i.e. $c_2 < c^*$) we can order:

$$\beta < u'(c_1) < u'(2Z - c_1) \quad \Rightarrow \quad c_1 > Z. \tag{10}$$

The last inequality on the right of (10) is the important result from the analysis. Agents who are affected by the first-wave of shocks receive more liquidity than agents who are affected by the second-wave. There is cross-insurance between agents. Intuitively, this is because the probability of the second-wave occurring is strictly smaller than that of the first-wave (or, equivalently, conditional on the first wave having taken place there is a chance the economy is spared of a second wave). Thus, when liquidity is scarce (small Z) it is optimal to allocate more of the limited liquidity to the more likely shock. On the other hand, when liquidity is plentiful (large Z) the liquidity allocation of each agent is not contingent on the *order* of the shocks. This is because there is enough liquidity to cover all shocks.

We summarize these results as follows:

PROPOSITION 1 *: The equilibrium in the benchmark economy with $K = 0$ has two cases:*

- *The economy has insufficient aggregate liquidity if $Z < c^*$. In this case,*

  $$c^* > c_1 > Z > c_2.$$

  Agents are partially insured against liquidity shocks. First wave liquidity shocks are more insured than second wave liquidity shocks.

- *The economy has sufficient aggregate liquidity if $Z \geq c^*$. In this case,*

  $$c_1 = c_2 = c^*$$

  and agents are fully insured against liquidity shocks.

Flight to quality effects, and a role for central bank intervention, arise only in the first case (insufficient aggregate liquidity). This is the case we analyze in detail in the next sections.

C. Implementation

There are two natural implementations of the equilibrium: financial intermediation, and trading in shock-contingent claims.

In the intermediation implementation, each agent deposits Z in an intermediary initially and receives the right to withdraw $c_1 > Z$, if he receives a shock in the first-wave. Since

shocks are fully observable, the withdrawal can be conditioned on the agents' shocks. Agents who do not receive a shock in the first wave own claims to the rest of the intermediary's assets ($Z - c_1 < c_1$). The second group of agents either redeem their claims upon incidence of the second wave of shocks, or at date T. Finally, if no shocks occur, the intermediary is liquidated at date T and all agents receive Z.

In the contingent claims implementation, each agent purchases a claim that pays $2(c_1 - Z) > 0$ in the event that the agent receives a shock in the first wave. The agent sells an identical claim to every other agent, paying $2(c_1 - Z)$ in case of the first wave shock. Note that this is a zero-cost strategy since both claims must have the same price.

If no shocks occur, agents consume their own Z. If an agent receives a shock in the first wave, he receives $2(c_1 - Z)$ and pays out $c_1 - Z$ (since one-half of the agents are affected in the first wave), to net $c_1 - Z$. Added to his liquidity of Z, this gives total liquidity of c_1. Any later agent has $Z - (c_1 - Z) = 2Z - c_1$ units of liquidity to either finance a second shock, or as date T consumption.

Finally, note that if there is sufficient aggregate liquidity either the intermediation or contingent claims implementation achieves the optimal allocation. Moreover, in this case, the allocation is also implementable by self-insurance. Each agent keeps his Z and liquidates $c^* < Z$ to finance a shock. The self-insurance implementation is not possible when $Z < c^*$, because the allocation requires each agent to receive more than his endowment of Z if the agent is hit first.

D. $K > 0$ Robustness Case

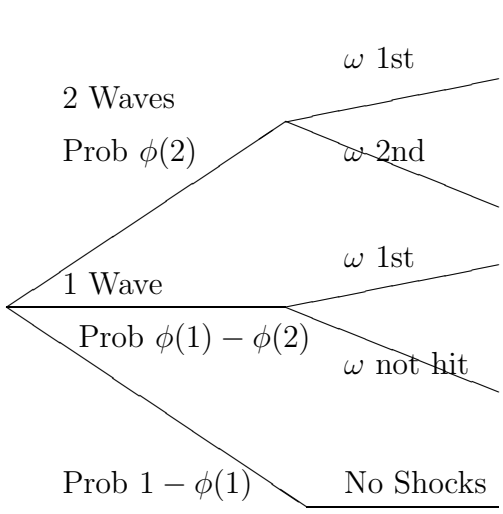
We now turn to the general problem when $K > 0$. Once again, we derive the equilibrium by solving a planning problem where the planner allocates the Z to agents as a function of shocks. When $K > 0$, agents make decisions based on a "worst-case" for the probabilities. This decision making process is encompassed in the planning problem by altering the planner's objective to,

$$\max_{\mathcal{C}} \min_{\theta^\omega_\omega \in \Theta} \sum p^{s,\omega} U(C^s) \tag{11}$$

The only change in the problem relative to the $K = 0$ case is that probabilities are based on the worst-case min rule.

Figure 2 redraws the event tree now indicating agent's worst-case probabilities. We use the notation that $\phi^\omega_\omega(2)$ is agent ω's worst-case probability of being hit second. In our setup, this assessment only matters when the economy is going through a two-wave event in which

the agent is unsure if other agents' shocks are going to occur before or after agent ω's.[9]



$s = (\#$ waves, ω's shock)	$p^{s,\omega}$
(2,1)	$\phi^{\omega}_{\omega}(1) - \frac{\phi(1)-\phi(2)}{2}$
(2,2)	$\phi^{\omega}_{\omega}(2)$
(1,1)	$(\phi(1) - \phi(2))/2$
(1,no)	$(\phi(1) - \phi(2))/2$
(0,no)	$1 - \phi(1)$

Figure 2. Robustness case. The tree on the left pictures the possible states realized for agent ω. The economy can go through zero (lower branch), one (middle branch), or two (upper branch) waves of shocks. In each of these cases, agent ω may or may not be affected. The first column lists the state, s, for agent ω corresponding to that branch of the tree. The second column lists the agent's *perceived* probability of state s occurring.

We simplify the problem following some of the steps of the previous derivation. $c_T^{0,no}$ must be equal to Z. Since the problem in the one-wave node is the same as in the previous case, we observe that $c_T^{1,1}$ and $c_T^{1,no}$ enter as sum in both objective and constraint and choose $c_T^{1,1} = 0$. The reduced problem is then,

$$V(\mathcal{C}; \theta^{\omega}_{\omega}) \equiv \max_{\mathcal{C}} \min_{\theta^{\omega}_{\omega} \in \Theta} \ \phi^{\omega}_{\omega}(1)u(c_1) + (\phi(2) - \phi^{\omega}_{\omega}(2))\,\beta c_T^{2,1} + \phi^{\omega}_{\omega}(2)\left(u(c_2) + \beta c_T^{2,2}\right) + \frac{\phi(1) - \phi(2)}{2}\beta c_T^{1,no} \tag{12}$$

The first two terms in this objective are the utility from the consumption bundle if the agent is hit first (either in the one wave or two wave event). The third term is the utility from the consumption bundle if the agent is hit second. The last term is the utility from the bundle when the agent is not hit in a one-wave event.

[9] We derive the probabilities as follows. $p^{2,2,\omega} = \phi^{\omega}_{\omega}(2)$ by definition. This implies that $p^{2,1,\omega} = \phi(2) - \phi^{\omega}_{\omega}(2)$ since the probabilities have to sum up to the probability of a two wave event ($\phi(2)$). We rewrite $p^{2,1,\omega} = \phi(2) - \phi^{\omega}_{\omega}(2) = \phi^{\omega}_{\omega}(1) - \frac{\phi(1)-\phi(2)}{2}$ using relation (5). The probability of ω being hit first is $\phi^{\omega}_{\omega}(1) = p^{2,1,\omega} + p^{1,1,\omega}$. Substituting for $p^{2,1,\omega}$, we can rewrite this to find that $p^{1,1,\omega} = \frac{\phi(1)+\phi(2)}{2}$. Finally, $p^{1,1,\omega} + p^{1,no,\omega} = \phi(1) - \phi(2)$, which we can use to solve for $p^{1,no,\omega}$.

The resource constraints for this problem are,

$$\begin{aligned} c_1 + c_T^{1,no} &\leq 2Z \\ c_1 + c_2 + c_T^{2,1} + c_T^{2,2} &\leq 2Z. \end{aligned}$$

The optimization is also subject to non-negativity constraints.

PROPOSITION 2 *: Let:*

$$\bar{K} \equiv \frac{\phi(1) - \phi(2)}{4} \left(\frac{u'(Z) - \beta}{u'(Z)} \right).$$

Then, the equilibrium in the robust economy depends on both K and Z as follows:

- *When there is insufficient aggregate liquidity, there are two cases:*
  - *For $0 \leq K < \bar{K}$, agents' decisions satisfy:*

$$\phi_\omega^\omega(1) u'(c_1) = \phi_\omega^\omega(2) u'(c_2) + \beta \frac{\phi(1) - \phi(2)}{2}. \tag{13}$$

    where, the worst-case probabilities are based on $\theta_\omega^\omega = K$:

$$\phi_\omega^\omega(1) = \frac{\phi(1)}{2} - K, \qquad \phi_\omega^\omega(2) = \frac{\phi(2)}{2} + K.$$

    In the solution,

$$c_2 < Z < c_1 < c^*$$

    with $c_1(K)$ decreasing and $c_2(K)$ increasing. We refer to this as the "partially robust" case.

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

$$c_1 = Z = c_2 < c^*.$$

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

- *When there is sufficient aggregate liquidity (Z), agents' decisions satisfy,*

$$c_1 = c_2 = c^* < Z.$$

The formal proof of the proposition is in the Appendix, and is complicated by the need to account for all possible consumption plans for every given θ_ω^ω scenario when solving the max-min problem. But, there is a simple intuition that explains the results.

We show in the Appendix that $c_T^{2,1}$, and $c_T^{2,2}$ are always equal to zero. Dropping these controls, the problem simplifies to:

$$\max_{c_1,c_2,c_T^{1,no}} \min_{\theta_\omega^\omega \in \Theta} \quad \phi_\omega^\omega(1)u(c_1) + \phi_\omega^\omega(2)u(c_2) + \frac{\phi(1)-\phi(2)}{2}\beta c_T^{1,no}.$$

For the case of insufficient aggregate liquidity, the resource constraints give:

$$c_2 = 2Z - c_1, \qquad c_T^{1,no} = 2Z - c_1.$$

Then the first order condition for the max problem for a given value of θ_ω^ω is,

$$\phi_\omega^\omega(1)u'(c_1) = \phi_\omega^\omega(2)u'(c_2) + \beta\frac{\phi(1)-\phi(2)}{2}.$$

In the benchmark case, the uncertain probabilities are $\phi_\omega^\omega(1) = \frac{\phi(1)}{2}$ and $\phi_\omega^\omega(2) = \frac{\phi(2)}{2}$, which yields the solution calling for more liquidity to whoever is affected by the first shock ($c_1 > c_2$). When $K > 0$, agents are uncertain over whether their shocks are early or late relative to other agents. Under the max-min decision rule, agents use the worst case probability in making decisions. Thus, they bias up the probability of being second relative to that of being first.[10] When K is small, agents' first order condition is,

$$\left(\frac{\phi(1)}{2} - K\right)u'(c_1) = \left(\frac{\phi(2)}{2} + K\right)u'(c_2) + \beta\frac{\phi(1)-\phi(2)}{2}.$$

As K becomes larger, c_2 increases toward c_1. For K sufficiently large, c_2 is set equal to c_1. This defines the threshold of $\bar{K}$. In this "fully robust" case, agents are insulated against their uncertainty over whether their shocks are likely to be first or second.

E. Flight to Quality

A flight to quality episode can be understood in our model as a comparative static across K. To motivate this comparative static within our model, let us introduce a date -1 as a

[10] In the solution, agents have distorted beliefs and in particular disagree: Agent ω thinks his $\theta_\omega = K$, but he also knows that $\int_{\omega\in\Omega}\theta_\omega d\omega = 0$. That is, a given agent thinks that all other agents on average have a $\theta_\omega = 0$, but the agent himself has the worst case θ. This raises the question of whether it is possible for the planner to design a mechanism that exploits this disagreement in a way that agents end up agreeing. We answer this question in the Appendix, and conclude that allowing for a fuller mechanism does not alter the solution.

contracting date for agents. Each agent has $Z < c^*$ units of the good at this date and has preferences as described earlier (only over consumption at date T and/or date $1, 2$). At date 0, a value of K is realized to be either $K = 0$ or $K > 0$. The $K > 0$ event is a low probability unusual event which may trigger a flight to quality. For example, the $K > 0$ event may be that the downgrade of a top name is imminent in the credit derivatives market. Today (i.e. date -1) market participants know that such an event may transpire and also are aware that in the event there will be considerable uncertainty over outcomes. At date -1, agents enter into an arrangement, where the terms of the contract are contingent on the state K, as dictated by Proposition 2. We can think of the flight to quality in comparing the contracts across the states.[11]

In this subsection we discuss three concrete examples of flight to quality events in the context of our model. Our first two examples identify the model in terms of the financial intermediation implementation discussed earlier, while the last example identifies the model in terms of the contingent claims implementation.

The first example is one of uncertainty-driven contagion and is drawn from the events of the fall of 1998. We interpret the agents of our model as the trading desks of an investment bank. Each trading desk concentrates in a different asset market. At date -1 the trading desks pool their capital with a top-level risk manager of the investment bank, retaining c_2 of capital to cover any needs that may arise in their particular market ("committed capital"). They also agree that the top-level risk manager will provide an extra $c_1 - c_2 > 0$ to cover shocks that hit whichever market needs capital first ("trading capital"). At date 0, Russia defaults. An agent in an unrelated market – i.e. a market in which shocks are now no more likely then before, so that $\phi^{\omega}_{\omega}(1) + \phi^{\omega}_{\omega}(2)$ is unchanged – suddenly becomes concerned that other trading desks will suffer shocks first and therefore the agent's trading desk will not have as much capital available in the event of a shock. The agent responds by lobbying the top-level risk manager to increase his committed capital up to a level of $c_2 = c_1$. As a result, every trading desk now has less capital in the (likelier) event of a single shock. Scholes (2000)

[11] An alternative way to motivate the comparative static is in terms of rewriting of contracts. Suppose that it is costless to write contracts at date -1, but that it costs a small amount ϵ to write contracts at date 0. Then it is clear that at date -1, agents will write contracts based on the $K = 0$ case of Proposition 2. If the $K > 0$ event transpires, agents will rewrite the contracts accordingly. We may think of a flight to quality in terms of this rewriting of contracts. Note that the only benefit in writing a contract at date -1 that is fully contingent on K is to save the ϵ of rewriting costs. In particular, if $\epsilon = 0$ it is not possible to improve the allocation based on signing contingent date -1 contracts. Agents are identical at both date -1 and at date 0, so that there are no extra allocational gains from writing the contracts early.

argues that during the 1998 crisis, the natural liquidity suppliers (hedge funds and trading desks) became liquidity demanders. In our model, uncertainty causes the trading desks to tie up more of the capital of the investment bank. The average market has less capital to absorb shocks, suggesting reduced liquidity in all markets.

In this example, the Russian default leads to less liquidity in other unrelated asset markets. Gabaix, Krishnamurthy, and Vigneron (2006) present evidence that the mortgage-backed securities market, a market unrelated to the sovereign bond market, suffered lower liquidity and wider spreads in the 1998 crisis. Note also that in this example there is no contagion effect if Z is large as the agents' trading desk will not be concerned about having the necessary capital to cover shocks when $Z > c^*$. Thus, any realized losses by investment banks during the Russian default strengthen the mechanism we highlight.

Our second example is a variant of the classical bank-run, but on the credit side of a commercial bank. The agents of the model are corporates. The corporates deposit Z in a commercial bank at date -1 and sign revolving credit lines that give them the right to c_1 if they receive a shock. The corporates are also aware that if banking conditions deteriorate (a second wave of shocks) the bank will raise lending standards/loan rates so that the corporates will effectively receive only $c_2 < c_1$. The flight to quality event is triggered by the commercial bank suffering losses and corporates becoming concerned that the two-wave event will transpire. They respond by preemptively drawing down credit lines, effectively leading all firms to receive less than c_1. Gatev and Strahan (2006) present evidence of this sort of credit-line run during periods when the spread between commercial paper and Treasury bills widens (as in the fall of 1998).

The last example is one of the interbank market for liquidity and the payment system. The agents of the model are all commercial banks who have Z Treasury bills at the start of the day. Each commercial bank knows that there is some possibility that it will suffer a large outflow from its reserve account, which it can offset by selling Treasury bills. To fix ideas, suppose that bank A is worried about this happening at 4pm. At date -1, the banks enter into an interbank lending arrangement so that a bank that suffers such a shock first, receives credit on advantageous terms (worth c_1 of T-bills). If a second set of shocks hits, banks receive credit at worse terms of c_2 (say, the discount window). At date 0, 9/11 occurs. Suppose that bank A is a bank outside New York City which is not directly affected by the events, but which is concerned about a possible reserve outflow at 4pm. However, now bank A becomes concerned that other commercial banks will need liquidity and that these needs may arise before 4pm. Then, bank A will renegotiate its interbank lending arrangements and

become unwilling to provide c_1 to any banks that receive shocks first. Rather, it will hoard its Treasury Bills of Z to cover its own possible shock at 4pm. In this example, uncertainty causes banks to hoard resources, which is often the systemic concern in a payments gridlock (e.g., Stewart, 2002, and McAndrews and Potter (2002)).

The different interpretations we have offered show that the model's agents and their actions can be mapped into the actors and actions during a flight to quality episode in a modern financial system. As is apparent, our environment is a variant of the one that Diamond and Dybvig (1983) study. In that model, the sequential service constraint creates a coordination failure and the possibility of a bad crisis equilibrium in which depositors run on the bank. In our model, the crisis is a rise in Knightian uncertainty rather than the realization of the bad equilibrium. The association of crises with a rise in uncertainty is the novel prediction of our model, and one which fits many of the flight to quality episodes we have discussed in this paper. Other variants of the Diamond and Dybvig model such as Rochet and Vives (2004) associate crises with low values of commercial bank assets. While our model shares this feature (i.e. Z needs to below c^*), it provides a sharper prediction through the uncertainty channel. Our model also offers interpretations of a crisis in terms of the rewriting of financial contracts triggered by the uncertainty increase, rather than the behavior of a bank's depositors. Of course in practice both the coordination failures that Diamond and Dyvbig highlight and the uncertainties we highlight are likely to be present, and possibly interact, during financial crises.

II. Collective Bias and the Value of Intervention

In this section, we study the benefits of central bank actions in the flight to quality episode of our model. We show that a central bank can intervene to improve aggregate outcomes. The analysis also clarifies the source of the benefit in our model.

A. Central Bank Information and Objective

The central bank knows the aggregate probabilities $\phi(1)$, $\phi(2)$, and knows that the ϕ_ω's are drawn from a common distribution for all ω. We have previously noted that this information is common knowledge, so we are not endowing the central bank with any more information than agents have. The central bank also understands that because of agents' ex-ante symmetry, all agents choose the same contingent consumption plan C^s. We denote $p_\omega^{s;CB}$ as the probabilities

that the central bank assigns to the different events that may affect agent ω. Like agents, the central bank does not know the true probabilities p^s_ω. Additionally, $p^{s,CB}_\omega$ may differ from $p^{s,\omega}_\omega$.

The central bank is concerned with the equally weighted ex-post utility that agents derive from their consumption plans:

$$\begin{aligned} V^{CB} &\equiv \int_{\omega\in\Omega} \sum p^{s,CB}_\omega U(C^s) d\omega \\ &= \sum p^s U(C^s). \end{aligned} \tag{14}$$

The step in going from the first to second line is an important one in the analysis. In the first line, the central bank objective reflects the probabilities for each agent ω. However, since the central bank is concerned with the aggregate outcome, we integrate over agents, exchanging the integral and summation, and arrive at a central bank objective that only reflects the aggregate probabilities p^s. Note that the individual probability uncertainties disappear when aggregating, and that the aggregate probabilities that appear are common knowledge (i.e., they can be written solely in terms of $\phi(1)$ and $\phi(2)$). Finally, as our earlier analysis has shown that only $c_1, c_2, c_T^{1,no} > 0$ need to be considered, we can reduce the objective to:

$$V^{CB} = \frac{\phi(1)}{2} u(c_1) + \frac{\phi(2)}{2} u(c_2) + \frac{\phi(1)-\phi(2)}{2} \beta c_T^{1,no}.$$

The next two subsections explain how a central bank that maximizes the objective function in (14) will intervene. For now, we note that one can view the objective in (14) as descriptive of how central banks behave: Central banks are interested in the collective outcome, and thus it is natural that the objective adopts the average consumption utility of agents in the economy. We return to a fuller discussion of the objective function in Section *D.* where we explain this criterion in terms of welfare and Pareto improving policies.

B. Collective Risk Management and Wasted Liquidity

Starting from the robust equilibrium of Proposition 2, consider a central bank that alters agents' decisions by increasing c_1 by an infinitesimal amount, and decreasing c_2 and $c_T^{1,no}$ by the same amount. The value of the reallocation based on the central bank objective in (14) is:

$$\frac{\phi(1)}{2} u'(c_1) - \frac{\phi(2)}{2} u'(c_2) - \frac{\phi(1)-\phi(2)}{2} \beta. \tag{15}$$

First, note that if there is sufficient aggregate liquidity, $c_1 = c_2 = c^* = u'^{-1}(\beta)$. For this case,

$$\frac{\phi(1)}{2}u'(c_1) - \frac{\phi(2)}{2}u'(c_2) - \frac{\phi(1)-\phi(2)}{2}\beta = 0$$

and equation (15) implies that there there is no gain to the central bank from a reallocation.

Turning next to the insufficient liquidity case, the first order condition for agents in the robustness equilibrium satisfies,

$$\phi^{\omega}_{\omega}(1)u'(c_1) - \phi^{\omega}_{\omega}(2)u'(c_2) - \beta\frac{\phi(1)-\phi(2)}{2} = 0$$

or

$$\left(\frac{\phi(1)}{2} - K\right)u'(c_1) - \left(\frac{\phi(2)}{2} + K\right)u'(c_2) - \beta\frac{\phi(1)-\phi(2)}{2} = 0.$$

Rearranging this equation we have that,

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

Substituting this relation into (15), it follows that the value of the reallocation to the central bank is $K(u'(c_1)+u'(c_2))$ which is positive for all $K > 0$. That is, the reallocation is valuable to the central bank because from its perspective agents are wasting aggregate liquidity by self-insuring excessively rather than cross-insuring risks.

Summarizing these results:

PROPOSITION 3 : *For any $K > 0$, if the economy has insufficient aggregate liquidity ($Z < c^*$), on average 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 agents' insurance toward the first shock.*

C. Is the Central Bank Less Knightian or More Informed than Agents?

In particular, are these the reasons the central bank can improve outcomes? The answer is no. To see this, note that any randomly chosen agent in this economy would reach the same conclusion as the central bank if charged with optimizing the expected ex-post utility of the collective of agents.

Suppose that agent $\tilde{\omega}$, who is Knightian and uncertain about the true values of θ_ω, is given such a mandate. Then this agent will solve,

$$\max_{c_1,c_2,c_T^{1,no}} \min_{\theta^{\tilde{\omega}}_{\tilde{\omega}} \in \Theta} \int \left(\phi^{\tilde{\omega}}_{\omega}(1)u(c_1) + \phi^{\tilde{\omega}}_{\omega}(2)u(c_2) + \frac{\phi(1)-\phi(2)}{2}\beta c_T^{1,no}\right) d\omega.$$

Since aggregate probabilities are common knowledge we have that:

$$\int \phi_{\omega}^{\tilde{\omega}}(1)\ d\omega = \frac{\phi(1)}{2}, \qquad \int \phi_{\omega}^{\tilde{\omega}}(2)\ d\omega = \frac{\phi(2)}{2}.$$

Substituting these expressions back into the objective, and dropping the min operator since now no expression in the optimization depends on $\theta_{\omega}^{\tilde{\omega}}$, yields:

$$\max_{c_1,c_2,c_T^{1,no}} \frac{\phi(1)}{2}u(c_1) + \frac{\phi(2)}{2}u(c_2) + \frac{\phi(1)-\phi(2)}{2}\beta c_T^{1,no},$$

which is the same objective as that of the central bank.

If it is not an informational advantage or the absence of Knightian traits in the central bank, what is behind the gain we document? The combination of two features drives our results: The central bank is concerned with aggregates *and* individual agents are "uncertain" (Knightian) not about aggregate shocks but about the impact of these shocks on their individual outcomes.

Since individual agents make decisions about their own allocation of liquidity rather than about the aggregate, they make choices that are *collectively biased* when looked at from the aggregate perspective. Let us develop the collective bias concept in more detail.

In the fully robust equilibrium of Proposition 2 agents insure equally against first and second shocks. To arrive at the equal insurance solution, robust agents evaluate their first order conditions (equation 13) at conservative probabilities:

$$\phi_{\omega}^{\omega}(1) - \phi_{\omega}^{\omega}(2) = \frac{\phi(1)-\phi(2)}{2}\left(\frac{u'(c^*)}{u'(Z)}\right) \tag{16}$$

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

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

The two in front of these expressions reflects the fact that only one-half of agents are affected by each of the shocks. Integrating equation (16) and using the definitions above, we find that agents' conservative probabilities are such that,

$$\bar{\phi}(1) - \bar{\phi}(2) = (\phi(1)-\phi(2))\left(\frac{u'(c^*)}{u'(Z)}\right) < \phi(1) - \phi(2).$$

The last inequality follows in the case of insufficient aggregate liquidity ($Z < c^*$).

Implicitly, these conservative probabilities overweight an agent's chances of being affected second in the two-wave event. Since each agent is concerned about the scenario in which he receives a shock last and there is little liquidity left, robustness considerations lead each agent to bias upwards the probability of receiving a shock later than the average agent. However, every agent cannot be later than the "average." Across all agents, the conservative probabilities violate the *known* probabilities of the first and second wave events.

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 become apparent the scenario that the collective of conservative agents are guarding against is impossible.

D. Welfare

We next discuss our specification of the central bank's objective in (14). Agents in our model choose the worst case among a class of priors when making decisions. That is, they are not rational from the perspective of Bayesian decision theory and therefore do not satisfy the Savage axioms for decision making. As Sims (2001) notes, this departure from rational expectations can lead to a situation where a maximin agent accepts a series of bets that have him lose money with probability one. The appropriate notion of welfare in models where agents are not rational is subject to some debate in the literature.[12] It is beyond the scope of this paper to settle this debate. Our aim in this section is to clarify the issue in the present context and offer some arguments in favor of objective (14).

At one extreme, consider a "libertarian" welfare criterion whereby agents' choices are by definition what maximizes their utility. That is, define,

$$\underline{V}^{CB} = \int_{\omega\in\Omega} \min_{\theta_\omega^\omega \in \Theta} \sum p_\omega^{s,\omega} U(C^s) d\omega.$$

This is an objective function based on each agent-ω's ex-ante utility which is evaluated using that agent's worst-case probabilities. The difference relative to the objective in (14) is that all utility here is "anticipatory." That is, the agent enjoys happiness *at date 0* from making a decision that avoids a worst-case outcome. Note that such a specification differs from standard expected utility whereby the agent only receives happiness at dates $1, 2$ and T when the agent actually consumes.

[12]The debate centers on whether or not the planner should use the same model to describe choices and welfare (see, e.g., Gul and Pesendorfer (2005) and Bernheim and Rangel (2005) for two sides of the argument). See also Sims (2001) in the context of a central bank's objective.

Under the objective $\underline{V}^{CB}$ the agent's choices are efficient and there is no role for the central bank. We can see this immediately because the planning problem in deriving Proposition 2 was based on the latter objective function.

The objective function we use in (14) is based on ex-post consumption utility, and assumes that agents do not receive any anticipatory utility. More generally, consider an objective function $\underline{\lambda}\,\underline{V}^{CB} + (1-\underline{\lambda})V^{CB}$ with $\underline{\lambda} \in [0,1]$. Then it is clear that as long as $\underline{\lambda} < 1$ – i.e. the welfare function places some weight on ex-post consumption utility – there is a role for the central bank. In this sense, the no-intervention case is an extreme one.

There is a further reason to restrict attention to the $\underline{\lambda} = 0$ case, as in (14). Consider the following thought experiment: Suppose that we repeat infinitely many times the liquidity episode we have described. At the beginning of each episode, agent ωdraws a $\theta_\omega \in \Theta$. These draws are i.i.d. across episodes, and the agent knows that on average his θ_ω will be zero. In each episode, since agent ω does not know the θ_ω for that episode, the agent's worst-case decision rule has him using $\theta_\omega = K$. V^{CB} is the average consumption utility of agent ω across all of these episodes.[13]

The preceding two arguments are ones in favor of using an ex-post consumption utility welfare criterion, where each agent is weighted equally. The last point we discuss is when equal weighting is appropriate. Thus far, since agents are ex-ante identical, a policy that improves the average agent's ex-post consumption utility also improves each agent's ex-post expected consumption utility. Suppose however that a fraction of the agents in the economy are Bayesian (i.e. rational) and moreover know that their true θ_ω is equal to K. For these agents, the worst-case probabilities are truly their own probabilities. Thus, define the welfare of the rational agents as

$$V^R = \int_{\omega \in \Omega^R} \sum p_\omega^{s,\omega} U(C^s) d\omega,$$

where Ω^R is the subset of Ω corresponding to the rational agents, and the probabilities $p_\omega^{s,\omega}$ are based on $\theta_\omega = K$.

The rest of the agents, $\omega \in \Omega \setminus \Omega^R$, are Knightian with θ_ω's such that the average θ_ω across both classes of agents is zero. We define V^K in a similar way to the objective in (14) as the average ex-post consumption utility of the Knightian agents.

We now have a situation where there is ex-ante heterogeneity among agents so that equal weighting is no longer appropriate. Suppose that the central bank cannot discriminate among

[13]Of course, in living through repeated liquidity events, an agent learns over time about the true distribution of θ_ω. However, it is still the case that along this learning path, K remains strictly positive (while shrinking) and hence the qualitative features of our argument go through for a small enough discount rate.

the two classes of agents. Is the intervention still Pareto improving?

The result in Proposition 3 still applies to the Knightian agents. The central bank will compute a first order gain in V^K from the reallocation intervention. Importantly, note that the envelope theorem implies that changing the rational agents' decisions results in only a second order utility loss to the rational agents. That is, a small intervention means that the loss in V^R is small compared to the gain in V^K. Thus, although the central bank's policy is not Pareto improving, it involves asymmetric gains to the Knightian agents. Camerer, et al., (2003) propose this type of asymmetric paternalism criterion in evaluating policies when some agents are behavioral.

E. Risk Aversion versus Uncertainty Aversion

We note previously that from a positive standpoint our model of uncertainty aversion predicts a flight to quality when there is a "new" shock, whereas a model with extreme risk aversion predicts conservative behavior in response to any negative shock, new or not. We close this section by noting that the normative implications of uncertainty aversion also differ from that of extreme risk aversion. Without collective bias, and regardless of the agent's degree of risk aversion, our central bank sees no reason to reallocate liquidity toward the first wave of shocks beyond the private sector's choices. We can see this because setting $K = 0$ in our model represents a model without uncertainty aversion. As we have imposed only weak requirements on $u(\cdot)$, the utility function can be chosen to represent extreme forms of risk aversion. However the results of Proposition 3 establish that there is a gain for the central bank only if $K > 0$ and $Z < c^*$.

We conclude that there is a role for the central bank only in situations of Knightian uncertainty and insufficient aggregate liquidity. Of course not all recessionary episodes exhibit these ingredients. But there are many scenarios where they are present, such as October 1987 and the Fall of 1998.

III. An Application: Lender of Last Resort

The abstract reallocation experiment considered in Proposition 3 makes clear that during flight to quality episodes the central bank will find it desirable to induce agents to insure less against second shocks and more against first shocks. In this section we discuss an application of this result and consider a lender of last resort (LLR) policy in light of the gain identified

in Proposition 3.

As in Woodford (1990) and Holmstrom and Tirole (1998), we assume the LLR has access to collateral that private agents 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 novel result is that, because of the reallocation benefit of Proposition 3, the value of the LLR exceeds the direct value of the intervention. Thus our model sheds light on a new benefit of the LLR.

The model also stipulates when the benefit is highest. As we have remarked previously, the reallocation benefit only arises in situations where $K > 0$ and $Z < c^*$. This carries over directly to our analysis of the LLR: the benefits are highest when $K > 0$ and $Z < c^*$. We also show that the LLR must be a *last-resort* policy. In fact, if liquidity injections take place too often, the reallocation effect works against the policy and reduces its value.

A. LLR Policy

Formally, the central bank credibly expands the resources of agents in the *two-shock event* by an amount Z^G. That is, agents who are affected second in the two-wave event ($s = (2, 2)$), will have their consumption increased from c_2 to $c_2 + 2Z^G$ (twice Z^G because one-half measure of agents are affected by the second shock). The resource constraints for agents (for the reduced problem) are:

$$c_1 + c_T^{1,no} \leq 2Z \tag{17}$$

$$c_1 + c_2 \leq 2Z + 2Z^G. \tag{18}$$

In practice, the central bank's promise may be supported by a credible commitment to costly ex-post inflation or taxation and carried out by guaranteeing, against default, the liabilities of financial intermediaries who have sold financial claims against extreme events. 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,

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

The anticipation of the central bank's second-shock insurance leads agents to reoptimize their insurance decisions. Agents reduce their private insurance against the publicly insured

second-shock and increase their first-shock insurance. The total benefit of the intervention includes both the direct benefit as well as any benefit from portfolio reoptimization:

$$V_{Z^G}^{CB,total} = \int_\Omega \left[\phi_\omega(1) u'(c_{1,\omega}) \frac{dc_{1,\omega}}{dZ^G} + \phi_\omega(2) u'(c_{2,\omega}) \frac{dc_{2,\omega}}{dZ^G} + \frac{\phi(1)-\phi(2)}{2} \beta \frac{dc_{T,\omega}^{1,no}}{dZ^G} \right] d\omega.$$

The first order condition for agent decisions, from (13), in the robust equilibrium gives,

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

We simplify the expression for $V_{Z^G}^{CB,total}$ by integrating through $\phi_\omega(1)$ and $\phi_\omega(2)$ and then substituting for $u'(c_1)$ from the first order condition. These operations yield,

$$V_{Z^G}^{CB,total} = \frac{\phi(2)}{2} u'(c_2) \left(\frac{dc_1}{dZ^G} + \frac{dc_2}{dZ^G} \right) + \beta \frac{\phi(1)-\phi(2)}{2} \left(\frac{dc_1}{dZ^G} + \frac{dc_T^{1,no}}{dZ^G} \right) + K(u'(c_1)+u'(c_2)) \frac{dc_1}{dZ^G}$$

Last, we differentiate the resource constraints (17) and (18) with respect to Z^G to find,

$$\frac{dc_1}{dZ^G} + \frac{dc_2}{dZ^G} = 2, \quad \frac{dc_1}{dZ^G} + \frac{dc_T^{1,no}}{dZ^G} = 0.$$

Then,

$$\begin{aligned} V_{Z^G}^{CB,total} &= \phi(2) u'(c_2) + K(u'(c_1) + u'(c_2)) \frac{dc_1}{dZ^G} \\ &= V_{Z^G}^{CB,direct} + K(u'(c_1) + u'(c_2)) \frac{dc_1}{dZ^G}. \end{aligned}$$

The additional benefit we identify is due to portfolio reoptimization: Agents cut back on the publicly insured second shock and increase first shock insurance, thereby moving their decisions closer to what the central bank would choose for them. In this sense, the LLR policy can help to implement the policy suggested in Proposition 3.

We also note that without Knightian uncertainty ($K = 0$), there is no gain (beyond the direct benefit) from the policy. Moreover, it is straightforward to see that if $Z > c^*$ then agents will not use the additional insurance to cover their liquidity shocks, but will reoptimize in a way as to use the insurance at date T. In this case there is no gain to offering the public insurance (since $\frac{dc_1}{dZ^G} = 0$). We summarize these results as follows:

PROPOSITION 4 *: For $K > 0$ and $Z < c^*$, the total value of the lender of last resort policy exceeds its direct value:*

$$V_{Z^G}^{CB,total} > V_{Z^G}^{CB,direct}.$$

It is important to note that under the LLR policy the central bank injects resources only rarely. As we associate the second-shock event with an extreme and unlikely event, in expectation the central bank does not promise many resources. This aspect of policy is similar to Diamond and Dybvig's (1983) analysis of a LLR. However there are a few important differences in the mechanism through which the policies work. As there is no coordination failure in our model, the policy does not work by ruling out a "bad" equilibrium. Rather, the policy works by reducing the agents' "anxiety" that they will receive a shock last when the economy has depleted its liquidity resources. It is this anxiety that leads agents to use a high $\phi_{\omega}^{\omega}(2)$ in their decision rules. From this standpoint, it is also clear that an important ingredient in the policy is that agents have to believe that the central bank will have the necessary resources in the two-event shock in order to reduce their anxiety. Credibility and commitment are central to the working of our LLR policy.[14]

B. Moral Hazard and Early Interventions

The policy we have suggested cuts against the usual moral hazard critique of central bank interventions. 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.[15] Public and private provision of insurance are complements in our model.

This logic suggests that interventions against first shocks may be subject to the moral hazard critique as agents' portfolio reoptimization would lead them toward more insurance against the second shock. To consider the "early intervention" case, suppose that the central bank credibly offers to increase the consumption of agents who are affected in the first shock from c_1 to $c_1 + 2Z^G$. The resource constraints for agents (for the reduced problem) are:

$$\begin{aligned} c_1 + c_T^{1,no} &\leq 2Z + 2Z^G \\ c_1 + c_2 &\leq 2Z + 2Z^G. \end{aligned}$$

[14]In this sense, the policy relates to the government bond policy of Woodford (1990) and Holmstrom and Tirole (1998) who argue that government promises are unique because they have greater collateral backing than private sector promises.

[15]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.

The *direct* benefit of intervention in the first shock is:

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

We compute the total benefit as previously except that we substitute agents' first order condition using,

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

Also, using the fact that,

$$\frac{dc_1}{dZ^G} + \frac{dc_2}{dZ^G} = 2, \quad \frac{dc_1}{dZ^G} + \frac{dc_T^{1,no}}{dZ^G} = 2.$$

we find that,

$$V_{Z^G}^{CB,total} = V_{Z^G}^{CB,direct,first} - K(u'(c_1)+u'(c_2))\frac{dc_1}{dZ^G} < V_{Z^G}^{CB,direct,first}$$

The expected cost of the early intervention policy is much larger than the second shock intervention, since the central bank rather than the private sector bears the cost of insurance against the (likely) single-shock event. Agents 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. We conclude that the lender of last resort facility, to be effective and improve private financial markets, has to be a *last* and not an intermediate resort.

C. 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 extend the model to consider multiple shocks. We assume the economy may experience $n = 1...N$ waves of shocks, each affecting $\frac{1}{N}$ of the agents. The probability of the economy experiencing n waves is denoted $\phi(n)$, with $\phi(n) < \phi(n-1)$. Also, each ω's probability of being affected in the nth wave satisfies $\int_{\omega\in\Omega}\phi_{\omega}(n)d\omega = \frac{\phi(n)}{N}$.

The LLR policy takes the following form: The central bank injects $\frac{1}{N-j+1}$ units of liquidity for all shocks after (and including) the jth wave ($j \leq N$). We also simplify our analysis by focusing on the fully robust case where c_n is the same for all n and by setting $\beta = 0$, thereby

assuring that $Z < c^*$ and allowing us to disregard effects on $c_T^{n,no}$. c_n rises to $c_n + \frac{N}{N-j+1}$ in the intervention (i.e. $\frac{1}{N-j+1}$ injected to a measure $\frac{1}{N}$ of agents).

The direct value of the intervention as a function of j is,

$$\begin{aligned} V_{Z^G}^{CB,direct} &= \frac{N}{N-j+1}\int_\Omega \sum_{n=j}^{N} \phi_\omega(n) u'(c_{n,\omega}) d\omega \\ &= u'(c_1)\frac{1}{N-j+1}\sum_{n=j}^{N}\phi(n). \end{aligned}$$

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:

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

From the resource constraint we have that

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

In the fully robust case, $c_{n,\omega}$ and $\frac{dc_{n,\omega}}{dZ^G}$ are the same for all n. Then,

$$V_{Z^G}^{CB,total} = u'(c_1)\frac{dc_1}{dZ^G}\frac{1}{N}\sum_{n=1}^{N}\phi(n) = u'(c_1)\frac{1}{N}\sum_{n=1}^{N}\phi(n) \tag{19}$$

Note that this expression is independent of the intervention rule j. In contrast, it is apparent that $V_{Z^G}^{CB,direct}$ is decreasing with respect to j since the $\phi(n)$'s are monotonically decreasing. Thus, the ratio:

$$\frac{V_{Z^G}^{CB,total}}{V_{Z^G}^{CB,direct}} = \frac{\frac{1}{N}\sum_{n=1}^{N}\phi(n)}{\frac{1}{N-j+1}\sum_{n=j}^{N}\phi(n)}$$

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

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.

The multiple shock model also clarifies another benefit of late intervention. As j 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 j rises, while the expected benefit remains constant.

In other words, as j 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 j increases.

Thus, while a well designed LLR policy may indeed have a direct effect only in highly unlikely events, the policy is not irrelevant for likely outcomes. Its main benefits come from unlocking private markets to insure more likely and less extreme events.

IV. Final remarks

We present a model of financial crises and the role of a lender of last resort that centers on Knightian uncertainty and liquidity shortages. While Knightian uncertainty is discussed in policy circles (e.g., see Greenspan's quote in the Introduction), it is not standard in academic analyses of crises, which instead emphasizes liquidation externalities (e.g., Diamond and Dybvig (1983)). Thus, rather than ending with a summary of findings, it is useful to take stock by considering points of similarity and departure between our Knightian uncertainty model and the standard liquidation externality model.

As we argue throughout the paper, an "uncertainty shock," distinct from —although possibly correlated with— a liquidation shock is an important element in many financial crises. The uncertainty shock, like the liquidation shock, reproduces the behavior of agents during a flight to quality episode. Agents move toward holding uncontingent and safe assets, financial intermediaries hoard liquidity, and investment banks and trading desks turn conservative in their allocation of risk capital so that capital becomes immobile across markets. On these points our model emphasizes a different mechanism but shares many of the predictions of liquidation models.

Despite this similarity, we think it is valuable to spell out the model of the uncertainty shock. At a minimum, doing so clarifies and expands the set of preconditions for crises. There may be crisis events where liquidation externalities are absent but in which there is a substantial amount of Knightian uncertainty. At the other extreme, the most severe crises in the US and abroad seem to have elements of both an uncertainty and a liquidation shock. In such cases our model helps to understand why these crises may have been so severe.

Furthermore, there are some events – e.g. the recent liquidation of the Amaranth hedge

fund's sizeable portfolio – for which the preconditions seem to fit the liquidation crisis model, but yet did not trigger crises. Our model suggests that the absence of significant uncertainty around these events may be one reason why (in the particular case of Amaranth, financial specialists had already learned from the LTCM experience). Thus our model also helps to shed light on the "dog that did not bark." More broadly, it suggests that an important precondition for crises is the presence of "new" shocks, perhaps surrounding new and untested financial innovations. This prediction is also useful to guide policymakers on where crises may arise.

From a policymaker's perspective, our model, like the liquidation externality model, shows that there are benefits to a lender of last resort facility during a crisis. However, in our model the central issue to determine the value of a LLR facility is not only the potential for coordination failure but also the extent of uncertainty in the marketplace. For example, we prescribe that a default by a hedge fund – even one that is large – should not elicit central bank reaction unless the default triggers considerable uncertainty in other market participants and hedge funds are financially weak.

Yet another subtle difference between the two models is in the incentive effects of policy intervention. In the liquidation model, moral hazard is often an important issue that tempers policy intervention. This is because private and public insurance are substitutes in this model. In contrast, our analysis reveals that in the uncertainty model there are dimensions in which private and public insurance are complements rather than substitutes, suggesting that the moral hazard issue is less important for uncertainty driven crises. Having said this, a different kind of moral hazard problem may arise in the uncertainty model if agents can invest resources to affect the degree of uncertainty they face.

In the liquidation model, ex-ante policy recommendations typically center on prudential risk management. For example, in many analyses there are externalities present that drive a wedge between private and social incentives to insure against a financial crisis episode. Then, ex-ante regulations to reduce leverage, increase liquidity ratios, or tighten capital requirements are beneficial. In our model, an important dimension of the crisis is that there is uncertainty about outcomes. Agents cannot refer to history to understand how a crisis will unfold because the historical record may not span the event space. In such a case it is unclear whether any entity, either private or public, can arrive at the appropriate ex-ante risk management strategy, calling into question the feasibility of these policy recommendations. Instead, in our uncertainty model, the most beneficial ex-ante actions are ones which help to reduce the extent of uncertainty should a crisis occur. In some cases, this may simply involve

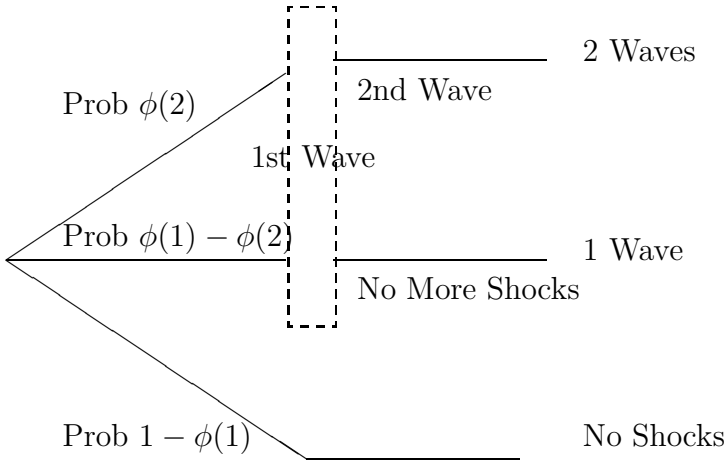
making common knowledge information that is known to subsets of market participants – for example, making common knowledge the portfolio positions of the major players in a market. In other cases, this may involve the central bank facilitating discussions among the private sector on how each party will react in a crisis scenario.

These points are pertinent, and can be illustrated, in the credit derivatives market. There is currently considerable uncertainty over how the downgrade of a top name will affect the credit derivatives market (see Geithner (2006)). Market participants are aware that such a downgrade will occur at some point, but given the lack of history, they are uncertain of how events will unfold. Will the market absorb such a shock without losing liquidity? Could such a shock result in a credit crunch that causes the corporate sector to suffer, and trigger a domino effect of downgrades? Are back-office and settlement procedures sufficient to handle such an event? Our model suggests that the central bank stands ready to act as the LLR in the event that a downgrade triggers these uncertainties. There have also been recent moves to increase transparency and risk assessment in this market, as well as streamline back-office settlement procedures. Our model suggests that such ex-ante actions may reduce uncertainty and be beneficial.

Finally, as we note, Knightian uncertainty may often be associated with financial innovations. This suggests that crises surrounding financial innovations may by fertile ground to look empirically for the effects we have modeled, and disentangle them from other more well understood effects. It also suggests a new perspective on the costs and benefits of innovation. For example, our model suggests that in a dynamic context with endogenous financial innovation, it is the pace of this innovation that inherently creates uncertainty and hence the potential for a flight to quality episode. Financial innovation facilitates risk sharing and leverage, but also introduces sources of uncertainty about the resilience of the new system to large shocks. This uncertainty is only resolved once the system has been tested by a crisis or near-crisis, after which the economy may enjoy the full benefits of the innovation. We are currently exploring these issues.

Appendix

Event Tree and Probabilities: The event tree is pictured below. The probability of two waves affecting the economy is $\phi(2)$; the probability of one wave affecting the economy is $\phi(1) - \phi(2)$; and, the probability of no waves affecting the economy is $1 - \phi(1)$. We used the dashed box around "1st wave" to signify that agents are unsure whether they are in the upper branch (one more wave will occur) or the middle branch (no further shocks).



Consider an agent ω who may be affected in these waves. Suppose that his probability of being affected by a shock when the event is the middle branch ("1 wave") is one-half. Suppose that his probability of being affected by a first shock when the event is the upper branch ("2 waves") is ψ_ω, while his probability of being affected by a second shock is $1 - \psi_\omega$. Moreover, suppose that the agent is uncertain about ψ_ω, which we interpret as the agent is uncertain about his likelihood of being first or second, in the case of a two wave event.

The agent's probability of being affected by a first shock is,

$$\phi_\omega(1) = \phi(2)\psi_\omega + (\phi(1) - \phi(2))\frac{1}{2}.$$

The agent's probability of being affected by a second shock is,

$$\phi_\omega(2) = \phi(2)(1 - \psi_\omega).$$

Note that,

$$\phi_\omega(1) + \phi_\omega(2) = \phi(2) + (\phi(1) - \phi(2))\frac{1}{2} = \frac{\phi(1) + \phi(2)}{2},$$

and,

$$\begin{aligned} \phi_\omega(1) - \frac{\phi(1)}{2} &= \phi(2)\psi_\omega - \frac{\phi(2)}{2} \\ \phi_\omega(2) - \frac{\phi(2)}{2} &= -\phi(2)\psi_\omega + \frac{\phi(2)}{2} \end{aligned}$$

These expressions show that the event tree is consistent with agents being certain about their probability of receiving a shock, but being uncertain about their relative probabilities of being first or second. In the text, we describe the uncertainty in terms of $\phi_\omega(2) - \frac{\phi(2)}{2}$ rather than in terms of ψ_ω.

Proof of Proposition 2: We focus on the case of insufficient aggregate liquidity ($Z < c^*$). The other case follows the same logic as the $K = 0$ case. We are looking for a solution to the problem in equation (12). We can describe this problem in the game-theoretic language often used in max-min problems. The agent chooses $\mathcal{C}$ to maximize $V(\mathcal{C}; \theta_\omega^\omega)$ anticipating that "nature" will choose θ_ω^ω to minimize $V(\mathcal{C}; \theta_\omega^\omega)$ given the agent's choice of $\mathcal{C}$.

The solution $(\bar{\theta}_\omega^\omega, \bar{\mathcal{C}})$ has to satisfy a pair of optimization problems. First, $\bar{\theta}_\omega^\omega \in argmin_{\theta_\omega^\omega} V(\bar{\mathcal{C}}; \theta_\omega^\omega)$. That is, nature chooses θ_ω^ω optimally given the agent's choice of $\bar{\mathcal{C}}$. Second, $\bar{\mathcal{C}} \in argmax_{\mathcal{C}} V(\mathcal{C}; \bar{\theta}_\omega^\omega)$. That is, the agent chooses $\mathcal{C}$ optimally given nature's choice of $\bar{\theta}_\omega^\omega$.[16]

We compute:

$$\frac{\partial V}{\partial \theta_\omega^\omega} = u(c_2) - u(c_1) + \beta(c_T^{2,2} - c_T^{2,1})$$

Let us first ask whether there exists a solution in which $\frac{\partial V}{\partial \theta_\omega^\omega} < 0$. If so, then clearly $\theta_\omega^\omega = +K$. Taking this value of θ_ω^ω let us consider the agent's problem in equation (12). First note that $c_T^{2,1} = 0$. To see this, suppose that $c_T^{2,1} > 0$. Then we can reduce $c_T^{2,1}$ by δ and increase $c_T^{2,2}$ by δ and produce a utility gain of $\delta(\phi_\omega^\omega(2) - \phi(2) + \phi_\omega^\omega(2)) > 0$ when $\theta_\omega^\omega > 0$.

With this knowledge, we rewrite the condition that $\frac{\partial V}{\partial \theta_\omega^\omega} < 0$ as,

$$u(c_2) + \beta c_T^{2,2} < u(c_1) \quad \Rightarrow \quad c_1 > c_2$$

If $c_1 > c_2$ and $Z < c^*$ it follows from the resource constraint that $c_2 < c^*$. But if $c_2 < c^*$ then from the agent's problem, we must have that $c_T^{2,2} = 0$ (i.e. do not save any resources for date T if these resources could be used earlier).

Thus, we only need to consider the agent's problem in (12) for values of $c_1, c_2, c_T^{1,no} > 0$. The first order condition for the agent at $\theta_\omega^\omega = +K$ is,

$$\left(\frac{\phi(1)}{2} - K\right) u'(c_1) = \left(\frac{\phi(2)}{2} + K\right) u'(c_2) + \beta \frac{\phi(1) - \phi(2)}{2}.$$

We note that for $K = 0$, the unique solution to the agent's problem is $c_1 > c_2$. Thus, for small values of K, a solution exists in which the agent chooses $c_1 > c_2$ and nature chooses $\theta_\omega^\omega = +K$. This is the partially robust solution given in the Proposition.

As K becomes larger, c_1/c_2 falls and at some point $c_1 = c_2 = Z$. This occurs when $\bar{K}$ solves,

$$\left(\frac{\phi(1)}{2} - \bar{K}\right) u'(Z) = \left(\frac{\phi(2)}{2} + \bar{K}\right) u'(Z) + \beta \frac{\phi(1) - \phi(2)}{2}.$$

[16]The fact that the agent chooses $\mathcal{C}$ before nature chooses θ_ω^ω does not affect our problem. To see this, note that choosing first only gives the agent an advantage if the agent can induce nature to choose a θ_ω^ω different than $\bar{\theta}_\omega^\omega$. Suppose the agent chooses $\mathcal{C} \neq \bar{\mathcal{C}}$ to increase $V(\cdot)$. Clearly this choice reduces V below $V(\bar{\theta}_\omega^\omega, \bar{\mathcal{C}})$. Thus, nature can always choose to set $\theta_\omega^\omega = \bar{\theta}_\omega^\omega$ and make the agent strictly worse off than at the choice $\mathcal{C} = \bar{\mathcal{C}}$.

which gives the expression for the value of $\bar{K}$ defined in the Proposition.

Note that if $K > \bar{K}$, the solution $\theta_\omega^\omega = +\bar{K}$ and $c_1 = c_2$ still solves both the agent's and nature's optimization problems. The agent's choice is uniquely optimal at $\theta_\omega^\omega = +\bar{K}$, while nature is indifferent over values of $\theta_\omega^\omega \in [-K, +K]$. This is the fully robust solution given in the Proposition.

We have thus far shown that considering the case where $\frac{\partial V}{\partial \theta_\omega^\omega} \leq 0$, the solution given in the Proposition is the only solution to the problem in (12). We conclude by showing that there are no other solutions to the problem. To do this, we only need to consider whether there exists a solution in which $\frac{\partial V}{\partial \theta_\omega^\omega} > 0$.

Suppose there does exist such a solution. If $\frac{\partial V}{\partial \theta_\omega^\omega} > 0$, then $\theta_\omega^\omega = -K$. We can go back through arguments similar to those previously offered to show that $c_T^{2,2}$ and $c_T^{2,1}$ must both be zero in this case. Then the condition that $\frac{\partial V}{\partial \theta_\omega^\omega} > 0$ is equivalent to,

$$c_1 < c_2$$

The first order condition for the agent is,

$$\left(\frac{\phi(1)}{2} + K\right) u'(c_1) = \left(\frac{\phi(2)}{2} - K\right) u'(c_2) + \beta \frac{\phi(1) - \phi(2)}{2}.$$

The solution to this problem is that $c_1 > c_2$, which is a contradiction. Thus there does not exist a solution in which $\frac{\partial V}{\partial \theta_\omega^\omega} > 0$.

The Expanded Planning Problem: The planning problem consided in Proposition 2 solved for allocations contingent on the number of shocks affecting the economy and the time at which a given agent is affected. In particular we did not allow allocations to be directly contingent on agent's beliefs about their shock probabilities. This Appendix considers an expanded planning problem which considers such allocations and shows that the solution is the same as that of Proposition 2. That is, allowing for more contingencies does not alter our results.

Why consider allocations contingent on beliefs? In our equilibrium, agents have distorted beliefs and in particular disagree: Agent ω thinks his $\theta_\omega = K$, but he also knows that $\int_{\omega \in \Omega} \theta_\omega d\omega = 0$. That is, a given agent thinks that all other agents on average have a $\theta_\omega = 0$, but the agent himself has the worst case θ. This raises the question of whether it is possible to construct a mechanism that exploits this disagreement in a way that agents end up agreeing.

Let us consider this issue formally as a mechanism design problem. Denote the "type" of an agent as θ_ω^ω and denote the report of the agent's type as $\hat{\theta}_\omega^\omega$. Also denote the set of reports

from all agents as $\hat{\Gamma} \in \bigcup_{\omega\in\Omega}[-K,+K]$. A consumption allocation to agent ω as a function of the reports is $\mathcal{C}(\omega,\hat{\Gamma})$.

The utility of ω from this consumption allocation given his type θ^{ω}_{ω} is:

$$U(\mathcal{C}(\omega,\hat{\Gamma});\theta^{\omega}_{\omega}) = \phi^{\omega}_{\omega}(1)u(c_1(\omega,\hat{\Gamma})) + \phi^{\omega}_{\omega}(2)u(c_2(\omega,\hat{\Gamma})) + \frac{\phi(1)-\phi(2)}{2}\beta c_T^{1,no}(\omega,\hat{\Gamma}).$$

The probabilities $\phi^{\omega}_{\omega}(1)$ and $\phi^{\omega}_{\omega}(2)$ are a function of θ^{ω}_{ω}. Also, we have written this problem for the simplified case where we only need to consider $c_1, c_2, c_T^{1,no} > 0$ (i.e., the reduced problem in Proposition 2).

The planner's problem is to choose an allocation function $\mathcal{C}$:

$$\max_{\mathcal{C}(\omega,\hat{\Gamma})} \int_{\omega\in\Omega} U(\mathcal{C}(\omega,\hat{\Gamma});\theta^{\omega}_{\omega}))\ d\omega, \tag{$A1$}$$

s.t. resource constraints:

$$\int_{\omega\in\Omega} \Big(c_1(\omega,\hat{\Gamma})) + c_2(\omega,\hat{\Gamma})\Big)\ d\omega \leq 2Z,$$

and,

$$\int_{\omega\in\Omega} \Big(c_1(\omega,\hat{\Gamma})) + c_T^{1,no}(\omega,\hat{\Gamma})\Big)\ d\omega \leq 2Z.$$

Finally the "type" of the agent is also a function of the allocation:

$$\theta^{\omega}_{\omega} \in argmin_{\theta^{\omega}_{\omega}\in\Theta} U(C(\omega,\hat{\Gamma});\theta^{\omega}_{\omega}). \tag{$A2$}$$

The problem as we have written it is quite general. It describes each agent's consumption allocation as a function of the entire set of reports of agent's types; where types are interpreted to be the agent's beliefs over his shock probabilities.

We argue that the current planning problem as described in Section II.D subsumes the one in ($A1$). The strategy for the proof is to consider a relaxed version of the problem in ($A1$) and show that the solution to that problem is the same as the solution to the planning problem described in the text. Given this result, we conclude that allowing for a more general allocation does not affect the results.

Suppose that the planner knows the agent types θ^{ω}_{ω} without having to rely on reports. Giving the planner more information allows the planner to implement strictly better allocations. As a result, this is a relaxed version of the problem we have written down in ($A1$). Thus, suppose that $\mathcal{C}(\omega,\Gamma)$, where $\Gamma \in \bigcup_{\omega\in\Omega}[-K,+K]$ is the set of agent types and $\mathcal{C}$ is the allocation to agent ω as a function of the agent types directly and not the reports of agents.

Note that θ^{ω}_{ω}, as given in $(A2)$, is a function of the consumption allocations of the planner to agent ω – i.e., given consumption allocations to agent ω, the planner can directly compute θ^{ω}_{ω}. This implies that we can drop the dependence in $\mathcal{C}$ on Γ, and the planner's choice is over $\mathcal{C}(\omega)$: the planner chooses numbers $c_1(\omega), c_2(\omega)$, and $c_T^{1,no}(\omega)$ for each ω. This problem is the one we have solved in the text of the paper. The result is given in Proposition 2. Thus we conclude that allowing for allocations that depend on the "types" of the agents does not alter the equilibrium.

REFERENCES

Allen, Franklin and Douglas Gale, 1994, Limited market participation and volatility of asset prices, *American Economic Review* 84, 933-955.

Allen, Franklin and Douglas Gale, 2004, Financial intermediaries and markets, *Econometrica* 72(4), 1023-1061.

Barro, Robert, 2006, Rare disasters and asset markets in the Twentieth century, *The Quarterly Journal of Economics* 121(3), 823-66.

Bernheim, Bernheim and Antonio Rangel, 2005, Behavioral public economics: Welfare and policy analysis with non-standard decision makers, NBER Working Paper # 11518.

Bhattacharya, Sudipto, and Douglas Gale, 1987, Preference shocks, liquidity, and central bank policy, in Barnett, William A. and Singleton, Kenneth J. eds.: *New approaches to monetary economics.* International Symposia in Economic Theory and Econometrics series: (Cambridge University Press).

Caballero, Ricardo and Arvind Krishnamurthy, 2003, Excessive dollar debt: Underinsurance and domestic financial underdevelopment, *Journal of Finance* 58(2), 867-893.

Calomiris, Charles, 1994, Is the discount window necessary? A Penn Central perspective, *Review*, Federal Reserve Bank of St. Louis.

Camerer, Colin, Samuel Issacharoff, George Loewenstein, Ted O'Donoghue, and Matthew Rabin, 2003, Regulation for conservatives: Behavioral economics and the case for 'asymmetric paternalism', *University of Pennsylvania Law Review* 151, 1211-1254.

Diamond, Douglas and Philip Dybvig, 1983, Bank runs, deposit insurance, and liquidity, *Journal of Political Economy* 91, 401-419.

Dow, James and Sergio Werlang, 1992, Uncertainty aversion, risk aversion, and the optimal choice of portfolio, *Econometrica* 60, 197-204.

Easley, David and Maureen O'Hara, 2005, Regulation and return: The role of ambiguity, working paper, Cornell University.

Epstein, Larry and Martin Schneider, 2004, Recursive multiple priors, *Journal of Economic Theory* 113(1), 1-31.

Epstein, Larry and Tan Wang, 1994, Intertemporal asset pricing under Knightian uncertainty, *Econometrica* 62, 283-322.

Gabaix, Xavier, Arvind Krishnamurthy, and Olivier Vigneron, 2007, Limits of arbitrage: Theory and evidence from the mortgage-backed securities market, *Journal of Finance* 62(2), 557-95.

Gatev, Evan, and Phil Strahan, 2006, Banks' advantage in hedging liquidity risk: Theory and evidence from the commercial paper market, *Journal of Finance* 61(2), 867-92.

Geithner, Timothy, 2006, Implications of growth in credit derivatives for financial stability, Remarks at the New York University Stern School of Business Third Credit Risk Conference.

Gilboa, Itzak and David Schmeidler, 1989, Maxmin expected utility with non-unique priors, *Journal of Mathematical Economics*18, 141-153.

Goldstein, Itay and Ady Pauzner, 2005, Demand deposit contracts and the probability of bank runs, *Journal of Finance* 60(3), 1293-1328.

Gorton, Gary and Andrew Winton, 2003, Financial intermediation, in G.M. Constantinides & M. Harris & R. M. Stulz ed.: *Handbook of the Economics of Finance* (Elsevier).

Greenspan, Alan, 2004, Risk and uncertainty in monetary policy, Remarks for the AEA Meetings.

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

Gul, Faruk and Wolfgang Pesendorfer, 2005, The case for mindless economics, mimeo, Princeton University.

Hansen, Lars and Thomas Sargent, 1995, Discounted linear exponential quadratic gaussian control, *IEEE Transactions on Automatic Control* 40, 968-971.

Hansen, Lars and Thomas Sargent, 2003, Robust control of forward-looking models, *Journal of Monetary Economics* 50(3), 581-604.

Hansen, Lars, Thomas Sargent, Gauhar Turmuhambetova, and Noah Williams, 2006, Robust control and model misspecification, *Journal of Economic Theory* 128(1), 45-90.

Holmstrom, Bengt and Jean Tirole, 1998, Private and public supply of liquidity, *Journal of Political Economy* 106(1), 1-40.

Knight, Frank, 1921, *Risk, uncertainty and profit* (Houghton Mifflin, Boston).

McAndrews, Jamie, and Simon Potter, 2002, The liquidity effcts of the events of September 11, 2001, *Federal Reserve Bank of New York Economic Policy Review* 8, 59-79.

Melamed, Leo, 1998, Comments to Nicholas Brady's: The crash, the problems exposed, and the remedies, in Robert E. Litan and Anthony M. Santomero, ed.: *Brookings-Wharton Papers on Financial Services* (The Brookings Institution Press).

Routledge, Bryan and Stanley Zin, 2004, Model uncertainty and liquidity, NBER Working Paper # 8683.

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

Scholes, Myron, 2000, Crisis and risk management, *American Economic Review* 90(2), 17-21.

Sims, Christopher, 2001, Pitfalls of a minimax approach to model uncertainty, *American Economic Review* 91(2), 51-54.

Skiadas, Costis, 2003, Robust control and recursive utility, *Finance and Stochastics* 7, 475-489.

Stewart Jr., Jamie, 2002, Challenges to the payments System following September 11, Remarks by First Vice President Jamie B. Stewart, Jr. before the Bankers' Association for Finance and Trade.

Woodford, Michael, 1990, Public debt as private liquidity, *American Economic Review, Papers and Proceedings* 80, 382-88.