Model Uncertainty and Liquidity*

Bryan R. Routledge†

and

Stanley E. Zin‡

June 2000
Revised: June 2001

Abstract

Extreme market outcomes are often followed by a lack of liquidity and a lack of trade. This market collapse seems particularly acute for derivative markets where traders rely heavily on a specific empirical model. Asset pricing and trading, in these cases, are intrinsically model dependent. Moreover, observed behavior of traders and institutions suggests that attitudes toward "model uncertainty" may be qualitatively different than Savage rationality would suggest. For example, a large emphasis is placed on "worst-case scenarios" through the pervasive use of "stress testing" and "value-at-risk" calculations. In this paper we use Knightian uncertainty to describe model uncertainty, and use Choquet-expected-utility preferences to characterize investors aversion to this uncertainty. We show that an increase in model uncertainty can lead to a reduction in liquidity as measured by the bid-ask spread set by a monopoly market maker. In addition, the non-standard nature of hedging model uncertainty can lead to broader portfolio adjustment effects like "flight to quality" and "contagion."

JEL Classification: G10, G13, G20

Keywords: Liquidity, Value-at-Risk, Knightian Uncertainty, Derivatives, Market Microstructure

---

*We thank Burton Hollifield, David Marshall, and seminar participants at Boston College, Cal-Tech, Carnegie Mellon University, the International Monetary Fund, CEF-2000 (Barcelona), and the Econometric Society World Congress (Seattle) for helpful comments.

†GSIA, Carnegie Mellon University.

‡GSIA, Carnegie Mellon University and NBER.

1 Introduction

In August 1998 an odd thing happened: the Russian government repudiated debt. While this event had a large effect on the value of Russian bonds, the event, by itself, is not odd. Long prior to this date, yields on Russian government bonds exhibited a significant premium over comparable U.S. Treasury securities, suggesting that default (or at least rescheduling) was not only possible but carried non-trivial probability. These bonds were undoubtedly *ex ante* risky. The *ex post* default and the change in the bond's price can be viewed simply as a realization from the distribution of possible payoffs (*i.e.*, "risk happens"). What is truly odd about the Russian-debt default and the subsequent collapse of the prominent hedge-fund Long Term Capital Management, was that during the crisis, markets for most emerging-markets debt exhibited a severe lack of liquidity. Anecdotal evidence suggests that people were unable to trade emerging market debt at *any price* following the Russian crisis. The lack of liquidity was not limited to the Russian debt market. The "flight-to-quality" made trading in most emerging market bonds, and even some corporate debt, difficult. In addition, several initial public offerings and below-investment-grade bond offerings were canceled.[1]

The extremely large credit spreads observed over this period had little, if any, historical precedent. Myron Scholes, a partner in Long-Term Capital Management (LTCM) at the time of the Russian debt crisis noted the improbability of the events of August 1998. One week after the Russian government default, the swap credit spread increased 20 basis points (treasury bonds versus AA rated debt). The increase is ten standard deviations above historic norms.[2] Many economic models can incorporate events like the Russian Crisis

[1] *The Wall Street Journal* reported on November 16th, 1998, (page A1) that "LTCM's partners... reported that their markets had dried up. There were no buyers, no sellers. It was all but impossible to maneuver out of large trading bets." On October 7th, they reported (page A1) on liquidity in an unrelated corporate bond market: "According to [Scott's Fertilizer Company's] lead investment bankers at Salomon Smith Barney, there is no bond market at any price."

[2] See Scholes (2000). Note that for a normal distribution, the probability of a 10 standard

as rare events, structural breaks, or changes in the risk premiums. From any of these perspectives, standard models would typically predict a capital loss by some, a capital gain by others, and perhaps a change in the market-price process. However, most models are unable to explain the drop in liquidity that accompanies the crisis. The puzzle is not the large change in financial prices, it is that people seem to stop trading.

The Russian debt crisis and market collapse is not unique. For example, Summers (2000) recounts the five other major international financial crises that took place during the 1990's involving economies in Mexico, Thailand, Indonesia, South Korea, and Brazil. Domestically, there have been many market collapses or crashes including the 1975 municipal bond crisis sparked by New York City's near default, various stock market crashes (1929, 1987, 1989,...), and the collapse of the high-yield debt market in the early 1990's. While all of these events have their unique features, they share two common features. First, crises are "unexpected." Almost by definition, a crisis involves a substantial change in financial prices. So the *ex ante* likelihood of the event is low. Second, crises are followed by a severe lack of liquidity. Following the various recent international and domestic crises, liquidity disappeared. People have difficulty executing trades for existing financial securities and new bond and equity offerings are postponed or canceled. In this paper we investigate the connection between these two features. We investigate whether a market break or severe reduction in liquidity can result from "model uncertainty." In particular, we focus on markets such as financial derivatives in which traders must rely on an empirical model for the stochastic cash-flow process of an underlying security. This is a setting where asset pricing and trading is intrinsically model-dependent. By specifying preferences that explicitly incorporate "model uncertainty" in a simple market-making setting, we show how uncer-

deviation tail event is on the order of 10^{-24} which is roughly the likelihood of winning the Powerball lottery three times in a row. It was also rumored at the time, that relative to some of LTCM's empirical models, spreads of this magnitude represented a 23 standard deviation event. The likelihood for a normal distribution of such an observation is on the order of 10^{-117} (for comparison, the number of atoms in the universe is on the order of 10^{78}).

tainty and liquidity are related.

To study the uncertainty-liquidity connection, we focus on a financial intermediary. The role of an intermediary is to facilitate trade. In well-developed liquid markets, the role of an intermediary is the straightforward matching of buyers and sellers (*e.g.*, the specialist at the NYSE). In contrast, in more specialized financial markets like "proprietary products," the intermediary participates directly in the transaction. For example, according to Scholes (2000), LTCM was in the "business of supplying liquidity." This type of intermediation requires an ability to value and hedge the financial contract that is being provided. Typically, firms attack this problem in two disjoint approaches. They use a model like Black and Scholes (1973) to calculate arbitrage bounds and hedge trades for a financial contract. However since the financial model is only an abstraction that is based on limited data, firms typically "stress test" their model to account for "model risk." For example, "Value at Risk" calculates the loss potential over a specified horizon for an arbitrarily specified probability. A portfolio resulting from the sale of a financial contract and an offsetting (perhaps dynamic) hedge position might have a 1% likelihood of losing $50 million over the next two weeks. Exactly how large a tail to measure and what distributional assumption to make are left to judgment. What is striking about the amount of attention paid to worst-case scenarios, stress testing and value-at-risk calculations is that trader attitudes towards uncertainty of the correctness of their model is distinct from the risk of stochastic prices. That is, their preferences do not adhere to the Savage (1954) axioms for expected utility rationality.

Savage rationality, in particular the independence or sure-thing axiom, implies that preferences should not depend on the source of the risk. Uncertainty about the appropriateness of a pricing model, "model uncertainty," is indistinguishable from the risk inherent in the assumed stochastic process. The Savage independence axiom implies that one can simply collapse the probability weighting across possible models ("uncertainty") with the probabilities for payoffs ("risk") to represent behavior with a single probability measure for states. However, in experimental settings, decision makers consistently violate the independence axiom. For example, Ellsberg (1961), demonstrated that individuals' decisions over lotteries could not be represented by an expected utility decision rule. People expressed (revealed) a preference to "know the odds" or an aversion to uncertainty. In the context of financial intermediation, not knowing the realization of an asset payoff (consumption risk) and not knowing the probability measure for payoffs (model uncertainty) have different behavioral implications. This distinction between risk and uncertainty, first described by Knight (1921), is axiomatized in Gilboa and Schmeidler (1989). The resulting decision rule that captures uncertainty aversion is represented by Choquet (1955) utility. Given a random variable $\omega \in \Omega$, an agent chooses the optimal action, $\theta \in \Theta$, according to

$$\max_{\theta \in \Theta} \left\{ \min_{\pi \in \Pi} E_{\pi}[u(\theta, \omega)] \right\}. \tag{1}$$

Uncertainty is captured by the set of probability measures Π. The aversion to uncertainty manifests itself in the "min" operator that appears after the action is chosen. Note that if the set Π is a singleton, then the decision rule is the standard Savage rationality of expected utility.[3] In this paper, we use the recursive intertemporal formulation of uncertainty aversion of Epstein and Wang (1994) and (1995). This specification facilitates dynamic programming and preserves dynamic consistency.[4] The robust control framework of Hansen, Sargent, and Tallarini (1999) is similar to the Epstein and Wang approach. In a linear-quadratic model the mean return, for example, is chosen by a malevolent nature. The result is the same "min" operator as in Choquet utility.

Our goal in the paper is to understand the relationship between model uncertainty and liquidity. The Choquet representation of uncertainty aversion is

[3]A closely related approach of Gilboa (1987) and Schmeidler (1989) models subjective prior beliefs to be non-additive. In a coin toss, uncertainty aversion is captured by $P(head) + P(tail) < 1$.

[4]Time inconsistent examples with uncertainty aversion is a concern since there is no restriction that conditional events have less uncertainty than unconditional. Seidenfeld and Wasserman (1993) define and provide examples of this dilation of beliefs.

well defined. However, at any level of generality, "liquidity" is difficult to define. Analogous to the vacuous distinction between unemployment and leisure in a perfect labor market, parties choosing not to trade in a frictionless financial market is not a lack of liquidity. Liquidity can only be defined relative to a market friction. Models of liquidity must include a market friction like an imperfectly competitive market or asymmetric information. Within the context of some market imperfection, liquidity is commonly measured as a "discount for immediacy" (*e.g.* Grossman and Miller (1988)) or the "price impact of a trade" (*e.g.* Kyle (1985)). In this paper we wish to study the relationship between liquidity and uncertainty rather than market microstructure *per se*. We, therefore, specify a rather simple market mechanism. We focus on the bid and ask prices for a proprietary derivative security. The market maker for this derivative is assumed to be a monopolist in that market. The market for the underlying security is assumed to be competitive. The market maker's preferences exhibit uncertainty aversion. We therefore treat the bid-ask spread and the associated probability that the market maker will make a trade, as a measure of liquidity in the market for this derivative security.[5]

Specifically, we consider a financial intermediary who makes a market for a propriety derivative security. This market maker chooses bid and ask prices for the derivative to optimally tradeoff the probability of attracting a seller or buyer with the current income and future utility implications implied from a trade in the derivative. When there is ambiguity about the appropriate probability distribution for the underlying security's cash flows, the market-maker is uncertain about these dynamic consequences, which we model with an Epstein-Wang uncertainty-averse utility function. We find that uncertainty increases the bid-ask spread and, hence, reduces liquidity. In addition, "hedge portfolios" for the market maker can look very different from those implied by a model without Knightian uncertainty.

In Section 2, we lay out the basic economic environment and describe the

[5]The setting we adopt here is similar to Ho and Stoll (1981) and related inventory-based microstructure models.

market-makers problem. In Section 3, we explore some simple two-period examples of the general model and in Section 4 extend these examples to an infinite time horizon. Section 5 concludes the paper.

2 The Model

The model we consider is that of a monopolist making a market in a derivative asset as well as choosing optimal portfolio and consumption. The market-maker sets a bid and ask price for a derivative whose payoff is $X(P_t) \geq 0$.[6] The demand for the derivative is summarized by the arrival of a random willingness-to-trade $\tilde{v}_t$. If $\tilde{v}_t$ is greater than or equal to the posted ask price, a_t, then a "buy order" is received and the market maker must go short one call (denoted as $d_t = -1$), at a price of a_t. If the willingness-to-trade $\tilde{v}_t$ is less than or equal to the posted bid price, b_t, then the market maker must go long one call ($d_t = 1$), at a price of b_t. If $\tilde{v}_t$ lies between the bid and ask prices, no trade takes place ($d_t = 0$). We assume the willingness to trade is an $i.i.d$ process with $\Phi(v) = \text{Prob}(\tilde{v} < v)$. The bid and ask prices determine the likelihood of trade in the derivative with $\text{Prob}(d_t = -1) = [1 - \Phi(a_t)]$, $\text{Prob}(d_t = 0) = [\Phi(a_t) - \Phi(b_t)]$, and, $\text{Prob}(d_t = 1) = \Phi(b_t)$.

The exogenous random arrival of a trade request is consistent with a number of deeper microstructure models. Since our focus is on uncertainty and the bid-ask spread, we will maintain this simple market specification throughout. After the arrival of the request to trade, the market maker chooses an optimal consumption and investment in a risky asset. This allows the market maker the opportunity to dynamically hedge the realized position in the derivative market.

Trades by the market maker are discrete long, short, or no-trade events, denoted $d_t \in \{-1, 0, 1\}$. The "size" of a trade can be incorporated into the

[6] In order to maintain the intuitive bid-ask relation, $0 < b < a$, we will only consider derivatives with non-negative payoffs.

definition of the derivative's payoff. For concreteness, our numerical examples focus on the case of a one-period call option $X(P_t) = s\max(P_t - x, 0)$. The parameter s determines the size or importance of each trade.

Investment Opportunities:

We denote as θ_{t-1} the asset holdings brought into period t, θ_t the assets purchased in period t to be carried into $t+1$, P_t as the *ex*-dividend price of the underlying risky asset, and δ_t as the period-t dividend paid by the risky asset.[7] The market for the underlying asset is assumed to be competitive. c_t is period-t consumption expenditures and ω_t is the income of the market-maker in period t. The period-t budget constraint is:

$$\theta_{t-1}(P_t + \delta_t) + \omega(y_t, v_t, d_{t-1}; a_t, b_t) = c_t + \theta_t P_t \ . \tag{2}$$

Total income, $\omega(y_t, v_t, d_{t-1}; a_t, b_t)$, includes both exogenous income and derivatives trading income. Exogenous period t income is denoted y_t. The market-making activity affects income both through the derivative position, d_{t-1}, carried into period t and through new trades in the derivative. The trading income in the current period depends on the choice of ask, a_t, and bid, b_t, and the realization of the willingness-to-trade, $\tilde{v}_t$.

$$\omega(y_t, v_t, d_{t-1}; a_t, b_t) = y_t + d_{t-1}X(P_t) + \begin{cases} a, & \text{if } v_t \geq a_t \\ 0, & \text{if } b_t < v_t < a_t \\ -b, & \text{if } v_t \leq b_t \end{cases} \tag{3}$$

The trading outcome also determines the position in the derivative d_t the market maker will carry forward into the next period. That is $d_t = -1$ if $v_t \geq a_t$, $d_t = 0$ if $b_t < v_t < a_t$, or $d_t = 1$ if $v_t \leq b_t$.

The timing of events implied by this notation is shown in Figure 1.

[7] The dividend on the risky asset is helpful in constructing the simple numerical example in Section 4. It is not needed for any of the analytical discussion or the two period example in Section 3.

Figure 1: Model Time-Line

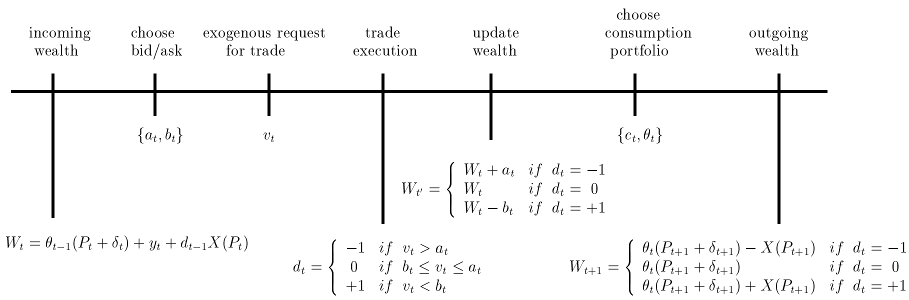

The market maker enters period t with holdings in the underlying security of θ_{t-1} which pay a dividend of δ_t and are liquidated at the price P_t. Holdings in the derivative of d_{t-1} have cash-flows realized of $X(P_t)$. Finally, he collects exogenous income of y_t. Given this information, he chooses a bid price, b_t, and an ask price, a_t. After the bid and ask are set, the exogenous request for a trade arrives, *i.e.,* $\tilde{v}_t$ is realized. Knowing the outcome of the trade in the derivative market, the market maker then chooses date t consumption, c_t, and investment in underlying risky security, θ_t.

Preferences:

The stochastic process governing the transition of the underlying security price and exogenous income is assumed to be Markov, with transition density given by

$$\text{Prob}\{P', \delta', y' \mid P, \delta, y\} = \pi(P, \delta, y) \ . \tag{4}$$

If the market maker has uncertainty or ambiguity about these probabilities, we will denote as Π the set of all distributions for future consumption paths that they believe are possible. Note that as in Epstein and Wang (1995), we assume that this set is time invariant. It should be thought of as part of the investor's preferences, rather than the physical environment.[8] Also following Epstein and Wang, we assume that preferences are given by the utility function, U that is the stationary, recursive specification of uncertainty aversion:

$$U(c_0, \tilde{c}_1, \tilde{c}_2, \ldots) = u(c_0) + \beta \min_{\pi \in \Pi} E_\pi U(\tilde{c}_1, \tilde{c}_2, \ldots) \,, \tag{5}$$

where $0 < \beta < 1$ is a utility discount factor and $u(c)$ is the single period utility derived from consumption c. Standard Savage preferences are included in this specification. If the set Π as a singleton, the agent adheres to Savage axioms.

Bellman Equation:

Combining the investment opportunity, the consumption implied by the budget constraint (2), and the specification of preferences, we can characterize this problem as a dynamic program. The Bellman equation associated with this program is given by:

$$\begin{aligned} V(\theta, d, P, \delta, y) = \max_{a,b} \Bigg\{ & [1 - \Phi(a)] \Big\{ \max_{\theta'} \Big[u(\theta(P+\delta) + y + dX(P) + a - \theta' P) \\ & \qquad + \beta \min_{\pi \in \Pi} E_\pi [V(\theta', -1, P', \delta', y')] \Big] \Big\} \\ & + [\Phi(a) - \Phi(b)] \Big\{ \max_{\theta'} \Big[u(\theta(P+\delta) + y + dX(P) - \theta' P) \\ & \qquad + \beta \min_{\pi \in \Pi} E_\pi [V(\theta', 0, P', \delta', y')] \Big] \Big\} \\ & + \Phi(b) \Big\{ \max_{\theta'} \Big[u(\theta(P+\delta) + y + dX(P) - b - \theta' P) \\ & \qquad + \beta \min_{\pi \in \Pi} E_\pi [V(\theta', 1, P', \delta', y')] \Big] \Big\} \Bigg\}. \end{aligned} \tag{6}$$

$V(\theta, d, P, \delta, y)$ is the value function. It depends on the five state variables: the position in the asset, the position in the derivative, the realized price for the asset, the realized dividend, and the realization for the exogenous income. The portfolio (and hence consumption), are chosen after the realization of $\tilde{v}$,

[8] Ambiguity is only relevant if the agent's preferences are averse to ambiguity.

which along with a and b, determines the outgoing position in the derivative. The outgoing position in the derivative, d', characterizes the future effect from the derivative trading (i.e., a, b, $\tilde{v}$ need not be included in the list of state variables).

With this formulation of the problem, it is straightforward to extend the model to include more assets and derivatives, other distributions of trade requests, or other market frictions.

Although closed-form solutions for the optimal policies of this dynamic program are unavailable, there are many computational algorithms that can be used to solve numerical versions of this model. When the set Π is not a singleton, the computational burden associated with solving this problem can be significantly greater than in the standard expected utility model. The additional non-linear program necessitated by the uncertainty averse preferences (minimizing over distributions) increases computation time. The Appendix outlines some new approaches to this problem that may yield significant computational gains.

3 Two-Period Model

To better understand the connection between uncertainty and liquidity, we first examine a simpler, two-period ($t = 0, 1$) version of the economy. Here, the market maker will make a market in a derivative of the single risky asset at period zero and the derivative will payoff at period one. The single risky asset, whose prices are P_0 and P_1, trades in a perfect market. In this section, the dividend is zero, $\delta_0 = \delta_1 = 0$, and exogenous income, y_0 and y_1 are non-stochastic.

Since the portfolio is chosen after the realization of trade in the derivative, we can consider the portfolio choice and the market-making activity separately. To do this, it is helpful to write the Bellman equation in two parts. Before determining, the optimal period zero bid and ask prices, consider the choice of the optimal consumption and portfolio. This is the inner maximization over θ' in equation (6). Note that a portfolio choice determines consumption via the budget constraint. That is, $c_0 = \omega_0 - \theta P_0$ and $c_1 = \omega_1 + \theta P_1$. Define the indirect total utility function $U(d, a, b)$ as follows:

$$U(d,a,b) = \max_{\theta} \left\{ u\left(\omega_0 - \theta P_0\right) + \beta \min_{\pi \in \Pi} E_\pi \left[u\left(\omega_1 + \theta P_1\right)\right] \right\} \tag{7}$$

The two-period version of income, equation (3), is

$$\omega_0 = y_0 + \begin{cases} -b & \text{if } d = 1 \\ 0 & \text{if } d = 0 \\ a & \text{if } d = -1 \end{cases} \qquad \text{and} \qquad \omega_1 = y_1 + dX(P_1) \tag{8}$$

Since $U(d, a, b)$ is conditional on the realization for the trade in the derivative market, the ask and bid prices have implications only for period-zero income. Once the trade occurs, they no longer affect the probability of a trade occurring. Therefore, $U(d, a, b)$ does not depend on a or b if $d = 0$. Similarly, it does not depend on ask price, a, if a "buy" order occurred and $d = 1$. Finally, U does not depend on the bid price b if a "sell" order $d = -1$ was received. Therefore, we can summarize equation (7) with three functions. Denote U_0 when there is no trade in the derivative ($d = 0$), U_b when the market maker has paid b and is long the derivative ($d = 1$), and U_a for when the market maker sold the derivative for a and holds a short position ($d = -1$).

Recall that the demand for the derivative asset is captured by the arrival of a trader with a valuation $\tilde{v}$ with distribution $\Phi(v) = \text{Prob}(\tilde{v} < v)$. The complete market-maker problem is

$$\max_{a,b} \left\{ \left[1 - \Phi(a)\right] U_a + \left[\Phi(a) - \Phi(b)\right] U_0 + \Phi(b) U_b \right\} \tag{9}$$

Portfolio Choice with Uncertainty Aversion

The portfolio problem in equation (7) can be studied independently of the market making activity. In particular, it is helpful to understand how Knightian uncertainty affects the optimal portfolio for a given income process ω_0 and $\tilde{\omega}_1$.

For concreteness, consider the portfolio choice for three different agents $i \in \{1, 2, K\}$. Two Savage individuals are captured with $\Pi^1 = \{\pi^1\}$ and $\Pi^2 = \{\pi^2\}$. An individual with an aversion to Knightian uncertainty is represented by $\Pi^K = \{\pi^1, \pi^2\}$. Obviously, to make this example interesting, $\pi^1 \neq \pi^2$.

For Savage individuals $i = 1, 2$, the optimal portfolio is characterized by the standard first order condition.

$$-u'(w_0 - \theta^i P_0)P_0 + \beta E_{\pi_i}\left[u'(w_1 + \theta^i P_1)P_1\right] = 0 \tag{10}$$

Assume that the optimal portfolio has $\theta^1 < \theta^2$.[9] This is the situation depicted in Figures 2 and 3. These figures plot total indirect utility (right-hand-side of (7)) against portfolio holdings, θ. The indirect utility is plotted for each distribution π^1 and π^2. For exposition, the utility function is quadratic and the distributions, π^i, are summarized by their mean and variance. The Knight agent's indirect utility, with $\Pi^K = \{\pi^1, \pi^2\}$, is the lower envelope of the two Savage agents.

There are three possibilities for how the optimal portfolio of the Knight individual is characterized. Denote θ^i as the optimal portfolio of the $i = 1, 2, K$ traders. First, it is possible that the that $\theta^K = \theta^1$ and the optimal portfolio satisfies equation (10). This is the case when $E_{\pi^1}\left[u\left(\omega_1 + \theta^1 P_1\right)\right] < E_{\pi^2}\left[u\left(\omega_1 + \theta^1 P_1\right)\right]$.[10] This situation is depicted in Figure 2. The second possibility is that the second distribution guides the portfolio choice and $\theta^K = \theta^2$. In both these cases, the aversion to uncertainty makes the Knight agent act according to the worst case probability distribution. The Knight agent simply looks like a pessimistic or "worst-case" Savage agent. Note that the definition

[9] For this discussion, we ignore the possibility that $\theta^1 = \theta^2$.

[10] This condition implies that $E_{\pi^2}\left[u\left(\omega_1 + \theta^2 P_1\right)\right] > E_{\pi^1}\left[u\left(\omega_1 + \theta^2 P_1\right)\right]$.

Figure 2: Optimal Portfolio: Quadratic Example 1

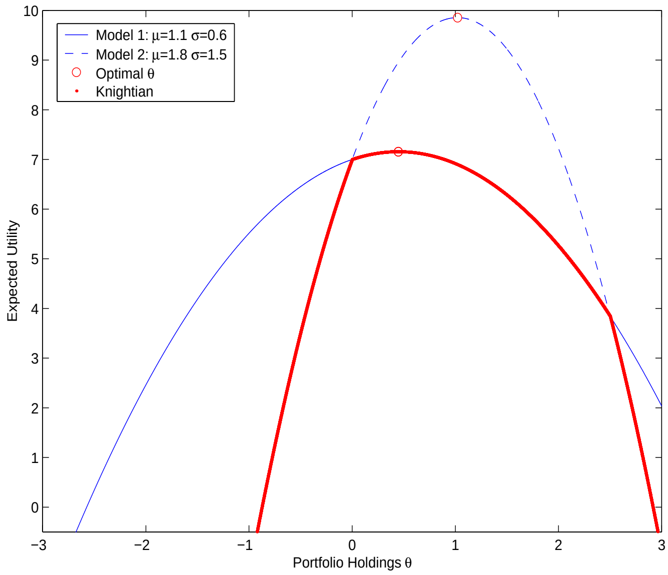

Note: The figure depicts the indirect utility from each of two distributions (low mean/low variance versus high mean/high variance) as a function of the investment in the risky asset θ. In this example, the Knightian-uncertainty portfolio choice is equivalent to assuming the low mean/low variance distribution and Savage expected utility.

of "worst case" depends in a complicated way on the portfolio position, θ, and on the endowment process. For example, for two distributions that differ only in the mean, the "worst case" depends on whether or not θ is positive or negative. Which distribution is considered the "worst case" will also depend on the correlation between period-one income and the asset payoff. Since, as we will explore below, market-making activity influences this correlation, the distribution that is considered as worst case may depend on the market-maker's position in the derivative.

The final possibility for the Knight agent's optimal portfolio is where $\theta^K \neq$

Figure 3: Optimal Portfolio: Quadratic Example 2

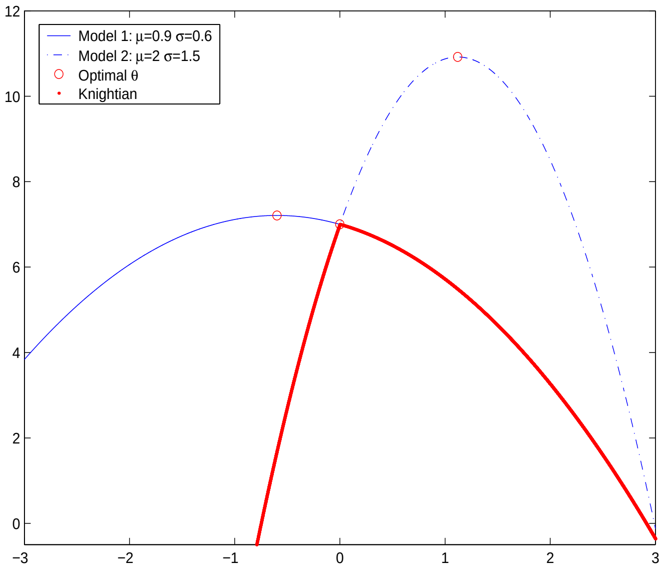

Note: The figure depicts the indirect utility from each of two distributions (low mean/low variance versus high mean/high variance) as a function of the investment in the risky asset θ. In this example, the Knightian-uncertainty portfolio choice differs significantly from the Savage expected-utility choice under either distribution.

θ^1 and $\theta^K \neq \theta^2$. This situation arises when

$$\begin{aligned} E_{\pi^1}\left[u\left(\omega_1+\theta^1 P_1\right)\right] &> E_{\pi^2}\left[u\left(\omega_1+\theta^1 P_1\right)\right] \quad \text{and} \\ E_{\pi^2}\left[u\left(\omega_1+\theta^2 P_1\right)\right] &> E_{\pi^1}\left[u\left(\omega_1+\theta^2 P_1\right)\right] \end{aligned} \tag{11}$$

When equation (11) holds, the optimal Knight portfolio, θ^K, solves

$$E_{\pi^1}\left[u\left(w_1+\theta^K P_1\right)\right] = E_{\pi^2}\left[u\left(w_1+\theta^K P_1\right)\right] \tag{12}$$

and the Knight agent looks like neither of the Savage agents. This case is illustrated in Figure 3. The optimal portfolio for the Knight individual occurs at the intersection of the intersection of the indirect utility calculated under the two distributions. In this case, it is not the case that the Knight agent is just pessimistic or focusing on a "worst case" distribution. The two Savage portfolios bound the optimal Knight portfolio (i.e., $\theta^1 < \theta^K < \theta^2$)[11] and there exists some (many) other Savage agents and beliefs, say π^3, such that $\theta^K = \theta^3$. However, since the Knight portfolio is not defined by a first order condition, it responds differently to changes in the state variables. In particular, consider a change in initial income, ω_0. In the situation where the portfolio is given by the first order condition, equation (10),

$$\frac{\partial \theta^i}{\partial \omega_0} = \frac{1}{P_0} \left(\frac{u'' P_0^2}{u'' P_0^2 + \beta E_{\pi^i} \left[u'' P_1^2 \right]} \right) \tag{13}$$

This implies that $0 < \frac{\partial \theta^i}{\partial \omega_0} < \frac{1}{P_0}$. This will be the case for Savage agents. It is also the case for a Knight trader when he acts as a "worst-case" Savage. However, in the case where equation (12) determines the optimal portfolio, $\frac{\partial \theta^K}{\partial \omega_0} = 0$ since, by inspection, ω_0 does not enter. The optimal portfolio of the Knight agent is insensitive to changes in initial wealth. This is true even for discrete changes in initial income. As long as equation (11) holds at the new level of income, ω_0, even large changes in initial income will not alter the optimal Knight holdings of the risky asset.[12]

The Market Maker Problem

The market maker problem for the two-period economy is contained in equation (9). In choosing the bid and ask prices, the tradeoff for the market maker is straightforward. Choosing a high value for the ask will generate more revenue should a high-value trader arrive. However, it lowers the probability of such a trade actually arriving. Likewise, choosing a low value for the bid

[11] For Savage agents the unique optimal portfolio is defined by the solution to equation (10). This is reflected in the "single-peaked" functions in Figures 2 and 3. An implication is that in the case where the Knight portfolio differs from both Savage portfolios, it will lie in an interval between the two Savage portfolios.

[12] The one case not considered here is where both equations (10) and (11) hold. In this situation, the peak for one of the Savage traders in Figure 2 or 3 occurs at the intersection with the other Savage agent. In this case the derivative $\frac{\partial \theta^K}{\partial \omega_0}$ will be zero in one direction and given by (13) in the other direction. Since this event is not likely (measure zero), it is not explored further.

will allow the market maker to obtain the future cash flows of the call for a low price should a low-value trader arrive, but it lowers the probability of such a trade actually arriving. The first order conditions for (9) are

$$\begin{aligned} U_b - U_0 &= \tfrac{\Phi(b)}{\phi(b)}(-U_b') \\ U_0 - U_a &= \tfrac{1-\Phi(a)}{\phi(a)}(U_a') \end{aligned} \tag{14}$$

where $U_b' = \partial U_b / \partial b$, $U_a' = \partial U_a / \partial a$, and $\phi(v) = \partial \Phi(v) / \partial v$.[13]

Market-making activity has implications for income at both date zero and date one (see equation (8)), so the analysis regarding the optimal portfolio in the previous section is helpful here. Denote the optimal portfolio from the solution equation (7) for each of the three agents $i \in \{1, 2, K\}$, as θ_0^i in the situation where there is no trade in the derivative ($d = 0$), θ_b^i when the market maker has paid b and is long the derivative ($d = 1$), and θ_a^i when the market maker sold the derivative for a and holds a short position ($d = -1$). Using this notation, for all three traders $i \in \{1, 2, K\}$, the derivative U_b' and U_a' exist and are given by

$$\begin{aligned} U_b' &= -u'\left(y_0 - b - \theta_b^i P_0\right) \\ U_a' &= u'\left(y_0 + a - \theta_a^i P_0\right) \end{aligned} \tag{15}$$

For Savage preferences, equation (15) is simply a statement of the envelope condition. For the Knight market-maker, the envelope condition holds when the optimal portfolio is given by (10); that is when $\theta_a^K = \theta_a^1$ or $\theta_a^K = \theta_a^2$ (similarly for the θ_b^K). In the case where θ_a^K is determined by (12), implying, $\theta_a^K \neq \theta_a^1$ and $\theta_a^K \neq \theta_a^2$, $\frac{\partial \theta_a^K}{\partial a} = 0$. This is the exact same calculation as $\frac{\partial \theta_a^K}{\partial \omega_0}$ since, conditioning on $d = -1$, a only affects current period income. Using

[13] To calculate the second order conditions, note that there are no cross products. That is the fist line of equation (14) is independent of a. The second order condition for b is

$$\phi'(b)\left(U_b - U_0\right) - 2\phi(b)u' + \Phi(b)u''$$

There is an analogous equation for the ask, a. This condition does not hold for all distributions. However, for a uniform distribution, used in the examples below, $\phi'(b) = 0$ and the second order condition is satisfied.

(15), we can write the optimal bid and ask price as solving:

$$\begin{array}{r}U_b - U_0 = \frac{\Phi(b)}{\phi(b)} u'\left(y_0 - b - \theta_b^i P_0\right) \\ U_0 - U_a = \frac{1-\Phi(a)}{\phi(a)} u'\left(y_0 + a - \theta_a^i P_0\right)\end{array} \tag{16}$$

From equation (16) and the preceding discussion, we can make a few general comments about spreads and the role of Knightian uncertainty. First, for both Savage and Knight market makers there is a preference to be long the derivative rather than short. That is, given optimal bids and portfolio choices, $U_b > U_0 > U_a$. This follows by noting the right-hand-side of equation (16) is strictly positive. This result is perhaps not surprising given that we are considering derivatives with risky, non-negative payoffs and risk averse market makers. Second, note that the right-hand-side of equation (16) is not directly affected by the min operator in preferences. Any non-differentiability induced by the min operator does not directly affect bid-ask spreads.[14] Knightian uncertainty does effect bid-ask spreads through two channels. First, channel is via the differences $U_b - U_0$ and $U_0 - U_a$. If Knightian uncertainty increases these differences, the bid-ask spread will be larger. The second channel is via the optimal portfolio. The previous section discussed how the portfolio of the Knight may look either like a "worst case" Savage trader or look like neither of the Savage market-makers. These two situations, as discussed in the numerical example below, produce different reactions in the portfolio from a long or short position in the derivative. In particular, the "hedge ratio" need not have the usual properties. In order to explore these features and properties, we next turn to a numerical example.

The Market Maker Problem - Two Period Numerical Example

To explore the effects of uncertainty on the bid-ask spread, we examine

[14] This is similar to a feature of Epstein and Wang (1994) and (1995). In the representative agent, endowment economy model these two papers present, the indeterminacy in equilibrium created by the min operator in uncertainty averse preferences is only important for economies without aggregate risk. In economies with aggregate risk and uncertainty only about aggregate endowment, the representative agent's Knightian utility is differentiable and prices are unique. See Epstein (2001).

some numerical examples of the two period economy. Preferences, u, in this example are quadratic and exogenous income, y_0 and y_1, is constant. The example considers a market-maker for a call option $X(P_t) = s\max(0, P_1 - 1.5)$ with $s = 1$ unless otherwise specified. The demand for the derivative is summarized by the arrival of a random willingness-to-trade $\tilde{v}$ where $\tilde{v}$ is distributed uniformly on the interval $[0.5, 1.5]$. The current price, $P_0 = 0.9$. The distribution(s) of the underlying asset's period one payoff, P_1, is assumed to be binomial with equally likely values of

$$\tilde{P}_1 = \begin{cases} \mu_m + \sigma_m \\ \mu_m - \sigma_m \end{cases} \tag{17}$$

which implies a mean of μ_m and a variance of σ_m^2. There are two possible distributions or models, π^m for $m \in \{1, 2\}$, for the underlying asset. Two Savage are captured by $\Pi^1 = \{\pi^1\}$ and $\Pi^2 = \{\pi^2\}$. The Knight market-maker with uncertainty aversion, is represented by $\Pi^K = \{\pi^1, \pi^2\}$. Model $m = 1$ has a mean and standard deviation variance of 0.9. Model 2 has a mean and standard deviation that we range from 0.7 to 1.8. For the different economies we consider, as the distribution in Model 2 moves further away from the distribution in Model 1, uncertainty increases.

Figure 4 depicts the effect on the bid and ask prices as uncertainty increases. The market maker with an aversion to Knightian uncertainty maintains a constant bid price, whereas a Savage market maker will allow the bid price to rise to reflect the higher value in the derivative. The affect that this has on the bid-ask spread is depicted in Figure 5, and the affect on the probability of a trade is depicted in Figure 5. As depicted in these figures, it is possible for model uncertainty to completely eliminate the willingness of a market maker who is uncertainty averse to provide liquidity, while a market maker who is not averse to this uncertainty but is merely a pessimistic expected utility maximizer will continue to provide liquidity.

From equation (16), spreads in this model are related to the change in utility from taking a long or short position in the derivative security. Therefore,

Figure 4: Bid and Ask as Uncertainty Increases

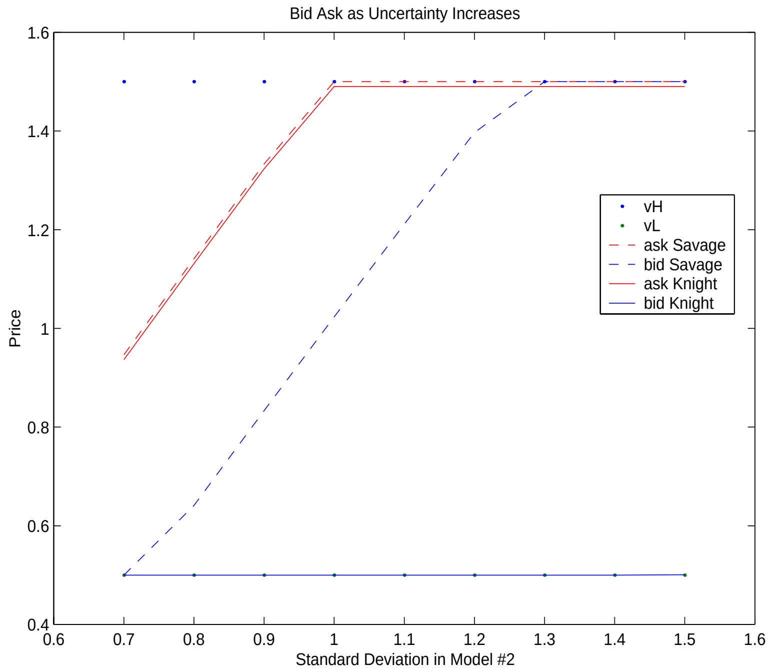

Note: The expected-utility market maker is willing to raise his bid price as volatility increases, whereas the market maker with an aversion to Knightian uncertainty does not.

Figure 5: Bid-Ask Spread as Uncertainty Increases

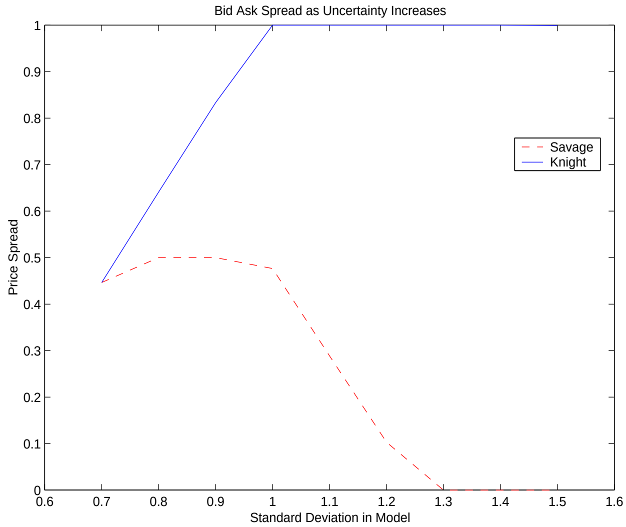

Figure 6: Probability of a Trade as Uncertainty Increases

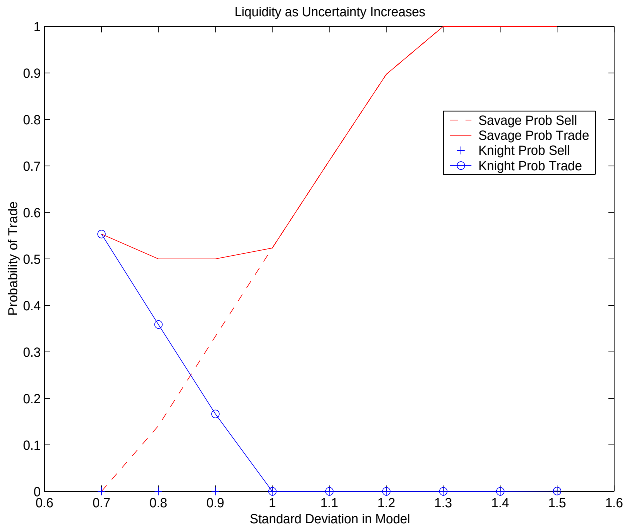

Figure 7: Total Indirect Utility as size of derivative payoff, s, Varies Continuously

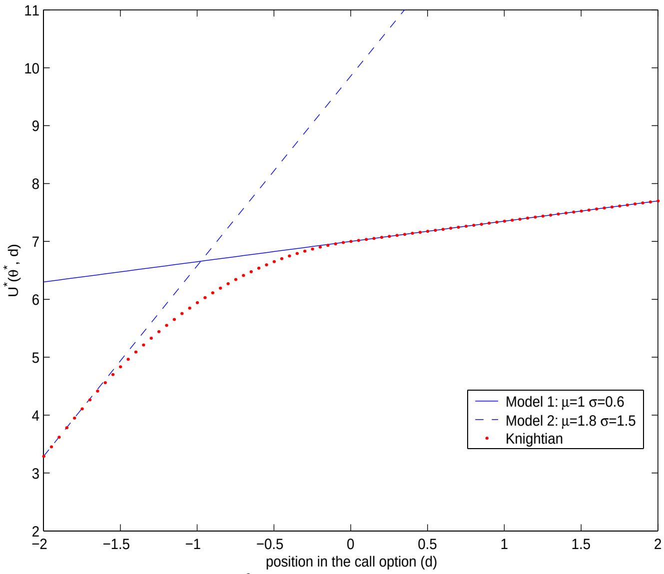

For fixed parameters, the size of the position in the derivative, $d \cdot s$, is varied. The payoff in the derivative is $X(P_1) = s\max(P_1 - 1.5, 0)$ and trades are discrete $d \in \{-1, 0, 1\}$. The total indirect utility shown is at the optimal portfolio.

bid-ask spreads are closely related to the concavity of the utility function at the optimal portfolio. Figure 7 plots the optimal total indirect utility as a function of the position in the derivative, d and size of derivative payoff s. Recall, $X(P_1) = s\max(P_1 - 1.5, 0)$. This plot lets us consider arbitrarily small long and short positions in the derivative. The plot shows the Savage expected utility preferences under the two possible distributions. The lower line represents the Knightian preferences. Since the Knightian preferences are more concave, they produce larger bid-ask spreads. This feature should also lead to larger spreads in other market specifications beyond the simple monopolist structure we consider here.

Figure 7 also highlights the necessity of a market friction in generating bid-ask spreads. Uncertainty aversion itself cannot be the source of a bid-ask spread. In a perfectly competitive situation where bid-ask spreads are set in a Bertrand competition, the spread will equal the change in utility from taking a long or short position in the derivative security. If trades in the derivative security are arbitrarily small (represented by arbitrarily small payoffs, i.e., $s \to 0$ from the left and right), a model with Savage expected-utility market makers will result in zero spreads since indirect utility is twice differential. For Knightian uncertain market makers in the same setting, this is true almost everywhere. Only at two points in Figure 7 is the Knightian market maker's indirect utility kinked. These kink points occur at the point where both (10) and (11) hold and the two Savage expected utilities intersect at an optimum. These two kinks are measure zero, so spreads even with Knightian uncertain market makers will be almost everywhere zero if markets are frictionless.[15] For example, to apply our model to the market collapse related to the 1998 Russian bond default discussed in the introduction, it important that both a market maker like LTCM is uncertainty averse as well as has some degree of market power. Given the important interaction between market frictions and Knightian uncertainty, we leave the question of an optimal market design given uncertainty to future research.

Hedging Derivatives Positions

An important facet of market making in derivatives is the ability to hedge the position in the underlying markets. The popularity of models like Black and Scholes (1973) is in their ability to provide an offsetting trade that hedges a position in the derivative. The ability to hedge positions is essential for a financial intermediary like LTCM to leverage their capital into large positions. In our two-period example, we can also look at the effect of model uncertainty on the trades used to hedge a position in the derivative. In our setting, the market is not complete, so market-makers cannot offset the full position in the derivative. However, we can look at how the optimal portfolio responds to a

[15] This is in contrast to comments in both Epstein and Wang (1995) and Dow and Werlang (1992). Knightain uncertainty, alone, is insufficient to generate bid-ask spreads.

Figure 8: Hedging a Short Call Position: "Natural"

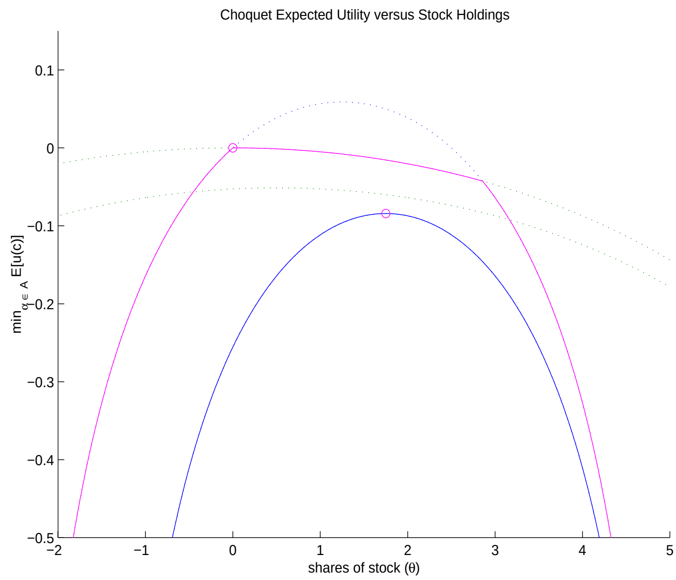

chance in the position in the derivative. For example, consider the optimal portfolio in the case where there is no derivative position relative to a short position in the derivative. Define the "hedge portfolio" induced by this short position in the call as $\theta_b^i - \theta_0^i$.[16] Figure 8 depicts this hedge portfolio for a short call position in a log-utility version of the two-period model. (The switch to logarithmic utility adds some asymmetry to the utility function relative to the quadratic case.) In this example, the short call position is hedged by buying more of the underlying asset. This is a natural hedging strategy and is consistent with the behavior of any Savage market maker. However, with a slightly different configuration of uncertainty, Figure 9 depicts a very strange situation. The short position in the call option is hedged by *reducing* investment in the underlying asset. The optimal portfolio, in response to a

[16] We are using the same notation as in the previous section: for traders $i \in \{1, 2, K\}$, θ_0^i is the optimal portfolio in the situation where there is no trade in the derivative ($d = 0$), θ_b^i when the market maker has paid b and is long the derivative ($d = 1$), and θ_a^i when the market maker sold the derivative for a and holds a short position ($d = -1$).

Figure 9: Hedging a Short Call Position: “Unnatural”

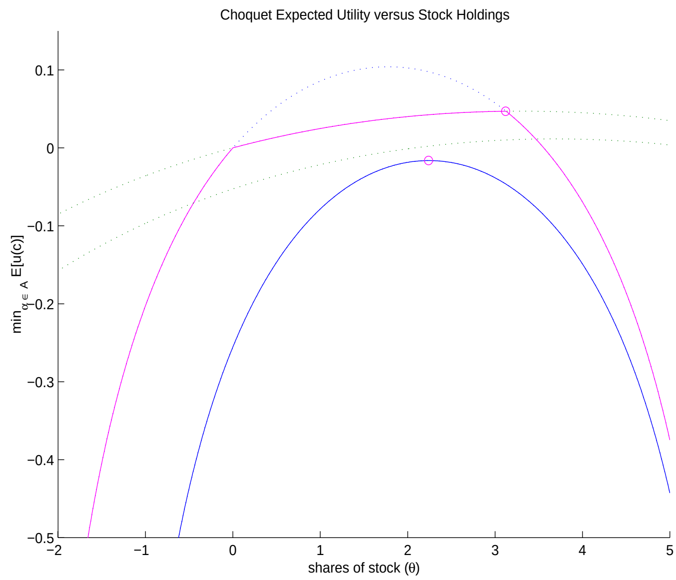

short position in the call option, has shifted left. The reason for this odd behavior is that, in this case, when $d = 0$, the optimal portfolio was not given the solution to the first condition as in equation (10). In this case, the optimal portfolio of the Knight trader does not resemble either of the Savage traders. Since the optimal portfolio is given by (12), it does not respond in a natural way to the derivative position. When the optimal portfolio is given by (12), the optimal hedge portfolio is not constrained to be positive or less than one as it would be for Savage market makers.

4 Infinite-Horizon Model

Building on our understanding of the two-period example, we now return to the infinite-horizon model summarized in equation (6). We focus on a relatively simple portfolio problem so that we can highlight the role of market making in the derivative. Assume that the underlying security price follows a two-state Markov process with $P_t \in \{0.75, 1.25\}$ with transition probabilities specified below. To ensure that the portfolio problem is well specified, it is necessary to assume the asset also carries a stochastic dividend. Without some additional cash-flow, the optimal portfolio is an arbitrarily large short sale position in the case $P_t = 1.25$. Therefore, we assume the underlying asset pays a dividend that also follows a two state Markov process, $\delta_t \in \{0, 0.4\}$. As in the previous section, we will consider two Savage market makers with beliefs, $\Pi^1 = \{\pi^1\}$ and $\Pi^2 = \{\pi^2\}$, and a Knight market maker with uncertainty represented by $\Pi^K = \{\pi^1, \pi^2\}$. For both possible models, π^1 and π^2, the states are *i.i.d.*. The two price and two dividend states produces a four state Markov process of (P_t, δ_t) $\in \{(0.75, 0), (0.75, 0.4), (1.25, 0), (1.25, 0.4)\}$ with transition probabilities,

$$\begin{aligned} \pi^1 &= \begin{bmatrix} 0.1875 & 0.5625 & 0.025 & 0.1875 \\ 0.1875 & 0.5625 & 0.025 & 0.1875 \\ 0.1875 & 0.5625 & 0.025 & 0.1875 \\ 0.1875 & 0.5625 & 0.025 & 0.1875 \end{bmatrix} \\ \pi^2 &= \begin{bmatrix} 0.1250 & 0.1250 & 0.3750 & 0.3750 \\ 0.1250 & 0.1250 & 0.3750 & 0.3750 \\ 0.1250 & 0.1250 & 0.3750 & 0.3750 \\ 0.1250 & 0.1250 & 0.3750 & 0.3750 \end{bmatrix} . \end{aligned} \tag{18}$$

Distribution π^1 is more pessimistic than distribution π^2. Other parameters used in the example are: utility is log, exogenous income is constant at $y_t =$ 12.750, and $\beta = 0.8$

Portfolio Choice with Uncertainty Aversion Example

The optimal portfolio in the case where there is no market-making activity, is shown in Figure 10 In each state, the Knight portfolio policy lies between the two Savage portfolio policies. This feature was discussed in the two period example. Unlike the Savage portfolio, the Knight portfolio policy has a region that is flat. In this region, the optimal portfolio, θ_t^K, is independent of the

Figure 10: Portfolio Policy for Two Savage Traders and a Knight Trader

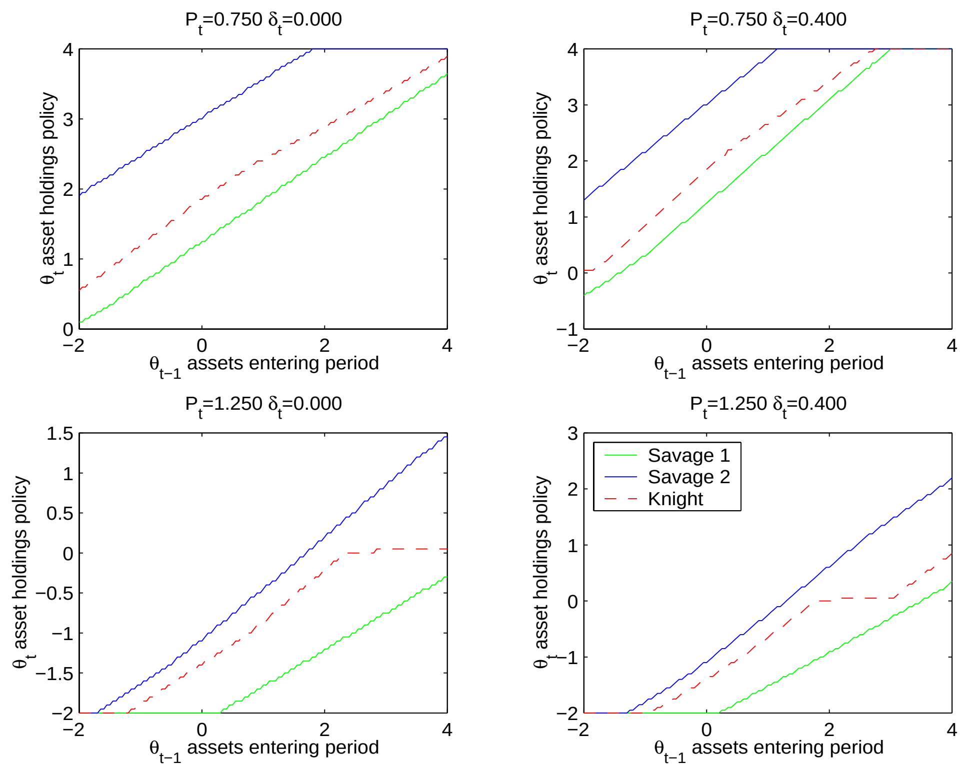

The optimal portfolio, θ^i, as a function of the previous asset holdings, is shown. Each sub-plot is a different value of the price-dividend state. The portfolio policy is shown for a Savage trader with beliefs $\Pi^1 = \{\pi^1\}$, a Savage trader with $\Pi^2 = \{pi^2\}$, and an uncertainty averse Knight trader with beliefs $\Pi^K = \{\pi^1, pi^2\}$. π^1 and π^2 are defined in equation (18).

asset holdings at the start of the period, θ_{t-1}^K. This occurs in the situation where the optimal portfolio for the Knight agent does not solve a first-order condition. Instead, the optimal portfolio is given by the infinite-horizon analog of equation (12). In this case, the optimal portfolio is not sensitive to initial wealth or, more precisely, the portfolio holdings brought into the period.

In the absence of market-making, Knightian uncertainty does not dramatically alter the portfolio behavior of a trader. While the optimal policy differs in the case of a Knight trader, it does not, by itself, have a dramatic effect on the time-series behavior of the portfolio holdings. To see this consider, Figure 11. The time-series behavior of the optimal Knight portfolio is constrained by the fact that it is bounded by the Savage-optimal portfolios state-by-state. Knight portfolio is never dramatically different than the Knight.

Market Maker Example - Policies

We use the same example to consider the infinite-horizon version of the market-maker problem. The derivative asset is a call option based on the *ex*-dividend price; that is $X(P_t) = s \max(P_t - x)$. We set the strike price at $x = 1.0$ and the derivative size at $s = 1.0$. The demand for the derivative is summarized by the arrival of a random willingness-to-trade $\tilde{v}$, where $\tilde{v}$ is distributed uniformly on the interval $[0.1, 0.2]$. Again, we consider the behavior of the two Savage market makers and an uncertainty averse, Knight, market maker.

Figures 12 and 13 summarize the bid and ask policy for the Savage market maker with beliefs π^1 and π^2 respectively. The figure shows the probability of a trade occurring, $1 - [\Phi(a_t) - \Phi(b_t)]$, as a function of the state variables. For the Savage market-maker with beliefs π^1, the bid and ask prices for the derivative are close to constant. The bid and ask prices are set such that the probability of trade is close to 0.5. For the Savage market-maker with the more optimistic beliefs of π^2 (see equation (18)), the probability for trade is slightly higher in the case where the underlying price is low ($P_t = 0.75$). Note,

Figure 11: Optimal Portfolio Time Series

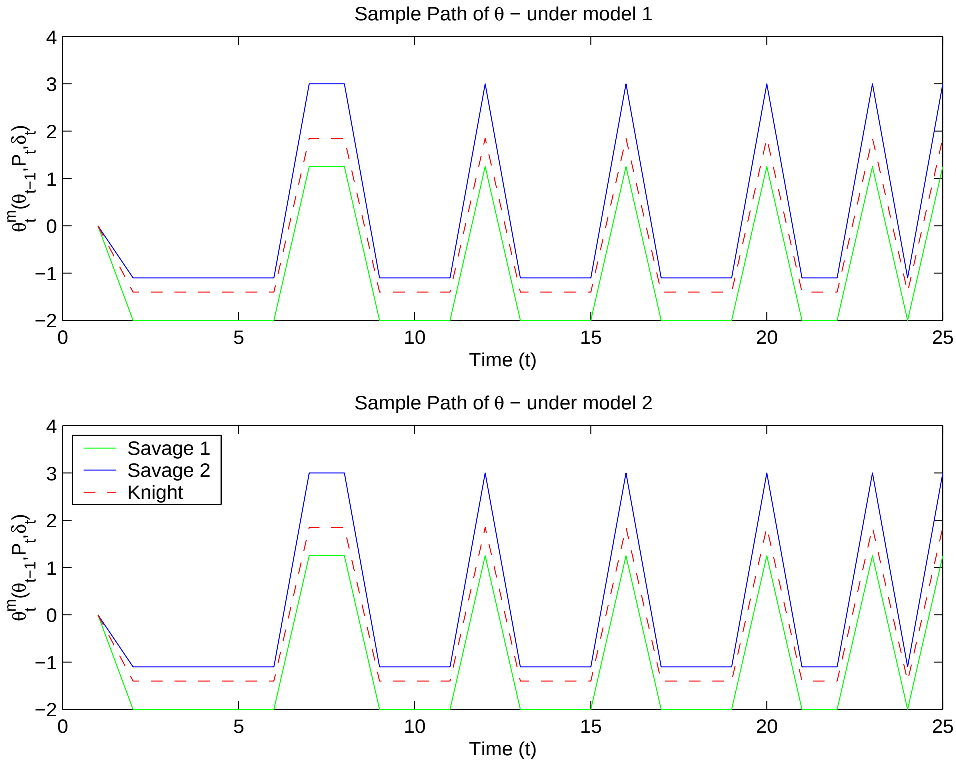

For a simulation of the economy, the optimal realized portfolio, θ_t^i, is shown for three traders: a Savage market maker with beliefs π^1, a Savage with beliefs π^2, and a Knight market maker with uncertainty averse beliefs represented by $\Pi^K = \{\pi^1, pi^2\}$. One period of the simulation consists of drawing a price P_t and a dividend δ_t The top panel is simulated under π^1 and the bottom panel is simulated under π^2 (see equation (18) for parameters).

Figure 12: Probability of Trade - Savage π^1

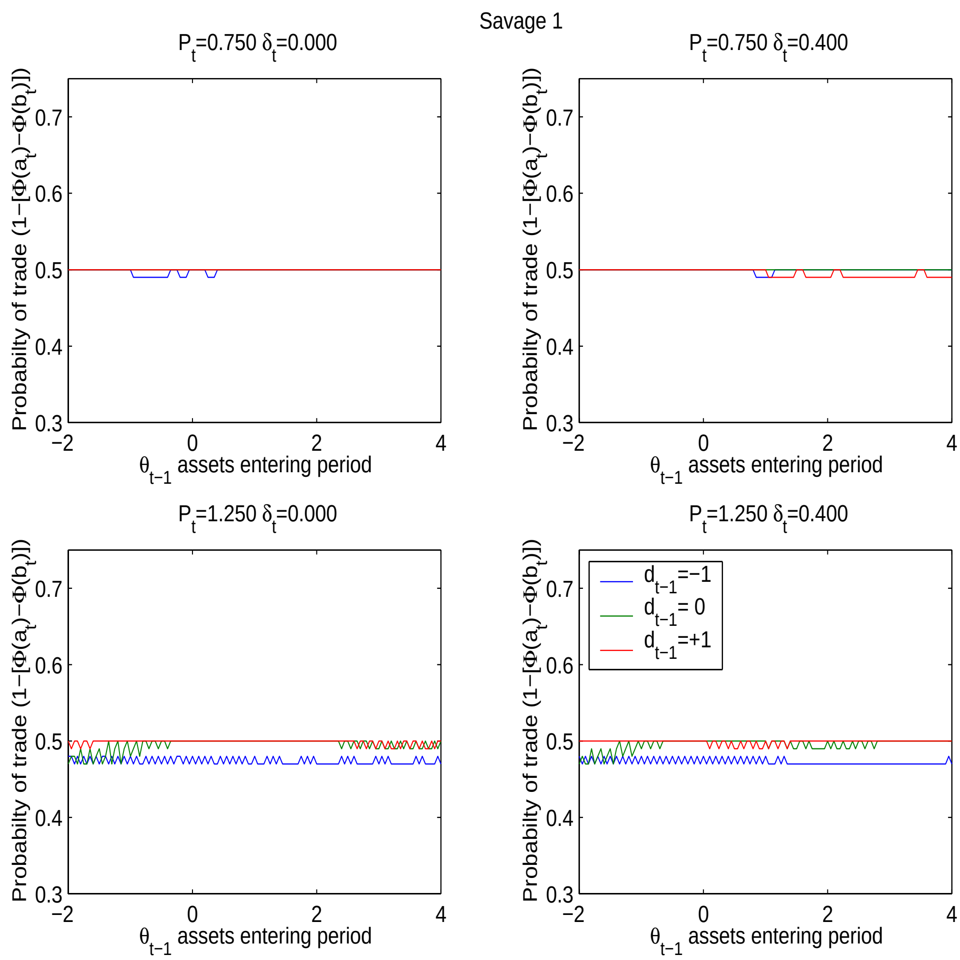

The figure shows the probability of a trade occurring given the bid and ask policy that solves equation (6) for the Savage market-maker with beliefs π^1 . The probability of a trade is calculated as $1-[\Phi(a_t)-\Phi(b_t)]$. The probability of as a function of the state variables: previous position in the derivative, d_{t-1} previous position in the portfolio, θ_{t-1},current asset price, P_t, and current asset dividend, δ_t.

Figure 13: Probability of Trade - Savage π^2

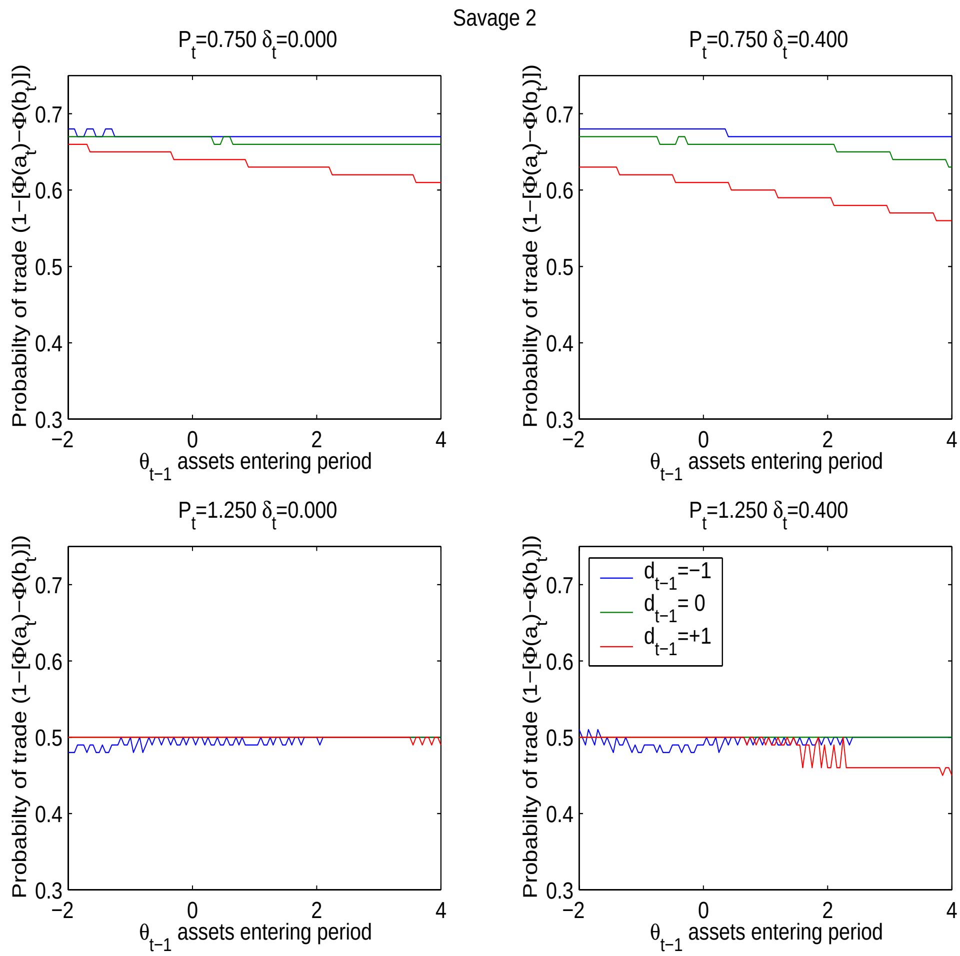

The figure shows the probability of a trade occurring given the bid and ask policy that solves equation (6) for the Savage market-maker with beliefs π^2 . The probability of a trade is calculated as $1-[\Phi(a_t)-\Phi(b_t)]$. The probability of as a function of the state variables: previous position in the derivative, d_{t-1} previous position in the portfolio, θ_{t-1},current asset price, P_t, and current asset dividend, δ_t.

for both Savage market makers, the optimal bid-ask policy is not that sensitive to the position in the underlying security, θ_{t-1}.

Figure 14 summarizes the bid and ask policy for the uncertainty averse Knight market-maker. Notice that the bid-ask behavior, reflected in the probability of trade, is much more sensitive to the incoming asset position, θ_{t-1}. It is also the case that the probability of trade can fall quite low (to 0.3). The low probability of trade for the Knight market-maker coincides with the case where the optimal portfolio is not sensitive to initial wealth since it is not characterized by a first-order-condition. Figure 15 shows the portfolio policy. Recall that the optimal asset position is chosen after the realization of the trade in the derivative and so depends on both previous, d_{t-1}, and current, d_t, position in the derivative. (The asset policy function for the two Savage traders is similar to that shown in Figure 10, so it is not repeated.) The regions where the probability for derivative trade is low occur, for example, when $P_t = 1.25$, $\delta_t = 0$, $d_{t-1} = 1$, and $\theta_t \approx 2$ (see lower-left panel of Figure 14). This situation corresponds to the lower three panels of Figure 15 where $d_{t-1} = 1$. In particular, the lower two lines. Note that the lack of liquidity is occurring at a point where the portfolio policy function is flat.

Market Maker Example - Time Series

To better understand the implications of these policy functions, it is helpful to simulate realizations for the economy. A simulation consists of drawing a price P_t, dividend δ_t, and willingness to trade, $\tilde{v}_t$ and applying the optimal policies for ask and bid prices and the portfolio. The result are for a simulation of 10,000 periods of the economy. Since there are two possible probability measures describing the evolution of price and dividend, the results show the simulation conducted under both distributions π^1 and π^2.

Figure 16 shows a realized path of bid and ask prices for the three market-makers. Fifty periods are shown. The two Savage traders differ in their beliefs about the likelihood of next period's price and dividend. The more optimistic market maker, π^1, tends to have a higher bid and ask price for the derivative.

Figure 14: Probability of Trade - Knight

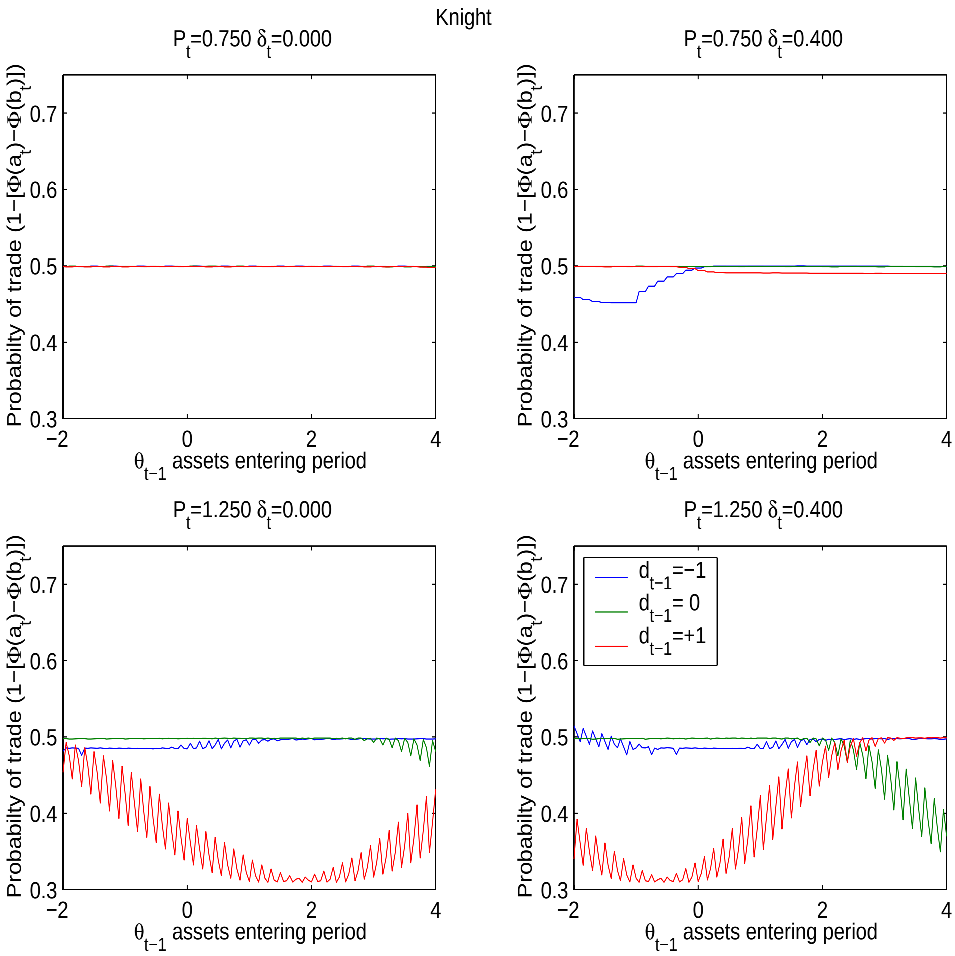

The figure shows the probability of a trade occurring given the bid and ask policy that solves equation (6) for the Knight market maker with uncertainty aversion represented by $\Pi^K = \{\pi^1, \pi^2\}$. The probability of a trade is calculated as $1 - [\Phi(a_t) - \Phi(b_t)]$. The probability of as a function of the state variables: previous position in the derivative, d_{t-1} previous position in the portfolio, θ_{t-1},current asset price, P_t, and current asset dividend, δ_t.

Figure 15: Optimal Portfolio - Knight π^2

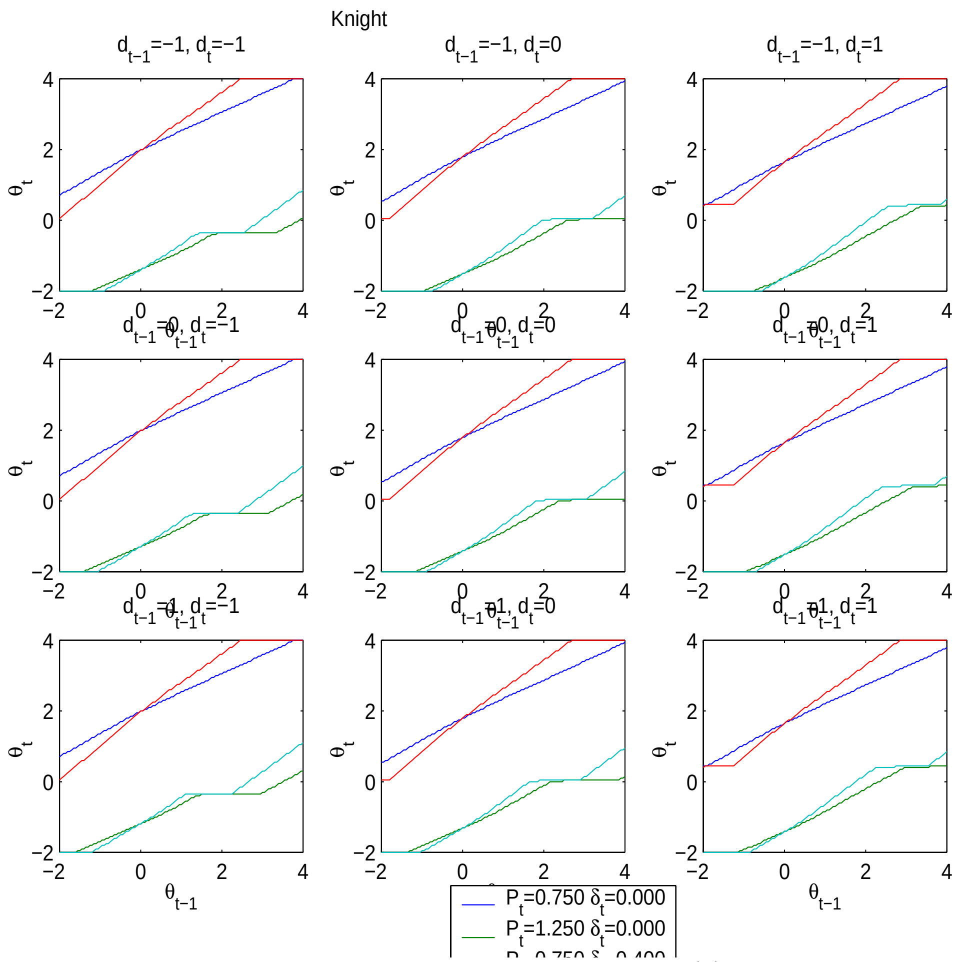

The figure shows the optimal portfolio that solves equation (6) for the Knight market maker with uncertainty aversion represented by $\Pi^K = \{\pi^1, \pi^2\}$. The optimal portfolio is a function of the state variables: current, d_t, and previous, d_{t-1} position in the derivative, previous position in the portfolio, θ_{t-1},current asset price, P_t, and current asset dividend, δ_t.

Figure 16: Simulated Bid and Ask Prices

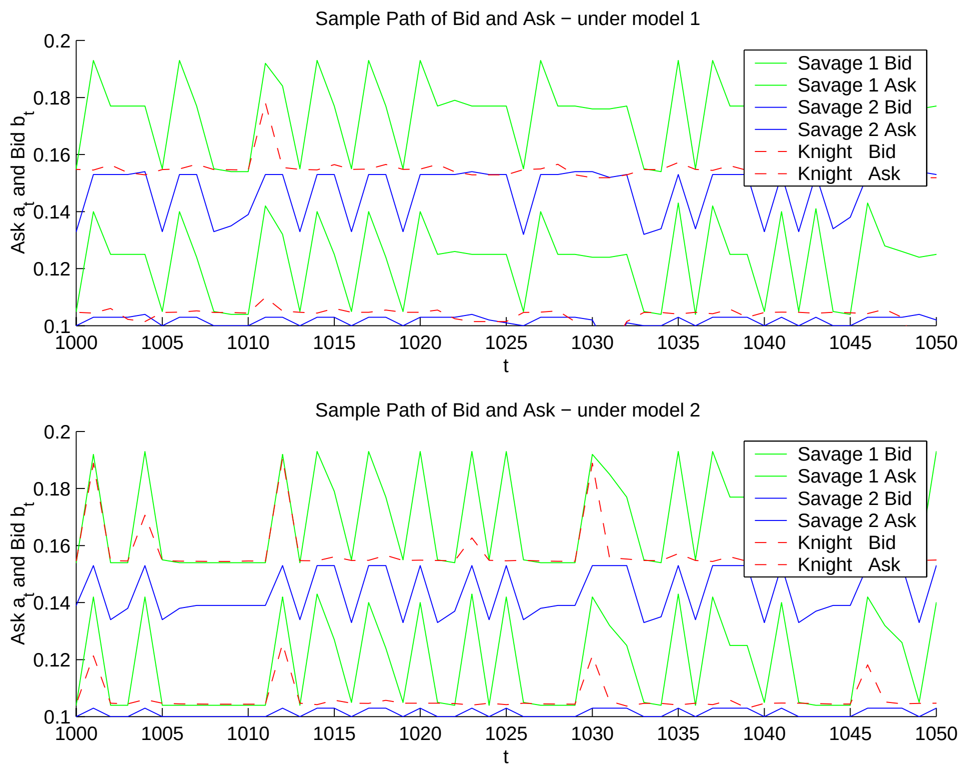

For a simulation of the economy, the optimal realized ask and bid prices are shown for three market-makers: a Savage market maker with beliefs π^1, a Savage with beliefs π^2, and a Knight market maker with uncertainty averse beliefs represented by $\Pi^K = \{\pi^1, pi^2\}$. One period of the simulation consists of drawing a price P_t, dividend δ_t, and willingness to trade, $\tilde{v}_t$. Of the 10,000 periods simulated, 50 periods are shown. The top panel is simulated under π^1 and the bottom panel is simulated under π^2 (see equation (18) for parameters.

Interestingly, it is not the case that the Knight market-maker simply adopts the "worst case" bid and ask. In other words, the derivative is not valued as a stand-alone investment according to the "worst case" distribution. Were this to be the case, the Knight market-maker would adopt the optimistic Savage's ask and the pessimistic Savage's bid.[17] However, as was discussed in the two-period portfolio choice case, uncertainty aversion manifests in behavior that is more complicated than simply "worst case."

Relative to the Savage market-makers, the uncertainty aversion of the Knight market-maker produces a less liquid market for the derivative in that the probability of trade is lower. Figure 17 shows the steady-state distribution for the likelihood of trade in any given period (based on a simulation). For all three traders, the median likelihood of trade is close to 0.5. The more optimistic Savage, π^2, has periods of higher liquidity. The Knight market-maker, has slightly lower median trade likelihood. Figure 18 shows the frequency of the position in the derivative. For all three market makers there is a higher frequency of short positions than long. This is specific to the parameterization of the example. As expected, the more pessimistic Savage market-maker, π^2, is less likely to take a long position in the derivative.

For the Knight market-maker, there is a small frequency of very low liquidity realizations. Figure 19 shows a sample path for the time-series of the probability of trade. With the Knight market-maker, the market experiences short, infrequent dips in liquidity where the probability of trade drops dramatically. As discussed previously, this drop in liquidity coincides with the case where the optimal portfolio is not described by a first-order condition. In these cases, the portfolio choice of the Knight and the bid-ask behavior are distinct from behavior under either of the Savage traders. In these times of crisis, the Knightian behavior is not representable by a "worst case" Savage trader

[17] If one is considering a long position in a call option, the worst case distribution has a low mean and low variance for the underlying asset. If one is going short a call, the worst case is a high mean and high variance.

Figure 17: Frequency for Probability of Trade in the Derivative

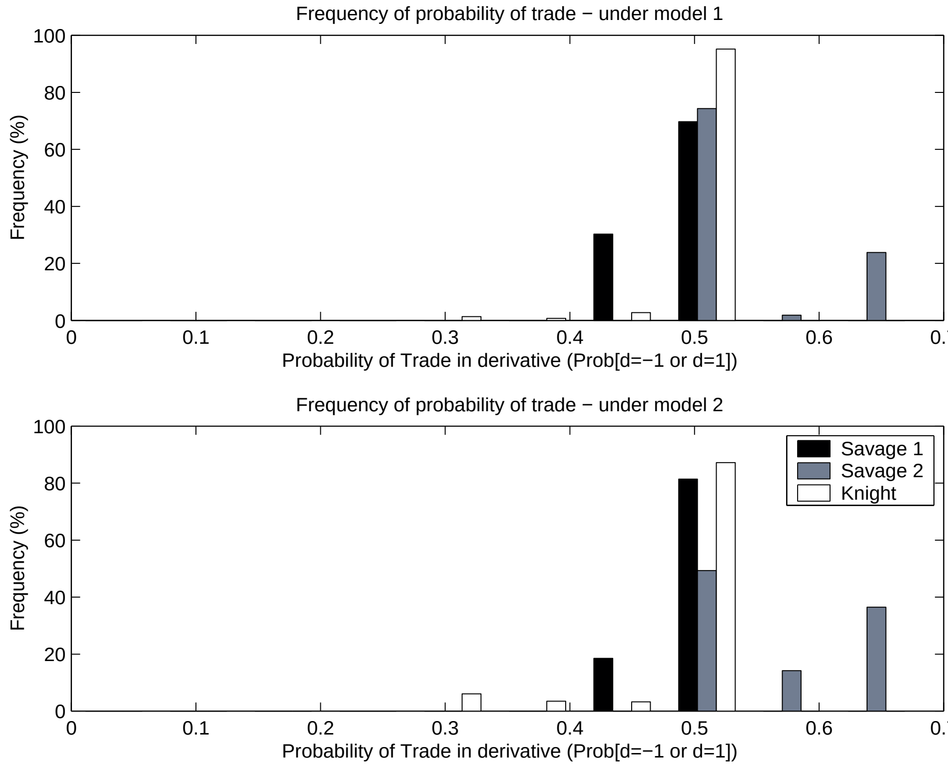

The figure is a histogram of the probability of trade based on 10,000 simulation periods. The probability of a trade is calculated as $1 - [\Phi(a_t) - \Phi(b_t)]$ and depends on the optimal ask and bid prices chosen by the three types of market makers. The three market makers are a Savage market maker with beliefs π^1, a Savage with beliefs π^2, and a Knight market maker with uncertainty averse beliefs represented by $\Pi^K = \{\pi^1, pi^2\}$. One period of the simulation consists of drawing a price P_t and a dividend δ_t The top panel is simulated under π^1 and the bottom panel is simulated under π^2 (see equation (18) for parameters.)

Figure 18: Frequency for Probability of Trade in the Derivative

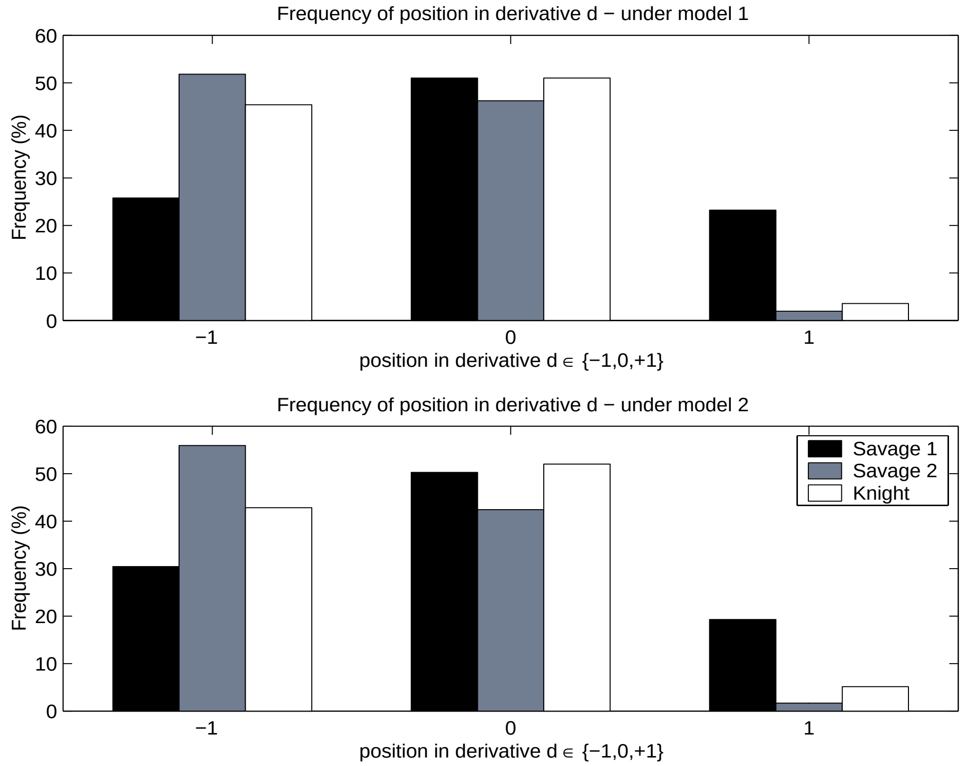

The figure is a histogram of the position in the derivative, d_t based on 10,000 simulation periods. The position in the derivative is the outcome of the optimal bid and ask prices the realization of the willingness to trade $\tilde{v}_t$. The frequency is shown for a Savage market maker with beliefs π^1, a Savage with beliefs π^2, and a Knight market maker with uncertainty averse beliefs represented by $\Pi^K = \{\pi^1, pi^2\}$. One period of the simulation consists of drawing a price P_t and a dividend δ_t The top panel is simulated under π^1 and the bottom panel is simulated under π^2 (see equation (18) for parameters.)

Figure 19: Simulated Probability of Trade in the Derivative

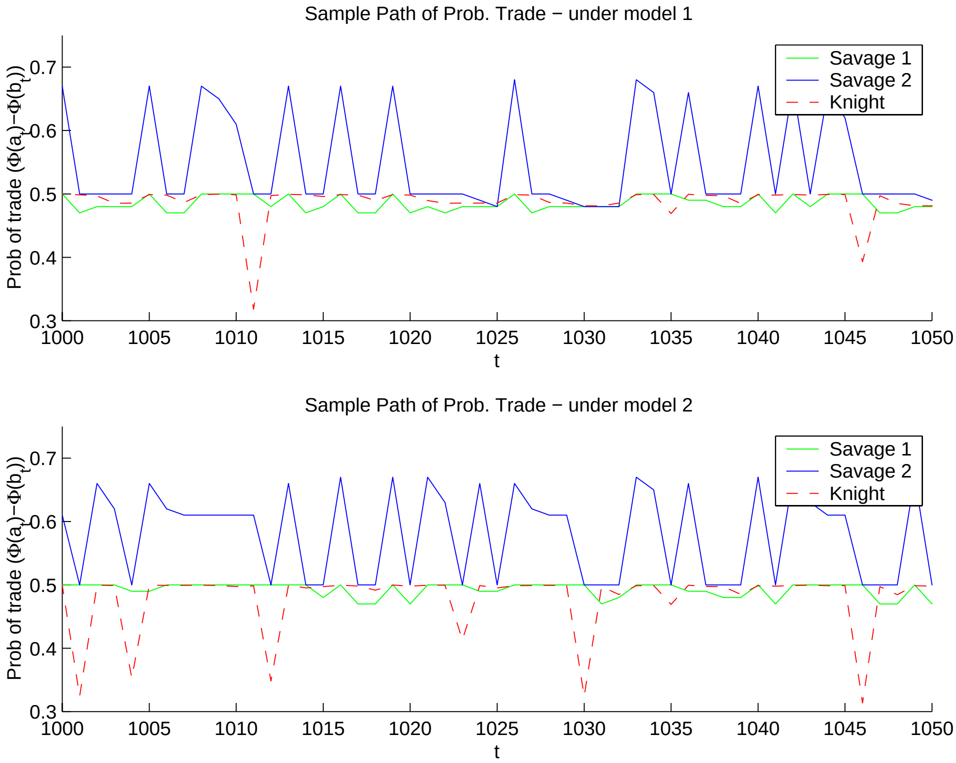

The figure is a realized path for the probability of trade. Of 10,000 simulation periods, a representative 50 periods are shown. The probability of a trade is calculated as $1 - [\Phi(a_t) - \Phi(b_t)]$ and depends on the optimal ask and bid prices chosen by the three types of market makers. The three market makers are a Savage market maker with beliefs π^1, a Savage with beliefs π^2, and a Knight market maker with uncertainty averse beliefs represented by $\Pi^K = \{\pi^1, pi^2\}$. One period of the simulation consists of drawing a price P_t and a dividend δ_t The top panel is simulated under π^1 and the bottom panel is simulated under π^2 (see equation (18) for parameters.)

In the simulation without any market-making activity shown in Figure 11, the realized path for the optimal portfolio of the Knight trader is bounded by the portfolio position of the two Savage traders. This is not surprising, given this relationship holds state-by-state (see Figure 10). In this setting, the realized path of the state variables P_t and δ_t is common across trader types. However, in the case where the trader is also a market maker, the optimal portfolio is not just a function of the exogenous state variables P_t and δ_t. It also depends on the position in the derivative, d_t. Since the bid and ask policies of the different market makers differ, the realized path in the derivative need not be common across all market-maker types. It is therefore not necessarily the case that the realized portfolio of the Knight market maker be bounded by the Savage portfolio. However, in the simulation, it is the case that $\theta_t^1 \leq \theta_t^K \leq \theta_t^2$. This is seen in Figure 20. In this entire simulation, as in the portion shown in the figure, the portfolio of the Knight market-maker is bounded by the two Savage portfolios. For completeness, we also show the frequency of portfolio holdings in Figure 21 Not surprisingly, the pessimistic Savage market-maker, π^1, typically has lower asset holdings and is more frequently short. The more optimistic market-maker, π^2, is more often long. The distribution for the Knight market-maker lies in between.

5 Conclusions

In a simple model of liquidity provision by a monopoly market maker, we have found that an aversion to Knightian uncertainty reduces liquidity. Future work will explore this connection further by extending the set of traded securities which will allow us to study spill-over effects like the "flight to quality" and "contagion" that these preliminary results suggest. We will also explore in more detail the empirical predictions of this model for financial market crises.

Figure 20: Simulated Portfolio Holdings

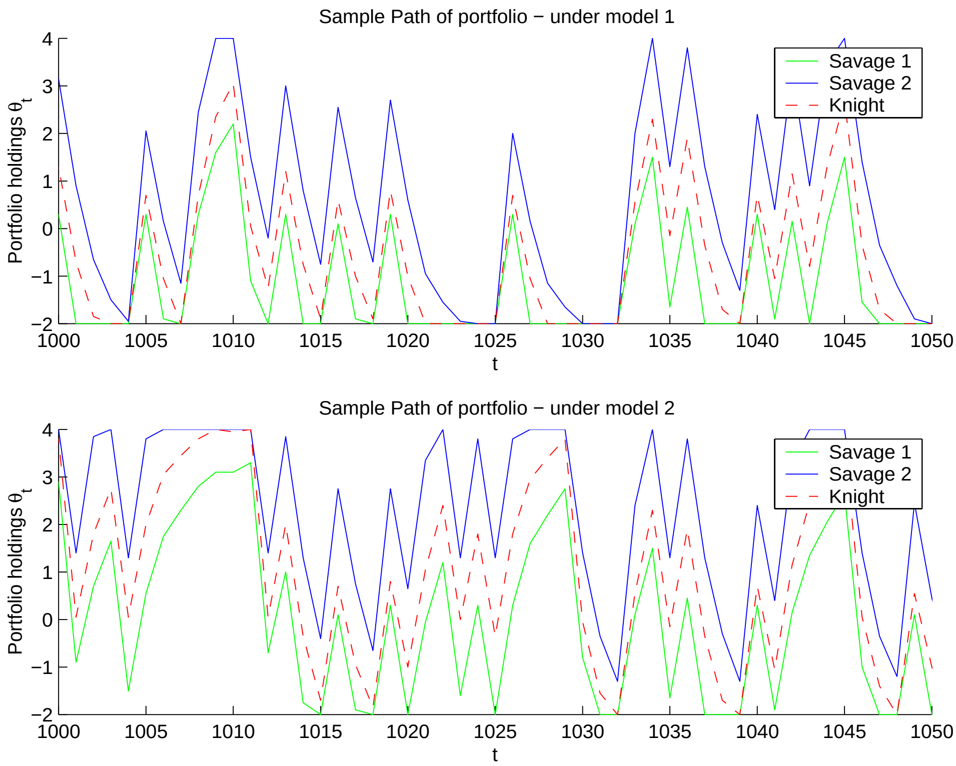

The figure is a realized path for the optimal portfolio. Of 10,000 simulation periods, a representative 50 periods are shown. The portfolio path is shown for three market makers are a Savage market maker with beliefs π^1, a Savage with beliefs π^2, and a Knight market maker with uncertainty averse beliefs represented by $\Pi^K = \{\pi^1, pi^2\}$. One period of the simulation consists of drawing a price P_t and a dividend δ_t The top panel is simulated under π^1 and the bottom panel is simulated under π^2 (see equation (18) for parameters.)

Figure 21: Frequency for Portfolio Holdings

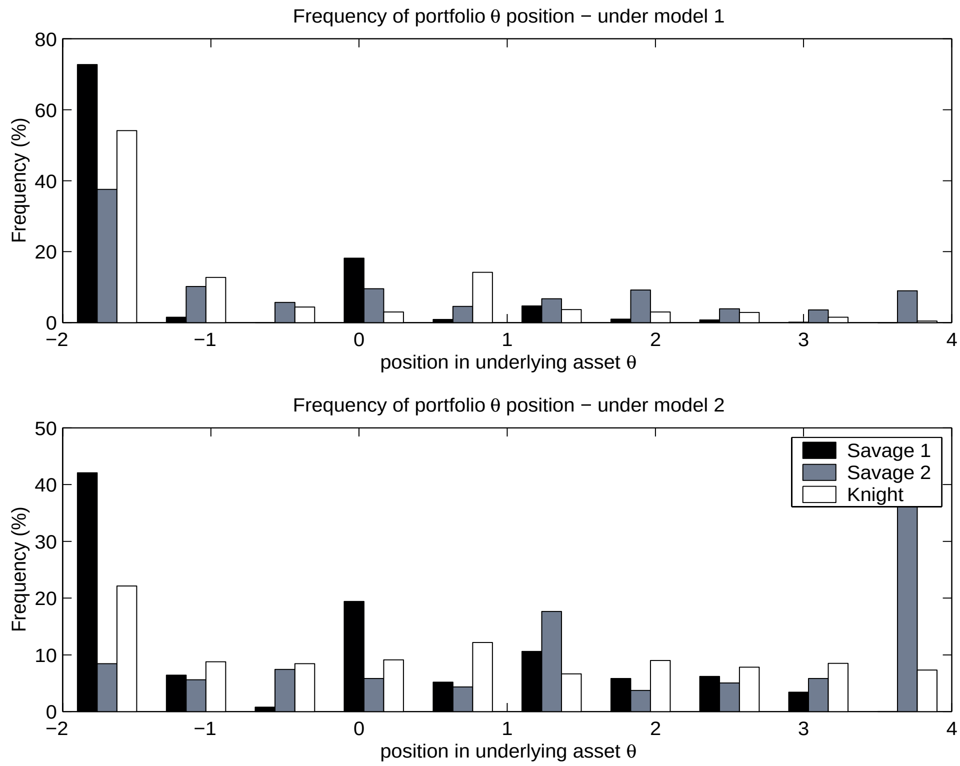

The figure is a histogram of the portfolio in the underlying asset, θ_t based on 10,000 simulation periods. The frequency is shown for a Savage market maker with beliefs π^1, a Savage with beliefs π^2, and a Knight market maker with uncertainty averse beliefs represented by $\Pi^K = \{\pi^1, pi^2\}$. One period of the simulation consists of drawing a price P_t and a dividend δ_t The top panel is simulated under π^1 and the bottom panel is simulated under π^2 (see equation (18) for parameters.)

References

Black, F., and M. Scholes (1973): "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637–659.

Choquet, G. (1955): "Theory of Capacities," *Annals de l'Institute Forier*, 5, 131–295.

Dow, J., and S. Werlang (1992): "Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio," *Econometrica*, 60(1), 197–204.

Ellsberg, D. (1961): "Risk, Ambiguity,and the Savage axioms," *Quarterly Journal of Economics*, 75, 643–669.

Epstein, L. G. (2001): "Sharing Ambiguity," *American Economic Review: Papers and Proceedings*, 91(2), 45–50.

Epstein, L. G., and T. Wang (1994): "Intertemporal Asset Pricing Under Knightian Uncertainty," *Econometrica*, 62, 283–322.

——— (1995): "Uncertainty, Risk-Neutral Measures and Security Price Booms and Crashes," *Journal of Economic Theory*, 67, 40–82.

Gilboa, I. (1987): "Expected Utility with Purely Subjective Non-Addative Probabilities," *Journal of Mathematical Economics*, 16, 65–88.

Gilboa, I., and D. Schmeidler (1989): "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18, 141–153.

Grossman, S. J., and M. H. Miller (1988): "Liquidity and Market Structure," *Journal of Finance (Papers and Proceedings)*, 43(3), 617–633.

Hansen, L. P., T. J. Sargent, and T. D. Tallarini (1999): "Robust Permanent Income and Pricing," *Review of Economic Studies*, 66, 873–907.

Ho, T., and H. Stoll (1981): "Opitmal Dealer Pricing Under Transactions and Return Uncertainty," *Journal of Financial Economics*, 9, 47–73.

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

Kyle, A. S. (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 53, 1315–1335.

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

Schmeidler, D. (1989): "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57(3), 571–587.

Scholes, M. S. (2000): "Crisis and Risk Management," *American Economic Review – Papers and Proceedings*, 90(2), 17–21.

Seidenfeld, T., and L. Wasserman (1993): "Dilation for Sets of Probabilities," *Annals of Statistics*, 21(3), 1139–1154.

Summers, L. H. (2000): "International Financial Crises: Causes, Prevention, and Cures," *American Economic Review – Papers and Proceedings*, 90(2), 1–17.

Appendix

In this appendix, we outline some new computational algorithms for solving portfolio-choice problems with aversion to Knightian uncertainty. We detail the algorithms for the standard stochastic growth model. the extension to the portfolio-choice problem in equation (6) is straightforward. Using standard notation, the model solves a Bellman equation:

$$\begin{aligned} v(k,\ z) &= \max_{k' \in \mathcal{A}(k,\ z)} \Big\{ u\left(zf(k) + (1-\delta)k - k'\right) \\ &+ \beta \min_{\mu(k,z) \in \mathcal{M}} \int v(k',\ z')\ d\mu(k,z) \Big\} . \end{aligned}$$

Our approach follow the following steps:

1. Discretize both k and z
   - k as *finely* as possible
   - z using quadrature
2. Define $\mathcal{M}$ using moment restrictions
   - $\Rightarrow$ linear in probabilities
3. Solve using linear programming
   - both optimizations are linear programs
   - use Trick and Zin (1995) constraint-generation algorithms
   - greatly reduce the number of value function evaluations

The discretized Bellman equation for this problem is given by:

$$v_{ij} = \max_{a \in \mathcal{A}_{ij}} \left\{ u_{ija} + \beta \min_{\pi_{j\cdot} \in \Pi_j} \sum_{l=1}^{n_z} \pi_{jl} v_{al} \right\}$$

1. Given probabilities $\Longrightarrow$

$$\min_v \sum_{ij} v_{ij}$$

subject to

$$v_{ij} \geq u_{ija} + \beta \sum_{l=1}^{n_z} \pi_{jl} v_{al}$$

for all i, j, and $a \in \mathcal{A}_{ij}$.

2. Given value function ordinates $\Longrightarrow$

$$\min_{\pi_{j\cdot}} \sum_{l=1}^{n_z} \pi_{jl} v_{al}$$

subject to $\pi_{j\cdot} \in \Pi_j$.

A complete algorithm to jointly solve these two optimizations simultaneously is as follows:

1. Find all of the extreme points of Π
2. Solve expected utility problem for each of these probabilities
3. Solve LP in π's for each of these value function solutions
4. Find the solutions that *agree*

The problems associated with this approach are:

- Huge number of extreme points in Π
- Full enumeration may be infeasible

We propose a heuristic algorithm based on these ideas as follows:

1. Guess value functions

- starting value tricks

2. Solve LP in π's for these value function costs
3. Given these π's solve expected utility Bellman equation by LP
4. Use these updated value function ordinates as costs and resolve the LP in π's
5. Iterate to convergence