

Model Uncertainty and Liquidity*

Bryan R. Routledge

Carnegie Mellon University

Stanley E. Zin

Carnegie Mellon University and NBER

June 2004

*We thank Burton Hollifield, Erzo Luttmer, David Marshall, and seminar participants at Boston College, Cal-Tech, Carnegie Mellon, the International Monetary Fund, Tulane, University of Vienna, University of Wisconsin-Madison, Computing in Economics and Finance, the Econometric Society World Congress, the NBER Summer Institute, Society for Financial Studies Conference Investments in Imperfect Capital Markets, and the Western Finance Association, for helpful comments. Correspondence: Bryan Routledge, Carnegie Mellon University, Pittsburgh, PA 15213-3890 <<http://sulawesi.gsia.cmu.edu>>

Model Uncertainty and Liquidity

Abstract

Extreme market outcomes are often followed by a lack of liquidity and a lack of trade. This market collapse seems particularly acute for markets where traders rely heavily on a specific empirical model such as in derivative markets. Moreover, the observed behavior of traders and institutions that places a large emphasis on “worst-case scenarios” through the use of “stress testing” and “value-at-risk” seems different than Savage expected utility would suggest. In this paper, we capture model-uncertainty using an Epstein and Wang (1994) uncertainty-averse utility function with an ambiguous underlying asset-returns distribution. To explore the connection of uncertainty with liquidity, we specify a simple market where a monopolist financial intermediary makes a market for a propriety derivative security. The market-maker chooses bid and ask prices for the derivative, then, conditional on trade in this market, chooses an optimal portfolio and consumption. We explore how uncertainty can increase the bid-ask spread and, hence, reduces liquidity. Our infinite-horizon example produces short, dramatic decreases in liquidity even though the underlying environment is stationary. We show how these liquidity crises are closely linked to the uncertainty aversion effect on the optimal portfolio. Effectively, the uncertainty aversion can, at times, limit the ability of the market-maker to hedge a position and thus reduces the desirability of trade, and hence, liquidity.

JEL Classification: G10, G13, G20

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

1 Introduction

A financial crisis, such as those that occurred in Russia, Mexico, Thailand, Indonesia, South Korea, and Brazil, is often associated with an increase in uncertainty and a decrease in liquidity. Almost by definition, a crisis involves a substantial change in financial prices so the *ex ante* likelihood of the event is low and the event is unusual. However, a crisis is also tends to be a period of higher uncertainty. For example, during the Asian financial crisis there was a pronounced increase in the dispersion of analysts' forecasts of economic growth (Prati and Sbracia (2001)). In the Russian financial crisis, a higher level of uncertainty could be attributed to the unusually large change in the swap credit spread,¹ questions about the solvency of some U.S.-based financial institutions, and the Federal Reserve Bank's unprecedented role in facilitating a recapitalization of the hedge fund Long Term Capital Management (LTCM). The second common feature of a crisis is a severe lack of liquidity. Following many financial crises, liquidity disappeared. Bid-ask spreads increased,² people had difficulty executing trades for existing financial securities,³ and new bond and equity offerings are postponed or cancelled.⁴

Many economic models can incorporate unusual "crisis" events as rare events, structural breaks, or changes in the risk premium. 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.

¹The swap credit spread increased 20 basis points (treasury bond yields versus AA rated debt yields). The increase is ten standard deviations above historic norms (Scholes (2000)).

²During the 1998 Russian crisis, bid-ask spreads on emerging market debt increased from 10-20 basis points to 60-80 basis points (International Monetary Fund (1998)). Becker, Chadha, and Sy (2000) documents a similar increase in document the increase in bid-ask spreads in foreign exchange and interbank rates following the 1997-1998 Asian crisis.

³Bank of International Settlements (1999) describe how a number of market-makers simply withdrew from trading following the 1998 Russian default and did not post quotations. Others in their survey reported that the market became "one-sided." Relatedly, *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."

⁴Most notably the Initial Public Offering of Goldman Sachs.

However, most models are unable to explain the drop in liquidity that accompanies the crisis.⁵ The puzzle is not the large change in financial prices, it is that people seem to stop trading. In this paper we investigate the connection between these uncertainty and liquidity. We investigate whether 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 uncertainty 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

⁵Marshall (2001) models a liquidity crisis as a “bad” equilibrium in coordination a game of asymmetrically informed lenders and borrowers.

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, the preferences expressed by this behavior 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 asset’s 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

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

Uncertainty and uncertainty aversion are captured by the set of probability measures Π and the “min” operator. If the set Π is a singleton, then the decision rule is the standard Savage rationality of expected utility.⁶ Note that the

⁶A 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 $P(head) + P(tail) < 1$.

set Π represents both uncertainty and uncertainty aversion. The choice axioms do not allow one to identify uncertainty and uncertainty aversion separately. In this paper, we use the recursive intertemporal formulation of uncertainty aversion of Epstein and Wang (1994) and (1995).⁷ This specification facilitates dynamic programming and preserves dynamic consistency.⁸

It may seem odd to model a market-making institution with axiomatically founded preferences. However, we proceed in this fashion for a number of practical reasons. The Gilboa-Schmeidler framework ensures that capturing a desire for model uncertainty does not introduce some other non-standard behavior. For example, Basak and Shapiro (2001) document that treating a value-at-risk concern as a constraint can lead to non-convex preferences. Second, the Gilboa and Schmeidler (1989) and Epstein and Wang (1994) approach facilitates a standard recursive formulation and is, therefore, analytically and numerically tractable. Lastly, as mentioned above, uncertainty averse preferences nest traditional expected utility preferences which are common in many other microstructure settings.⁹

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

⁷The robust control framework of Hansen, Sargent, and Tallarini (1999) is similar to the Epstein and Wang approach. In the robust control setting 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.

⁸See Epstein and Schneider (2003) for more on time consistency in models with uncertainty aversion.

⁹The setting we adopt here is similar to Ho and Stoll (1981) and related inventory-based microstructure models with a risk averse market-maker.

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 while the market for the underlying security is frictionless. 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.

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. More interestingly, our infinite-horizon example produces short, dramatic decreases in liquidity even though the underlying environment is stationary. The mechanism through which model uncertainty leads to low liquidity is the interaction of the portfolio policy and market-making activity. In some situations, model uncertainty manifests itself simply as pessimism. That is, an uncertainty averse individual is identical to a standard Savage expected utility individual with a pessimistic beliefs. However, at other times the uncertainty aversion is of first-order importance and the behavior is distinct from any expected utility agent. In this case the optimal portfolio policy is invariant to small changes in the income process like those created by a derivative position. This rigidity of the portfolio policy limits the market-maker’s ability to hedge a position in the derivative. Since taking on a naked or un-hedged position in a derivative is less appealing, the uncertainty averse market maker adjusts the

bid and ask price to reduce the likelihood of a derivative trade.

In Section 2, we lay out the basic economic environment and describe the 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. There is a frictionless market for an underlying security whose price is P_t . The market-maker sets a bid and ask price for a derivative whose payoff is $X(P_t) \geq 0$.¹⁰ Trades by the market maker are discrete short, no-trade, or long events, denoted $d_t \in \{-1, 0, 1\}$. The “size” of a trade can be incorporated into the definition of the derivative’s payoff. For concreteness, our numerical examples focus on the case of a one-period call option $X(P_t) = s \max(P_t - x, 0)$ (with x as the strike price). The parameter s determines the size or importance of each trade.

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 $\text{Prob}(d_t = -1) = [1 - \Phi(a_t)]$,

¹⁰In order to maintain the intuitive bid-ask relation, $0 < b < a$, we will only consider derivatives with non-negative payoffs.

$\text{Prob}(d_t = 0) = [\Phi(a_t) - \Phi(b_t)]$, and, $\text{Prob}(d_t = 1) = \Phi(b_t)$. For simplicity, we assume $\Phi(v)$ is atomless and define $\phi(v) = \partial\Phi(v)/\partial v$.

The exogenous random arrival of a trade request is consistent with a number of deeper microstructure models. This simple specification of the market structure for the derivative lets us focus on the effect of return-uncertainty on liquidity. Note that there is no uncertainty about the distribution governing trade arrival. To explore uncertainty about the market microstructure itself, a more detailed specification of the market would be needed.

After the arrival of the request to trade, the market maker chooses an optimal consumption and investment in a risky asset. The timing of consumption is not particularly important in our setting. We allow for consumption each period, rather than focus on the utility of terminal wealth, primarily to facilitate the stationary infinite-horizon in section 4. The timing of trade is important. The investment in the risky asset after observing trade in the derivative allows the market maker the opportunity to (at least partially) hedge the realized position in the derivative market. The effect of uncertainty aversion on the ability to hedge a derivative position turns out to be central to the main result. Lastly, note that the derivative position can never be perfectly hedged since we consider only a single risky asset available for trade. We could, of course, consider a more complicate multi-asset portfolio problem. However, for uncertainty aversion to play a role in the bid-ask of a derivative it is necessary the derivative cannot be perfectly hedged. Obviously, if the derivative can be hedged then positions in the derivative are free of uncertainty and there is no “model risk.” The incomplete market we study here captures the notion that some of the model risk institutions face is uncertainty about what constitutes the appropriate hedge.

2.1 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.¹¹ The market for the underlying asset is assumed to be frictionless. c_t is period- t consumption and ω_t is the income of the market-maker in period t . The period- t budget constraint is:

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

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 .

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

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. 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 .

¹¹The 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.

2.2 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

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

If the market maker has uncertainty or ambiguity about these probabilities, we will denote as Π the set of all such distributions. 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 since the uncertainty is only relevant if the agent's preferences are averse to ambiguity. 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:

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

where $0 < \beta < 1$ is a utility discount factor and $u(c)$ is the single period utility derived from consumption c . Note, if the set Π as a singleton, the agent adheres to Savage axioms and expected utility.

2.3 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:¹²

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

$V(\theta, d, P, \delta, y)$ 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} , 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).

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 makes a market in a derivative of the single risky asset at period zero and the derivative payoffs at period one. The single risky asset, whose prices are P_0 and P_1 , trades in a perfect market. In this section, the

¹²As in Epstein and Wang (1994), we use the recursive intertemporal formulation to preserve dynamic consistency.

dividend is zero, $\delta_0 = \delta_1 = 0$, and exogenous income, y_0 and y_1 are non-stochastic.

3.1 Portfolio Choice with Uncertainty Aversion

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, write the Bellman equation in two parts. The full market-maker problem is discussed in Section 3.2. However, before characterizing the optimal period zero bid and ask prices, we first consider the choice of the optimal consumption and portfolio. How Knightian uncertainty affects the optimal portfolio turns out to be very important to understanding the market-making problem.

The portfolio problem 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:

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

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

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

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. 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 ($d = -1$).

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, $\pi^1 \neq \pi^2$. (We consider a convex set for uncertainty below.)

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

$$-u'(w_0 - \theta^i P_0)P_0 + \beta E_{\pi^i} [u'(w_1 + \theta^i P_1)P_1] = 0 \quad (9)$$

Figures 2 and 3 plot total utility (right-hand-side of (7)) against portfolio holdings, θ , given income and prices. The utility is plotted for each distribution π^1 and π^2 and the Savage-optimal portfolios, θ^1 and θ^2 , are indicated. For exposition, the utility function is quadratic and the distribution, π^i , is summarized by its mean and variance.

The Knight agent's utility in Figures 2 and 3 is the lower envelope of the two Savage agents. There are two possibilities for how the optimal portfolio of the Knight individual is characterized. First, the aversion to uncertainty can make the Knight agent act according to the worst-case probability distribution. Here, the Knight agent simply looks like a pessimistic or "worst-case" Savage agent. Distribution π^1 is the worst-case near θ^1 in Figure 2 since $E_{\pi^1} [u(\omega_1 + \theta^1 P_1)] < E_{\pi^2} [u(\omega_1 + \theta^1 P_1)]$.¹³ In this case, $\theta^K = \theta^1$ and the optimal portfolio satisfies the first-order condition in equation (9). The definition of worst-case depends 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.

¹³This condition implies that $E_{\pi^2} [u(\omega_1 + \theta^2 P_1)] > E_{\pi^1} [u(\omega_1 + \theta^2 P_1)]$.

The definition of the worst-case distribution will also depend on the portfolio. In Figure 2 very large long or short positions in the asset switch the characterization of worst-case to π^2 (note that π^2 has a larger variance than π^1). However, for portfolios in the neighborhood of θ^K (which is equal to θ^1), distribution π^1 is always the worst-case. In contrast, Figure 3 is an example where, at the optimal Knight portfolio, there is no portfolio-free characterization of the worst-case. This occurs when

$$\begin{aligned} E_{\pi^1} [u(\omega_1 + \theta^1 P_1)] &> E_{\pi^2} [u(\omega_1 + \theta^1 P_1)] \quad \text{and} \\ E_{\pi^2} [u(\omega_1 + \theta^2 P_1)] &> E_{\pi^1} [u(\omega_1 + \theta^2 P_1)] \quad . \end{aligned} \tag{10}$$

In this situation, there is not a clear worst case distribution and the Knight agent acts like neither of the Savage agents $\theta^K \neq \theta^1$ and $\theta^K \neq \theta^2$. At θ^K , a marginal change in the portfolio alters which distribution is considered worst-case and uncertainty is of first-order importance.

When equation (10) holds, the optimal Knight portfolio is not characterized by a first-order condition. Instead, the optimal portfolio for the Knight individual occurs at the intersection of the utility calculated under the two distributions. That is, θ^K , solves

$$E_{\pi^1} [u(w_1 + \theta^K P_1)] = E_{\pi^2} [u(w_1 + \theta^K P_1)] \tag{11}$$

and θ^K lies between θ^1 and θ^2 .¹⁴

The behavior of the Knight trader in these two cases is distinct. Consider, for example, how the Knight's optimal portfolio responds to a change in initial income, w_0 . In the situation where the Knight's portfolio is identical to a pessimistic Savage portfolio it is defined by the first order condition (see Figure 2). Differentiating equation (9) implies that $0 < \frac{\partial \theta^K}{\partial w_0} < \frac{1}{P_0}$. However, in the case where there is no portfolio-free characterization of the worst-case as in Figure 3, equation (11) determines the optimal portfolio and $\frac{\partial \theta^K}{\partial w_0} = 0$. Due to

¹⁴Equation (11) may have multiple solutions. However, the optimal portfolio is the only solution that lies in the interval between θ^1 and θ^2 . This is because, for Savage individuals, the first-order condition is necessary and sufficient.

the first-order nature of the uncertainty, 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 (10) 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.

3.1.1 Convex Uncertainty

Characterizing uncertainty with a discrete set, $\Pi^K = \{\pi^1, \pi^2\}$ is without loss of generality. To see why, consider a case where Knightian uncertainty is defined by the convex set $\Pi^C = \{\alpha\pi^1 + (1 - \alpha)\pi^2 | \alpha \in [0, 1]\}$. First note that since expectations are linear in probabilities, adding the convex hull to the set of possible distributions does not alter behavior. In particular, the optimal portfolio, θ^C , for a Knight trader with this convex uncertainty, Π^C , will be identical to a Knight trader with the discrete set of models, Π^K ; that is $\theta^C = \theta^K$. However, since Π^C is convex, by the min-max theorem,

$$\max_{\theta} \left\{ u(c_0) + \beta \min_{\pi^\alpha \in \Pi^C} E_{\pi^\alpha} [u(c_1)] \right\} = \min_{\pi^\alpha \in \Pi^C} \max_{\theta} \left\{ u(c_0) + \beta E_{\pi^\alpha} [u(c_1)] \right\} \quad (12)$$

Therefore, there exists a Savage individual with beliefs $\pi^\alpha = \alpha\pi^1 + (1 - \alpha)\pi^2$ whose optimal portfolio is identical to the Knight's; that is $\theta^C = \theta^\alpha$. Figure 4 is identical to Figure 3 except utility is plotted for additional distributions. In this case the optimal Knight portfolio, θ^C , is the same as that of a Savage individual whose beliefs are given by $\alpha = 0.1$. One can not distinguish between a Savage individual with beliefs $\pi^{\alpha=0.1}$ and Knight trader with uncertainty of Π^C simply by observing the portfolio choice. However, the two traders are distinguished by other behavior. For example, the Savage individual reacts in the usual way to a change in initial income, $0 < \frac{\partial \theta^\alpha}{\partial \omega_0} < \frac{1}{P_0}$. This is not true of the Knight portfolio. As before, $\frac{\partial \theta^C}{\partial \omega_0} = 0$. When initial income changes, the optimal Knight portfolio can still be represented by a particular Savage agent's choice. However, it will be a different-looking Savage. That is, when $0 < \alpha < 1$, a change in endowment requires a changes in the beliefs if one is

to mimic the Knight portfolio with a Savage individual. As in Figure 3, the uncertainty is of first-order importance. Since the relevant aspects of Knightian uncertainty are present in the simple discrete-set characterization we focus on this specification in the remainder of the paper.

3.2 The Market Maker Problem

Given a characterization of the portfolio problem and the resulting indirect utility, the market-maker problem at period zero in the two-period model is given by

$$\max_{a,b} \left\{ [1 - \Phi(a)] U_a + [\Phi(a) - \Phi(b)] U_0 + \Phi(b) U_b \right\}. \quad (13)$$

Recall that the demand for the derivative asset is captured by the arrival of a trader with a valuation \tilde{v} with distribution $\Phi(v) = \text{Prob}(\tilde{v} < v)$. 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 will allow the market maker to obtain the future cash flows of the derivative for a low price should a low-value trader arrive, but it lowers the probability of such a trade actually arriving. In both cases, the period-zero income effect that results from a trade is offset by the period-one income effect of the derivative's payoff.

3.2.1 Example

To explore the effects of uncertainty on the bid-ask spread, we examine a numerical example 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 $s = 1$. The demand for the derivative is summarized by the arrival of a random willingness-to-trade \tilde{v} with $\Phi(v)$ as uniform 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

$$\tilde{P}_1 = \begin{cases} \mu_m + \sigma_m \\ \mu_m - \sigma_m \end{cases} \quad (14)$$

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 market-makers 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 5 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 of the derivative. The affect that this has on the bid-ask spread is depicted in Figure 6, and the affect on the probability of a trade is depicted in Figure 7. As depicted in these figures, it is possible for model uncertainty to completely eliminate the willingness of a uncertainty-averse market-maker 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.

3.2.2 Determinants of Spreads

In the preceding example, the liquidity provided by a Knight market-maker was always less than or equal to the liquidity provided by a market maker with Savage preferences. To explore why this occurs, consider the first-order

conditions for the market-maker problem in equation (13).

$$\frac{1 - \Phi(a)}{\phi(a)} \left(\frac{\partial U_a}{\partial a} \right) = U_a - U_0 \quad (15a)$$

$$\frac{\Phi(b)}{\phi(b)} \left(-\frac{\partial U_b}{\partial b} \right) = U_b - U_0. \quad (15b)$$

Denote the optimal portfolio from the solution equation (7) as θ_a^i when the market maker has received a and is short the derivative ($d = -1$) and define θ_b^i and θ_0^i analogously (for $i \in \{1, 2, K\}$). For both Savage and Knight traders $\frac{\partial U_a}{\partial a}$ and $\frac{\partial U_b}{\partial b}$ are given by

$$\frac{\partial U_a}{\partial a} = u' (y_0 + a - \theta_a^i P_0) \quad (16a)$$

$$\frac{\partial U_b}{\partial b} = -u' (y_0 - b - \theta_b^i P_0). \quad (16b)$$

For Savage preferences, equation (16) is the envelope condition. For the Knight market-maker, the envelope condition holds when the optimal portfolio is identical to a worst-case Savage and is defined by the first-order condition (9). When there is no portfolio-free characterization for the worst-case distribution as in equation (10) and Figure 3, equation (16) holds since $\frac{\partial \theta_a^K}{\partial a} = 0$ and $\frac{\partial \theta_b^K}{\partial b} = 0$. This simplifies the characterization of the optimal bid and ask prices since Knightian uncertainty only effects the right-hand-side of equation (15) indirectly through the optimal portfolio. Equation (16) is a useful property since it facilitates decomposing the ask-bid first-order condition to isolate the effects of Knightian uncertainty.

Market-making activity has implications for income at both date zero and date one (see equation (8)). Denote $U_{b=0}$ as the total indirect utility from a long position in the derivative acquired at zero cost. Similarly, we can denote $U_{a=0}$ as the total indirect utility from a short position in the derivative when no compensating period-zero ask is received. Using this notation and equation (16), we can re-write the equation (15) to decompose the first-order condition

to highlight period-zero and period-one income effects of derivative trade.

$$-\frac{1 - \Phi(a)}{\phi(a)} u' (y_0 + a - \theta_a^i P_0) + U_a - U_{a=0} = U_0 - U_{a=0} \quad (17a)$$

$$\frac{\Phi(b)}{\phi(b)} u' (y_0 - b - \theta_b^i P_0) + U_{b=0} - U_b = U_{b=0} - U_0 \quad (17b)$$

$$\underbrace{\hspace{10em}}_{\substack{\text{trade} \\ \text{arrival}}} \quad \underbrace{\hspace{3em}}_{\substack{\text{period 0} \\ \text{income}}} \quad \underbrace{\hspace{3em}}_{\substack{\text{period 1} \\ \text{income}}}$$

Consider the optimal ask price given by equation (17a). The left-hand side of (17a) captures the basic trade-off in terms of period zero income. The first term shows that a larger ask price decreases the likelihood of a trade arrival (and decreases marginal utility of period-zero consumption if a trade arrives). The second term, $U_a - U_{a=0}$, captures the benefit of a larger ask price has on utility if a trade does occur. The left-hand side of (17a) is increasing in a . The right-hand side of (17a), $U_0 - U_{a=0}$, captures only the disutility from the effect that a short position in the derivative has on period-one income. Since $X(P) \geq 0$ a short position, given no adjustment in period zero income ($a = 0$), makes the market-maker worse off ($U_0 - U_{a=0} > 0$). Since the right-hand side includes only the period-one income effect of the derivative cash-flow, it is independent of a . The decomposition for equation (17b) has an analogous interpretation. The left-hand side of (17b) increasing in the bid price, b , and the right-hand side independent of the bid, b .

Based on equation (17), all else equal, ask prices are higher when $U_0 - U_{a=0}$ is large and bid prices are lower when $U_{b=0} - U_0$ is small. In this case, spreads are large and the market is less liquid. Figure 8 plots utility as a function of the investment in the risky asset θ and is analogous to Figures 2 and 3. The solid lines depict the utility under the case where the market-maker is long the derivative at cost $b = 0$ (solid line at left), has no position in the asset (solid line in center), is short the derivative at ask price of zero, $a = 0$ (solid line

at right). The distance between the peaks determines the right-hand-side of equation (17). If the distance $U_0 - U_{a=0}$ is large, the market-maker's optimal ask will be large to offset the large utility cost of the short-position at period one. At the optimal ask, U_a (peak of dashed line on right) must lie above U_0 . At the optimal ask, $U_a - U_0 > 0$ since the market maker can choose to set the ask price arbitrarily large and drive the probability of the short position to zero.¹⁵ Similarly, if the distance between $U_{b=0}$ and U_0 is small, the benefit from being long the derivative is small and the market-maker has little room to bid aggressively for the derivative. Again, note that at the optimal bid U_b (peak of dashed line on left) must lie above U_0 .

The decomposition in (17) is also helpful since it highlights the effect of Knightian uncertainty on bid-ask spreads. Since the left-hand-sides of (17) are concerned with the impact of ask and bid prices on period zero income, uncertainty about the distribution of period one income plays no direct role. The only way the left-hand-sides of (17) can behave differently for a Savage market-maker relative to a Knight market-maker is through differences in the optimal portfolio. In contrast, since the right-hand-sides of (17) concerns period one income, uncertainty has a direct effect.

Knightian uncertainty will lower all three levels of utility in Figure 8. That is relative to the Savage market makers, $U_a^K < U_a^i$, $U_0^K < U_0^i$, and $U_b^K < U_b^i$, for $i = 1, 2$. Consider, for example, the case where the optimal portfolio for the Knight and the Savage market-makers, with no position in the derivative, yields consumption that is close to riskless, and, hence, without ambiguity. In this case, $U_0^K \approx U_0^i$ and a Knightian market-maker bid-ask spread will be larger than a Savage. That is: $a^K \geq a^i$ and $b^K \leq b^i$ (for $i = 1, 2$ Savage market makers). However, if U_0^i is affected by uncertainty, then it is possible that the Knight market-maker posts a bid or ask that is more aggressive. This occurs when the derivative position "hedges ambiguity." In particular, the difference $U_0^K - U_{a=0}^K$ may be smaller for a Knight market-maker than a Savage if the worst-case distribution in the U_0^K case differs from the case in $U_{a=0}^K$. In this

¹⁵Substitute equation (16) into (15) and not the left-hand side is positive.

case, the optimal Knight market-maker may post a more aggressive (lower) ask.

3.2.3 Market Structure

Non-zero bid-ask spreads are, of course, directly related to our assumption of a monopolist market-maker. Interestingly, the difference between bid-ask spread of a Knight market maker relative to a Savage market-maker is also closely linked to the market structure. For example, consider the effect of making a small trade in a call option. Define $X(P_1) = s \max(P_1 - x, 0)$ and let $s \rightarrow 0$.¹⁶ Since period-one income is continuous in s , indirect utility in equation (7) is continuous in s even under Knightian uncertainty. Therefore the right-hand-sides of equation (17) are both zero as $s \rightarrow 0$. Therefore, Knightian uncertainty can have no direct impact on bid-ask spreads. It will only affect spreads indirectly through the effect on the optimal portfolio.

More generally, uncertainty aversion itself cannot be the source of a bid-ask spread. Consider the case where, in addition to arbitrarily small trades in the derivative, the market-maker faces Bertrand competition. In Bertrand competition, the ask prices will simply be the compensating period zero income that offsets the disutility from a short position in the derivative security. That is the ask price, a , solves $U_a - U_0 = 0$ and the bid price, b , solves $U_b - U_0 = 0$. Figure 9 plots the indirect utility (at the optimal portfolio) as a function of the position in the derivative, d , and the size of derivative payoff, s , with $X(P_1) = s \max(P_1 - 1.5, 0)$ (the same call option as in the previous example). This plot lets us consider arbitrarily small long and short positions in the derivative. The plot shows Savage expected utility preferences under the two possible distributions. For the Savage market-maker, the indirect utility is differentiable in s . In particular, small positive and small negative positions in the derivative have an equal (but opposite) effect on utility. Therefore,

¹⁶Ask and bid prices need to be redefined to be “per unit” so that the market maker pays $s \cdot b$ to go long and gets $s \cdot a$ to go short.

bid-ask spreads are zero (i.e., $a = b$). For the Knightian-uncertain market-maker in the same setting, this is true almost everywhere. The lower line in Figure 9 represents the indirect utility of the Knightian market-maker. For small and large values of s , the Knightian indirect utility is the same as the worst case Savage (as in Figure 2). For intermediate values of s , the Knightian indirect utility is strictly lower as in the case when the uncertainty aversion is of first-order importance (see Figure 3). However, only at two points in Figure 9 is the Knightian market maker's indirect utility kinked (non-differentiable) and, hence, only at these two points does the ask price exceed the bid. These kink points occur where both (9) and (10) simultaneously occur. Since these two points are unlikely to occur (of zero measure), when markets are frictionless, Knightian uncertain preferences are not sufficient to generate a bid-ask spread.¹⁷

To apply our model to the market collapse related to the 1998 Russian bond default, for example, it is not sufficient that a market-maker like LTCM is uncertainty averse. It is also important the trades in the derivative be large and that the market-maker has some degree of market power in order for Knightian uncertainty to effect market liquidity. Interestingly, the Bank of International Settlements (1999) lists the increased concentration of market-making activity as one of the factors that amplified the crisis. Given the important interaction between market frictions and Knightian uncertainty, we leave the question of an optimal market design given uncertainty to future research.

¹⁷This is consistent with Dow and Werlang (1992a) and (1992b) that point out that in a frictionless economy, Knightian uncertainty leads to a bid-ask spread only if the asset is in zero net supply. Similarly, the indeterminacy in equilibrium described in Epstein and Wang (1994) and (1994) does not occur in representative-agent economies with an uncertain aggregate endowment (See Epstein (2001)).

3.2.4 Hedging Derivatives Positions

Figures 2 and 3 helped characterize the optimal Knight portfolio independent of the market-making activity. However, it is interesting to consider how the optimal portfolio responds to a change in the position in the derivative. Looking back at equation (17), the optimal portfolio appears in most of the terms. Therefore, Knightian uncertainty has an effect on bid-ask spreads indirectly through the optimal portfolio. Moreover, hedging derivative positions in the underlying markets is an important facet of market making. 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 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, the change in the optimal portfolio due to a trade in the derivative captures the hedging behavior of the market-maker. 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 $\theta_a^i - \theta_0^i$.

Figure 10 depicts the hedge portfolio for a short call position. For simplicity, consider the case where the short position is granted at an ask of zero ($a = 0$ as in Figure 8). In Figure 10, 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, since the position in the derivative changes which model is viewed as worst-case scenario, the hedge trade to offset the risk of a single short call position is larger than one share of the underlying asset. Moreover, with a slightly different configuration of uncertainty, Figure 11 depicts an odd situation. The short position in the call option is hedged by *reducing* investment in the underlying asset. The optimal portfolio, in response to a 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 by the solution to the first condition as in equation (9) and uncertainty is of first-order importance. 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 (11), it does not respond in a natural way to the derivative position. When the optimal portfolio is given by (11), the optimal hedge portfolio is not constrained to be either 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, $\delta_t \in \{0, 0.4\}$. The two price and two dividend states produces a four state Markov process of $(P_t, \delta_t) \in \{(0.75, 0), (0.75, 0.4), (1.25, 0), (1.25, 0.4)\}$. 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 $\Pi^K = \{\pi^1, \pi^2\}$. Both possible models, π^1 and π^2 , the are *i.i.d.* with transition probabilities for π^1 of (0.1875, 0.5625, 0.0625, 0.1875) and for π^2 of (0.1250, 0.1250, 0.3750, 0.3750). These distribution produce asset returns that are state-dependent. For distribution π^1 , a more pessimistic view, the mean (standard deviation) asset return is 0.63 (0.36) when $P_t = 0.75$ and -0.02 (0.22) when $P_t = 1.25$. The π^2 distribution is more optimistic with mean (standard deviation) asset return of 0.73 (0.40) in the $P_t = 0.75$ states

and 0.04 (0.24) in the $P_t = 1.25$ states. Other parameters used in the example are: utility is log, exogenous income is constant at $y_t = 12.750$, and $\beta = 0.8$.

4.1 Portfolio Choice

The optimal portfolio in the case where there is no market-making activity, is shown in Figure 12. In each state, the Knight portfolio policy lies between the two Savage portfolio policies. This feature is similar to the discussion of the two period example of Figures 2 and 3. Unlike the Savage portfolio, the Knight portfolio policy has a region that is flat. In this region, the uncertainty is of first-order importance since there is no portfolio-free characterization of the worst-case distribution and the optimal portfolio, θ_t^K , is independent of the asset holdings at the start of the period, θ_{t-1}^K .

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 the simulation presented in Figure 13. 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 . 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.

4.2 Market-Maker 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 14 and 15 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 , the probability for trade is slightly higher in the case where the underlying price is low ($P_t = 0.75$). Note, for both Savage market makers, the optimal bid-ask policy is not that sensitive to the position in the underlying security, θ_{t-1} (not sensitive to current wealth).

Figure 16 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 there is no portfolio-free worst-case distribution and the portfolio is not characterized by a first-order-condition. Figure 17 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 12, 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 16). This situation corresponds to the lower three panels of Figure 17 where $d_{t-1} = 1$. In particular, the lower two lines.

The interaction of the portfolio policy and market-making activity highlights the mechanism by which model uncertainty leads to low liquidity in our

model. When the uncertainty is of first-order importance, the optimal portfolio policy is invariant to changes in income. The rigidity of the portfolio policy limits the market-maker's ability to hedge a position in the derivative. Since taking on a naked or un-hedged position in a derivative is less appealing than a hedged position, the market maker adjusts the bid and ask price to reduce the likelihood of a derivative trade and, in the case of trade, generate higher compensation.

4.3 Time Series Behavior

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 results are for a simulation of 10,000 periods of the economy under both of the possible probability measures, π^1 and π^2 , that describe the evolution of prices and dividends.

Figure 18 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 Savage market-maker, π^1 , tends to have a higher bid and ask price for the derivative. 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.¹⁸ However, as was discussed in the two-period portfolio choice case, uncertainty aversion manifests in behavior that is more complicated than simply the worst case since the characterization of the worst case can depend on the portfolio choice.

¹⁸If 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.

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 19 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 20 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 21 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 uncertainty is of first-order importance and there is no portfolio-free worst-case distribution. 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.

In the simulation without any market-making activity shown in Figure 13, 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 12). 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 22. 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 23. 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, an aversion to Knightian uncertainty can significantly reduce liquidity. The market-maker in our model chooses bid and ask prices for a derivative security to optimally tradeoff the current and future income of a particular derivatives position against the probability of attracting a trade. When there is ambiguity about the appropriate probability distribution for the underlying security's future cash-flows, the market maker is uncertain about the dynamic consequences of their derivatives trading. This uncertainty increases the market-maker's bid-ask spreads and reduces liquidity. For example, in an infinite-horizon model with a stationary environment, the optimal behavior of the market maker can produce short and dramatic decreases in liquidity in situations where an expected-utility market maker would exhibit no such decrease. This occurs when the uncertainty aversion is of first-order importance. In this case the optimal portfolio policy is invariant to small changes in the income process and effectively limits the market maker's ability to hedge a derivative position. Since an un-hedged position is unattractive, the market-maker reduces the likelihood of derivative trade. Surprisingly, the uncertainty aversion does not produce markets that are always less liquid. Much of the time, the uncertainty aversion manifests itself in the form of the "worst-case" distribution. In these

situations, the bid-ask behavior is similar to a Savage market-maker – albeit a Savage with pessimistic beliefs. It is only when the uncertainty is of first-order importance, and there is no obvious characterization of the worst-case distribution, that liquidity is much smaller than with a Savage market-maker.

Much of the understanding of the infinite-horizon example is generated from a simple two-period version of the model. By characterizing the optimal bid and ask prices into their effect on trade arrival, current income, and future income, we isolate the effect of Knightian uncertainty. For example, model uncertainty and uncertainty aversion alone are not sufficient to reduce liquidity. Significant liquidity effects require an interaction between uncertainty aversion and the market structure. Non-competitive market-making, relatively large discrete trades, and an aversion to Knightian uncertainty are all necessary to generate liquidity effects that are significantly different from standard Savage expected utility models. Finally, “hedge portfolios” for the market-maker, which are an important component to understanding spreads and liquidity, can look very different from those implied by a model without Knightian uncertainty.

References

- Bank of International Settlements, 1999, “A Review of Financial Market Events in Autumn 1998,” Committee on the global financial system, Bank of International Settlements, Basel, Switzerland.
- Basak, S., and A. Shapiro, 2001, “Value-at-Risk Based Risk Management: Optimal Policies and Asset Prices,” *Review of Financial Studies*, 14, 371–405, 371–405.
- Becker, T., B. Chadha, and A. Sy, 2000, “Foreign Exchange Bid-Ask Spreads in the Asian Crisis,” Working Paper, International Monetary Fund.
- 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’Institut Fourier*, 5, 131–295.
- Dow, J., and S. Werlang, 1992a, “Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio,” *Econometrica*, 60, 197–204.
- Dow, J., and S. Werlang, 1992b, “Excess Volatility of Stock Prices and Knightian Uncertainty,” *European Economic Review*, 36, 00–00.
- 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, 45–50.
- Epstein, L. G., and M. Schneider, 2003, “Recursive multiple-priors,” *Journal of Economic Theory*, 113, 1–31.
- Epstein, L. G., and T. Wang, 1994, “Intertemporal Asset Pricing Under Knightian Uncertainty,” *Econometrica*, 62, 283–322.
- Epstein, L. G., and T. Wang, 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-Additive 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, 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, "Optimal Dealer Pricing Under Transactions and Return Uncertainty," *Journal of Financial Economics*, 9, 47–73.
- International Monetary Fund, 1998, "World Economic Outlook and International Capital Markets: Interim Assessment December 1998," World economic and financial surveys, International Monetary Fund, Washington, DC.
- 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.
- Marshall, D. A., 2001, "The Crisis of 1998 and the Role of the Central Bank," *Economic Perspectives, Federal Reserve Bank of Chicago*, 25, 2–23.
- Prati, A., and M. Sbracia, 2001, "Currency Crises and Uncertainty About Fundamentals," Working Paper, International Monetary Fund.
- Savage, L., 1954, *The Foundations of Statistics*, Wiley, New York, NY.
- Schmeidler, D., 1989, "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57, 571–587.
- Scholes, M. S., 2000, "Crisis and Risk Management," *American Economic Review – Papers and Proceedings*, 90, 17–21.

Figure 1: Model Time-Line

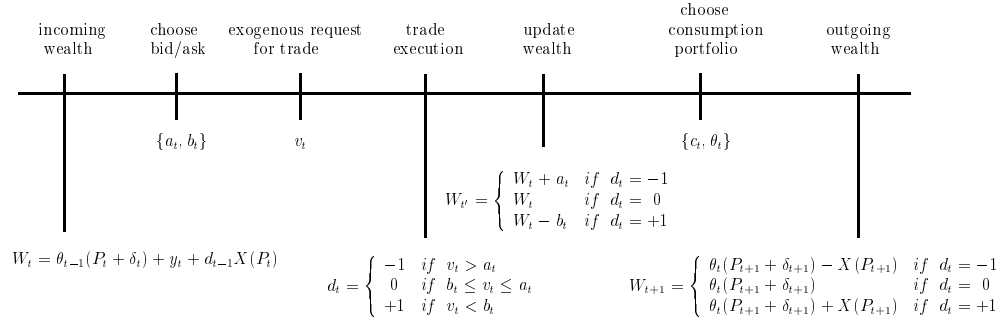
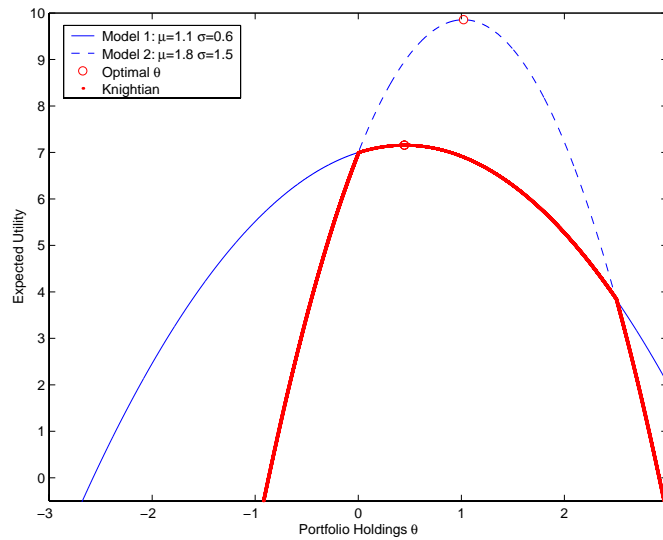



Figure 2: Optimal Portfolio: Quadratic Example 1




Note: The figure depicts the 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.

Figure 3: Optimal Portfolio: Quadratic Example 2




Note: The figure depicts the 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.

Figure 4: Optimal Portfolio: Convex Uncertainty



Note: The figure depicts the utility under several of the possible distributions $\Pi^C = \{\alpha\pi^1 + (1 - \alpha)\pi^2 | \alpha \in [0, 1]\}$ as a function of the investment in the risky asset θ . In this case, the optimal Knight portfolio, $\theta^C = \theta^K$ is the same as a Savage portfolio with $\alpha = 0.1$, $\theta^{\alpha=0.1}$

Figure 5: 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 6: Bid-Ask Spread as Uncertainty Increases

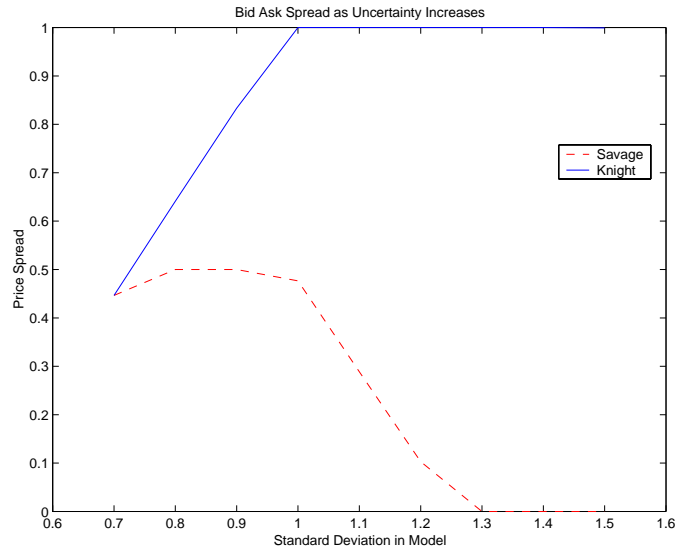


Figure 7: Probability of a Trade as Uncertainty Increases

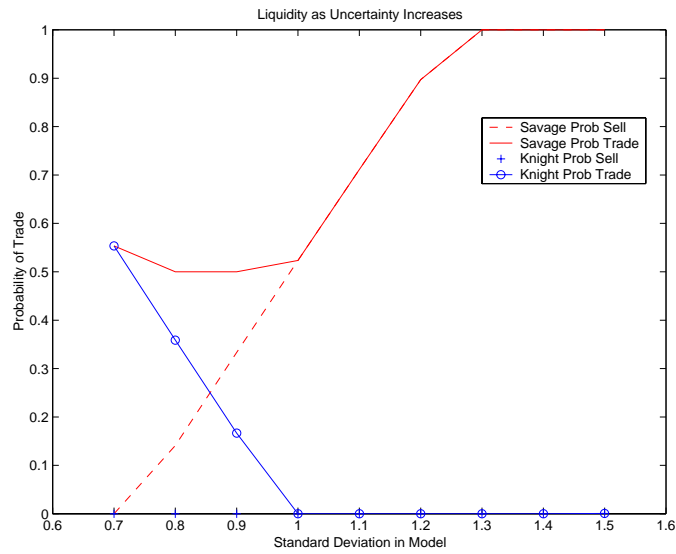




Figure 8: Utility-Based Spreads




Note: The figure depicts the utility as a function of the investment in the risky asset θ . The solid lines depict the utility under the case where the market-maker is long the derivative at cost $b = 0$ (left), has no position in the asset (center), is short the derivative at ask price of zero, $a = 0$ (right). The peak of each of these three lines determines the right-hand-side of equation (17). The dashed lines show the utility in the long (left) and short (right) cases given the optimal bid and ask price.

Figure 9: 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 utility shown is at the optimal portfolio.

Figure 10: Hedging a Short Call Position: “Natural”




Note: The figure depicts the utility from each of two distributions as a function of the investment in the risky asset θ . π^1 is the solid line, π^2 is the dashed line. For each distribution, two situations are considered $d = 0$ no position and $d = -1$ a short position in the call option. The lower envelope for each case ($d = 0, -1$) is the utility of the Knightian market-maker. The optimal portfolio θ_0^K in the case of $d = 0$, and given a short position in the call option, $\theta_{a=0}^K$ is indicated. The difference, $\theta_{a=0}^K - \theta_0^K$, is the hedge portfolio.

Figure 11: Hedging a Short Call Position: “Unnatural”




Note: The figure depicts the utility from each of two distributions as a function of the investment in the risky asset θ . π^1 is the solid line, π^2 is the dashed line. For each distribution, two situations are considered $d = 0$ no position and $d = -1$ a short position in the call option. The lower envelope for each case ($d = 0, -1$) is the utility of the Knightian market-maker. The optimal portfolio θ_0^K in the case of $d = 0$, and given a short position in the call option, $\theta_{a=0}^K$ is indicated. The difference, $\theta_{a=0}^K - \theta_0^K$, is the hedge portfolio.

Figure 12: Portfolio Policy for Two Savage Traders and a Knight Trader



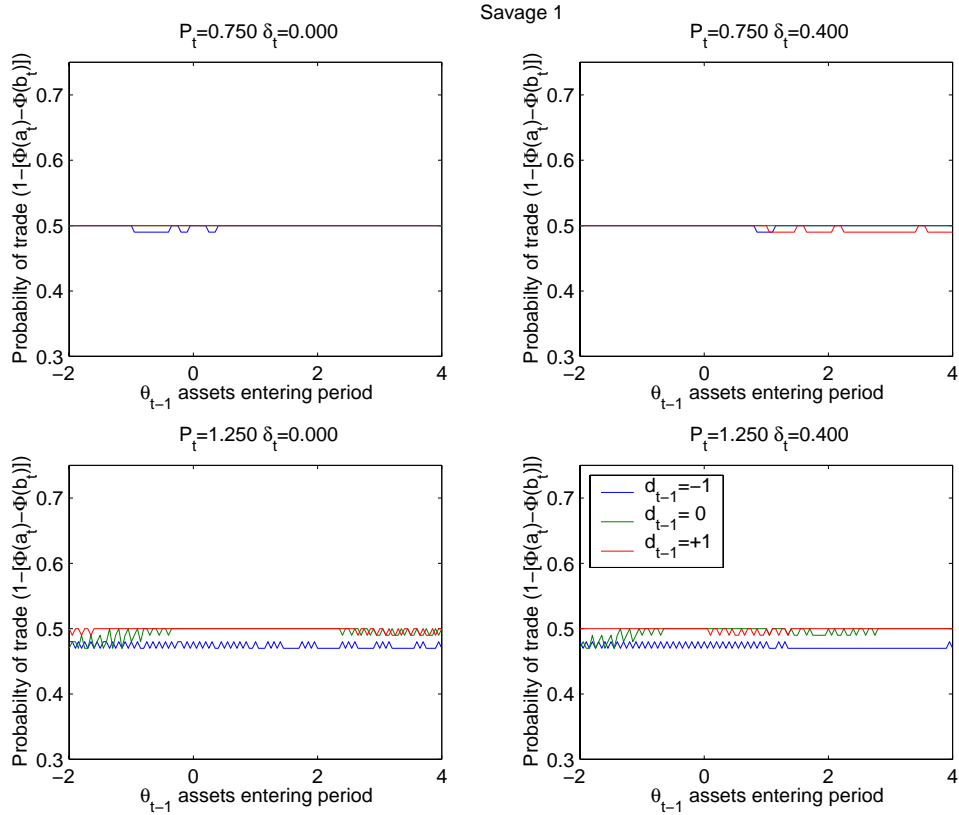
The optimal portfolio, θ^i , as a function of the previous asset holdings, is shown. Each sub-plot is a different value of the price-dividend state. The portfolio policy is shown for a Savage trader with beliefs $\Pi^1 = \{\pi^1\}$, a Savage trader with $\Pi^2 = \{p^i\}$, and an uncertainty averse Knight trader with beliefs $\Pi^K = \{\pi^1, \pi^2\}$. Beliefs are *i.i.d.* with $\pi^1 = (0.1875, 0.5625, 0.0625, 0.1875)$ and $\pi^2 = (0.1250, 0.1250, 0.3750, 0.3750)$.

Figure 13: 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 .

Figure 14: Probability of Trade - Savage π^1



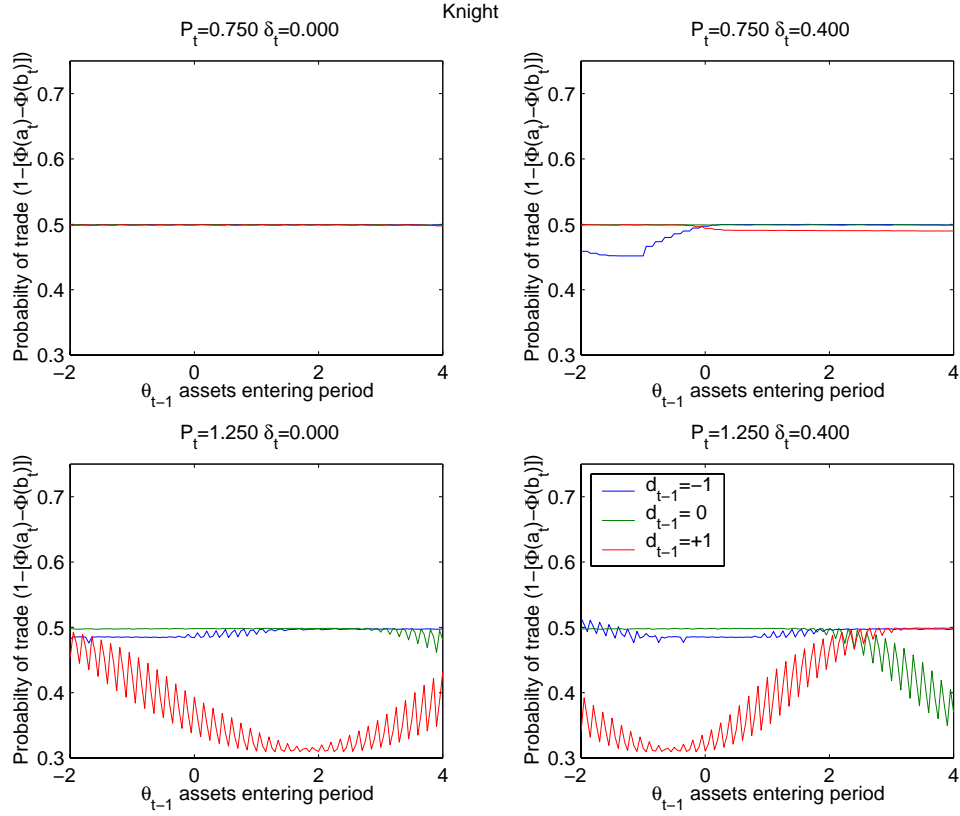
The figure shows the probability of a trade occurring given the bid and ask policy that solves equation (6) for the Savage market-maker with beliefs π^1 . The probability of a trade is calculated as $1 - [\Phi(a_t) - \Phi(b_t)]$. The probability of a trade is a function of the state variables: previous 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: Probability of Trade - Savage π^2



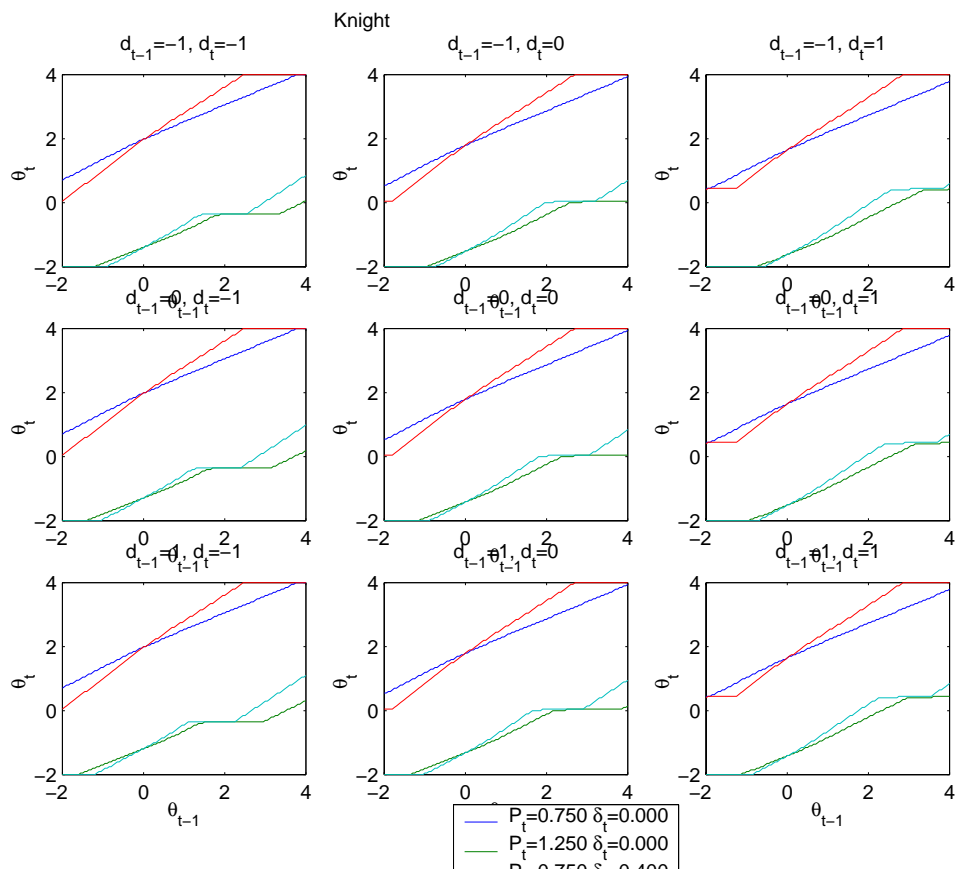
The figure shows the probability of a trade occurring given the bid and ask policy that solves equation (6) for the Savage market-maker with beliefs π^2 . The probability of a trade is calculated as $1 - [\Phi(a_t) - \Phi(b_t)]$. The probability of a trade is a function of the state variables: previous 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 16: Probability of Trade - Knight



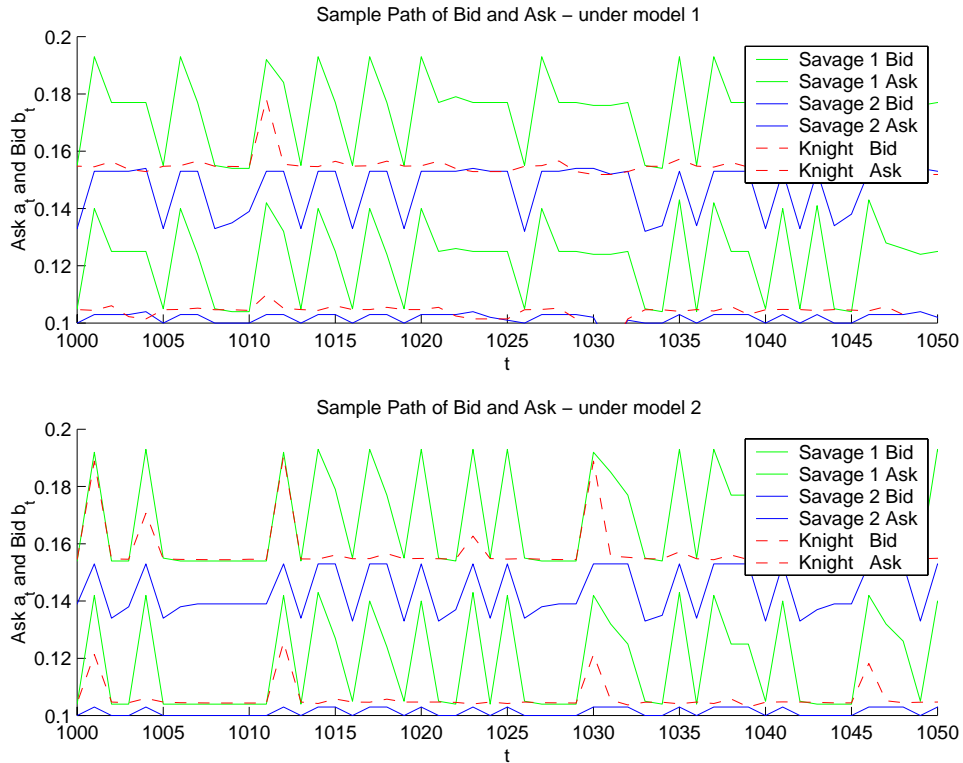
The figure shows the probability of a trade occurring given the bid and ask policy that solves equation (6) for the Knight market maker with uncertainty aversion represented by $\Pi^K = \{\pi^1, \pi^2\}$. The probability of a trade is calculated as $1 - [\Phi(a_t) - \Phi(b_t)]$. The probability of a trade is 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 17: Optimal Portfolio - Knight π^2



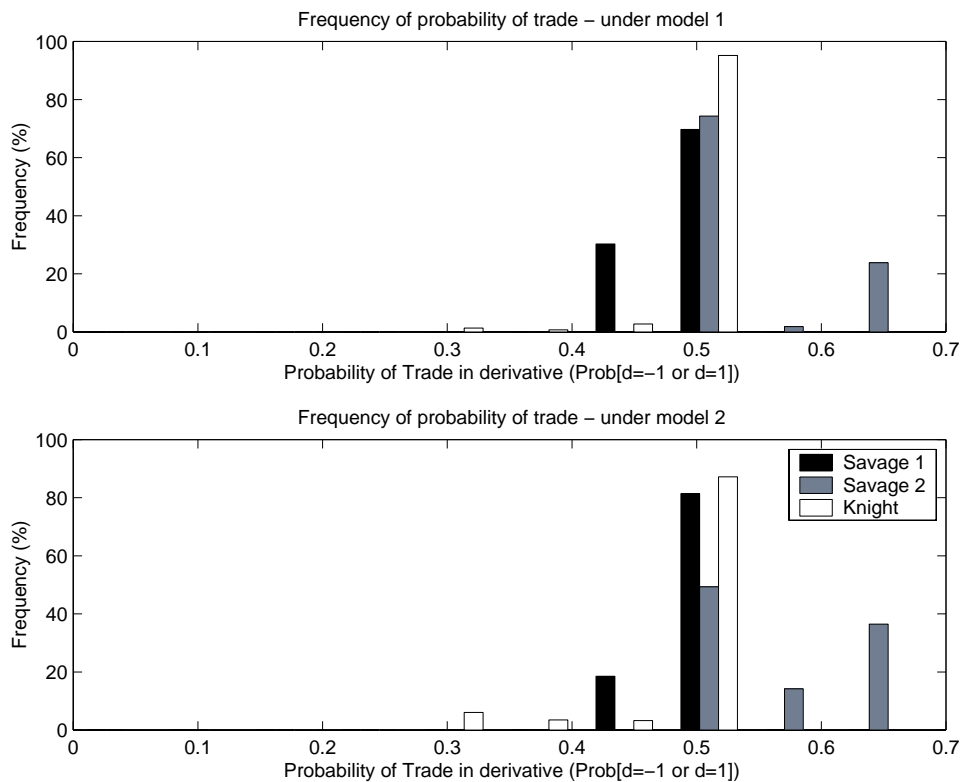
The figure shows the optimal portfolio that solves equation (6) for the Knight market maker with uncertainty aversion represented by $\Pi^K = \{\pi^1, \pi^2\}$. The optimal portfolio is a function of the state variables: current, d_t , and previous, d_{t-1} position in the derivative, previous position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

Figure 18: 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 .

Figure 19: Frequency for Probability of Trade in the Derivative



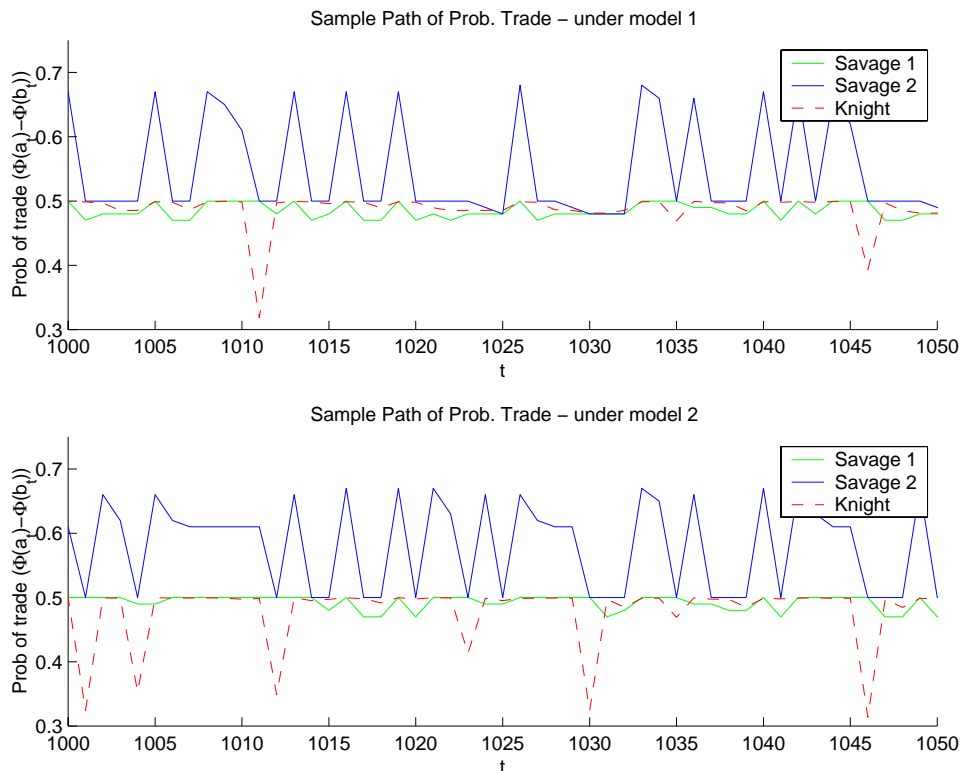
The figure is a histogram of the probability of trade based on 10,000 simulation periods. The probability of a trade is calculated as $1 - [\Phi(a_t) - \Phi(b_t)]$ and depends on the optimal ask and bid prices chosen by the three types of market makers. The three market makers are a Savage market maker 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 .

Figure 20: Frequency for Probability of Trade in the Derivative



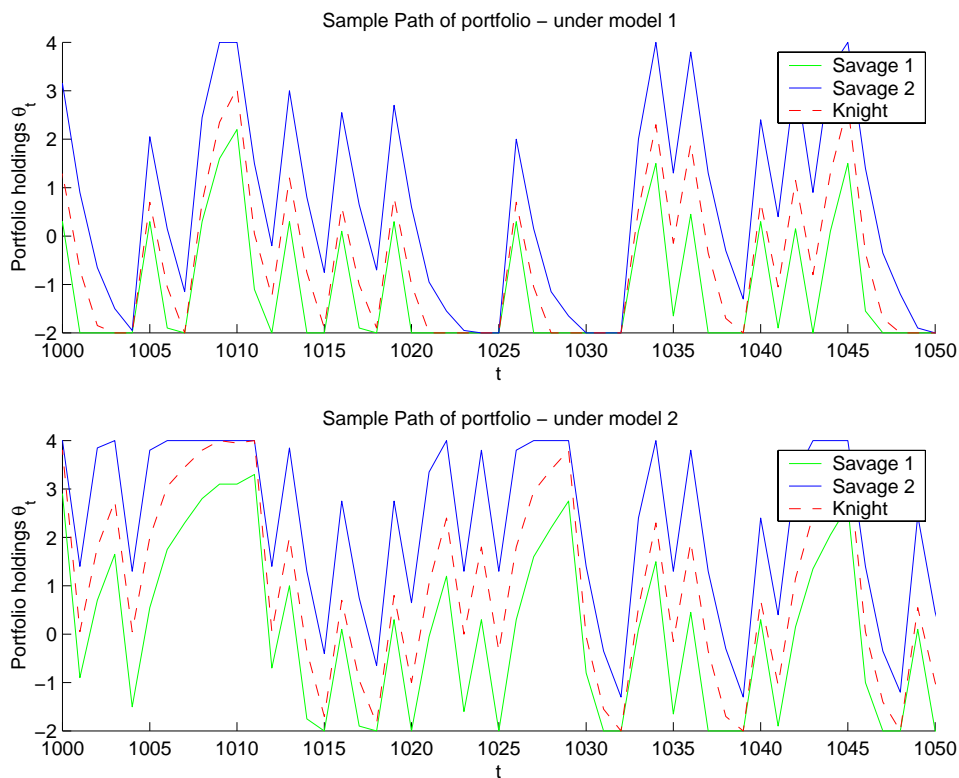
The figure is a histogram of the position in the derivative, d_t based on 10,000 simulation periods. The position in the derivative is the outcome of the optimal bid and ask prices the realization of the willingness to trade \tilde{v}_t . The frequency is shown for a Savage market maker with beliefs 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 .

Figure 21: 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 .

Figure 22: 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 .

Figure 23: 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 .