

NBER WORKING PAPER SERIES

COLLECTIVE RISK MANAGEMENT IN A FLIGHT TO QUALITY EPISODE

Ricardo J. Caballero
Arvind Krishnamurthy

Working Paper 12896
<http://www.nber.org/papers/w12896>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2007

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, William Hawkins, Burton Hollifield, Bengt Holmstrom, Dimitri Vayanos, Ivan Werning and 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. 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. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2007 by Ricardo J. Caballero and Arvind Krishnamurthy. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Collective Risk Management in a Flight to Quality Episode
Ricardo J. Caballero and Arvind Krishnamurthy
NBER Working Paper No. 12896
February 2007
JEL No. E30,E44,E5,F34,G1,G21,G22,G28

ABSTRACT

We present a model of optimal intervention in a flight to quality episode. The reason for intervention stems from a collective bias in agents' expectations. Agents in the model make risk management decisions with incomplete knowledge. They understand their own shocks, but are uncertain of how correlated their shocks are with systemwide shocks, treating the latter uncertainty as Knightian. We show that when aggregate liquidity is low, an increase in uncertainty leads agents to a series of protective actions -- decreasing risk exposures, hoarding liquidity, locking-up capital -- that reflect a flight to quality. However, the conservative actions of agents leave the aggregate economy over-exposed to negative shocks. Each agent covers himself against his own worst-case scenario, but the scenario that the collective of agents are guarding against is impossible. A lender of last resort, even if less knowledgeable than private agents about individual shocks, does not suffer from this collective bias and finds that pledging intervention in extreme events is valuable. The intervention unlocks private capital markets.

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

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

1 Introduction

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

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

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

of Russian government bond spreads, Brazilian spreads, and U.S. Treasury bond spreads was a surprise to even sophisticated market participants. These high correlations rendered standard risk management models obsolete, leaving financial market participants searching for new models. Agents responded by making decisions using “worst-case” scenarios and “stress-testing” models. 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* and not merely an increase in risk. The emphasis on tail outcomes and worst-case scenarios in agents’ decision rules suggests uncertainty aversion rather than simply risk aversion. It is also noteworthy 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 recent default of the Amaranth hedge fund barely caused a ripple in financial markets. The one-of-a-kind aspect of flight to quality episodes supports the view that uncertainty is an important ingredient in these episodes.

Section 2 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. Knightian uncertainty leads agents to make decisions based on worst-case scenarios. 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 flexibly move it across markets.

The main results of our paper are in Sections 3 and 4. As indicated by Greenspan’s

²See 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 Ashcraft and Duffie (2006) on 9/11.

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 as they exacerbate the aggregate shortage of liquidity. Central bank policy can be designed to alleviate the over-conservatism. We show that a lender of last resort (LLR), even one facing the same incomplete knowledge that triggers agents' Knightian responses, can unlock capital and improve outcomes. It does so by committing to intervene during extreme events when agents' liquidity is depleted. Importantly, because these extreme events are unlikely, the expected cost of this intervention is minimal. If credible, the policy derives its power from a private sector multiplier: each pledged dollar of public intervention in the extreme event is matched by a comparable private sector reaction to free up capital. In this sense, the LTCM bailout was important not for its direct effect, but because it served as a signal of the Fed's readiness to intervene should conditions worsen. The signal unlocked *private* capital markets.

The value of the LLR facility is usually analyzed in terms of Diamond and Dybvig's (1983) model of bank runs. The LLR rules out the "bad" run equilibrium in their model. While our environment is a variant of Diamond and Dybvig's, it does not include the sequential service constraint that leads to a coordination failure and informs their discussion of policy. The benefit of the LLR in our model derives from a different mechanism, which is the collective bias in agent decisions caused by an increase in Knightian uncertainty. Our model also provides a clear prescription to central banks on when to intervene in financial markets: The benefit of the LLR is highest when there is *both* insufficient aggregate liquidity and Knightian uncertainty. 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.

Holmstrom and Tirole (1998) study how a shortage of aggregate collateral limits private liquidity provision (see also Woodford, 1990). Their analysis suggests that a credible government can issue government bonds which can then be used by the private sector for liquidity provision. The key difference between our paper and those of Holmstrom and Tirole, and Woodford, is that we show how even 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.³ 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 and asks to facilitate trade of an asset. Their model captures an important aspect of flight to quality: uncertainty aversion can lead to a sudden widening of the bid-ask spread, causing agents to halt trading and reducing market liquidity. Both our paper and Routledge and Zin share the emphasis on financial intermediation and uncertainty aversion as central ingredients in flight to quality episodes. But each paper captures different aspects of flight to quality. Easley and O'Hara (2005) study a model where ambiguity (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.

Finally, in our model agents are only Knightian with respect to systemic events. Epstein (2001) explores the home bias in international portfolios in a related setup, where agents are more uncertain about foreign than local markets.⁴ As Epstein points out, this modelling also highlights the difference between high risk aversion and aversion to Knightian uncertainty. Moreover, our modelling shows that max-min preferences interact with macroeconomic conditions in ways that are not present in models with an invariant amount of risk aversion. We show that when aggregate liquidity is plentiful, Knightian and standard agents behave identically. However when there is an aggregate liquidity shortage, the actions of these agents differ, leading to flight to quality in the Knightian model.

³There 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.

⁴Epstein's model is closely related to our model in Caballero and Krishnamurthy (2005).

2 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 unexpected 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.

2.1 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,

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

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,

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

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

We note that the sequential shock structure means that,

$$\phi(1) > \phi(2) > 0. \quad (3)$$

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

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

With $\alpha_1 = 1, \alpha_2 = 0$ 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.”

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

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:⁵

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

We define,

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

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 uncertainty aversion and write:

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

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 Ashcraft and Duffie, 2006). 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.

⁵For further clarification of the structure of shocks and agents' uncertainty, see the event tree that is detailed in the Appendix.

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.

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.

2.2 A benchmark

We begin by analyzing the problem when $K = 0$. This case will clarify 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 $s = (\# \text{ of waves, } \omega\text{'s shock})$ 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.

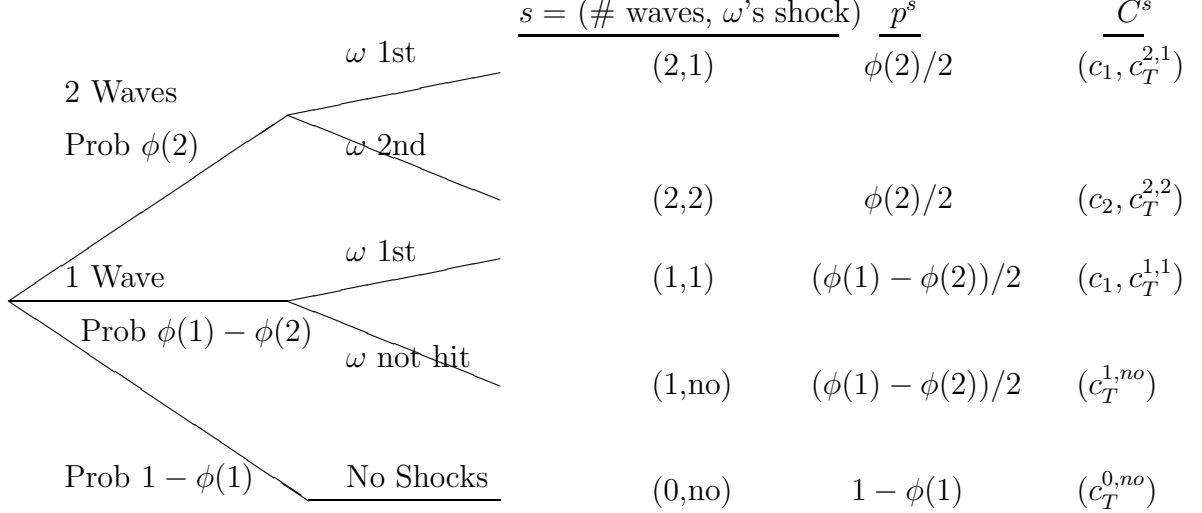


Figure 1: Benchmark Case

The planner's problem is to solve,

$$\max_c \sum p^s U^\omega(C^s)$$

subject to resource constraints that for every shock realization, the promised consumption amounts are not more than the total endowment of Z :

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

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:

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

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:

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

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 .

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,

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

with first order condition,

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

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

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

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.

2.3 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.

2.4 $K > 0$ 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 planners objective to,

$$\max_c \min_{\theta_\omega \in \Theta} \sum p^{s,\omega} U(C^s) \quad (11)$$

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.⁶

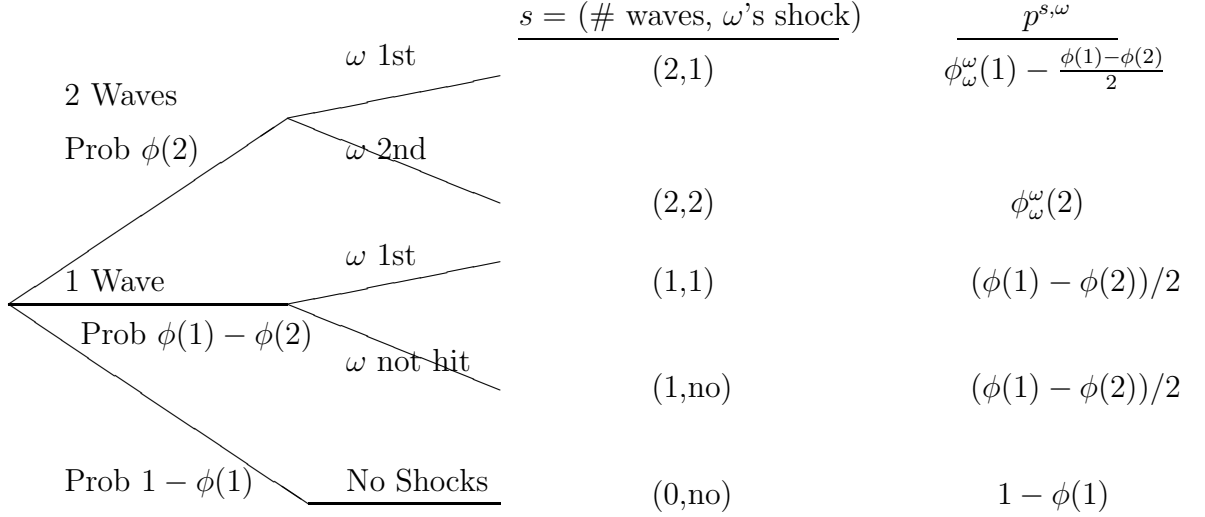


Figure 2: Robustness Case

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,

$$V(\mathcal{C}; \theta_\omega^\omega) \equiv \max_c \min_{\theta_\omega^\omega \in \Theta} \phi_\omega^\omega(1)u(c_1) + (\phi(2) - \phi_\omega^\omega(2)) \beta c_T^{2,1} + \phi_\omega^\omega(2) (u(c_2) + \beta c_T^{2,2}) + \frac{\phi(1) - \phi(2)}{2} \beta c_T^{1,no} \quad (12)$$

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.

⁶We 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:*

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

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

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

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:

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

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

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

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

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

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. 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.

2.5 Flight to quality

A flight to quality episode can be understood in our model as a comparative static across K or Z . Let us fix a value of K and Z and suppose that at a date -1 , agents enter into a contractual arrangement as dictated by Proposition 2. At date 0, there is an unanticipated (non-contracted) event that increases K (or reduces Z) and leads agents to rewrite contracts. We may think of a flight to quality in terms of this rewriting of contracts.

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 than 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) 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 be 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 Ashcraft and Duffie, 2006).

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 an unanticipated rise in Knightian uncertainty rather than the realization of the bad equilibrium. 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 creditors.⁷

⁷There is an underlying connection between the crisis in our model and that of Diamond and Dybvig (1983). As Diamond and Dybvig emphasize, a deposit contract implements an optimal shock-contingent allocation of

3 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. Since the agents of the model are not standard expected utility maximizers (i.e. adhering to the Savage axioms), there are delicate issues that arise when defining the central bank’s objective. On the one hand, if the central bank’s objective is the min-max objective of the planner in (11), then since we have previously derived the equilibrium as the solution to a planner’s problem, the central bank cannot improve on the allocation. On the other hand, throughout this section, we argue for a paternalistic central bank objective which does not incorporate agents’ worst-case probability assessments. We argue that there are “natural” paternalistic policies that arise in our setting.⁸

3.1 Collective bias

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:

$$\phi_{\omega}^{\omega}(1) - \phi_{\omega}^{\omega}(2) = \frac{\phi(1) - \phi(2)}{2} \left(\frac{u'(c^*)}{u'(Z)} \right) \quad (14)$$

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

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

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 (14) and using the definitions above, we find that liquidity. In the fully robust equilibrium of our model, agents choose an allocation that is non-contingent on each other’s shocks instead of the contingent liquidity allocation of the benchmark economy. In this sense, robust agents’ preference for shock-independent liquidity allocations is related to the behavior of panicked depositors in a bank run.

⁸The appropriate notion of welfare in models where agents are not rational is subject to some debate in the literature. 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.

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, implying that agents' conservative probabilities are *collectively biased*.

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

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

3.2 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. 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_ω^s . Additionally, $p_\omega^{s,CB}$ may differ from $p_\omega^{s,\omega}$.

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

$$\begin{aligned} V^{CB} &\equiv \int_{\omega \in \Omega} \sum p_\omega^{s,CB} U(C^s) \\ &= \sum p^s U(C^s) \end{aligned} \tag{15}$$

where the second line follows from exchanging the integral and summation, and using the fact that the aggregate probabilities are common knowledge.

One can view this objective as descriptive of how central banks behave: As noted above, 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.

The objective can also be seen as corresponding to agents' welfare. Two potential issues arise when making the latter identification. First, since the objective is based on ex-post consumption utility, it ignores any utility costs that agents may suffer because of date 0 "anxieties" about ex-post outcomes. Second, if we introduce some rational agents into the economy, they may suffer from the type of policies we discuss in this and the next section. We address the second concern below. In contrast, we ignore the first one, which is what makes ours a paternalistic criterion.

3.3 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 is:

$$\frac{\phi(1)}{2}u'(c_1) - \frac{\phi(2)}{2}u'(c_2) - \frac{\phi(1) - \phi(2)}{2}\beta. \quad (16)$$

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 (16) implies that 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 (16), 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 $K > 0$.

Summarizing these results:

Proposition 3 *For any $K > 0$, if the economy has insufficient aggregate liquidity ($Z < c^*$), agent decisions are collectively biased. Agents choose too much insurance against receiving shocks second relative to receiving shocks first. A central bank that maximizes the expected (ex-post) utility of agents in the economy can improve outcomes by reallocating agents' insurance toward the first shock.*

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.

The central bank reaches this conclusion requiring only knowledge of aggregate probabilities. As we have remarked, the central bank may be more confused than individual agents about individual θ_ω s. In this sense, the central bank may be the *least informed* agent of the economy. The important point is that the central bank does not suffer from collective bias.⁹ In fact, going beyond local perturbations, the linearity with respect to individual probabilities of the central bank objective in (15) implies that the central bank solves the $K = 0$ problem of the planner in Section 2, regardless of how uninformed it may be about individual θ_ω s.

3.4 Welfare and paternalism

To what extent does the central bank's objective correspond to welfare? We offer two further arguments in favor of our paternalistic welfare criterion in this subsection. First, suppose that a fraction of the agents in the economy are rational and had probabilities such that $\theta_\omega = +K$; i.e. these agents know that the worst-case probabilities are truly their own probabilities. Suppose the rest of the agents are Knightian agents with θ_ω 's such that the average θ_ω across all agents is zero. In this case, reallocating insurance from second to first shocks hurts the rational agents while it helps the average Knightian agent. However, note that the envelope theorem implies that the utility cost to the rational agents is second order, while, since the

⁹If the central bank is uncertain about the values of $\phi(1)$ and $\phi(2)$, then we could overturn the result. In particular, we may imagine a situation in which the central bank is uncertain about these probabilities, and its objective function overweights liquidity crises (i.e. the incidence of both shocks occurring). In this case, the central bank will also be subject to the "overinsurance" bias of agents. However, this "bias" is of a different nature than the one we emphasize as it would not be collectively inconsistent with conditional probabilities.

envelope theorem does not apply to the average Knightian agent, the policy results in a first order utility gain to the Knightian agents. Thus, although the central bank's policy is not Pareto improving, it involves asymmetric gains to the Knightian agents. Our policy satisfies the asymmetric paternalism criterion that Camerer, et. al., (2003) propose for evaluating policy when some agents are behavioral.

Another justification is through the following thought experiment: Suppose that we repeat the liquidity episode we have described infinitely many times. At the beginning of each episode, agent ω draws 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 will have him using $\theta_\omega = +K$. The central bank's ex-post utility criterion corresponds to the expected consumption utility of an agent across all of these episodes, where the expectation is taken using agents' known probabilities.¹⁰

3.5 Risk aversion versus uncertainty aversion

Given the centrality of collective bias in Proposition 3, it should be apparent that while increased risk aversion may generate positive implications (flight to quality) that are similar to those of Knightian uncertainty, its normative implications are not. Without collective bias, and regardless of the agent's degree or change in 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 make this point precise by returning to the agents' first order condition in the $K = 0$ case of Proposition 1. Equation (9) simplified by setting $\beta = 0$ is,

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

Suppose that $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$. Then the first order condition is,

$$\left(\frac{c_1}{c_2}\right)^{-\gamma} = \frac{\phi(2)}{\phi(1)}.$$

If we think of a flight to quality event in terms of increased risk aversion then it is clear that c_1 falls and c_2 rises as γ rises. However, since $K = 0$, there is no role for the central bank in this case.

¹⁰In living through repeated liquidity events, an agent will learn 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 will go through for a small enough agent discount rate.

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.

4 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 will be 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.

4.1 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:

$$V_{Z^G}^{CB,total} = \int_{\Omega} \left[\phi_{\omega}(1) u'(c_{1,\omega}) \frac{dc_{1,\omega}}{dZ^G} + \phi_{\omega}(2) u'(c_{2,\omega}) \frac{dc_{2,\omega}}{dZ^G} + \frac{\phi(1) - \phi(2)}{2} \beta \frac{dc_T^{1,no}}{dZ^G} \right] d\omega.$$

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 yield,

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

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 also 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.¹¹

4.2 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

¹¹In 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.

shock. However, when flight to quality is the concern, the reallocation improves (ex-post) outcomes on average.¹² 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:

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

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.

¹²Note 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.

4.3 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 $n = 1 \dots N$ 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 n th 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 j th 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,

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

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

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

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 N th shock. Instead, it suggests that for any given amount of resources available for intervention, the LLR should first pledge resources to the N th 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.

5 Final remarks

Flight to quality is a pervasive phenomenon that exacerbates the impact of recessionary shocks and financial accelerators. The starting point of this paper is a model of this phenomenon based on Knightian uncertainty aversion among financial specialists. We show that when aggregate liquidity is plentiful, agents' uncertainty does not affect the equilibrium. However, when there is both Knightian uncertainty and agents think that aggregate liquidity is scarce, they take a set of protective actions to guarantee themselves safety, but which leave the aggregate economy overexposed to negative shocks.

In this context, a Lender of Last Resort policy is valuable when used to support rare events. Our model prescribes that the benefit of the LLR is highest when there is *both* insufficient

aggregate liquidity and Knightian uncertainty. Many of the examples we have discussed in this paper (1987 market crash, Fall of 1998, 9/11) satisfy these criteria. But just as important are times when the “dog did not bark.” For example, a recession in which the ingredients are not both central, does not call for central bank action. Likewise, 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. The model also suggests that it is important for the central bank to have an understanding of whether a given market dynamic is driven by risk or uncertainty.

There is currently considerable uncertainty over how the downgrade of a top name will affect the credit derivatives market. Our model suggests that ex-ante actions to reduce the extent of uncertainty during a flight to quality episode are valuable. For example, recent moves to increase transparency and risk assessment in this market as well as streamline back-office settlement procedures can be viewed in this light (Geithner, 2006).

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

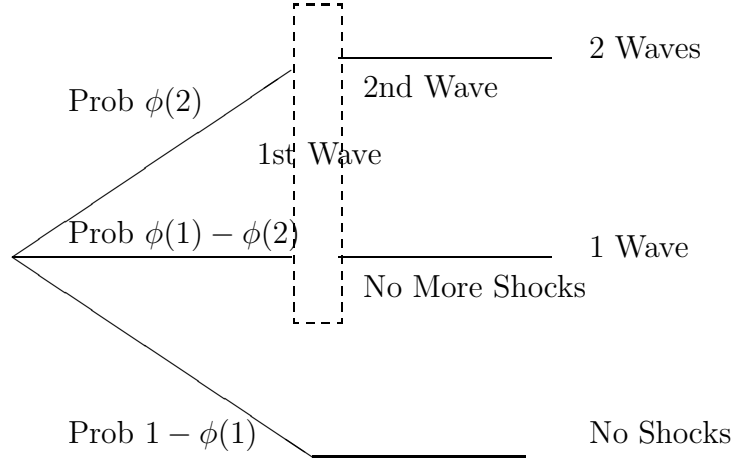
References

- [1] Ashcraft, Adam and Darrell Duffie, 2006, "Systemic Dynamics in the Federal Funds Market," Mimeo, Stanford University.
- [2] Barro, R.J., 2005, "Rare Events and the Equity Premium," Mimeo, Harvard University.
- [3] Bernheim, B.D. and A. Rangel, 2005, "Behavioral Public Economics: Welfare and Policy Analysis with Non-Standard Decision-Makers," Stanford Mimeo, June.
- [4] Caballero, R.J. and A. Krishnamurthy, 2005, "Financial System Risk and Flight to Quality," NBER WP # 11996.
- [5] Calomiris, Charles W. 1994, "Is the Discount Window Necessary? A Penn Central Perspective," Review, Federal Reserve Bank of St. Louis (May/June), pp. 31-55
- [6] Camerer, C., S. Issacharoff, G. Loewenstein, T. O'Donoghue, and M. Rabin, "Regulation for Conservatives: Behavioral Economics and the Case for 'Asymmetric Paternalism'," University of Pennsylvania Law Review, v151, 1211-1254, 2003.
- [7] Diamond, D.W. and P.H. Dybvig, 1983, "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy* 91, pp. 401-419.
- [8] Dow, J. and S. Werlang, 1992, "Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio," *Econometrica* 60, p. 197-204.
- [9] Easley, David and Maureen O'Hara, 2005, "Regulation and Return: The Role of Ambiguity," working paper, Cornell University.
- [10] Epstein, L., 2001, "Sharing Ambiguity," *American Economic Review*, 91(2), 45-50.
- [11] Epstein, L. and M. Schneider, 2004, "Recursive Multiple Priors," *Journal of Economic Theory*, 113(1), 1-31.
- [12] Epstein, L. and T. Wang, 1994, "Intertemporal Asset Pricing Under Knightian Uncertainty," *Econometrica* 62, 283-322.
- [13] Gabaix, Xavier, Arvind Krishnamurthy, and Olivier Vigneron, 2006, "Limits of Arbitrage: Theory and Evidence from the Mortgage-Backed Securities Market," forthcoming, *Journal of Finance*.

- [14] 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.
- [15] 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.
- [16] Gilboa, I. and D. Schmeidler, 1989, "Maxmin Expected Utility with Non-unique Priors," *Journal of Mathematical Economics*, 18, 141-153.
- [17] Greenspan, A, 2004, "Risk and Uncertainty in Monetary Policy," Remarks for the AEA Meetings, January.
- [18] Gul, F. and W. Pesendorfer, 2005, "The Case for Mindless Economics," Princeton Mimeo, November.
- [19] Hansen, L. and T. Sargent, 1995, "Discounted Linear Exponential Quadratic Gaussian Control," *IEEE Transactions on Automatic Control* 40, 968-971.
- [20] Hansen, L.P. and T.J. Sargent, 2003, "Robust Control of Forward-looking Models," *Journal of Monetary Economics* 50(3), 581-604.
- [21] Hansen, L.P., T.J. Sargent, G.A. Turmuhambetova, N. Williams, 2004, "Robust Control and Model Misspecification," April, Mimeo.
- [22] Holmstrom, Bengt and Jean Tirole, "Private and Public Supply of Liquidity," *Journal of Political Economy*, 106(1), 2/98, pp. 1-40.
- [23] Melamed, L., 1998, "Comments to Nicholas Brady's: The Crash, the Problems Exposed, and the Remedies," in *Brookings-Warton Papers on Financial Services*, Brookings, D.C.
- [24] Routledge, Bryan and Stanley Zin, 2004, "Model Uncertainty and Liquidity," working paper, Carnegie Mellon University.
- [25] Scholes, Myron S. 2000, "Crisis and Risk Management," *American Economic Review*, Vol. 90, No. 2, Papers and Proceedings of the One Hundred Twelfth Annual Meeting of the American Economic Association (May), pp. 17-21

- [26] Sims, Christopher A., 2001, "Pitfalls of a Minimax Approach to Model Uncertainty," working paper, Princeton University.
- [27] Skiadas, Costis, 2003, "Robust Control and Recursive Utility," *Finance and Stochastics* 7, 475-489.
- [28] Stewart Jr., Jamie B., 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.
- [29] Woodford, M., "Public Debt as Private Liquidity," *American Economic Review, Papers and Proceedings* 80, May 1990, pp. 382-88.

A Event Tree and Probabilities



The event tree is pictured above. 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 .

B Proof or 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, $\bar{\theta}_\omega^\omega \in \operatorname{argmin}_{\theta_\omega^\omega} V(\bar{\mathcal{C}}; \theta_\omega^\omega)$. That is, nature chooses θ_ω^ω optimally given the agent’s choice of $\bar{\mathcal{C}}$. Second, $\bar{\mathcal{C}} \in \operatorname{argmax}_{\mathcal{C}} V(\mathcal{C}; \bar{\theta}_\omega^\omega)$. That is, the agent chooses \mathcal{C} optimally given nature’s choice of $\bar{\theta}_\omega^\omega$.¹³

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,

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

¹³The 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 $\mathcal{C} = \bar{\mathcal{C}}$.

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,

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

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 $\frac{\partial V}{\partial \theta_\omega^\omega} > 0$.