PRICING SYSTEMIC CRISES:

Monetary and Fiscal Policy When Savers Are Uncertain*

Andreas Lehnert
Board of Governors of the Federal Reserve System
Mail Stop 93
Washington DC, 20551
(202) 452-3325
alehnert@frb.gov

Wayne Passmore
Board of Governors of the Federal Reserve System
Mail Stop 93
Washington DC, 20551
(202) 452-6432
wpassmore@frb.gov

Version information:
Revision: 6.8
Last Revised: August 12, 1999

*The views expressed in this paper are ours alone and do not necessarily reflect those of the Board of Governors of the Federal Reserve System or its staff. We thank Wen-Fang Liu, Darrel Cohen, John Curran, Eric Engen, Stephen Oliner, and seminar participants at the Federal Reserve Board for helpful discussions. Any remaining errors are our own.

PRICING SYSTEMIC CRISES:

Monetary and Fiscal Policy When Savers Are Uncertain

Abstract

In our model, the return on assets depends on the joint behavior of all savers; if all savers sell the asset simultaneously then there will be a financial "Armageddon" and the return will be quite low. We assume that individual savers are risk-neutral but uncertainty-averse and cannot form precise estimates of the behavior of the market (all other savers). We determine equilibrium investment using the modern theory of *Knightian uncertainty*, that is, decision-making with multiple subjective prior distributions, which shows that savers will act to maximize their minimum payoff in the presence of uncertainty. Savers divide their portfolios between two assets: A risky asset, representing a fully-diversified share in economy-wide production, and a riskless asset, representing government bonds. Government bonds are backed by the tax authority of the state, and so always pay off. In times of high uncertainty, savers' demand for government bonds will increase, and spreads between the return on bonds and the expected return on risky assets will widen. As a result, investment in the risky asset will decrease. If the government responds with a purely *monetary* policy of reducing the risk-free rate, it will make the bond less attractive and force savers to hold more of the risky asset. We show that such monetary expansions are Pareto-improving, but that they fail to recapture the optimal allocation. To restore investment and savings to their optimal levels, the government must also use a *fiscal* policy of distortionary taxes to discourage current consumption and leisure. *Journal of Economic Literature* classification numbers: G2, E5, E44, D8.

Keywords: Knightian uncertainty, financial crisis, monetary and fiscal policy

1 Introduction

During times of market turmoil, savers are unable to form precise estimates of the distribution of returns of their portfolios. Such times are also often marked by unusually high interest-rate spreads between riskless and risky assets, unusually high levels of asset return volatility, and an attendant difficulty in obtaining project financing. The distribution of an asset's returns is difficult to forecast, at least in part, because it depends on the likely actions of other market participants. If worried savers liquidate part of their portfolios, the value of all risky assets can fall dramatically. To combat such crises of confidence, governments are urged, on the one hand, to take dramatic action to restore confidence in financial markets with a mix of an expansionary monetary policy and explicit or implicit guarantees to producers or to financial institutions. On the other hand, governments are cautioned not to provoke inflation or promote moral hazard. In this paper, we characterize the government's optimal policy responses to such an episode.

Financial assets typically carry a promised rate of return, but this return is subject to a host of so-called "risks." These range from the well understood and predictable (such as *actuarial risk*) to the exotic (such as "*fat-finger*" *risk*) or unpredictable (such as *political risk*).[1] Somewhere in the middle

[1] "Fat-finger" risk is the risk of mis-keying an order into an electronic terminal. For a representative "taxonomy of risks" see Rahl (1998).

lies so-called *model risk*, the risk of having mis-modelled the underlying stochastic process of interest. This "risk" is largely what we have in mind when we discuss "uncertainty," although, of course, all risks carry some uncertainty with them. Uncertainty occurs when the information available about a particular event is too sparse or too vague to form a precise estimate of its probability; the sparser or vaguer the information, the greater the uncertainty.

Uncertainty has a qualitatively different effect on financial markets than risk. For example, on November 3, 1998 the Chicago Mercantile Exchange (CME) listed a futures contract based on the Quarterly Bankruptcy Index (QBI).[2] If savers were truly risk-neutral, there should have been a brisk trade in the contract, given the disparity in estimates of the distribution of bankruptcies, and the fact that the CME waived its contract fee. However, as of mid-April, 1999, not a single contract had been traded. While there are several possible explanations for this, market participants often cite high levels of model risk for staying away from contracts based on the CME-QBI, suggesting that uncertainty indeed plays a role.[3]

Similarly, in the late summer and fall of 1998 spreads between risky and riskless assets widened dramatically, while asset return volatility also in-

[2] The CME-QBI is a count, measured in thousands, of all new bankruptcy case filings in U.S. bankruptcy courts. For more information, see Chicago Mercantile Exchange (1999).

[3] As a further example of this kind of uncertainty, Christofferson, Diebold, and Schuermann (1998) analyze the typical volatility forecast underpinning the popular "Value at Risk" risk-management model and find that, beyond a ten-day trading horizon, the volatility forecast is of very poor quality.

creased. Market participants reported that they were dissatisfied with their risk-management models. There was a widespread sense that savers were less willing to accept risk, which reduced liquidity in financial markets and, as a result, jeopardized the continued expansion of real economic output. The Federal Reserve responded by cutting the target Fed Funds rate by 75 basis points. The return on risky assets also declined, although spreads remained elevated, and the real consequences of the financial turmoil never materialized. We argue that the turmoil was a direct consequence of uncertainty: savers began to doubt their estimates of the distribution of portfolio returns, in part because these returns depend on the behavior of other savers. In particular, savers' estimates of the probabilities of extreme financial events became more imprecise, causing them to flee risky assets in favor of riskless assets.[4]

In our model savers are able to invest in two assets: A risky asset, representing a share in economy-wide production and a riskless asset, representing a government bond. The risky asset carries only *aggregate* risk because we assume that the saver has diversified away all of the idiosyncratic, project-specific, risk. We model financial crises by assuming that if total investment in the risky asset falls below a critical level, its per-share return will plunge. Above the critical level, the return is smoothly decreasing in total investment.

[4]Formally, we model uncertainty following Liu (1999) and Epstein and Wang (1994), where uncertainty and uncertainty-aversion are summarized by a scalar parameter.

Government bonds, in contrast, pay a return that varies neither with the state of nature nor with the level of aggregate investment. They are ultimately backed by the tax authority of the state and so will be paid off even if a bad state of nature is realized or if investment falls below the critical level [see Burnham (1989)]. Uncertainty-averse savers will demand an uncertainty premium over the return on these bonds to hold a share in the productive, though risky, asset. In periods of greater uncertainty these spreads will widen, investment will decline, and total output will be lower.

We show that, in times of increased uncertainty, savers will "flee to quality" and demand too much of the government's safe bond relative to the risky asset. We further show that a purely *monetary* policy of reducing the risk-free rate is Pareto improving, because it makes the bond less attractive and induces savers to hold more of the risky asset, although it also causes savers to save less in total. We also discuss the consequences of a policy that holds the stock of riskless bonds constant, so that uncertainty drives down the equilibrium risk-free rate. When there is a bad technology shock (so that labor's product is low), we show that the scope for monetary expansions will be sharply limited; and we further show that if uncertainty is very high, monetary policy may be wholly ineffective. To recapture the optimum–that is, to restore total savings, labor effort and consumption to their optimal levels–the government must use an additional *fiscal* policy of distortionary taxes to discourage current consumption and leisure.

The paper is organized as follows: In section 2 we present the model without uncertainty, in section 3 we present a condensed discussion of the modern theory of uncertainty and show how it affects outcomes in our model and in section 4 we derive optimal monetary and fiscal policies. In section 5 we provide conclusions and a brief discussion of the implications of this work for the conduct of monetary policy. Proofs of propositions are relegated to an appendix.

2 The Model

2.1 Savers

Many individual savers, whose population is normalized to one, are born in each period t and live for two periods. Savers consume an amount c_0^t while young and c_1^t while old. (We will generally denote generations with superscripts and age with subscripts.) In addition, young savers are endowed with a single unit of time which they may split between leisure, ℓ_0^t, and labor effort, $1 - \ell_0^t$. Savers have preferences given by:

(1) $$u(c_0^t, \ell_0^t) + c_1^t, \text{ where } u_c, u_\ell > 0,\ u_{cc}, u_{\ell\ell} < 0 \text{ and } u_{c\ell}^2 < u_{cc}u_{\ell\ell}.$$

Note that savers are risk-neutral with respect to consumption while old. This allows us to isolate the pure effect of uncertainty aversion. Finally, because all savers are identical, we will define equilibria using a representative saver.

Labor effort earns a certain real wage of W_t, which will depend on the state (see below). Thus at the end of the first period of life, agents have available to them an amount $s_t = (1 - \ell_0^t)W_t - c_0^t$ to save for consumption while old. Their savings portfolio is divided between investment in an economy-wide risky asset, x_t and holdings of the safe asset, b_t. The safe asset will pay a gross return in the following period of r_t. The risky asset will pay a state-contingent gross return in the following period, $t + 1$, conditional on the total amount invested (see below). For now, assume that the risky asset pays a gross per-share return of $R(\omega_{t+1}, X_t)$.

Aggregate states ω, which represent a technology shock, must lie in the set of allowed states Ω. There is a known, true, distribution over these states $\pi^\star$, which does not change over time.

The wage earned by young agents will depend on the state in the current period, $W_t = W(\omega_t)$. Assume that the wage function $W(\omega)$ satisfies:

(A1) $$W(\omega, X) > 0, \text{ all } \omega;$$

This assumption guarantees that young workers will always earn a positive wage.

Savers must pay an excise tax rate of τ_c on consumption while young and a labor income tax rate of τ_ℓ. In addition, the government makes lump-sum transfers (taxes if negative) in each period of life of H_0^t and H_1^t. The saver's budget constraints in each period of life are:

$$(1 + \tau_c)c_0^t + (1 - \tau_\ell)W_t\ell_0^t + x_t + b_t \leq (1 - \tau_\ell)W_t + H_0^t, \tag{2}$$

$$c_1^t(\omega_{t+1}) = R(\omega_{t+1}, X_t)x_t + rb_t + H_1^t. \tag{3}$$

Savers take as given the return to the risky asset in each state, $R(\omega_t)$.

2.2 Producers

The risky asset stands for a fully diversified investment share of a large number of risky projects. These projects are owned and operated by two-period-lived entrepreneurs who consume only in the second period of life, are risk-neutral and have no other productive assets. As a result, they will be willing to accept a loan to finance their projects if the expected value of output net of loan costs is greater than or equal to zero. Projects' ex ante risk characteristics and ex post outputs are costlessly observed by all parties in the economy. There is a continuum of producers and their associated projects; these are named n and are distributed uniformly in the range $[0, I_1]$. Projects of type n *potentially* produce a return to their owners (the producers) of $\rho(n, \omega)$ in state ω, where $\rho(n, \omega)$ is twice continuously differentiable in n for each ω. Assume that these projects are ordered so that:

$$\frac{\partial \rho(n, \omega)}{\partial n} \equiv \rho_n(n, \omega) < 0, \text{ all } \omega \text{ in } \Omega.$$

In other words, the most productive projects have the lowest name. In addition, assume that $\rho(0, \omega) < \infty$, so that even the most productive project in the best possible state has a finite return. Finally, define the potential expected return to projects of type n as $\rho(n)$:

$$\rho(n) \equiv \int_{\Omega} \rho(n, \omega) \; d\pi^{\star}.$$

When a project type is funded by savers, it generates an external social benefit. If not enough projects are funded, then the projects that are funded do not achieve their potential, and produce a very low return, ρ_0. Let this critical level of investment be denoted as X_L. Thus if some total amount X is invested in the risky projects (i.e. the project type named X and all better project types $n < X$ are financed), total expected economy-wide production is:

$$f(X) = \begin{cases} \rho_0 X & \text{if } X \le X_L, \\ \int_0^X \rho(n)\, dn & \text{if } X > X_L. \end{cases}$$

Finally, assume that the return when aggregate investment is too low, ρ_0, is smaller than the return earned by any project in any state:

$$\rho_0 < \rho(n, \omega), \text{ all } n \in [0, I_1], \omega \text{ in } \Omega.$$

An expected return function satisfying these characteristics is displayed in figure 1.

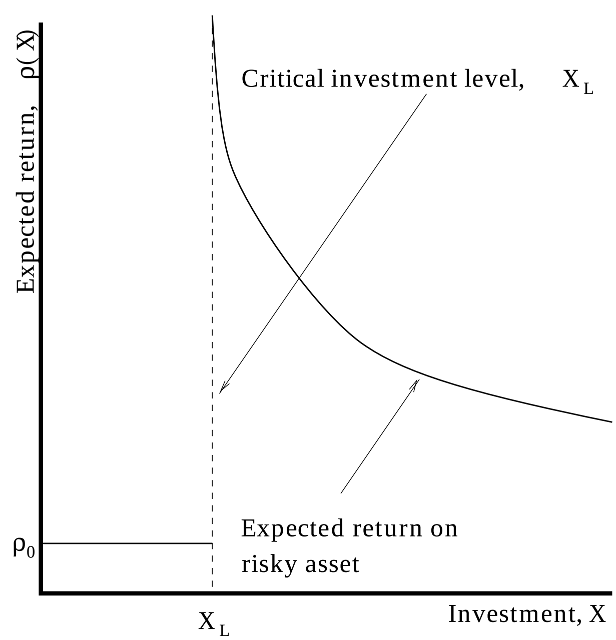

Figure 1: Expected returns to the risky asset as a function of total investment, $\rho(X)$.

This form of production captures the sense of recent research into financial crises. As total investment decreases, the projects that are financed may suffer because they do not have access to the products of other projects, as in the "race for the exit" model of Subrahmanyam and Titman (1999). Further, the projects themselves may be illiquid, so that a decrease in investment from one period to the next requires the premature (and costly) liquidation of fundamentally illiquid assets, as in Allen and Gale (1998) and Diamond and Dybvig (1983). Finally, failure to invest in one project may cause the default of another, as in Lagunoff and Schreft (1998,1999).

Note that all of our results are unaffected if the critical aggregate investment level is zero, $X_L = 0$. All that we require is that the returns realized by the very first project to be funded are quite low if that project is the only one funded. As a result, a saver who believes that no other savers intend to finance projects will not find it profitable to fund a project.

Savers place their investments x_t with a zero-cost, zero-profit economy-wide competitive intermediary that allows them to own a fully-diversified share of the projects that are funded. Both savers and producers will act as price takers with respect to the interest rate charged on individual loans. All projects will be charged the same rate in each state. If a total quantity $X > X_L$ is committed to the intermediary, the owners of the last project type that is financed must be just indifferent between accepting their loans and undertaking production and not. Thus savers realize per-share returns of $\rho(X, \omega)$. If an amount below the critical level is invested, $X < X_L$, we assume that all production is surrendered to the savers by the producers, so that the per-share return is ρ_0.

Thus if an amount $X > X_L$ is invested, total expected producer's surplus (the sum of all expected utilities) is given by:

$$\text{PS}(X | X \geq X_L) = \int_0^X [\rho(n) - \rho(X)] \, dn. \tag{4}$$

If less than X_L is invested, we assume that producers make zero surplus and all output is divided among savers. Although we do not explicitly consider them here, we also assume that the government has available to it a complete array of lump-sum taxes and transfers that it can levy on producers.

2.3 The Government Bond

The government receives goods in period t from the young of that period in exchange for tokens which may be redeemed in period $t + 1$ for some known amount. In this paper, we assume that the government has a monopoly on the safe asset, which is backed by a storage technology earning a unit return. Thus by manipulating the return to the only truly safe asset in the economy, the government will be able to directly affect the risk-free rate of return.

If the government pays a gross return $r_t < 1$ on its bonds from period t to $t+1$, it realizes net revenues in period $t+1$, while if it pays a gross return $r_t > 1$, it must make a net outlay in period $t+1$. In the first case, we imagine that this “seigniorage” revenue is refunded lump-sum to the generation that pays it, while in the latter case, we imagine that the government levies taxes on either generation to meet its obligation. Thus government bonds are special because they are backed by the tax authority of the government and because the return on bonds is totally independent to the quantity of bonds sold.

The government's period-t budget constraint is thus:

$$(5) \qquad -H_0^t - H_1^{t-1} + \tau_\ell(1-\ell_0^t)W_t + \tau_c c_0^t + (1-r_{t-1})B_{t-1} + B_t \geq 0.$$

Recall that superscripts denote generations and subscripts denote age. In period t the government removes from storage the bond purchases of the current old generation, B_{t-1}. The government promised a gross return r_{t-1} on its bonds, and so realizes revenue of $(1-r_{t-1})B_{t-1}$ on its bonds. Notice that the government can transfer resources between the two generations alive in each period with lump-sum transfers (or taxes) on the current young (H_0^t) and the current old (H_1^{t-1}). We assume (see below) that the government treats each generation as essentially independent from its neighbors, and so does not favor one over the other. Because we assume that the government has no expenses of its own the inequality in (5) holds with equality.

2.4 Equilibrium Without Knightian Uncertainty

In the case without Knightian uncertainty we look for an equilibrium in each period t, conditional on a realization of the shock term ω_t (and thus the wage rate W_t), in which the government does nothing but provide a zero-cost storage asset paying a gross return $r_t > 0$.

Assume that the risk-free rate lies between the maximum and minimum expected returns on the risky asset:

$$\min_{X \geq 0} \rho(X) \leq r_t \leq \max_{X \geq 0} \rho(X).$$

Because savers are risk neutral, they are not concerned with the distribution of returns, and will demand a common expected return of r_t on elements of their savings portfolio (if both assets are held in positive quantities). As a consequence of this, we can characterize a saver's optimal policy, the solution to maximizing (1) subject to the constraints (2) and (3), as a savings function $s[W(\omega_t), r_t]$ or $s(\omega_t, r_t)$. Note that a saver's total savings is increasing in the risk-free rate r_t.

Definition 1 (Equilibrium Without Knightian Uncertainty)

Given (a) An announced choice for the risk-free rate r_t; (b) A realization of the aggregate state ω_t; and (c) A wage rate $W_t = W(\omega_t)$, an additive probability *measure Q_t (with support on the interval $[0, s(W_t, r_t)]$) over the investment levels x of other savers and an aggregate level of investment X_t are a Nash equilibrium if:*

1. *The representative saver is indifferent, under Q_t, among all levels of risky investment, $0 \leq x_t \leq s(\omega_t, r_t)$.*

2. *The level of aggregate investment in the risky asset is given by:*

$$X_t = \int_X x dQ_t(x).$$

   Further, the aggregate quantity of savings, S_t, is given by the representative saver's choice of savings, $S_t = s(\omega_t, r_t)$; and the aggregate quantity of bond holdings, B_t, is given by $S_t - X_t$.

Here, Q_t is the probability distribution over levels of investment in the risky asset that the representative saver ascribes to all other savers existing at time t. If the representative saver is to be indifferent among all levels of investment, x_t, then it must be the case that there is expected rate-of-return equality between assets:

$$\rho(E_Q\{X_t\}) = r_t,$$

where $E_Q\{\cdot\}$ is the mathematical expectations operator under distribution Q. There are many possible equilibria to the investment game in each period, associated with many possible distributions Q, including one in which the saver expects no-one else to invest, and so does not himself. For positive quantities of investment, aggregate investment will be $X^{\star}(r)$, defined as:

$$\rho[X^{\star}(r)] = r. \tag{6}$$

The representative saver born each period has some probability distribution Q_t in mind for the investment behavior of his fellows, where Q_t must only satisfy equation (6). We interpret Q_t as a purely subjective distribution, and determine the level of aggregate investment from equation (6).

As shown by figure 2, when the risk-free rate is $r_t = 1$ and there is no uncertainty, total savings is $s(W_t, 1)$ and equilibrium investment is $X^{\star}(1)$ (we assume that for all realizations of the technology shock ω_t, total savings exceed $X^{\star}(1)$ when $r_t = 1$). The remaining $S_t - X^{\star}(1)$ of aggregate savings is allocated to the riskless asset. Clearly, when labor's product W_t, and hence savings, is low, bond sales will also be low. As we shall see in the next section, this will limit the government's ability to use monetary policy.

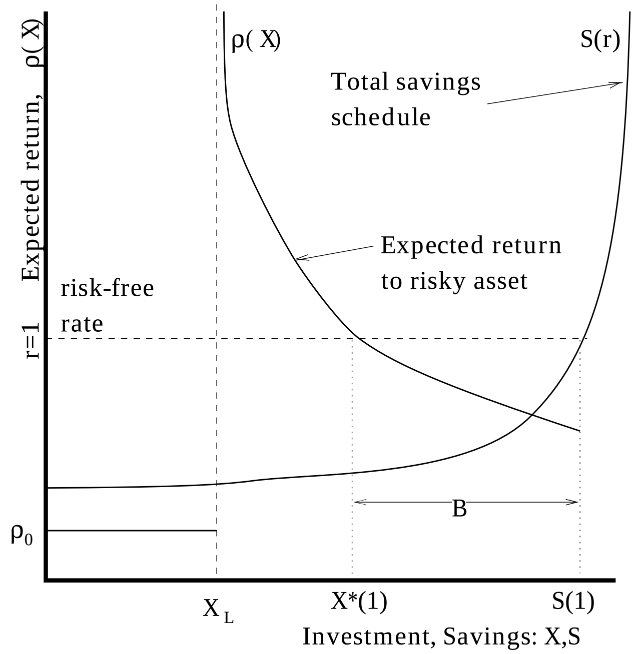

Figure 2: Equilibrium without uncertainty.

3 Analysis With Knightian Uncertainty

3.1 Description of Knightian Uncertainty

Knight (1921) made the original distinction in economics between risk and uncertainty that has since become conventional. For this reason, uncertainty in the sense of multiple prior distributions has become known as Knightian uncertainty. Keynes [see Glahe (1991)] famously argued that uncertainty opened the door for "animal spirits" to affect financial markets. However, the modern economic analysis of choice under randomness, due to Savage (1954), explicitly requires a unique subjective prior distribution. The formal analysis of choice when there are multiple priors begins with Gilboa (1987), Schmeidler (1989) and Gilboa and Schmeidler (1989). The Schmeidler-Gilboa analysis assumes, as does the classic Savage analysis, a series of axioms of choice. To accommodate the presence of multiple priors, they weaken the axiom of independence and add an axiom of uncertainty aversion. With these additions, preferences can be represented using a "maxmin" functional form—agents will act to maximize their payoff under the most pessimistic probability distribution.

The theoretical advances of Schmeidler-Gilboa over the standard Savage economic theory are necessarily complex and beyond the scope of this article. Gilboa and Schmeidler prove that expected utility may be represented by either a single but non-additive prior or a set of additive priors. (Non-additive distributions are subjective probability distributions that do not sum to one, and are also known as *capacities*.) In simple cases, like the ones considered here, the two approaches yield identical results. [See especially Gilboa and Schmeidler (1993) and Wakker (1989) for a description of the connection between the two approaches.] Because information about extreme events is relatively vague, Knightian uncertainty is particularly apposite for modeling such events [see Epstein and Wang (1995) and Dow and Werlang (1992a)].

A classic example of uncertainty is one of appraising the value of a (purported) Ming dynasty vase [see Shafer (1976) and Dempster (1968)]. The vase in question is worth $1000 if it really is from the Ming dynasty, but zero if it is a fake. A risk-neutral saver who is not an expert in Chinese pottery examines the vase and concludes that it is as likely real as fake. A strictly Bayesian approach would then require us to ascribe a subjective probability distribution of {0.5,0.5} to the states {real, fake}. The agent would then buy the vase for any price below $500 and sell it for any price above $500. (Imagine that if the agent sells the vase, he receives the sale price with certainty and then loses the revealed true value of the vase. The agent is "shorting" the vase asset.) This description seems to miss something crucial about how agents, even risk-neutral agents, react to incomplete information.

If the agent's behavior is instead described by the non-additive subjective probability distribution {0.4,0.4} the two events are still equiprobable (so that the agent still believes that the vase is as likely real as fake), but their sum is less than one. An amount 0.2 of the total probability over states has been "lost" to uncertainty. Expected values of random variables under non-additive probabilities may be calculated using the Choquet integral [see Dow and Werlang (1992b) and Gilboa (1987)]. If the agent buys the vase for some price p, and it turns out to be fake, his realized utility is $-p$. If it turns out to be real, the agent's utility is $1000 - p$, an improvement of $1000. Thus his expected utility from buying the vase is $-p + 0.4(1000)$ or $400 - p$. The agent would buy the vase for any price below $400. If instead the agent sells the vase for some price p, and it turns out to be real, his utility is $p - 1000$. If it turns out to be fake, his utility is p, an improvement of \$1000. Thus his expected utility from selling the vase is $p - 1000 + 0.4(1000)$ or $p - 600$. The agent would sell the vase for any price above \$600. At prices between \$400 and \$600, the agent would neither buy nor sell the vase.

More generally, if agents have a subjective probability distribution over events of $P_\epsilon = (1 - \epsilon)Q$, where Q is an additive probability distribution over events in the σ-algebra $\mathfrak{S}$ constructed from the discrete set of states Ω (so that P_ϵ is a "uniform squeeze" of Q), we say that agents have constant uncertainty aversion of degree ϵ. If an agent has a utility in each state of $u(\omega)$, then "expected utility" is formed as:

(7) $$E_{P_\epsilon}\{u(\omega)\} = \epsilon \min_{\omega} u(\omega) + (1 - \epsilon) E_Q\{u(\omega)\}.$$

Thus uncertain savers behave as if they were pessimistic, in the sense that they ascribe all of the missing probability to the worst-case scenario.

In related work, Gilboa and Schmeidler (1989) and Wakker (1989) demonstrate formally that, under certain assumptions, agents with many additive prior distributions behave as if they had a single but non-additive probability distribution. Such agents will also have the maxmin form of the utility function. Assume that an agent has the collection of additive priors $\mathbb{P}_\epsilon$, defined as:

$$\mathbb{P}_\epsilon \equiv \{(1-\epsilon)Q + \epsilon m : m \in \mathcal{M}\}, \tag{8}$$

where the constant, $0 \leq \epsilon \leq 1$ indexes uncertainty, Q is the reference or true additive probability distribution and $\mathcal{M}$ is the set of *all* probability measures defined on the support of Q. Then expected utility is again formed using the maxmin formulation, equation (7), above.

Notice the interesting parallel with certain types of difficult-to-quantify financial risks, e.g. model risk. Savers have many different competing hypotheses about the probability distribution of the returns to some financial asset. These correspond with the many distributions lying in $\mathbb{P}_\epsilon$. Such savers, according to the theory of uncertainty, will act to maximize utility, equation (7). They will, as a result, be disproportionately concerned with the worst-case scenario. The more the competing hypotheses differ, i.e. the more uncertain agents are, the greater the weight on the worst-case.

3.2 Nash Equilibrium With Knightian Uncertainty

In this section, we extend the definition of a Nash equilibrium in the investment game, played by the representative saver against the market, to the case when the saver is uncertain about the market's behavior. In doing so, we use the extension of Nash equilibrium due to Dow and Werlang

(1994). Notice that uncertainty will not be over the aggregate state, which follows the known i.i.d. process $\pi^{\star}$, but rather over the behavior of other savers.

Investor \ Market	L	H
l	0,0	0,a
h	-b,0	a,a

Figure 3: A simplified version of the investment game.

To motivate our definition of a symmetric Nash equilibrium with Knightian uncertainty, we provide the following example. Figure 3 displays a stylized version of the investment game that the representative saver plays against "the market" (all other savers taken together) in which both sides are allowed only two simple actions. The individual saver can play either a high ("h") or a low ("ℓ") level of investment in the risky asset. The market can do the same, playing either "H" or "L". The behavior of the individual (small) saver will not affect the payoff of all other savers taken together. If the market invests the low amount, it realizes the risk-free rate, here normalized to zero, and if it invests the high amount, it realizes some increment $a > 0$ over that. If the individual follows the market, he realizes the same return as the market. However, if the individual invests the high amount when the market invests the low amount, he suffers a loss of $-b < 0$. This corresponds to the case of being the last saver out of a "rush to the exit" model, or a depositor at the back of the line in a bank run model. Although the saver can eliminate L as a strategy of the market (because it is strictly dominated by H), if the market (for whatever reason) did play L, and b was large, then the saver would be exposed to a large loss.

We use the augmented definition of a Nash equilibrium due to Dow and Werlang (1994). Assume that the saver's beliefs about the play of the market can be described by the set of probabilities $\{p_L, p_H\}$. Assume further that the saver ascribes probability zero of the market playing L, but, because of Knightian uncertainty, a probability less than unity of the market playing H, $p_H < 1$. We can form expected utilities as described in equation (7) above, and see that the individual saver invests the high amount if and only if $p_H > b/(a+b)$. Thus if the loss exposure b is small, the individual saver will invest even if p_H is quite small, while if b is large, p_H must be relatively close to unity. This is equivalent to assuming that the saver has an uncertainty parameter of $\epsilon = 1 - p_H$.

With this example in mind, we generalize the definition of a Nash equilibrium in the period-t investment game to include the case when agents have Knightian uncertainty. Because he is uncertain about the behavior of the market, the representative saver will have beliefs about aggregate investment that can be represented by a non-additive probability distribution. As a result, the saver will demand a premium for holding the risky asset. He is guarding against the possibility of a financial Armageddon–when aggregate investment falls below the critical level, X_L.

Definition 2 (Equilibrium With Knightian Uncertainty)

Given (a) An announced choice for the risk-free rate r_t; (b) A realization of the aggregate state ω_t; (c) A wage rate $W_t = W(\omega_t)$; and (d) A realization of the uncertainty parameter ϵ_t, an additive probability measure Q_t (with support on the interval $[0, s(W_t, r_t)]$) over the investment levels x of other savers and an aggregate level of investment X_t are a Nash equilibrium with Knightian uncertainty if:

1. *The representative saver is indifferent, under the non-additive probability measure $P_{\epsilon_t}(x) = (1 - \epsilon_t)Q_t(x)$, among all levels of risky investment, $0 \leq x_t \leq s(\omega_t, r_t)$.*

2. *The level of aggregate investment in the risky asset is given by:*

$$X_t = \int_X x dQ_t(x).$$

   The level of aggregate savings S_t and bond holdings, B_t are given by the representative saver's choices: $S_t = s(\omega_t, r_t)$ and $B_t = S_t - X_t$.

We further assume that the uncertainty parameter ϵ_t is drawn i.i.d. each period from a known distribution on the interval $[0,1]$, and that it is costlessly observed by all agents before they make their consumption, effort and investment decisions in a period. We make no assumptions about the covariance (if any) between the preference shock ϵ_t and the technology shock ω_t.

3.3 Investment Under Knightian Uncertainty

For a saver with uncertainty parameter ϵ_t to be indifferent between all portfolio divisions, the rate of return equality condition (6) must be altered to reflect uncertainty. Thus aggregate investment will be given by $X^\star_{\epsilon_t}(r_t)$, defined implicitly by:

$$(1-\epsilon_t)\int_\Omega \rho[\omega, X^\star_{\epsilon_t}(r_t)]\; d\pi^\star + \epsilon_t \rho_0 = r_t. \tag{9}$$

Because $\rho(X)$ is decreasing in X, the uncertainty-contaminated level of investment in the risky asset in period t, $X^\star_{\epsilon_t}(r_t)$, is below the no-uncertainty level of investment $X^\star(r_t)$, as shown in figure 4. Recall from section 2.2 that the return $\rho(X)$ is finite; as a result, if the uncertainty parameter ϵ_t grows large enough, investment in the risky asset falls to zero because even the project with the highest return, when corrected for the saver's fear of a financial Armageddon, does not yield more than the risk-free rate.

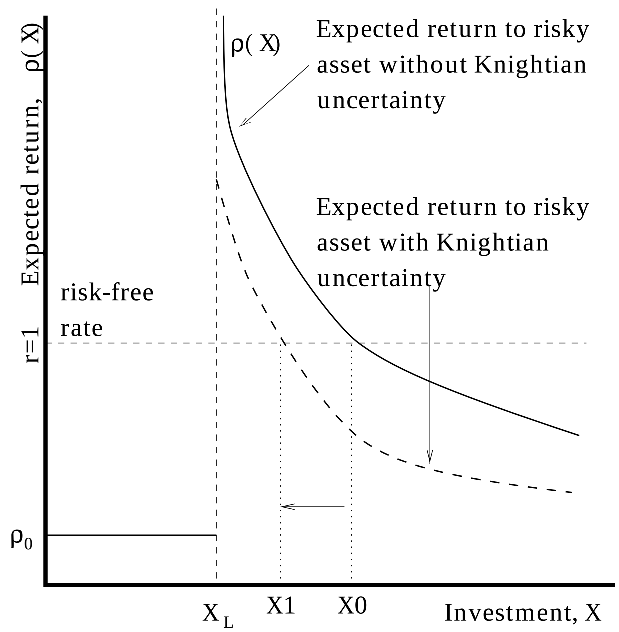

Figure 4: Effect of uncertainty. Notice that aggregate investment in the risky asset falls from X_0 to X_1.

The following proposition formalizes the effect of increases in the uncertainty parameter ϵ_t. It is quite close in spirit to proposition 1 from Liu (1998).

Proposition 1 (Effect of Uncertainty)

If no government action is taken, a generation j born with an uncertainty parameter ϵ_j will invest less than all generations i born with smaller uncertainty parameters, $\epsilon_i < \epsilon_j$. In addition, in times of high uncertainty, spreads between risky assets and the riskless asset will widen.

Proposition 2 (Financial Armageddon)

In each period, given an announced risk-free rate r_t, if $r_t > \rho_0$ there is some level of uncertainty $\epsilon^\star < 1$ such that, if $\epsilon_t \geq \epsilon^\star$, no equilibrium with positive investment exists.

4 Optimal Monetary and Fiscal Policy

In periods of certainty (when generations are born with uncertainty parameters $\epsilon_t = 0$) there is no need for government intervention. In such periods, the government sets the interest rate on bonds to the storage rate $r_t = 1$ and does not levy any taxes on the young. In periods of uncertainty (when $\epsilon_t > 0$) the government may choose a mix of fiscal and monetary policies to undo the distortions–decreased investment and possibly a complete financial Armageddon–described in section 3.3 above. We take as monetary policy a choice for the return on government bonds, r_t; and as fiscal policy a choice of distortionary tax policy, $\{\tau_c^t, \tau_\ell^t\}$. An expansionary monetary policy is one which pushes the gross return on government bonds below its natural level of one (the storage return).

The government will take as its problem that of maximizing a social welfare function formed by the equally-weighted sum of the expected utility of the representative saver plus the producers' surplus. Notice that because savers and producers are risk-neutral with respect to consumption while old, an increase in this social welfare function means that the government, by using a lump-sum tax and transfer policy in the second period of life, could make both producers and savers (at least weakly) better off. Thus we identify increases in the social welfare function with potential Pareto improvements, with the caveat that the government may have to lump-sum transfer resources from one class of agents to another.

Governments often use monetary policy to respond to (at least the initial stages of) a financial crisis, perhaps because fiscal policy is costly to change. In this model, because government bonds are backed by a riskless storage technology, the optimal monetary policy when there is no uncertainty is to set the gross return on bonds to the technologically-determined return of unity. When we introduce Knightian uncertainty, a solely monetary policy (that is, a monetary policy without an associated fiscal policy) will have two countervailing effects. First, it causes agents to work less and consume more while young, pushing down total savings. Second, it will, by the portfolio balance equation (9), cause them to hold more of the risky asset, which has a true expected rate of return greater than the risk-free rate. The second, good, effect will dominate the first, bad, effect for at least small decreases in the risk-free rate, if uncertainty is not too large.

The effect of a monetary expansion will depend on the current realization level of labor productivity $W(\omega_t)$, and the uncertainty level ϵ_t. Consider the case, displayed in figure 5, of "good times," when labor productivity, W_t, is high. In such periods, the savings schedule $s(W_t, r_t)$ lies relatively far to the right–for any given interest rate, savers will save in total more when W_t is high. If the saver had no uncertainty, then the equilibrium would be the optimum presented in section 2.4 above. At the point marked "C" in the figure, the saver saves a total amount $s(W_t, 1)$, of which an amount $X^{\star}(1)$, marked "A," is invested in the risky asset. The difference, the interval marked "B," is devoted to bonds. The uncertainty realization ϵ_t shifts the apparent (to the saver) return on the risky asset down, from the solid line to the dashed line. If the risk-free rate is still held at one, then the saver's portfolio now contains an amount $X_a < X^{\star}(1)$ of the risky asset, where the dashed line and the solid line intersect at the point marked "D." The government (in this figure) responds with a large monetary expansion, forcing down the risk-free rate from its initial level of one (the solid line) to $r^{\star} < 1$ (the dashed line). Now the saver's portfolio contains

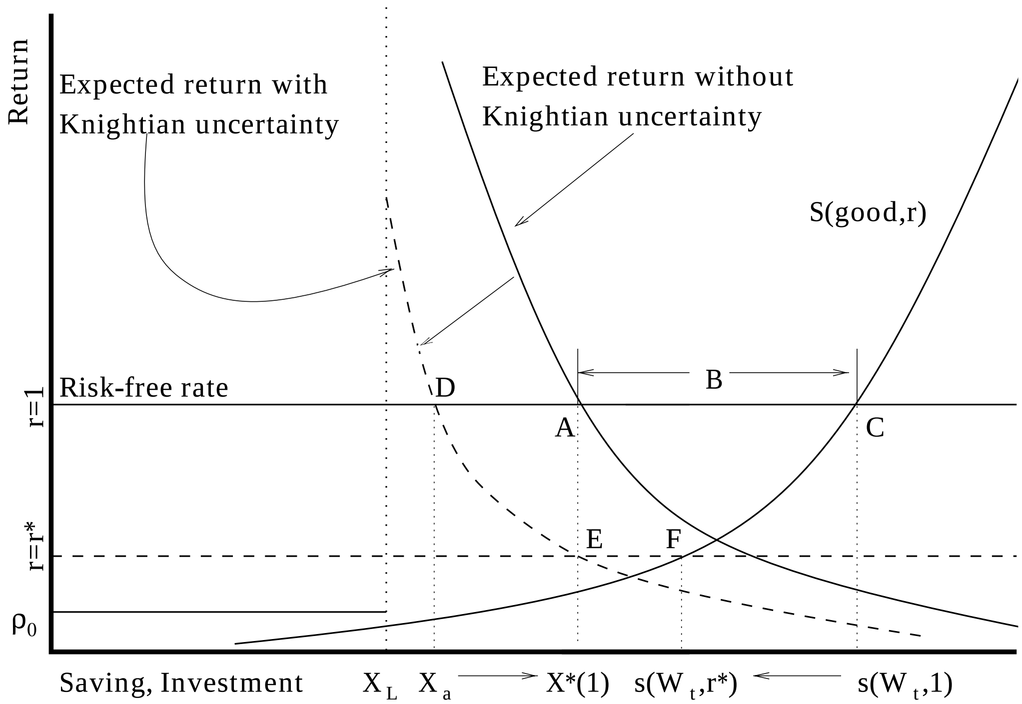

Figure 5: Effect of monetary policy in relatively good times. Investment in the risky asset returns to its optimal level, but total savings falls.

just the right amount, $X^{\star}(1)$, of the risky asset again, at the point marked "E." However, the amount of total savings fall to $s(W_t, r^{\star}) < s(W_t, 1)$, at the point marked "F."

In contrast, when labor's productivity is low, i.e. in "bad times," as shown in figure 6, the scope for monetary expansions is limited. If there is no uncertainty, aggregate investment in the risky asset is still $X^{\star}(1)$ (at the point marked "A") and aggregate bond holdings are still given by the interval marked "B." Because the savings schedule is shifted quite far to the left, the demand for government bonds when there is no uncertainty is quite small. We assumed, in section 3.2 above, that without uncertainty, there is always a positive demand for bonds at $r_t = 1$.

When uncertainty is present in addition to lower labor productivity, the apparent expected return to the risky asset is given by the dashed line in figure 6. As before, when the government decreases r_t total savings fall and the portfolio level of risky assets rises, so the demand for bonds falls. When the interest rate is $\underline{r}$, the demand for bonds falls to zero (at the point marked "E") and hence further decreases in r_t do not fuel further increases in holdings of the risky asset. The maximum level of investment in the risky asset that can be generated from monetary policy alone is $X_c < X^\star(1)$.

Now consider the effect of the uncertainty parameter ϵ_t on the effectiveness of monetary policy. As uncertainty increases, the apparent (to the saver) rate of return on risky assets falls. If uncertainty is high enough, the highest possible apparent return to the risky asset (when aggregate investment is X_L) would not stimulate enough savings to cover the minimum investment level X_L. Here monetary policy is completely ineffective and the saver will hold a portfolio made up entirely of government bonds. As shown on figure 7, if the uncertainty-adjusted apparent rate of return schedule on risky assets (the dashed line) at X_L falls below the point marked S_a, then savers would never be willing to save enough to

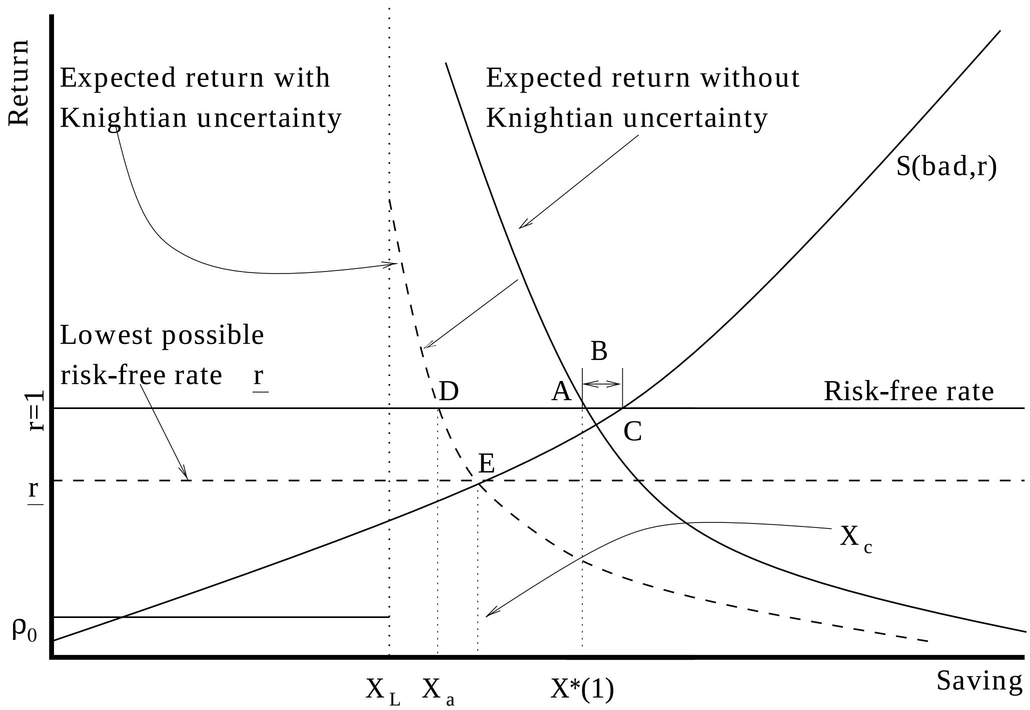

Figure 6: Effect of monetary policy in bad times. Note the lower bound on the effective risk-free rate and the (suboptimally low) upper bound on aggregate investment in the risky asset.

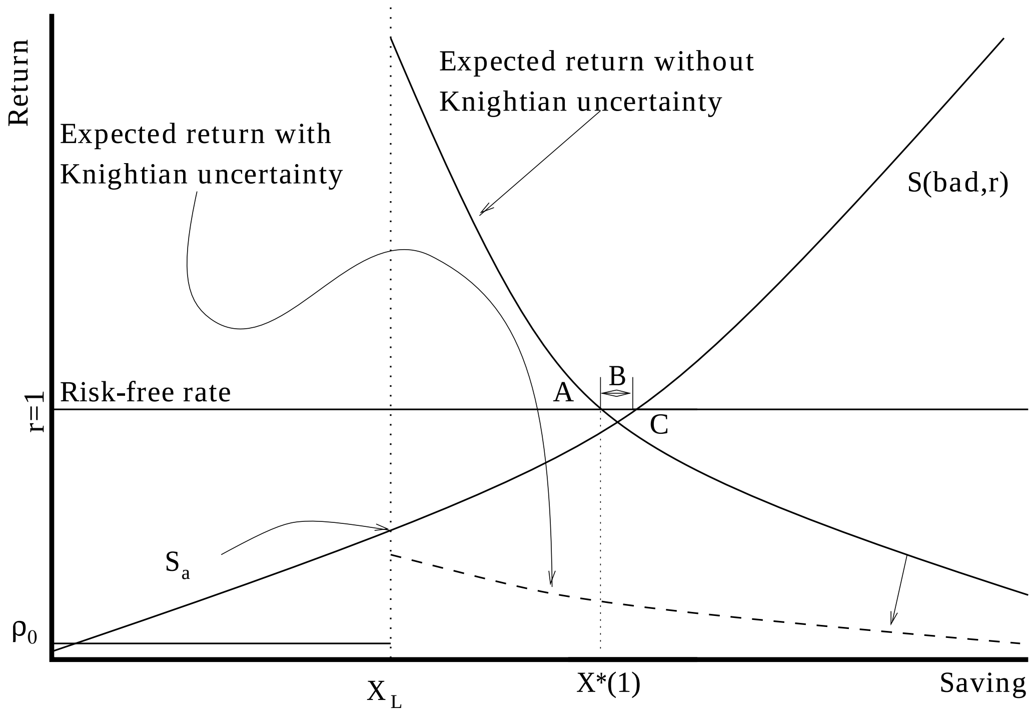

Figure 7: Effect of very high uncertainty. No monetary policy can stimulate investment; without fiscal policy, portfolios are made up entirely of the safe asset.

cover the critical investment level X_L. In contrast, in figure 5, because the savings schedule is shifted quite far to the right, even at very low risk-free rates the saver is willing save enough to cover the critical level X_L. Thus there is further interaction between the technology shock and the effectiveness of monetary policy: When labor's product is low, there may be some level of uncertainty so large that no monetary policy can avoid a financial Armageddon.

We now formalize these ideas.

Proposition 3 (Pareto Improving Role of Monetary Expansions)

In times of moderate uncertainty, when $0 < \epsilon_t \leq \epsilon_t^{\star\star}$*, where* $\epsilon^{\star\star}$ *is defined as:*

$$s[W_t, (1 - \epsilon_t^{\star\star})\rho(X_L) + \epsilon_t^{\star\star}\rho_0] = X_L,$$

decreasing the risk-free rate r_t *from* $r_t = 1$ *increases the social welfare function. Monetary expansions alone cannot restore the optimal allocation. If there is no uncertainty the optimal risk-free rate is one.*

Proposition 4 (Limits to Monetary Expansions)

There is some minimum effective risk-free rate, $\underline{r}_t(W_t)$*, available to the government. The level of aggregate risky investment at this rate may be below the optimum. The minimum rate* $\underline{r}_t(W_t)$ *is higher when labor's productivity,* W_t*, is low.*

To fix ideas, consider figure 8 below, which shows the relationship of $\epsilon_t^{\star\star}$ to labor productivity, W_t. When labor is productive, it is more likely that at least a small decrease in the risk-free rate will have an effect. However, even within the region marked as effective, as an economy approaches the border with ineffectiveness (from a combination of high uncertainty and low labor productivity), the minimum possible level for the risk-free rate, $\underline{r}_t$, will be increasing.

So far we have imagined that the government has conducted monetary policy by choosing a risk-free rate and accommodating the resulting demand for bonds. In this view, the demand for bonds is irrelevant, as long as it is positive. However, we might imagine that the government's ability to vary bond sales was limited either because the government has some minimum financing needs, or because its monetary policy is simply to sell the same number of bonds each period, regardless of the level of labor's product or uncertainty. Figure 9 below plots the demand schedule for bonds under different realizations of labor's product $W(\omega_t)$ and uncertainty ϵ_t. The two solid schedules correspond to different realizations of labor's product $W(\omega_t)$ when uncertainty is zero, and the dashed schedule shows the effect of uncertainty. For the reasons discussed above, bond demand drops to zero at $\underline{r}_t(W_t)$. This level is higher in bad times and lower when uncertainty is high. Demand for bonds rises smoothly with the rate they pay, until the critical level at which bonds pay a higher return than the largest possible return, correcting for Knightian uncertainty, that the risky asset can pay. If bonds pay a return greater than this critical level, demand for bonds jumps up as savers adjust their portfolios to contain only bonds. Bonds have completely crowded out investment in the risky asset. If uncertainty exceeds $\epsilon_t^{\star\star}$ as defined in proposition 3 above, then this critical level is at or below $\underline{r}_t$, so no equilibrium with positive investment is possible.

If the government followed a policy of always selling (if possible) some amount $B^{\star} > 0$ of bonds, and adjusting the risk-free rate to accommodate this, then the return would vary depending on labor's product, $W(\omega_t)$ and uncertainty, ϵ_t. In bad times, when labor's product was low, the equilibrium risk-free rate would be higher than in good times, when labor's product was high, even if there was no uncertainty. In times of increased uncertainty, as the demand schedule for bonds shifts out, the equilibrium risk-free rate would be lower. This immediately raises the possibility that, if the government adopted a constant bond-sale policy, monetary policy would be self-stabilizing, with decreases in the risk-free rate in times of high uncertainty. However, this turns out not to be the case:

Corollary 1

There is no level of bonds, $B^\star \geq 0$, such that, if the government maintained bond sales at $B^\star$ for all realizations of ω_t and ϵ_t, altering bond sales would not produce a Pareto improvement.

This is a simple application of propositions 3 and 4. First, notice that if the government pursues a constant bond-sale policy, the risk-free rate will vary in response to shocks to labor's product $W(\omega_t)$. We know from proposition 3 that the optimal risk-free rate is the same for all levels of $W(\omega_t)$. Second, from proposition 4 we know that in times of high uncertainty and a realization of labor's product, $W(\omega_t)$, the optimal monetary policy will be to sell zero bonds. If $B^\star > 0$, then this condition can never be met; while if $B^\star = 0$, then too few bonds will be sold in periods of low uncertainty.

Next, we describe how a fiscal policy can encourage savings by discouraging current consumption and leisure, thus allowing the government to recapture the optimum. By increasing the tax on current consumption and decreasing the labor income tax (that is, actually subsidizing labor effort), the optimal tax system rewards saving. As a result, for any prevailing risk-free interest rate, savers facing such a tax system will save more than savers who face no distortionary taxes. The government can then depress the risk-free rate, recovering the optimal level of risky investment through the portfolio-balance effect, without causing an over-all decrease in savings. The additional effect of a fiscal policy is illustrated in figure 10 below. A combined monetary and fiscal policy move the aggregate investment level to A' and total savings to C'. Note that these are the same as the optimal levels, at points A and C. Such a policy is always feasible, because we assumed that $S(\omega_t, 1) > X^\star(1)$ for all ω_t.

Proposition 5 (Optimal Fiscal and Monetary Policy)

In times of uncertainty, when $\epsilon_t > 0$, if the government sets the risk-free rate to $r^\star(\epsilon_t) < 1$ satisfying $X^\star_{\epsilon_t}(r^\star_t) = X^\star(1)$; the tax rates on consumption while young and labor income to: $\tau^\star_c = -1 + 1/r^\star$; and $\tau^\star_\ell = 1 - 1/r^\star$; and the lump-sum transfers to satisfy the government's budget constraint, then the representative saver will consume, work, save and invest exactly as if there were no uncertainty. In particular, risky investment in period t will be at its optimal level: $X_t = X^\star(1)$.

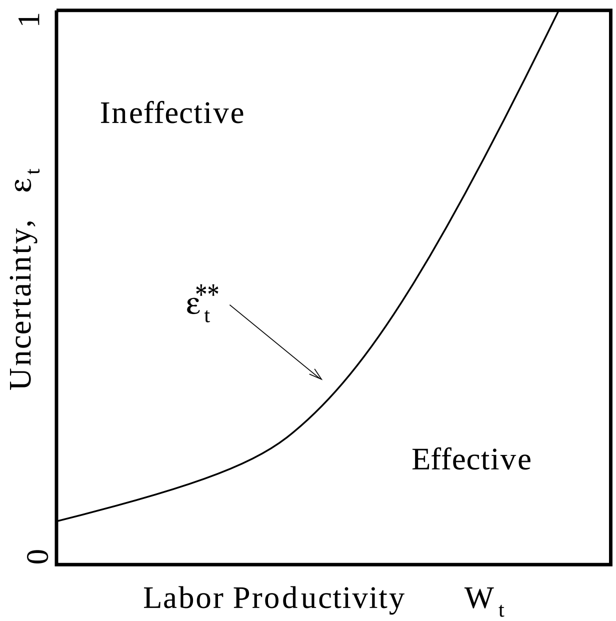

Figure 8: Region in which monetary policy is effective.

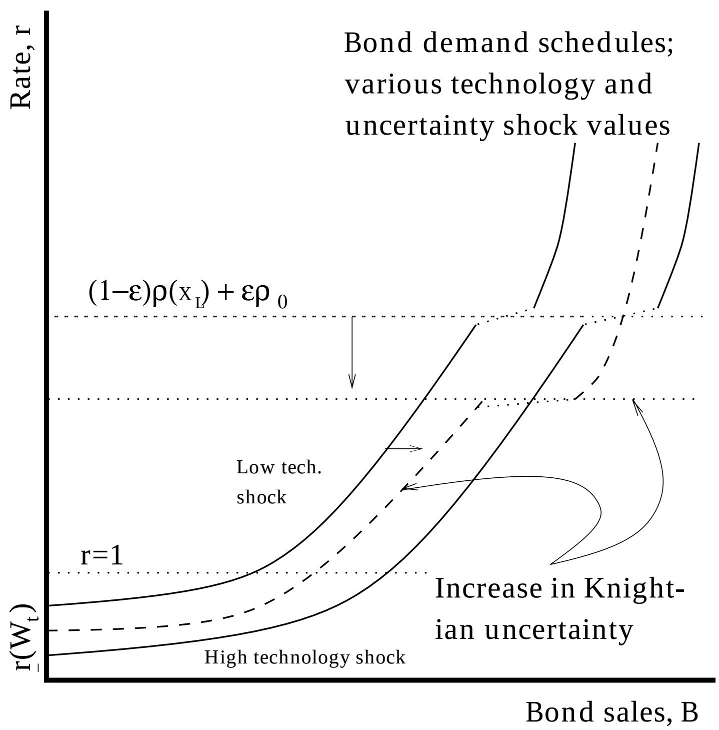

Figure 9: Bond demand as a function of technology shock and Knightian uncertainty. The dashed line indicates an increase in Knightian uncertainty.

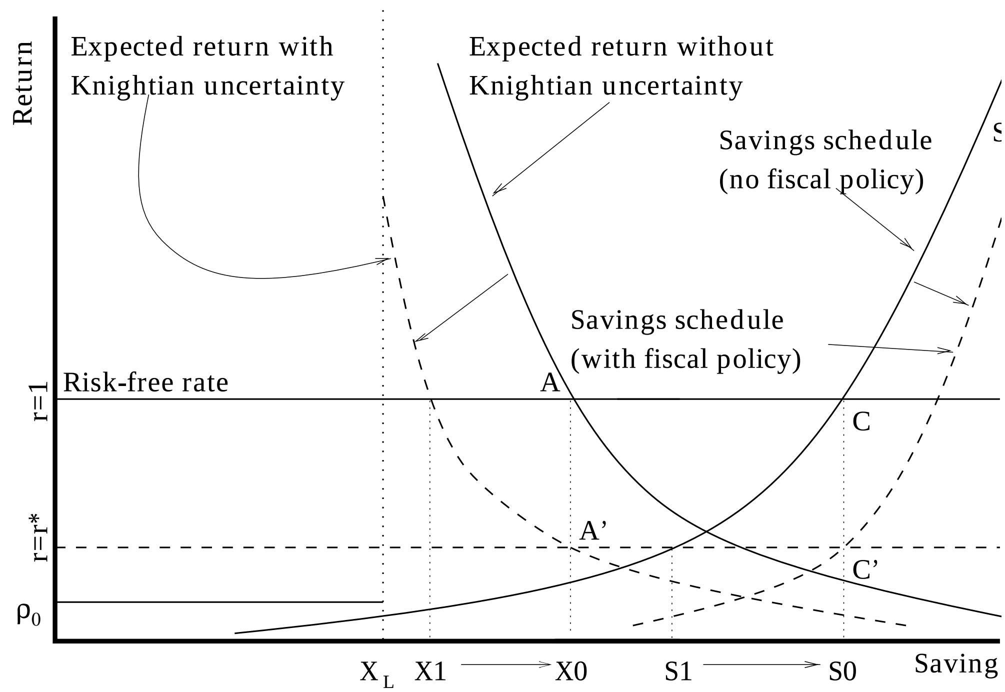

Figure 10: Effect of fiscal and monetary policy in good or bad times. The appropriate fiscal policy manipulates aggregate savings so that monetary policy is always effective.

5 Conclusion

In this paper we used the relatively new theory of choice under Knightian uncertainty to study the equilibrium effects of saver uncertainty and the optimal government policy response. Because uncertain savers behave like pessimists, they will underinvest in the risky asset, starving the economy of productive capital, and overinvest in the risk-free asset, the government bond. The optimal monetary policy response to uncertainty is one that decreases the rate of return to bonds, inducing savers to hold risky assets. A strategy of combating uncertainty with a purely monetary policy is shown to be Pareto improving, but to be unable to recapture the optimal allocation. To recapture the optimum, monetary policy must be combined with a fiscal policy of taxing current consumption and subsidizing current labor effort, so that total savings does not fall.

This analysis allows us to draw several conclusions about the market turmoil of the late summer and fall of 1998. First, it implies that portfolio adjustments may have been the primary reason that a monetary expansion undid some of the ill-effects of the financial turmoil. Second, it implies that interest rate spreads between riskless and risky assets are in large part determined by the level of uncertainty in the economy. As a result, spreads remained elevated even after the monetary expansion, indicating that savers continued to be uncertain. Third, our analysis implies that it is easier for the government to counteract uncertainty-driven financial crises in good times than in bad times. The fact that the U.S. economy was particularly healthy in the fall of 1998 allowed it to use monetary policy alone to calm financial markets. Finally, the model provides an explanation for the underlying cause of the sudden increase in uncertainty. If savers perceived that other savers were withdrawing from risky investments, they would have had no incentive to maintain their own risky investments.

More generally, our analysis points to rules for conducting monetary policy in the face of shocks to labor income and uncertainty. Roughly speaking there are three levels of the uncertainty parameter: zero (or very low), moderate, and extremely high. In the same way, we can consider the two extremes of labor's product: low and high. If uncertainty is close to zero, then there is no role for monetary policy in our model (proposition 3). This result derives fundamentally from the fact that financial markets are assumed to be perfect; that is, they are zero-cost, fully transparent, and free of any underlying moral hazard problems. When there is moderate uncertainty, spreads between risky and riskless asset will widen, and there will be a general flight to safety (proposition 1). An expansionary monetary policy will be Pareto improving, but will be limited by the current level of labor's product (propositions 3 and 4). If labor's product is high, demand for bonds will be relatively strong, and the government will be able to push the risk-free rate fairly low. If labor's product is low, demand for bonds will be relatively weak, and the government will not be able to push the risk-free rate very far away from its natural level. The monetary expansion undoes, to a certain extent, the flight to quality. Even after the decrease in the risk-free rate, spreads will remain elevated, although returns to the risky asset will fall. Thus although spreads are useful in signalling increases in uncertainty, monetary policy should not target them directly, instead targeting the returns on risky assets.

Finally, if uncertainty is extremely high, investors may perceive that aggregate investment will fall below the critical level, and thus choose to prudently invest nothing in the risky asset, with disastrous consequences for output (proposition 2). Further, at some levels of uncertainty, monetary policy may be ineffective, and so monetary policy alone will be unable to stem a catastrophic decline in investment (proposition 4). In such cases, only a combination of monetary and fiscal policies will be able to avert this financial Armageddon (proposition 5).

Appendix

Proof of Proposition 1

Given that the government is not conditioning variables on the signal, we must show that investment in the risky asset is decreasing in ϵ for all values of the risk-free rate $r > 0$ and realizations of the production shock, ω:

$$X^\star_{\epsilon_j}(r) \geq X^\star_{\epsilon_i}(r), \text{ all } r > 0, \epsilon_j > \epsilon_i.$$

Here $X^\star_\epsilon(r)$ is determined by the portfolio balance equation (9). This equation implies that the expected return to the risky asset is increasing in the uncertainty parameter, ϵ, as savers demand an uncertainty premium for holding the risky asset. Because the production function $\rho(\omega, X)$ is decreasing in aggregate investment X, it must be the case that equilibrium investment in the risky asset declines as uncertainty increases. The effect of the production shock ω is to shift the total savings curve $s(\omega, r)$. For any given level of total savings, higher values of ϵ are associated with a decreased portfolio holding of the risky asset.

To think about the spread between a risky and the riskless asset, consider an asset that pays a state-contingent return of $R(\omega)$. Assume for convenience's sake that the return can take on only two values: $R_1 > 0$, when aggregate investment lies above X_L, and zero, when aggregate investment falls below X_L (imagine a AAA-rated corporate bond that repays in all states of the world, and defaults only if there is a financial "Armageddon"). Uncertainty-contaminated rate-of-return equality then requires that the expected return to the bond, formed under the non-additive probability measure P_ϵ, must satisfy:

$$(1-\epsilon)R_1 = r,$$

where r is the prevailing risk-free rate. Thus for risk-neutral (but uncertainty-averse) savers to hold positive quantities of this asset, its return must satisfy:

$$R_1 \geq \frac{1}{1-\epsilon} r.$$

The spread is then:

(A.1.1) $$R_1 - r = r\frac{\epsilon}{1-\epsilon}.$$

This is increasing in ϵ. ■

Proof of Proposition 2

The highest possible expected return to the risky asset occurs when aggregate investment just equals X_L. We can then define $\epsilon^\star$ from the portfolio balance equation (9) as that level of uncertainty at which, even if savers expect the risky technology to pay off at its highest possible level (that is, if aggregate investment is expected to be just X_L), they are just indifferent between the risky and the riskless asset:

$$(1-\epsilon^\star)\rho(X_L) + \epsilon^\star\rho_0 = r.$$

Or, manipulating:

$$\epsilon^\star = \frac{\rho(X_L) - r}{\rho(X_L) - \rho_0}.$$

Notice the importance of the assumption of $r \geq \rho_0$. In the next proposition, we discuss why the risk-free rate might be bounded from below. ■

Proof of Proposition 3

In this section we consider the problem of a benevolent government constrained to combat uncertainty with a purely monetary policy. We construct a social welfare function and show that, in times of uncertainty, it is increasing in an expansionary monetary policy (decreasing in the risk-free rate). Because increases in the risk-free rate can make savers better off by capturing some monopoly rents from the producers, we have to consider both the welfare of savers and producers. The government will have available to it a lump-sum tax on (transfer to) producers which it uses as a transfer to (tax on) savers.

Note first that monetary policy will never influence portfolio decisions if uncertainty is greater than a critical level $\epsilon_t^{\star\star}$, defined implicitly as:

$$s[W_t, (1-\epsilon_t^{\star\star})\rho(X_L) + \epsilon_t^{\star\star}\rho_0] = X_L.$$

If uncertainty ϵ_t is greater $\epsilon_t^{\star\star}$, the Knightian-uncertainty contaminated expected rate of return schedule never intersects the savings schedule at any level of savings greater than the critical level, X_L. Note that $\epsilon_t^{\star\star}$ is increasing in W_t. If W_t is large enough, there may be no $\epsilon_t^{\star\star}$ less than unity, in which case monetary policy is effective no matter how uncertain savers are. Thus when the production shock, ω_t, takes on a bad value monetary policy is more likely to be ineffective.

Now we turn to the question of the effect of monetary policy when uncertainty is not too great: $\epsilon_t < \epsilon_t^{\star\star}$.

From equation (4) above, if the prevailing risk-free rate is r_t and savers invest an amount $X_{\epsilon_t}^{\star}(r_t)$ in the productive (risky) technology, the representative producer has an expected surplus of:

$$\text{PS}_t(r_t) = \int_0^{X_{\epsilon_t}^{\star}(r_t)} \{\rho(n) - \rho[X_{\epsilon_t}^{\star}(r_t)]\} \; dn.$$

Note that the price paid for loans is not r_t, but rather an uncertainty-premium over r_t. Recall from equation (9) that the uncertainty adjusted expected rate of return on the risky asset satisfies, in period t:

$$\rho[X_{\epsilon_t}^{\star}(r_t)] = \frac{r_t - \epsilon_t \rho_0}{1 - \epsilon_t}.$$

Substituting back into the expression for $\text{PS}_t(r_t)$ produces:

$$\text{(A.3.1)} \qquad \text{PS}_t(r_t) = \int_0^{X_{\epsilon_t}^{\star}(r_t)} \rho(n) \; dn - \frac{r_t - \epsilon_t \rho_0}{1 - \epsilon_t} X_{\epsilon_t}^{\star}(r_t).$$

Appealing to Leibnitz's rule, the derivative is:

$$\text{PS}_t'(r_t) = \left\{ \rho[X_{\epsilon_t}^{\star}(r_t)] - \frac{r_t - \epsilon_t \rho_0}{1 - \epsilon_t} \right\} \frac{\partial X_{\epsilon_t}^{\star}(r_t)}{\partial r} - \frac{1}{1 - \epsilon_t} X_{\epsilon_t}^{\star}(r_t).$$

Again substituting in for the uncertainty premium from equation (9) produces:

$$\text{(A.3.2)} \qquad \text{PS}_t'(r_t) = -\frac{1}{1 - \epsilon_t} X_{\epsilon_t}^{\star}(r_t).$$

To compute the saver's indirect utility, it will be convenient to abstract from his choice of consumption and leisure while young, c_0^t and ℓ_0^t. Consider the partial value function $\phi(W_t - s_t)$:

$$\phi(W_t - s_t) \equiv \max_{c_0^t, \ell_0^t} \; u(c_0^t, \ell_0^t) \text{ subject to: } c_0^t + W_t \ell_0^t - H_0^t \leq W_t - s_t.$$

If the saver faces a risk-free rate of r_t, his problem may be expressed as:

$$\max_{s_t} \; \phi(W_t - s_t) + r_t s_t.$$

As a direct consequence, notice that $\phi'(\cdot) = r$. This problem induces a savings relation $s(\omega_t, r_t)$ in the usual way. Given that there is a continuum of savers of mass unity, aggregate savings are $S_t = s(\omega_t, r_t)$.

Begin by considering the representative saver's expected utility, formed using maxmin preferences:

$$V_t^{\text{KU}}(r_t) = \phi[W_t - S_t(r_t)] + r_t S_t(r_t) + H_1^t.$$

The agent expects a rate of return of r_t on all parts of the total savings portfolio, including holdings of the risky asset (although he will earn, in reality, a higher return on the risky asset). Assume that the government refunds (taxes) lump-sum any seigniorage revenue (cost) derived from manipulating the rate paid on storage, r_t:

$$H_1^t = (1 - r_t)[S_t - X_{\epsilon_t}^{\star}(r_t)].$$

Now the saver's indirect utility function becomes:

$$V_t^{\text{KU}}(r_t) = \phi[W_t - S_t(r_t)] + S_t - (1 - r_t) X_{\epsilon_t}^{\star}(r_t). \tag{A.3.3}$$

The slope with respect to r_t is:

$$\frac{dV_t^{\text{KU}}(r_t)}{dr} = S_t'(\cdot)[1 - \phi'(\cdot)] - (1 - r_t)\frac{\partial X_{\epsilon_t}^{\star}(r_t)}{\partial r} + X_{\epsilon_t}^{\star}(r_t). \tag{A.3.4}$$

Notice immediately that when the paid return on government bonds equals the technological rate of return unity, this slope becomes:

$$\left.\frac{dV_t^{\text{KU}}(r_t)}{dr}\right|_{r_t = 1} = X_{\epsilon_t}^{\star}(r_t).$$

This term is positive because, if the risk-free rate increases, the saver captures some surplus from producers.

So far we have concentrated on the representative saver's expected utility under Knightian uncertainty. We may also be interested in the true expected utility of the saver, thus explicitly recognizing that the risky asset will pay off more than r_t in expected value. The representative saver's indirect utility function is now:

$$V_t^{\text{true}}(r_t) = \phi[W_t - S_t(r_t)] + S_t(r_t) - X^\star_{\epsilon_t}(r_t) + \rho[X^\star_{\epsilon_t}(r_t)]X^\star_{\epsilon_t}(r_t).$$

The representative saver divides his total savings S_t into a portfolio of the risky asset, $X^\star_{\epsilon_t}(r_t)$, and the safe asset, $S_t - X^\star_{\epsilon_t}(r_t)$. The saver will earn a certain return of unity on savings placed in the safe asset (because any difference between the prevailing risk-free rate and unity will be refunded lump-sum in H_1^t). However, the saver will earn a *true* expected return of $\rho[X^\star_{\epsilon_t}(r_t)] > 1$ on investments in the risky asset. Thus by decreasing the risk-free rate, the government can stimulate greater investment in the high-return risky asset that is shunned by pessimistic savers.

Substituting out for the uncertainty premium, the representative saver's true indirect utility over r_t becomes:

$$V_t^{\text{true}}(r_t) = \phi[W_t - S_t(r)] + S_t(r_t) - X^\star_{\epsilon_t}(r_t) + \frac{r_t - \epsilon_t\rho_0}{1-\epsilon_t}X^\star_{\epsilon_t}(r_t). \tag{A.3.5}$$

Taking the derivative with respect to r produces:

$$\begin{aligned} V_t^{\text{true}\prime}(r_t) = &-\phi'[W_t - S_t(r_t)]S_t'(r_t) + S_t'(r_t) \\ &- \frac{\partial X^\star_{\epsilon_t}(r_t)}{\partial r}\left[\frac{1-\epsilon_t - r_t + \epsilon_t\rho_0}{1-\epsilon_t}\right] + \frac{1}{1-\epsilon_t}X^\star_{\epsilon_t}(r_t). \end{aligned} \tag{A.3.6}$$

The final term again represents the monopoly rents captured by savers when the risk-free rate increases. It will be exactly offset by a corresponding decrease in the surplus of the producers.

The government (in its role as a social planner) chooses a purely monetary policy, that is, a level of the risk-free interest rate, to solve the social welfare problem. Here we assume that producers and savers are assigned equal Pareto weights, so the social welfare function becomes:

$$\max_{r_t}\ \mathcal{F}_t(r_t) = V_t^{\text{KU}}(r_t) + \text{PS}_t(r_t),\ \text{subject to: } X^\star_{\epsilon_t}(r_t) \le S_t(r_t).$$

The value functions PS_t and V_t^{KU} are given by equations (A.3.1) and (A.3.4), respectively. By assumption, the constraint that $S_t \ge X^\star_{\epsilon_t}$ does not bind when $r_t = 1$, for any level of the uncertainty parameter $\epsilon_t \geq 0$.

Using equations (A.3.2) and (A.3.4), the slope of the social welfare function is:

$$\frac{d\mathcal{F}_t(r_t)}{dr} = S_t'(r_t)[1 - \phi_t'(\cdot)] - (1 - r_t)\frac{\partial X^\star_{\epsilon_t}(r_t)}{\partial r} - \frac{\epsilon_t}{1 - \epsilon_t} X^\star_{\epsilon_t}(r_t).$$

Because $\phi_t'(\cdot) = r_t$, when $\epsilon_t = 0$ the solution to the social welfare problem is to set $r_t = 1$:

$$\left.\frac{d\mathcal{F}_t(r_t)}{dr}\right|_{r_t=1,\epsilon_t=0} = 0.$$

If $\epsilon_t > 0$ and $r_t = 1$ then the slope of the social welfare function is negative:

$$\left.\frac{d\mathcal{F}_t(r_t)}{dr}\right|_{r_t=1,\epsilon_t>0} = \frac{\epsilon_t}{1 - \epsilon_t}\frac{\partial X^\star_{\epsilon_t}(r_t)}{\partial r}.$$

Thus, even using the Knightian uncertainty contaminated expected utility of the representative saver, a decrease in the risk-free rate increases the total surplus available to be divided between producers and savers. The reason is that Knightian uncertainty produces a wedge between what the producer must pay for a loan and what the saver expects to realize on it. A decrease in the risk-free allows more production (worthy projects are funded), and thus an increase in producer's surplus greater than the decrease in saver's surplus. With appropriate transfers between the two parties, savers can be made better off without hurting producers. In a democracy, all parties would vote for such a system.

This implies that an expansionary monetary policy, that is, a policy of depressing the risk-free rate below its natural rate of unity, increases the social welfare function, and is thus potentially Pareto improving, when $\epsilon_t > 0$. However, as the risk-free rate decreases, the constraint that investment in the risky asset not exceed total savings will begin to bind, so that we cannot derive exactly the optimal pure monetary policy. Further, as we shall see, this constraint is more likely to bind in bad times (when labor's product W_t is low) than in good times. With the addition of the fiscal policy instruments of distortionary taxes, the government can completely manipulate the savings schedule and restore the optimum.

Finally, notice that the true expected consumption of a saver while old is:

$$E\{c_1^t(r_t)\} = S_t(r_t) - X^\star_{\epsilon_t}(r_t) + \frac{r_t - \epsilon_t\rho_0}{1 - \epsilon_t} X^\star_{\epsilon_t}(r_t).$$

Taking the derivative with respect to r_t:

$$\frac{\partial E\{c_1^t(r_t)\}}{\partial r_t} = S_t'(r_t) + \frac{1}{1-\epsilon_t}X^\star_{\epsilon_t}(r_t) + \frac{\epsilon_t}{1-\epsilon_t}(1-\rho_0)\frac{\partial X^\star_{\epsilon_t}(r_t)}{\partial r}.$$

This may be either positive or negative, depending on the sensitivity of the portfolio level of investment to the prevailing interest rate, which in turn depends on the sensitivity of the production function $\rho(\cdot, X)$ to investment.

Proof of Proposition 4

Monetary policy faces another limit beyond that imposed by the maximum level of uncertainty which it can combat, $\epsilon_t^{\star\star}$. Even if uncertainty is below this critical level, there is a lower bound on the risk-free rate, $\underline{r}_t(W_t)$. When the risk-free rate reaches this bound, demand for bonds will drop to zero and further changes in r_t will not affect portfolio or savings decisions.

This critical return $\underline{r}_t(W_t)$ is defined implicitly from:

$$s[W_t, \underline{r}_t(W_t)] = X^\star_{\epsilon_t}[\underline{r}_t(W_t)].$$

If $r_t < \underline{r}_t(W_t)$ then savings S_t fall below investment in the risky asset, $X^\star_{\epsilon_t}(r_t)$. In such a situation, the non-negativity constraint on the storage technology is binding: Savers would like to bring forward assets from the future to finance investment in the present. In section 2 we saw that in this situation savers would invest their portfolio entirely in the risky asset, so that further changes in r_t would not affect portfolios. Notice that because the savings schedule $s(W_t, r_t)$ is increasing in W_t, that $\underline{r}_t$ is therefore also increasing in W_t. In bad times, the government will find that monetary policy is less effective. ■

Proof of Proposition 5

Assuming that the saver's savings schedule could be manipulated to any desired level (see below), the right choice of risk-free rate is the one that induces a portfolio with the optimal quantity of risky investment:

$$r^\star : \; X^\star_{\epsilon_t}(r^\star) = X^\star(1).$$

Even if $r^\star$ is below $\underline{r}_t$, it will still be achievable, because the government will use fiscal policies to manipulate the savings schedule. Further, by our assumption that $S(\omega_t, 1)$ is always greater than $X^\star(1)$, this will be feasible.

Take as given for the moment the saver's choice of portfolio conditional on the risk-free rate $r^\star$, and consider his optimization problem, that of maximizing (1) subject to the budget constraints (2) and (3), by choice of consumption while young, leisure and total savings. The first-order conditions from this problem are:

$$u_c = r_t(1+\tau_c^t), \text{ and:}$$
$$u_\ell = r_t W_t(1-\tau_\ell^t).$$

Thus a decrease in r_t will increase consumption and leisure while young. However, this decrease can be offset by an optimal choice of τ_c^t and τ_ℓ^t. For any $0 < r_t \le 1$:

(A.5.1) $$\tau_c^{\star,t} = -1 + 1/r_t^\star, \text{ and:}$$

(A.5.2) $$\tau_\ell^{\star,t} = \quad 1 - 1/r_t^\star.$$

Note that by (A.5.2), the optimal tax rate on labor income will be negative.

The taxes and transfers levied on the young, $\tau_c^{\star,t}$ and $\tau_\ell^{\star,t}$, may cause the government to run a net surplus or deficit. The seigniorage revenue realized on bond holdings $(1 - r_t^\star)$ will also leave the government a surplus in the second period of each generation's life. The differences are made up with lump-sum transfers (taxes if negative) so that:

(A.5.3) $$H^{0\star} = \tau_c^{\star,t} c_0^\star + \tau_\ell^{\star,t}(1-\ell_0^\star),$$

(A.5.4) $$H^{1\star} = (1-r_t^\star)[1 - X^\star(1)].$$

Here $c_0^\star$ and $\ell_0^\star$ denote the saver's optimal choices of consumption and leisure while young, given the government's fiscal and monetary policy choices. ■

References

Allen, F. and D. Gale (1998). Financial contagion. Working Paper, The Wharton School, University of Pennsylvania.

Burnham, D. (1989). *A Law Unto Itself*. New York: Vintage Books.

Chicago Mercantile Exchange (1999). CME Quarterly Bankruptcy Index Futures and Options. `http://www.cme.com/qbi`.

Christofferson, P. F., F. X. Diebold, and T. Schuermann (1998). Horizon problems and extreme events in financial risk management. *Federal Reserve Bank of New York Economic Policy Review 4*(3), 109–116.

Dempster, A. P. (1968). A generalization of Bayesian inference. *Journal of the Royal Statistical Society, Series B 30*(2), 205–247.

Diamond, D. W. and P. H. Dybvig (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy 91*(3), 401–19.

Dow, J. and S. R. d. C. Werlang (1992a). Excess volatility of stock prices and Knightian uncertainty. *European Economic Review 36*(2–3), 631–638.

Dow, J. and S. R. d. C. Werlang (1992b). Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica 60*(1), 197–204.

Dow, J. and S. R. d. C. Werlang (1994). Nash equilibrium under Knightian uncertainty: Breaking down backward induction. *Journal of Economic Theory 64*(2), 305–324.

Epstein, L. G. and T. Wang (1994). Intertemporal asset pricing under Knightian uncertainty. *Econometrica 62*(2), 283–322.

Epstein, L. G. and T. Wang (1995). Uncertainty, risk-neutral measures and security price booms and crashes. *Journal of Economic Theory 67*(1), 40–82.

Gilboa, I. (1987). Expected utility with purely subjective non-additive probabalities. *Journal of Mathematical Economics 16*(1), 65–88.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics 18*(2), 141–53.

Gilboa, I. and D. Schmeidler (1993). Updating ambiguous beliefs. *Journal of Economic Theory 59*(2), 33–49.

Glahe, F. R. (1991). *Keynes's The General Theory of Employment, Interest, and Money: A Concordance*. Savage, Md.: Rowman & Littlefield.

Knight, F. H. (1921). *Risk, Uncertainty and Profit*. Boston: Houghton Mifflin.

Lagunoff, R. and S. Schreft (1998). A model of financial fragility. Working Paper, Department of Economics, Georgetown University.

Lagunoff, R. and S. Schreft (1999). Financial fragility with rational and irrational exuberance. Department of Economics, Georgetown University. Forthcoming in *Journal of Money, Credit, and Banking*.

Liu, W. (1998). Savings, portfolio decisions and investment: Observable implications of Knightian uncertainty. Manuscript, Department of Economics, University of Washington, Seattle.

Liu, W. (1999). Heterogeneous agent economies with Knightian uncertainty. Manuscript, Department of Economics, University of Washington, Seattle.

Rahl, L. (1998). Galaxy of risks. Capital Market Risk Advisors, Inc. `http://www.cmra.com`.

Savage, L. (1954). *The Foundations of Statistics*. New York: John Wiley and Sons.

Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica 57*(3), 571–587.

Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton, New Jersey: Princeton University Press.

Subrahmanyam, A. and S. Titman (1999). Real effects of financial market trading: Market crises and the going public process. Manuscript, College of Business Administration, University of Texas at Austin.

Wakker, P. (1989). Continuous subjective expected utility with non-additive priors. *Journal of Mathematical Economics 18*(1), 1–27.