STEPHEN MORRIS

RISK, UNCERTAINTY AND HIDDEN INFORMATION

ABSTRACT. People are less willing to accept bets about an event when they do not know the true probability of that event. Such “uncertainty aversion” has been used to explain certain economic phenomena. This paper considers how far standard private information explanations (with strategic decisions to accept bets) can go in explaining phenomena attributed to uncertainty aversion. This paper shows that if two individuals have different prior beliefs about some event, and two sided private information, then each individual’s willingness to bet will exhibit a bid ask spread property. Each individual is prepared to bet for the event, at sufficiently favorable odds, and against, at sufficiently favorable odds, but there is an intermediate range of odds where each individual is not prepared to bet either way. This is only true if signals are distributed continuously and sufficiently smoothly. It is *not* true, for example, in a finite signal model.

KEY WORDS: Subjective probability, uncertainty aversion, asymmetric information, willingness to bet.

> The old-fashioned way of measuring a person’s beliefs is to propose a bet and see what are the lowest odds which he will accept. This method I regard as fundamentally sound. ... (But) *the proposal of a bet may inevitably alter his state of opinion*.[1]

> Another shortcoming of the definition – or of the device for making it operational – is the possibility that people accepting bets against our individual have better information than he has (or know the outcome of the event considered). This would bring us to game theoretic considerations.[2]

1. INTRODUCTION

In Savage’s (1954) theory of choice under uncertainty, rational individuals behave as if they are maximizing the expected value of some utility function over outcomes, with respect to some probabilities. We are not encouraged to ask where the probabilities come from, or if the individual ‘knows’ the ‘true’ probabilities. This is not relevant to the argument.

Theory and Decision **42:** 235–270, 1997.

Much criticism of this *subjective expected utility* (SEU) approach has focused on the idea that an individual's knowledge of true probabilities does matter. In particular, it is an implication of the SEU approach that our individual should be prepared to assign a probability to any event and accept a bet either way on the outcome of that event at odds actuarially fair given his probability of that event. Yet both introspection and some experimentation[3] suggest that most people are prepared to do so only if they know the true probability. This distinction is potentially important in explaining economic phenomena. Indeed, there is an old tradition making a distinction between *risk* (when objective probabilities are known) and *uncertainty* (when objective probabilities are not known).[4] More recently, generalizations of expected utility theory have been proposed to capture the distinction.[5] 'Uncertainty aversion', over and above 'risk aversion', has been used to explain a number of apparently puzzling economic phenomena, including aspects of 'entrepreneurship',[6] financial asset pricing,[7] and non-rationalizable play in games.[8]

It is argued in this paper that the phenomena attributed to 'uncertainty', or unknown objective probabilities, are consistent with the traditional subjective expected utility approach with private information. Naturally, our individual is hesitant to bet either way given his subjective beliefs, if he believes that the person he is betting against possesses some superior information.[9] Notice that if an individual does not know the true probability, there is certainly some potentially valuable information which he does not possess. This is tautologically true: he does not know the true probability. Notice also that if the individual is making a real bet, there is someone who stands to gain at our individual's expense, if she had access to that information. Thus all we need to argue, in order to rationalize apparent uncertainty aversion in betting as a consequence of private information, is that our individual assign some probability to the person he is betting against knowing more about the true probability than he does. This does not seem to be a very stringent requirement.

The motivation for this reinterpretation is discussed in detail in Section 3. It is argued that proponents of subjective expected utility have always understood that a 'bid–ask spread' in rational individuals' willingness to bet is consistent with SEU maximization in the presence of private information. It is true that it is possible to imagine environments where strategic considerations are ruled out, and our individual nonetheless displays uncertainty aversion. However, it is argued that such situations are unlikely to be economically relevant.

The core of the paper (in Section 2) is a formal model of bid-ask spreads in betting prices for subjective expected utility maximizers with two sided private information. Suppose two individuals with different prior beliefs about some event each observe some signal correlated with the event. Suppose some outside agent specifies a 'price', or betting odds, at which the individuals may bet with each other (the choice of price does not reveal any information). Given that their different priors and the distribution of the signals are common knowledge, and that they both know that a bet will be implemented only if both accept the bet, at which odds will they bet against each other? In particular, is it the case that each individual is prepared to accept a bet for an event at any odds sufficiently favorable to him, is prepared a bet against at any odds sufficiently favorable to him, but is not prepared to accept a bet either way at some intermediate range of odds?

This paper gives a condition on the distribution of individuals' signals which ensures an affirmative answer to this question. But it is useful to see first what might go wrong. Suppose I am considering betting for some event. As the odds move in my favor, I should, other things being equal, be keener to accept the bet. But, with two sided private information, as the odds move in my favor, my opponent will be more cautious about accepting a bet with me (remember the *distributions* of signals were taken to be common knowledge). In particular, she will accept the bet only on signals which make the event I am betting for unlikely. Thus I must trade off my improving odds against the increasingly bad news implicit in my opponent's willingness to accept a bet against me.

In Section 2.3, I give an example with finite signals where the bad news implicit in improving odds outweighs the improved odds themselves. In fact, for any finite distribution of signals, we can find priors for the individuals such that this phenomenon occurs. There will always be a situation where an infinitesimal improvement in my odds will lead my opponent to stop accepting the bet on some signal more favorable to me than the other signals where he is accepting the bet. The discrete change in his willingness to accept the bet will outweigh my infinitesimally improved odds.

Individuals will exhibit the bid ask spread behavior that is said to characterize uncertainty aversion if signals are continuous and sufficiently smooth. The rate of change of the increasing *marginal* likelihood ratio must be no greater than the rate of change of either of two measures of the *average* likelihood ratio. In Section 2.2, I show that the bid and ask prices, which exist under the smoothness condition, satisfy the comparative static properties we might expect. Thus, as my prior probability of an event increases, the less favorable odds I require to accept a bet for and the more favorable odds I require to accept a bet against. As *your* prior probability of an event increases, the more favorable odds I require to accept a bet for, since in that case there is more bad news implicit in your willingness to bet against me. Conversely, as your prior probability of an event increases, the less favorable odds I require to accept a bet against, since there is then less bad news implicit in your willingness to bet against me (and for the occurrence of the event). Finally, it is shown that under natural conditions, improved information for either you or me will tend to widen my bid ask spread in both directions.

There are a number of papers which address issues related to those in this paper. Leamer (1986) discusses a number of alternative explanations for observed bid ask spreads for subjective probabilities. If individuals are actually choosing the prices of the bets (unlike in this paper), there will be strategic reasons to report different buying and selling prices for bets, even if they are in fact indifferent between betting for or against at some odds. He also discusses a 'winner's curse' in betting which is implicit in the model studied in this paper.

Morris (1993) provides a more general model of when two sided private information leads to bid ask spreads. Existing informational models of bid ask spreads in financial asset markets – e.g. Glosten and Milgrom (1985) – consider environments where all transactions take place between an uninformed market maker and possibly informed traders, so that the model in effect deals with only one sided private information. Shin (1991, 1992, 1993) uses a framework similar to Glosten and Milgrom in an explicitly betting environment: the horse racing market in the U.K. Again, the assumption is that there is one sided private information and noise traders. In all these papers, comparative static properties of the bid ask spread can be studied. In Shin (1992), it is shown that the bid ask spread will be higher for low probability bets, because a punter's willingness to bet for a low probability event conveys more information than his willingness to bet for a high probability event.

This paper describes a particular institution by which individuals with heterogeneous prior beliefs and private information might bet with each other in ways representable by bid ask spreads. A more general approach to this problem (allowing for more general trading institutions) is studied in Morris (1994).

In this paper, the game where individuals decide whether to accept or reject a bet typically has multiple equilibria. In order to carry out comparative statics, or even identify the bid ask spread, it is necessary to select among the different equilibria of the betting game. But because there are *strategic complementarities* in individuals' willingness to accept or reject a bet, there is a natural 'largest' Nash equilibrium, which is used in the analysis of this paper. Morris (1992) shows that strategic complementarities allow us to carry out comparative statics in a more general class of 'acceptance games'.

The outline of the paper is as follows. The general informational model of bid ask spreads for subjective probabilities with two sided private information is presented in Section 2.2. In Section 2.1, an example is presented where it is possible to explicitly solve for bid ask spreads as a function of all the parameters. This section enables the game theoretic solution concept to the problem to be introduced. Section 3 contains a more complete discussion of how important hidden information is in explaining phenomena attributed to uncertainty aversion, or unknown probabilities. Section 4 concludes with a discussion of the significance of the assumptions driving the results of this paper.

2. AN INFORMATIONAL MODEL OF BID–ASK SPREADS FOR SUBJECTIVE PROBABILITIES

There is a 'bid–ask' spread for subjective probabilities for some individual if [1] he is prepared to accept a bet for an event at some odds sufficiently favorable to him, [2] he is prepared to bet against at some odds sufficiently favorable to him, but [3] there is also a range of odds in between where he is not prepared to accept a bet either way. A formal model explaining the existence of such bid ask spreads in the presence of two sided private information is presented in this section. The format of the section is as follows. The basic environment of two individuals, with possibly different prior beliefs, considering a bet about some event, is introduced first. Then in Section 2.1 an example is presented where each individual observes a particular continuous signal. This enables us to motivate the basic conceptual framework of the paper by means of an example where we can explicitly derive the bid–ask spread as a function of the parameters. Section 2.2 presents the model for general continuous signals. The results of Section 2.2 rely on continuous signals. As argued in the introduction, if there were only a finite set of signals, then generically the bid and ask probabilities of our individuals will be badly behaved. Moving odds in favor of an individual may make him *less* willing to bet, because of informational externalities. An example where this occurs is presented in Section 2.3.

Suppose two individuals 1 and 2 have possibly different beliefs about the outcome of some event – say, whether a Republican will win the U.S. Presidential election in 2008. Individual *i*'s *prior* belief is $\pi_i \in (0,1)$. They are considering making a bet about the possible outcome. Say that one individual, 'F', is considering betting for a Republican victory, while the other individual, 'A', will bet against. Let us represent the odds at which they are to bet by a probability $p \in (0,1)$. If a Republican wins, individual A must pay individual F \$(1-p), while if a Republican fails to win, individual F must pay individual A \$$p$. Notice that an alternative interpretation is that \$$p$ is the 'price' of \$1, contingent on a Republican victory. In the absence of any private information, and thus any strategic considerations, an individual will be prepared to accept the bet for if $\pi_i \geq p$ and against if $\pi_i \leq p$. Thus a bet will occur between individual F and individual A if $\pi_F \geq p \geq \pi_A$.

It is useful to introduce additional notation for the *ratios* of probabilities. Write π_i^* for $\pi_i/(1-\pi_i)$ and $p^* = p/(1-p)$. Thus 'prior odds ratio' π_i^* is the ratio of likelihoods that individual i assigns to a Republican victory and a non-Republican victory. It represents the ratio of marginal value to him of dollars in those states. The 'betting odds ratio' p^* is the relative price of \$1 contingent on Republican victory with respect to \$1 contingent on a non-Republican victory. Thus the uninformed individual will be prepared to bet for Republican victory if $\pi_i^* \geq p^*$. Throughout the paper, x^* is used to represent $x/(1-x)$ for any variable, x, ranging from 0 to 1.

2.1. *A Continuous Signals Example*

Now suppose each individual observes some private information. Each individual observes an independent but identically distributed signal in the interval [0,1]. The density of the signal on the interval [0,1] is $2s$, conditional on a Republican victory, while it is $2(1-s)$ conditional on a non-Republican victory.

Let us calculate when the individuals would accept the bet, if they failed to take into account the other individual's willingness to bet. Individual i's probability of Republican victory, divided by the probability of a non-Republican victory, conditional on observing signal s, would be $\pi_i^* s/(1-s)$. Thus individual F would accept the bet for, after observing signal s, if $\pi_F^* s/(1-s) \geq p^*$. Individual A would accept the bet against, after observing signal s, if $\pi_A^* s/(1-s) \leq p^*$.

But in practise a bet will be implemented only if both individuals are prepared to accept it. Suppose each individual must decide simultaneously (and irrevocably) whether to accept. If both accept, the bet is sealed. If one of them rejects the bet, no bet is implemented (and there is no renegotiation). Formally, this is a game of incomplete information, where each individual's strategy is to accept or reject the bet, contingent on his signal. The terms of the bet (who is for and against, and the betting odds ratio p^*) are taken as given in the analysis of the game.[10]

Any non-trivial Nash equilibrium of this game will have the property there is some critical pair of signals $\bar{s}$ and $\underline{s}$, such that individual F will accept the bet only if he observes signal $\bar{s}$ or higher, while individual A will accept the bet only if he observes signal $\underline{s}$ or lower.[11] Therefore, in searching for Nash equilibria, we can focus attention on simple strategies of this form.

So suppose that individual A will accept the bet only if he observes signal $\underline{s}$ or lower. What is individual F's best response? We need to calculate individual F's probability that a Republican will win the election, divided by the probability that a non-Republican will win, conditional on F observing signal s_F and A observing a signal $s_A \leq \underline{s}$. Individual F will accept the bet if this ratio is at least as great as the betting odds ratio p^*, i.e. if

$$(1) \qquad \frac{\text{prob(Republican win)}}{\text{prob(non} - \text{Republican win)}} =$$

$$\frac{\pi_F 2 s_F \int_0^{\underline{s}} 2 s_A ds_A}{(1-\pi_F) 2(1-s_F) \int_0^{\underline{s}} 2(1-s_A) ds_A} = \pi_F^* \frac{s_F}{1-s_F} \frac{\underline{s}}{2-\underline{s}} \geq p^*.$$

This expression can be rewritten as

$$(2) \qquad s_F \geq \frac{1}{1 + (\pi_f^*/p^*) \cdot [(\underline{s}/2 - \underline{s})]}.$$

Thus we can represent F's *reaction function* as follows:

$$(3) \qquad \rho_F(\pi_F^*, p^*, s) = \frac{1}{1 + (\pi_F^*/p^*) \cdot [(\underline{s}/2 - \underline{s})]}$$

with the interpretation that F accepts the bet if and only if his signal $s_F \geq \rho_F(\pi_F^*, p^*, \underline{s})$. Notice that ρ_F is decreasing in $\underline{s}$, so that the more willing A is to accept the bet, the more willing F is to accept it too. If A is always prepared to accept the bet (i.e. if $\underline{s} = 1$), $\rho_F = p^*/(p^* + \pi_F^*)$, so that F is prepared to accept 'naively', based on his own observation alone. If $\underline{s}$ is small, so A is prepared to accept only on very bad news about Republican victory, F will be prepared to accept only if he has observed very good news, i.e. ρ_F is close to 1. As $\underline{s}$ tends to 0, ρ_F tends to 1.

We can perform the symmetric calculation for individual A betting against Republican victory: suppose that individual F will accept the bet only if he observes signal $\bar{s}$ or higher. What is individual A's best response? We calculate individual A's probability that a Republican will win the election, divided by the probability that a non-Republican will win, conditional on A observing signal s_A and F observing a signal $s_F \geq \bar{s}$. Individual A will accept the bet if this ratio is no higher than the betting odds ratio p^*, i.e. if

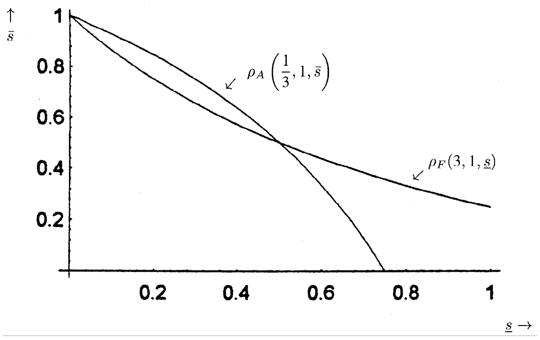

Fig. 1.

$$(4) \qquad \frac{\text{prob(Republican win)}}{\text{prob(non} - \text{Republican win)}} =$$

$$\frac{\pi_A 2 s_A \int_{\bar{s}}^{1} 2 s_F ds_F}{(1-\pi_A) 2 (1-s_A) \int_{\bar{s}}^{1} 2(1-s_F) ds_F} = \pi_A^* \frac{s_A}{1-s_A} \frac{1+\bar{s}}{1-\bar{s}} \leq p^*.$$

This expression can be re-written as

$$(5) \qquad s_A \leq \frac{1}{1 + (\pi_A^*/p^*) \cdot [(1+\bar{s})/(1-\bar{s})]}.$$

Individual A's *reaction function* is:

$$(6) \qquad \rho_A(\pi_A^*, p^*, \bar{s}) = \frac{1}{1 + (\pi_A^*/p^*) \cdot [(1+\bar{s})/(1-\bar{s})]}$$

with the interpretation that *A* accepts the bet if and only if his signal $s_A \leq \rho_A(\pi_A^*, p^*, \bar{s})$.

A Nash equilibrium of this game will now be a pair $(\underline{s}, \bar{s})$ satisfying $\rho_F(\pi_F^*, p^*, \underline{s}) = \bar{s}$ and $\rho_A(\pi_A^*, p^*, \bar{s}) = \underline{s}$. Notice first that there is *always* a Nash equilibrium where $\underline{s} = 0$ and $\bar{s} = 1$.[12] This is essentially equivalent to the Nash equilibrium where neither individual ever accepts the bet (since it is a zero probability event that $s_F = 1$ and $s_A = 0$). But does there exist any other Nash equilibrium where betting does occur? Figure 1 plots the reaction functions for the case where $\pi_F^* = 3$, $p^* = 1$, $\pi_A^* = 1/3$. Here there is a Nash equilibrium when $\bar{s} = \underline{s} = 1/2$. *F* accepts the bet if he observes a signal greater than or equal to 1/2; *A* accepts if he observes a signal less than or equal to 1/2.

It can be verified that if $\pi_F^* \leq 4\pi_A^*$, there does not exist a Nash equilibrium where betting occurs (i.e. a pair $(\bar{s}, \underline{s}) \neq (1, 0)$ solving the mutual best response property). Not only must π_F^* be greater than π_A^* in order for betting to occur, it must be sufficiently greater (i.e. by a factor of four) to overcome the asymmetry of information. If $\pi_F^* > 4\pi_A^*$, there is exactly one such Nash equilibrium where,

$$\text{(7)} \qquad \bar{s} = \frac{p^* + 2\pi_A^*}{\pi_F^* + p^* - 2\pi_A^*}, \qquad \underline{s} = \frac{\pi_F^* - 4\pi_A^*}{\pi_F^* - 2\pi_A^* + \frac{\pi_F^*\pi_F^*}{p^*}}.$$

Let us adopt the convention that if such a Nash equilibrium exists, it is *this* Nash equilibrium which is played. This gives a unique prediction of play in the game:

$$\text{(8)} \qquad \bar{s}(\pi_F^*, \pi_A^*, p^*) = \begin{cases} \frac{p^*+2\pi_A^*}{\pi_F^*+p^*-2\pi_A^*}, & \text{if} \quad \pi_F^* > 4\pi_A^* \\ 1, & \text{otherwise} \end{cases}$$

$$\text{(9)} \qquad \underline{s}(\pi_A^*, \pi_F^*, p^*) = \begin{cases} \frac{\pi_F^*-4\pi_A^*}{\pi_F^*-2\pi_A^*+\frac{\pi_F^*\pi_A^*}{p^*}}, & \text{if} \quad \pi_A^* < \frac{1}{4}\pi_F^* \\ 0, & \text{otherwise.} \end{cases}$$

Notice that these functions have all the comparative static properties we might expect. Thus $\bar{s}$ is decreasing in π_F^*: the more confident I am that a Republican will win, the lower will be the critical signal at which I start betting; $\bar{s}$ is increasing in π_A^*: the more confident you are that a Republican will win, the more serious a signal your willingness to bet against me is; $\bar{s}$ is increasing in p^*: as the betting odds move in my favor, the more willing I am to bet.

It may improve understanding to consider the limiting case here. As $\pi_F^* \to 0$, $\bar{s}$ tends to 1: as I become more and more confident of Republican victory, I am less prepared to bet; conversely, as $\pi_F^* \to \infty$, $\bar{s}$ tends to 0. As $\pi_A^* \to 0$, $\bar{s}$ tends to $p^*/[p^* + \pi_F^*]$: as you become more and more confident in Republican failure, the less and less information is conveyed by your willingness to bet. In the limit, I bet on the basis of my ‘naive’ posterior beliefs. On the other hand, as $\pi_A^* \to \infty$, $\bar{s}$ tends to 1: your willingness to bet against a Republican victory, *despite* your extreme prior confidence in a Republican victory must be very discouraging to me. If we consider variations in the betting odds ratio, p^*, we must consider two cases. If $\pi_F^* \leq 4\pi_A^*$, then $\bar{s}$ is 1 (and no betting occurs always). If $\pi_F^* > 4\pi_A^*$, then as p^* tends to 0, $\bar{s}$ tends to a constant $2\pi_A^*/[\pi_F^* - 2\pi_A^*]$. In this case, the improving betting odds ratio is being exactly canceled out in the limit by the bad news implicit in *A*'s willingness to bet. Finally, as $p^* \to \infty$, $\bar{s}$ tends to 1, as might be expected.

It would be reasonable to finish our analysis here, with this satisfactory characterization of the betting game for each p^*. But it is useful to be able to turn the solution concept around in order to see when the individuals will exhibit bid-ask spread behavior. Suppose you (individual 2) and I (individual 1) are considering making a bet. Our prior probabilities are π_1 and π_2, and we have each observed an independent signal, generated as above. All this is common knowledge. Suppose I have observed signal s_1. At which betting odds ratios would I be prepared to bet with you (either way)? Using the solutions above, we see that I, individual 1, will be prepared to accept a bet for Republican victory at betting odds p^* if $s_1 \geq \bar{s}\,(\pi_1^*, \pi_2^*, p^*)$; I will be prepared to accept a bet against Republican victory at betting odds p^* if $s_1 \leq \underline{s}(\pi_1^*, \pi_2^*, p^*)$. Because $\bar{s}$ and $\underline{s}$ are both (weakly) increasing in p^*, we can give a simply characterization of the set of betting odds ratios where I will be prepared to bet: I will bet for Republican victory if betting odds p^* are less than or equal to $\bar{p}^*(\pi_1^*, \pi_2^*, s_1)$; I will bet against Republican victory if betting odds p^* are greater than or equal to $\underline{p}^*(\pi_1^*, \pi_2^*, s_1)$; where functions $\bar{p}^*$ and $\underline{p}^*$ are given by

$$(10)\qquad \begin{aligned} &\bar{p}^*(\pi_i^*, \pi_j^*, s) \\ &\quad = \max\{p^* \in \mathcal{R}_+ \cup \{\infty\} \mid s \geq \bar{s}(\pi_i{}^*, \pi_j{}^*, p^*)\} \\ &= \begin{cases} \frac{1}{1-s}[(\pi_i^* - 2\pi_j^*)s - 2\pi_j^*], & \\ \qquad\qquad \text{if } \pi_i^* > 4\pi_j^* \text{ and } s > \frac{2\pi_j^*}{\pi_i^* - 2\pi_j^*} & \\ 0, & \text{otherwise} \end{cases} \end{aligned}$$

$$(11)\qquad \begin{aligned} &\underline{p}^*(\pi_i^*, \pi_j^*, s) \\ &\quad = \min\{p^* \in R_+ \cup \{\infty\} \mid s \leq s(\pi_i^*, \pi_j^*, p^*)\} \\ &\quad = \begin{cases} \dfrac{s}{\left\{\frac{1-s}{\pi_i^*} - \frac{4-2s}{\pi_j^*}\right\}}, & \text{if } \pi_i^* < \frac{1}{4}\pi_j^* \text{ and } s < \frac{\pi_j^* - 4\pi_i^*}{\pi_j^* - 2\pi_j^*} \\ \infty, & \text{otherwise.} \end{cases} \end{aligned}$$

It is useful to compare $\bar{p}^*$ and $\underline{p}^*$ with the price at which a naive individual (who ignored the information implicit in his opponent's willingness to bet) would be prepared to bet (either way): $\pi_i^* s/(1-s)$.

Observe that

$$(12) \quad \bar{p}^*(\pi_i^*, \pi_j^*, s) \leq \pi_i^* \frac{s}{1-s} \leq \underline{p}^*(\pi_i^*, \pi_j^*, s), \text{ for all } \pi_i^*, \pi_j^*, s.$$

Thus $\bar{p}^*$ and $\underline{p}^*$ do indeed generate the same prediction as the minmax approach: there is a highest and lowest price at which the individual is prepared to bet. His 'true' willingness to bet (in the absence of private information) is in between those two prices.

We can confirm also that $\bar{p}^*$ and $\underline{p}^*$ have all the comparative static and limit properties we might expect. Thus $\bar{p}^*$ (the maximum price individual *i* is prepared to pay for a bet for Republican victory) is increasing in *i*'s prior odds ratio, π_i^*, decreasing in his opponent's prior odds ratio, π_j^*, and increasing in his signal.

2.2. *The General Continuous Signals Case*

The framework in this section is exactly as in the previous section, except that here we allow for more general information signals. Here we will write *E* for the event (Republican victory) that the individuals are betting about.

Each individual observes a signal whose distribution depends on whether *E* has occurred or not. The signal can be distributed on any closed interval on the real line, and for notational convenience we will let that interval be [0,1].[13] There is a smooth distribution function (with support [0,1]) for individual *i*'s signal: $g_i(\cdot|E)$ if *E* has occurred, $g_i(\cdot| \sim E)$ otherwise. Both individuals are assumed to know the true distribution of their own and each other's signals. Thus they disagree about the likelihood of E, but not about the conditional distribution of the signals. The signals are conditionally independent. Thus we have:

$$(13) \quad \begin{aligned} prob_i[E, s_i, s_j] &= \pi_i g_i[s_i \mid E] g_j[s_j \mid E] \\ prob_i[\sim E, s_i, s_j] &= (1-\pi_i) g_i[s_i \mid\sim E] g_j[s_j \mid\sim E]. \end{aligned}$$

We assume that the signals have been ordered so that each likelihood ratio $r_i(s)$ is strictly increasing in s, where

$$(14) \quad r_i(s) = \frac{g_i(s \mid E)}{g_i(s \mid\sim E)}.$$

Now I introduce notation for cumulative distributions given those acceptance strategies. Let $G_i(\cdot|E)$ and $G_i(\cdot| \sim E)$ be the cumulative distribution function of individual i's signal conditional on E and not E, respectively, and let R_i be the ratio of the cumulative distributions. Thus

$$(15) \quad \begin{aligned} G_i(s \mid E) &= \int_0^s g_i(t \mid E)dt, \\ G_i(s \mid\sim E) &= \int_0^s g_i(t \mid\sim E)dt, \\ \text{and} \quad R_i(s) &= \frac{G_i(s \mid E)}{G_i(s \mid\sim E)}. \end{aligned}$$

Also let $\underline{G}_i(\cdot|E)$ and $\underline{G}_i(\cdot|\sim E)$ be the cumulative distribution function of individual i's signal from s to 1, conditional on E and not-E respectively, and let $\underline{R}_i$ be the ratio of those cumulative distributions. Thus

$$(16) \quad \begin{aligned} &\underline{G}_i(s \mid E) = \int_s^1 g_i(t \mid E)dt, \\ &\underline{G}(s \mid\sim E) = \int_s^1 g_i(t \mid\sim E)dt, \quad \text{and} \\ &\underline{R}_i(s) = \frac{\underline{G}_i(s \mid E)}{\underline{G}_i(s \mid\sim E)}. \end{aligned}$$

Observe that both R_i and $\underline{R}_i$ are strictly increasing (because r_i is strictly increasing); $R_i(s) \to r_i(0)$ as $s \to 0$; $R_i(1) = 1$; $\underline{R}_i(0) = 1$; and $\underline{R}_i(s) \to r_i(1)$ as $s \to 1$. It is convenient in what follows to *define* $R_i(0)$ to be equal to $r_i(0)$ and $\underline{R}_i(1)$ to be equal to $r_i(1)$. Also observe that $\underline{R}_i(s) \geq r_i(s) \geq R_i(s)$ and $\underline{R}_i(s) \geq 1 \geq R_i(s)$, for all signals s.

Thus, in the previous section, we had $g_i(s \mid E) = 2s$ and $g_i(s|\sim E) = 2(1-s)$, for *both* individuals. Then we would have $r_i(s) = s/(1-s)$,

$G_i(s|\ E) = s^2$, $G_i(s|\sim E) = s(2-s)$, $R_i(s) = s/(2-s)$, $G_i\ (s|\ E) = 1 - s^2$, $G_i(s|\sim E) = (1-s)^2$, $R_i(s) = (1+s)/(1-s)$.

Following the logic of the previous section, suppose one individual, F, is considering betting for, and the other individual, A, is considering betting against the occurrence of event E. Suppose that individual A will accept the bet only if he observes signal $\underline{s}$ or lower. What is individual F's best response? We need to calculate individual F's probability that a Republican will win the election, divided by the probability that a non-Republican will win, conditional on F observing signal s_F and A observing a signal $s_A \leq \underline{s}$. Individual F will accept the bet if this ratio is at least as great as the betting odds ratio p^*, i.e. if

$$(17)\quad \frac{prob(E)}{prob(\sim E)} = \frac{\pi_F g_F(s_F \mid E) G_A(\underline{s} \mid E)}{(1-\pi_F) g_F(\underline{s}_F \mid\sim E) G_A(\underline{s} \mid\sim E)} = \pi_F^* r_F(s_F) R_A(\underline{s}) \geq p^*.$$

This is equivalent[14] to

$$(18)\quad s_F \ \geq\ r_F^{-1}\left\{\frac{p^*}{\pi_F^* R_A(\underline{s})}\right\},\quad \text{if}\quad \frac{p^*}{\pi_F^* r_F(1)} \ \geq\ R_A(\underline{s}) \geq \frac{p^*}{\pi_F^* r_F(0)}.$$

Thus we can represent F's *reaction function* as follows:

$$(19)\quad \rho_F(\pi_F^*, p^*, \underline{s}, g_F, g_A) = \begin{cases} 0, & \text{if}\quad R_A(\underline{s}) \geq \frac{p^*}{\pi_F^* r_F(0)} \\ 1, & \text{if}\quad R_A(\underline{s}) \leq \frac{p^*}{\pi_F^* r_F(1)} \\ r_F^{-1}\left\{\frac{p^*}{\pi_F^* R_A(\underline{s})}\right\}, & \text{otherwise} \end{cases}$$

with the interpretation that F accepts the bet if and only if his signal $s_F \geq \rho_F(\pi_F^*, p^*, \underline{s}, g_F, g_A)$.[15] We can perform the symmetric calculation for individual A betting against E: suppose that individual F will accept the bet only if he observes signal $\bar{s}$ or higher. What is individual A's best response? We calculate individual A's probability that a Republican will win the election, divided by the probability that a non-Republican will win, conditional on A observing signal s_A and F observing a signal $s_F \geq \bar{s}$. Individual A will accept the bet if this ratio is no higher than the betting odds ratio p^*, i.e. if

$$(20) \quad \frac{prob(E)}{prob(\sim E)} = \frac{\pi_A g_A(s_A \mid E) G_F(\bar{s} \mid E)}{(1-\pi_A) g_A(s_A \mid\sim E) G_A(\bar{s} \mid\sim E)} = \pi_A^* r_A(s_A) R_F(\bar{s}) \leq p^*.$$

This is equivalent to

$$(21) \quad s_A \leq r_A^{-1}\left\{\frac{p^*}{\pi_A^* \underline{R}_F(\bar{s})}\right\}, \quad \text{if} \quad \frac{p^*}{\pi_A^* r_A(1)} \geq R_F(\bar{s}) \geq \frac{p^*}{\pi_A^* r_A(0)}.$$

Individual A's *reaction function* is:

$$(22) \quad \rho_A(\pi_A^*, p^*, \bar{s}, g_A, g_F) = \begin{cases} 0, & if \quad R_F(\bar{s}) \geq \frac{p^*}{\pi_A^* r_A(0)} \\ 1, & if \quad R_F(\bar{s}) \leq \frac{p^*}{\pi_A^* r_A(1)} \\ r_A^{-1}\left\{\frac{p^*}{\pi_A^* \underline{R}_F(\bar{s})}\right\}, & otherwise \end{cases}$$

with the interpretation that A accepts the bet if and only if his signal $s_A \leq \rho_A(\pi_A^*, p^*, \bar{s}, g_A, g_F)$.

LEMMA 1. *Properties of the Reaction Functions*:

[1] ρ_F *is increasing in* p^*, *decreasing in* π_F^*, *and decreasing in* $\underline{s}$.

[2] ρ_A *is increasing in* p^*, *decreasing in* π_A^*, *and decreasing in* $\bar{s}$.[16]

Proof. Follows from the explicit forms of ρ_F and ρ_A given above, and the strict monotonicity of r_i, R_i and $\underline{R}_i$ for each i.

A Nash equilibrium of this game will now be a pair $(\underline{s}, \bar{s})$ satisfying $\rho_F(\pi_F^*, p^*, \bar{s}, g_F, g_A) = \bar{s}$ and $\rho_A(\pi_A^*, p^*, \bar{s}, g_A, g_F) = \underline{s}$. In contrast to the example of the previous section, there is no general, simple, characterization of when there exists a Nash equilibrium where betting occurs with positive probability (i.e. $(\bar{s}, \underline{s}) \neq (1,0)$). Moreover, there may exist many Nash equilibria where betting occurs (unlike in the example). But notice that since both reaction functions are continuous and negatively sloped, there exists a Nash equilibrium which reflects the largest possible amount of betting. Formally, define the set of Nash equilibria as follows:

$$(23)\quad N(\pi_F^*, \pi_A^*, p^*, g_F, g_A) = \{(\bar{s}, \underline{s}) \mid \rho_F(\pi_F^*, p^*, \underline{s}, g_F, g_A) = \bar{s}, \rho_A(\pi_A^*, p^*, \bar{s}, g_A, g_F) = \bar{s}\}.$$

Because both reaction functions are negatively sloped and continuous, we know that there exists a unique 'largest' $(\bar{s}, \underline{s}) \in N(\pi_F^*, \pi_A^*, p^*, g_F, g_A)$ such that for every $(\bar{s}', \underline{s}') \in N(\pi_F^*, \pi_A^*, p^*, g_F, g_A)$, $\bar{s} \leq \bar{s}'$ and $\underline{s} \geq \underline{s}'$. We want to assume that this 'largest' equilibrium is always played. We let $\bar{s}(\pi_F^*, \pi_A^*, g_F, g_A, p^*)$ and $\underline{s}(\pi_A^*, \pi_F^*, g_A, g_F, p^*)$ be that 'largest' equilibrium. Thus we have 'acceptance functions' predicting exactly what the individuals will do in every situation.[17]

In order to study the comparative static properties of these acceptance functions, we require a natural restriction on the information signals, and a property of the best response functions along the acceptance functions.

DEFINITION. Individual *i*'s information signal satisfies the *smoothness* condition if

$$(24)\quad \frac{r_i'(s)}{r_i(s)} \geq \max\left\{\frac{R_i'(s)}{R_i(s)}, \frac{\underline{R}_i'(s)}{\underline{R}_i(s)}\right\}.$$

This condition requires that the rate of change of the (marginal) likelihood ratio is at least as great as each measure of the rate of change of the average likelihood ratio.[18] Thus it is a requirement that the likelihood ratio changes sufficiently smoothly. We can check that the example of the previous section satisfies this condition:

$$(25)\quad r_i(s) = \frac{s}{1-s}, \quad R_i(s) = \frac{s}{2-s}, \quad \underline{R}_i(s) = \frac{1-s}{1+s}, \quad so$$
$$r_i'(s) = \frac{1}{(1-s)^2}, \quad R_i'(s) = \frac{2}{(2-s)^2},$$

$$\underline{R}_i'(s) = \frac{2s}{(1+s)^2},$$
$$\textit{and } \frac{r_i'(s)}{r_i(s)} = \frac{1}{s(1-s)}, \quad \frac{R_i'(s)}{R_i(s)} = \frac{2}{s(2-s)},$$
$$\frac{\underline{R}_i'(s)}{\underline{R}_i(s)} = \frac{2s}{(1-s)(1+s)}.$$

In the next section, a finite signals example is presented and it is shown how any continuous approximation to a finite signals distribution would fail this property.

Before stating the main result, we require a property of the best response functions which is guaranteed to hold along the acceptance functions.

LEMMA 2. *Suppose* $\bar{s} = \bar{s}(\pi_F^*, \pi_A^*, g_F, g_A, p^*)$ *and* $\underline{s} = \underline{s}(\pi_A^*, \pi_F^*, g_A, g_F, p^*)$ *are both in the open interval (0,1). Then at* $(\bar{s}, \underline{s})$,

$$\frac{\partial \rho_F}{\partial \bar{s}} \frac{\partial \rho_A}{\partial \bar{s}} < 1. \tag{26}$$

Proof. Recall that both derivatives are strictly negative. Now, if the condition of the lemma failed, then by continuity and monotonicity of the reaction functions, there would exist another Nash equilibrium $(\bar{s}', \underline{s}')$, 'larger' than $(\bar{s}, \underline{s})$, contradicting the construction of the acceptance functions.

Finally, let us say that there is *no betting* if $(\bar{s}, \underline{s}) = (1, 0)$ and there is *always betting* if $(\bar{s}, \underline{s}) = (0, 1)$. Then we have the following characterizations of the acceptance functions.

THEOREM 1. *Properties of the acceptance functions*:

[1] *If* $\pi_F^* \le \pi_A^*$, *then there is no betting.*
[2] *If* $r_F(0)\, \pi_F^* \ge p^* \ge r_A(1)\, \pi_A^*$, *then there is always betting.*
[3] $\bar{s}$ *is decreasing in* π_F^* *and increasing in* π_A^*; *if A's information satisfies the smoothness condition, then* $\bar{s}$ *is increasing in* p^*.
[4] $\underline{s}$ *is decreasing in* π_A^* *and increasing in* π_F^*; *if F's information satisfies the smoothness condition, then* $\underline{s}$ *is increasing in* p^*.

Part (1) states that if *A* thinks event E at least as likely before observing information, then no odds p^* will induce that individual to bet on any signals. Thus, in particular, if the individuals have a common prior ($\pi_F^* = \pi_A^*$), there is no betting. This was shown by Sebenius and Geanakoplos (1983) and is closely related to the 'no trade' theorem of Milgrom and Stokey (1982). Part (2) states that if F thinks E sufficiently likely, and A thinks E sufficiently unlikely, relative to the betting odds ratio p^*, then they will bet whatever signal has been observed. Notice that for some information structures (such as the example in the previous section) this condition need never be satisfied (because $r_F(0) = 0$ and $r_A(1) = \infty$). Parts (3) and (4) ensure that the acceptance functions have the natural monotonicity properties. Notice that A's information must satisfy smoothness, in order for F's acceptance function to be increasing in p^* (so that F bets less often as the odds move against him).

Proof. (1) and (2) follow directly from reaction functions. I will give a proof of (3) and (4) for an interior equilibrium (i.e. where $\bar{s}$ and $\underline{s}$ are in the interval (0,1)). If one of $\bar{s}$ and $\underline{s}$ are equal to 0 or 1, the argument is easier in the sense that changes in p^* do not in general affect one individual's willingness to bet at the margin. But there are many cases that have to be considered. At an interior equilibrium, the following equilibrium conditions hold:

$$(27) \quad \pi_F^* r_F(\bar{s}) R_A(\underline{s}) = p^* = \pi_A^* r_A(\underline{s}) R_F(\bar{s}).$$

Totally differentiating with respect to p^* gives:

$$(28) \quad \begin{bmatrix} \pi_F^* r_F'(\bar{s}) R_A(\underline{s}) & \pi_F^* r_F(\bar{s}) R_A'(\underline{s}) \\ \pi_A^* r_A(\underline{s}) R_F'(\bar{s}) & \pi_A^* r_A'(\underline{s}) R_F(\bar{s}) \end{bmatrix} \begin{bmatrix} \frac{d\bar{s}}{dp^*} \\ \frac{d\underline{s}}{dp^*} \end{bmatrix} = M \begin{bmatrix} \frac{d\bar{s}}{dp^*} \\ \frac{d\underline{s}}{dp^*} \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

Now Lemma 2 gives

$$(29) \quad \frac{\partial \rho_F}{\partial \underline{s}} \frac{\partial \rho_A}{\partial \bar{s}} = \left\{ -\frac{r_F(\bar{s}) R_A'(\underline{s})}{r_F'(\bar{s}) R_A(\underline{s})} \right\} \left\{ -\frac{r_A(\underline{s}) R_F'(\bar{s})}{r_A'(\underline{s}) R_F(\bar{s})} \right\} < 1$$

so

$$(30) \quad \det(M) = \pi_F^* \pi_A^* \left\{ \begin{array}{c} r_F'(\bar{s}) R_A(\underline{s}) r_A'(\underline{s}) R_F(\bar{s}) \\ -r_F(\bar{s}) R_A'(\underline{s}) r_A(\underline{s}) R_F'(\bar{s}) \end{array} \right\} > 0,$$

$$(31)\quad \begin{bmatrix} \frac{d\bar{s}}{dp^*} \\ \frac{d\underline{s}}{dp^*} \end{bmatrix} = M^{-1}\begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

$$= \frac{1}{\det[A]}\begin{bmatrix} \pi_A^* r'_A(\underline{s})R_F(\bar{s}) & -\pi_F^* r_F(\bar{s})R'_A(\underline{s}) \\ -\pi_A^* r_A(\underline{s})R'_F(\bar{s}) & \pi_F^* r'_F(\bar{s})R_A(\underline{s}) \end{bmatrix}\begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

$$= \frac{p^*}{\det[A]}\begin{bmatrix} \frac{r'_A(\underline{s})}{r_A(\bar{s})} & -\frac{R'_A(\underline{s})}{R_A(\underline{s})} \\ \frac{R'_F(\bar{s})}{R_F(\bar{s})} & \frac{r'_F(\bar{s})}{r_F(\bar{s})} \end{bmatrix}\begin{bmatrix} 1 \\ 1 \end{bmatrix},$$

by substituting
equilibrium conditions (27)

$$= \frac{p^*}{\det[A]}\begin{bmatrix} \frac{r'_A(\underline{s})}{r_A(\bar{s})} - \frac{R'_A(\underline{s})}{R_A(\underline{s})} \\ \frac{r'_F(\bar{s})}{r_F(\bar{s})} - \frac{R'_F(\bar{s})}{R_F(\bar{s})} \end{bmatrix} \geq \begin{bmatrix} 0 \\ 0 \end{bmatrix},$$

by smoothness assumption.

Thus if *A*'s signal satisfies smoothness, $\bar{s}$ is increasing in p^*, and if *F*'s signal satisfies smoothness, $\underline{s}$ is increasing in p^*. The signs of derivatives with respect to prior probabilities can be obtained by totally differentiating (27) with respect to π_F^* and π_A^*:

$$(32)\quad \begin{bmatrix} \frac{d\bar{s}}{d\pi_F^*} \\ \frac{d\underline{s}}{d\pi_F^*} \end{bmatrix} = M^{-1}\begin{bmatrix} -r_F(\bar{s})R_A(\underline{s}) \\ 0 \end{bmatrix}$$

$$and \quad \begin{bmatrix} \frac{d\bar{s}}{d\pi_A^*} \\ \frac{d\underline{s}}{d\pi_A^*} \end{bmatrix} = M^{-1}\begin{bmatrix} 0 \\ -r_A(\underline{s})R_F(\bar{s}) \end{bmatrix}.$$

The functions $\bar{s}$ and $\underline{s}$ determine what each individual should do for a given betting odds ratio p^* and specification of who is to bet for, and who against. Now we do the translation into specifying for each individual, having observed some signal s, at which betting odds ratios he would be prepared to accept a bet for, and at which betting odds ratios he would be prepared to accept a bet against. Thus let

$$(33)\quad \bar{P}^*(\pi_i^*, \pi_j^*, g_i, g_j, s)$$
$$= \left\{p^* \in \mathcal{R}_+ \cup \{\infty\} | s \geq \bar{s}(\pi_i^*, \pi_j^*, , g_i, g_j, p^*)\right\}$$

$$(34)\qquad \bar{P}^*(\pi_i^*, \pi_j^*, g_i, g_j, s) = \left\{p^* \in \mathcal{R}_+ \cup \{\infty\} | s \leq \underline{s}(\pi_i^*, \pi_j^*, g_i, g_j, p^*)\right\}.$$

In the example of the previous section, the sets of betting odds where acceptance (for and against) takes place had a natural structure. If a bet for was accepted at betting odds p^*, it was also accepted at any betting odds less than p^*. While if a bet against was accepted at p^*, it was also accepted at any betting odds greater than p^*. This is a natural property because, other things being equal, an individual prefers to bet for at a lower p^*. But, in equilibrium, a higher p^* will imply that the other individual's willingness to bet becomes a worse signal. Whether the $\bar{P}^*$ and $\underline{P}^*$ sets are well behaved in the sense described above depends on how $\bar{s}$ and $\underline{s}$ vary with p^* which in turn depends, as the previous theorem showed, on the smoothness restriction on signals.

THEOREM 2. *Existence and properties of bid-ask probabilities: Suppose that individuals' signals satisfy the smoothness condition. Then there exist functions* $\overline{p}^*$ *and* $\underline{p}^*$ *such that*

[1] *The bid ask spread property is satisfied:*

$$(35)\qquad \begin{aligned} &\bar{P}^*(\pi_i^*, \pi_j^*, g_i, g_j, s) \\ &\quad = \left\{p^* \in \mathcal{R}_+ \cup \{\infty\} | p^* \leq \bar{p}^*(\pi_i^*, \pi_j^*, g_i, g_j, s)\right\} \\ &\underline{P}^*(\pi_i^*, \pi_j^*, g_i g_j, s) \\ &\quad = \left\{p^* \in \mathcal{R}_+ \cup \{\infty\} \mid p^* \geq \underline{p}^*(\pi_i^*, \pi_j^*, g_i, g_j, s)\right\}. \end{aligned}$$

[2] *The betting interval includes naive posterior beliefs:*

$$(36)\qquad \bar{p}^*(\pi_i^*, \pi_j^*, g_i, g_j, s) \leq \pi_i^* r_i(s) \leq \underline{p}^*(\pi_i^*, \pi_j^*, g_i, g_j, s).$$

[3] *Monotonicity:* $\overline{p}^*$ *and* $\underline{p}^*$ *are both increasing in s, increasing in* π_i^**, and decreasing in* π_j^*.

Proof. [1] is true because $\bar{s}$ and $\underline{s}$ are increasing in p^*. Naive posterior beliefs reflect your willingness to bet if you assumed the other individual would bet everywhere. Given that $\bar{s}$ and $\underline{s}$ are increasing in p^*, you can only be less willing to bet if the other individual is not betting everywhere. This argument implies [2]. [3] follows from the corresponding monotonicity conditions on $\bar{s}$ and $\underline{s}$.

Theorem 2 captures everything we would like to say about willingness to bet except that we have not considered the impact of changing the information structure. In particular, it is natural to expect that individuals will become less willing to bet as the accuracy of their independent information improves. But how to measure the precision of the information? One approach would be to use the partial order on information signals described by Blackwell (1951). One signal is more valuable than another if the latter is a 'noisy' version of the former. For our purposes, however, it seems sufficient to study a natural one parameter class of information signals.

We will consider a class of signals that is parameterized by a number $\lambda \in (0,1]$. Suppose that, for all λ, the signal is uniformly distributed, conditional on E *not* occurring, i.e. $g^\lambda(s|\sim E) = 1$, for all s,λ; whereas if E occurs, the signal is a weighted sum of a uniform distribution and a strictly increasing distribution $g^1(s|E)$, i.e. $g^\lambda(s|E) = 1 - \lambda + \lambda g^1(s|E)$. Then it can be easily verified that $r^\lambda(s) = 1 - \lambda + \lambda r^1(s)$, $R^\lambda(s) = 1 - \lambda + \lambda R^1(s)$, $\underline{R}^\lambda(s) = 1 - \lambda + \lambda \underline{R}^1(s)$.[19] As λ tends to zero, the signal becomes valueless. Now we can ask what happens as λ varies, if some individual i observes signal g_i^λ.

What happens as information quality improves? There are two effects. First of all, any given signal conveys more information for the individual receiving it. Thus if $r_i^\lambda(s) < 1$, individual i will be less willing to bet for and more willing to bet against, as the quality of information (λ) increases. Whereas if $r_i^\lambda(s) > 1$, individual i will be more willing to bet for and less willing to bet against, as the quality of information increases. But secondly, whether $r_i^\lambda(s)$ is greater than or less than to 1, i's improved information will make j less willing to bet against him, which makes i less willing (in equilibrium) to accept the bet. If these effects work against each other, we cannot predict the effect on the bid ask spread. The following theorem applies to the case where they move in the same direction.

THEOREM 3. *Suppose information satisfies the smoothness condition. Then* $\bar{p}^*(\pi_i^*, \pi_j^*, g_i^\lambda, g_j^\mu, s)$ *is decreasing in* λ *if* $r_i^\lambda(s) \leq 1$ *and decreasing in* μ *if* $r_j^\mu[\underline{s}(\pi_j^*, \pi_i^*, g_j^\mu, g_i^\lambda, \bar{p}^*(\pi_i^*, \pi_j^*, g_i^\lambda, g_j^\mu, s))] \geq 1$*; and* $\underline{p}^*(\pi_i^*, \pi_j^*, g_i^\lambda, g_j^\mu, s)$ *is increasing in* λ *if* $r_i^\lambda(s) \geq 1$ *and increasing in* μ *if* $r_j^\mu[\bar{s}(\pi_j^*, \pi_i^*, g_j^\mu, g_i^\lambda, \underline{p}^*(\pi_i^*, \pi_j^*, g_i^\lambda, g_j^\mu, s))] \leq 1$.[20]

To understand exactly what the theorem says, consider first the case where $r_i^\lambda(s) = 1$, so that individual *i*'s naive posterior probability of *E* is the same as his *ex ante* probability. Then any increase in the quality of his own information (λ) will imply that he requires better odds to bet either way, so that his bid ask spread range widens in both directions. Intuitively, this is because while there is no direct effect of the change in quality of *i*'s information, the other individual *j* will become less willing to accept a bet (in either direction) given that *i* is better informed. Allowing r_i^λ(s) to be less than one reinforces *i*'s decreasing willingness to bet for as information quality improves. Allowing r_i^λ(s) to be greater than one reinforces *i*'s decreasing willingness to bet against as information quality improves. What is the effect of changes in the quality of *j*'s information (μ) on *i*'s bid-ask spread for betting? The complicated expression, $r_j^\mu[\underline{s}(\pi_j^*, \pi_i^*, g_j^\mu, g_i^\lambda, \bar{p}^*(\pi_i^*, \pi_j^*, g_i^\lambda, g_i^\mu, s))]$ represents the likelihood ratio of the highest signal at which *j* would be prepared to bet against event *E*, if *s* was the lowest value at which *i* would be prepared to bet for *E*, and if the price p^* was such that these were the 'largest' equilibrium of the acceptance game. If this likelihood were at least as great as one, then the critical signal at which *j* is accepting a bet *against* is one which makes *j* less likely. Thus *j* will be less willing to bet, and this will lead to *i* being less willing to bet.

Proof. Again totally differentiate equilibrium conditions (27), now with respect to information quality parameters λ_F and λ_A, where $r_F = 1 - \lambda_F + \lambda_F r_F^1$ and $r_A = 1 - \lambda_A + \lambda_A r_A^1$, gives

$$\text{(37)} \qquad \begin{bmatrix} \frac{d\bar{s}}{d\lambda_F} \\ \frac{d\underline{s}}{d\lambda_F} \end{bmatrix} = M^{-1} \begin{bmatrix} \pi_F^*(1 - r_F^1(\bar{s}))R_A^1(\underline{s}) \\ \pi_A^*(1 - R_F^1(\bar{s}))r_A^1(\underline{s}) \end{bmatrix}$$

$$\text{and} \quad \begin{bmatrix} \frac{d\bar{s}}{d\lambda_A} \\ \frac{d\underline{s}}{d\lambda_A} \end{bmatrix} = M^{-1} \begin{bmatrix} \pi_F^*(1 - R_A^1(\bar{s}))r_F^1(\bar{s}) \\ \pi_A^*(1 - r_A^1(\underline{s}))R_F(\underline{s}) \end{bmatrix}$$

but $R^1_F(s) \geq 1$ for all s. So if $r^1_F(\bar{s}) \leq 1$, we have $d\bar{s}/d\lambda_F \geq 0$ and $d\underline{s}/d\lambda_F \leq 0$ also $R_A(s) \leq 0$, for all s, So if $r^1_A(\underline{s}) \geq 1$, we have $d\bar{s}/d\lambda_A \geq 0$ and $d\underline{s}/d\lambda_A \leq 0$. These properties of acceptance rules translate into the theorem's conditions on bid-ask spreads.

2.3. *A Finite Signals Example*

The smoothness condition was required in the previous section to ensure that the existence of a bid-ask spread in individuals' willingness to bet. The smoothness condition automatically fails in a finite signal setting. We use the simplest possible example, where each individual observes one of two signals, to illustrate what goes wrong.

Suppose individual 1 has prior probability 5/6 of event *E*, while individual 2 has prior probability 1/2 of event *E*. Each individual observes a signal with two possible realizations, 'good' or 'bad'. Each individual's signal has the same distribution. The good signal is correlated with *E*, the bad with not *E*. Specifically, there is a 2/3 chance of the good signal if *E*, or of the bad signal if not-*E*. As before, there is some betting probability *p* at which the individuals may bet.

Rather than formally re-defining the solution concepts of the previous section for this finite case, I will assume it is clear how to make the translation. First, note that at no *p* will we have 1 betting against *E* and 2 betting for *E*. So we can restrict attention to 1 betting for, 2 against. I first check the values of *p* for which certain strategy profiles are Nash equilibria of the betting game. Then I translate these results into acceptance rules, and show that those acceptance rules do not generate bid-ask behavior.

We first check when there is Nash equilibrium of the betting game when both individuals always accept the bet whatever the signal. This would require that 1's posterior probability of event *E* after observing a bad signal (5/7) is greater than *p*. We also require that 2's posterior probability of event *E* after observing a good signal (2/3) is less than *p*. Thus this is a Nash equilibrium only if $2/3 \leq p \leq 5/7$.

Let us now check when there is Nash equilibrium of the betting game when individual 1 accepts a bet for only after observing a good signal and individual 2 accepts a bet against only after observing a bad signal. This requires that 1's posterior probability of event E, conditional on 2 observing a bad signal and 1 observing a good signal (5/6) is greater than p, while his posterior probability of event E, conditional on 2 observing a bad signal and 1 observing a bad signal (5/9) is less than p. It is also required (for this Nash equilibrium) that 2's posterior probability of event E, conditional on 2 observing a bad signal and 1 observing a good signal (1/2) is less than p, while her posterior probability of event E, conditional on 2 observing a good signal and 1 observing a good signal (4/5) is greater than p. Thus this is a Nash equilibrium if $5/9 \leq p \leq 5/6$ and $1/2 \leq p \leq 4/5$, and thus if $5/9 \leq p \leq 4/5$.

It can further be checked that there is a Nash equilibrium where 1 always accepts and 2 accepts only if she observes a bad signal, if $1/3 \leq p \leq 5/9$. There is a Nash equilibrium where 2 always accepts and 1 accepts only if he observes a good signal if $4/5 \leq p \leq 10/11$.

For Nash equilibria where some individual never accepts the bet, the 'perfection' issue that arose in the continuous signals case arises in a very similar form. Any strategy is a best response to the other individual never accepting. It can be addressed in a very similar way. Let us restrict our attention to Nash equilibria where one individual never accepts to those where, if one individual does accept on some signal, it is a best response to his most pessimistic conjecture of when the other individual might accept (i.e. 1 conjectures that 2 accepts only on the bad signal; 2 conjectures that 1 accepts only on the good signal). The strategy profile where 1 always accepts and 2 never accepts is such a Nash equilibrium if $0 \leq p \leq 1/3$. The strategy profile where 2 always accepts and 1 never accepts is such a Nash equilibrium if $10/11 \leq p \leq 1$.

This covers all pure strategy Nash equilibria of the betting game. We translate this information into acceptance rules by the assumption that the 'largest' Nash equilibrium is always played. Thus if $0 \leq p < 1/3$, 1 always accepts and 2 never accepts. If $1/3 \leq p \leq 5/9$, 1 always accepts and 2 accepts only if she observes a bad signal. If $5/9 < p < 2/3$, 1 accepts only if he observes a good signal, 2 accepts only if she observes a bad signal. If $2/3 \leq p \leq 5/7$, 1 always accepts and 2 always accepts. If $5/7 < p < 4/5$, again we have 1 accepting only if he observes a good signal, 2 accepting only if she observes a bad signal. If $4/5 \leq p \leq 10/11$, 1 accepts only on a good signal, 2 accepts always. If $10/11 < p \leq 1$, 1 never accepts and 2 always accepts.

We can translate this information into acceptance rules for each individual. Thus individual 1 will never accept a bet against event E, whatever signal he observes and whatever the odds. If he observes the good signal, he will accept a bet for event E if and only if $p \leq 10/11$. If he observes the bad signal, he will accept a bet for event E if and only if *either* $0 \leq p \leq 5/9$ *or* $2/3 \leq p \leq 5/7$.

Individual 2 will never accept a bet for event E, whatever signal she observes and whatever the odds. If she observes a bad signal, she will accept a bet against event E if and only if $p \geq 1/3$. If she observes a good signal, she will accept a bet against event E if and only if *either* $4/5 \leq p \leq 1$ *or* $2/3 \leq p \leq 5/7$.

Thus there is no bid-ask spread. Consider individual 1's decision to bet for if he has observed a bad signal, and what happens as the betting probability moves in his favor (i.e. from 1 to 0). He starts betting once p is smaller than 5/7. But when the odds move even more in his favor, he stops being prepared to bet, because once p falls below 2/3, individual 2 will be deterred from accepting the bet against if she has observed a good signal. But when the odds move yet further in his favor, to 5/9, he is prepared to accept, despite have observed a bad signal, and knowing that individual 2 is accepting only when she has a bad signal. A symmetric pattern holds for individual 2.

The logic of this example can be used to show that whenever individuals observe a finite distribution of signals, we can choose a ratio of prior beliefs for the two individuals, π_1/π_2, such that the bid ask spread property fails to hold. It may useful to relate this finite signals failure to the smoothness condition in the continuous signals case. Imagine making a continuous signal approximation to a finite signals example. Then at some value of s, we would require $r(s)$ to jump suddenly, so that $r'(s)/r(s)$ would have to be high at s for any approximation. Thus the smoothness condition would fail.

3. UNCERTAINTY OR HIDDEN INFORMATION?

The informational model of the previous section showed that it is *possible* to rationalize bid ask spreads in willingness to bet by strategic concerns about hidden information possessed the person on the other side of the bet. It remains to argue that this rationalization is going to be correct. In this section, I first review the emphasis which authors on subjective probability have always put on informational asymmetries in explaining bid ask spreads. I argue that it is tautological to claim that if decision makers do not 'know' the 'true' probabilities, then there is some information they do not possess, which is in some form available to others. Thus there is always some relevant information in existence which potentially explains the bid ask spread. Nonetheless, it could be that decision makers exhibit 'uncertainty aversion' even in situations where there is not in fact any possibility of this information being used strategically against them. It is suggested that such uncertainty aversion may occur because boundedly rational decision makers misapply a very sensible heuristic: "if someone asks you for odds at which you are prepared to bet either way, they are probably about to exploit you".

The earliest proponents of subjective expected utility always recognized that the presence of private information will lead to bid-ask spreads in willingness to accept bets. Ramsey and de Finetti, who both anticipated Savage's development of subjective expected utility, were explicit about this issue, as shown by the quotations preceding the introduction to this paper. De Finetti notes that a game theoretic solution will be required. Even earlier, Borel can be interpreted as proposing the creation of game theory precisely to deal with the problem of the information implicit in others' willingness to bet!

> The problem can be put in a ... simple form, which is, however, complex enough to contain all its difficulty, by a consideration of a game like poker, where each player bets on his own play against the play of the adversary. It the adversary proposes a large bet, this tends to make people believe that he has a good hand, or at least that he is not bluffing. *Consequently the fact alone that the bet is made modifies the judgement which the bet is about.* The deep study of certain games will perhaps lead to a new chapter in the theory of probabilities, a theory whose origins go back to the study of games of chance of the simplest kind.[21]

A more recent text on Bayes Theory is equally clear on why we don't expect subjective expected utility maximizers to bet either way at the same odds:

> Suppose you are in a room full of knowledgeable meteorologists, and you declare the probability it will rain tomorrow is .95. They all rush at you waving money.

Don't you modify the probability? We may not be willing to bet at all if we feel others know more. Why should the presence of others be allowed to affect our probability?[22]

The meteorologist example raises the question of what we mean by information. It would be true but no doubt not very helpful to maintain that your unwillingness to bet either way at a probability of .95 at your subjective probability is because you do not know the true probability. Even if you do not think meteorologists know the true probability, and even if the meteorologists have already been forced to share all 'objective' information with you, we are inclined to think that you would want to alter your beliefs in the light of their willingness to bet. Presumably there is some hidden 'information' which is implicit in the fact that they are experts, and you are not.

Conversely, it is useful to consider what could be meant by 'knowing the true probability'. It is useful not because we care about the semantics, but because there is a strong intuition around that knowing the true probability matters, so to get to the root of the intuition, we need to understand what it means. Borel gave one definition

Observe however that there are cases where it is legitimate to speak of the probability of an event: these are the cases where one refers to the probability which is common to the judgements of all the best informed persons, that is to say, the persons possessing all the information that it is humanly possible to possess at the time of the judgements.[23]

This surely captures exactly our intuition of what we mean by the true probability of an event. I know that the true probability of the dime I am holding coming up heads is 1/2. I can imagine that someday someone might invent a machine which, by physical tests of the coin, could determine to one hundred decimal places the 'true' probability of that coin coming up heads on a 'fair' toss (which might not be exactly 1/2). Once that machine is invented, or even if I thought such a machine might have been invented, I would not claim to know the 'true' probability, even if 1/2 was in fact the *ex ante* expectation unconditional on the machine's test results.

An alternative representation of uncertainty aversion uses non-additive expected utility functions, and a 'hidden information' interpretation is common here too, going all the way back to Dempster (1967, 1968). Thus Sarin and Wakker (1992) report (p. 1255) with approval Keynes' (1921) view that "ambiguity in the probability of such events may be caused, for example, by a lack of available information relative to the amount of conceivable information". Recent 'expanded state space' justifications for non-additive expected utilities include Gilboa and Schmeidler (1994) and Mukerji (1994a). Kelsey and Milne (1993) show how a hidden moral hazard problem induces non-additive expected utilities.

But even granted that 'not knowing the true probability' means the same thing as 'someone may know something I don't', this does not imply that the person I'm betting against might be the person who might know something I don't. Furthermore, we can design decision problems in the face of uncertainty in such a way that strategic issues should not arise. It is useful to illustrate this issue by the classic example of Ellsberg (1961) which is regularly used to motivate the notion of 'uncertainty aversion' in the face of unknown probabilities.

Consider the following thought experiment. Urn I contains 50 red and 50 black balls. Urn II contains 100 balls, some red and some black. The true proportion is not known to you. We are interested in your preferences among the following four contingent claims:

1. \$100 if a red ball is drawn from urn I
2. \$100 if a black ball is drawn from urn I
3. \$100 if a red ball is drawn from urn II
4. \$100 if a black ball is drawn from urn II

A typical preference pattern might have you indifferent between (1) and (2), indifferent between (3) and (4), but strictly preferring either (1) or (2) to either (3) or (4). You strictly prefer \$100 with a known probability of 1/2 to \$100 with unknown probability, even though SEU theory apparently suggests that the decision maker ought to be assigning probability at least 1/2 to either to a red ball or to a black ball being drawn from urn II. More explicitly, expected utility theory (assuming the decision maker assigns probability 1/2 to a red ball being drawn from urn I) implies that one of the following three holds:

(i) probability of red ball being drawn from urn II is more than 1/2 and $(3) \succ (1) \sim (2) \succ (4)$;

(ii) probability of red ball being drawn from urn II is 1/2 and $(1) \sim (2) \sim (3) \sim (4)$;

(iii) probability of red ball being drawn from urn II is less than 1/2 and $(4) \succ (1) \sim (2) \succ (3)$.

The apparent uncertainty aversion may initially seem to be easily explained by private information. It is certainly conceivable that someone knows the exact number of red balls in urn II. If someone offers to sell me option (3), say, I would certainly typically infer that the person offering to sell it to me knows something I don't, in particular that there are not too many red balls in urn II.

This explanation will not go through, however, if I am really offered the menu of choices (1) through (4). It is true that the person offering me a choice of (1) through (4) quite likely has some information about (3) and (4) that I don't. Therefore, she may have preferences over which of the options I take. But the thought experiment is set up in such a way that I cannot infer any of that information, and so I ought to behave non-strategically (i.e. satisfy the Savage axioms *despite* the presence of private information). Yet the evidence is that many apparently rational decision makers still display uncertainty aversion, even when the strategic aspects are ruled out.[24]

One explanation might go as follows. In most situations where you must make a choice between uncertain outcomes and where somebody may know something which you don't, there is some information implicit in the offer made to you which should lead you to not buy and sell at the same price. In those rare situations where it is not the case, you nonetheless continue to use this heuristic. Psychologists Kahneman and Tversky (1974) identified a number of heuristics which decision makers use in making choices under uncertainty. These rules make sense in simple environments (explaining how they might have come about in an evolutionary setting), but their failures in certain more complicated circumstances explain many anomalies in observed choices under uncertainty. Observed 'uncertainty aversion', like Kahneman and Tversky's 'representativeness', 'availability' and 'adjustment' heuristics, may be a decision rule which works well in many, but not all, circumstances.

Some authors have argued that not only do the three heuristics identified by Kahneman and Tversky imply violations of the expected utility assumption in certain circumstance, but they do so in economically significant ways.[25] Can the same be said for uncertainty aversion? Presumably, we would want to offer uncertainty aversion as an explanation for some observed behavior only if a standard private information explanation fails. Yet economic environments are precisely those where private information explanations are often going to be most plausible. Among all decisions under uncertainty, it is a particular characteristic of those we make in economic life that someone with more information than us stands to gain from our decision. For example, using uncertainty aversion to explain financial asset pricing[26] seems especially redundant, since clearly financial asset trades represent near zero-sum transactions and there is plenty of private information around.

Dow and Werlang (1994) explain the breakdown of backwards induction in perfect information games by the existence of uncertainty aversion. The argument can be seen as an alternative formalization of the reputation story of Kreps and Wilson (1982) and Milgrom and Roberts (1982), which explicitly assumes incomplete information ('crazy' types) rather than uncertainty aversion to motivate an intuitive alternative to the standard backwards induction outcome. The argument of this paper has been that incomplete information about payoffs (and/or agents' rationality) is what lies at the heart of our intuitive rejection of the backwards induction outcome. 'Uncertainty aversion' is generated by the incomplete information, either directly (as in the model of Section 2), or indirectly through an uncertainty aversion heuristic. Even the work of Knight (1921) is consistent with asymmetric information interpretation of uncertainty aversion. LeRoy and Singell (1987) have recently argued that

> Knight shared the modern view that agents can be assumed always to act as if they have subjective probabilities. We document our contention that by uncertainty Knight meant situations in which insurance markets collapse because of moral hazard or adverse selection.

When can we be sure that there are no strategic considerations in decision making under uncertainty, so that apparent uncertainty aversion cannot be rationalized by asymmetric information? The most obvious case is when the uncertainty concerns nature rather than the actions of other economic actors. Bewley (1989) – in one interpretation of Knight (1921) – has argued that entrepreneurs are exactly those people who are prepared to bet on events for which true probabilities are not known. Entrepreneurs are presumably at least partly engaged in betting against nature (will I be able to develop some new technology?), although in a market economy, they may also be making indirect bets against other economic actors (potential competitors, for example). On the other hand, it is not clear that standard private information stories do not explain apparent uncertainty aversion to entrepreneurial activities. If I am considering a new project which no one else has yet undertaken (and do not know the 'true' probability), there is clearly a winner's curse aspect. Information is revealed by the fact that others have not undertaken the project.

This appeal to unmodelled hidden information as an alternative to uncertainty aversion may appear to verge on the tautologous. The following may be an operational way to distinguish between hidden information and true 'uncertainty aversion'. Suppose it is the case that any particular decision of an apparently uncertainty averse agent can be plausibly rationalized by some hidden information, but that each decision requires the decision maker to be imputing *different* information to those he is interacting with. Then we have compelling evidence (which does not rely on the assumption that the decision maker does not perceive the situation to be strategic) that there is true uncertainty aversion. Dow and Werlang (1994) and Mukerji (1994) examine implications of uncertainty aversion which would require changing information for a hidden information explanation.

4. CONCLUSION

The main contribution of this paper was to show that in the face of two sided asymmetric information, with sufficiently smooth signals, rational individuals' willingness to bet will be characterized by well-behaved bid ask spreads.

Let us briefly see how this result depended on the exact assumptions of this paper. It was assumed that the priors and independent signal distributions were common knowledge and that the decision makers, while disagreeing on the prior probability of the event they were betting on, agreed on the posterior probability of their signals conditional on whether the event occurred or not. Each of these assumptions could be weakened while maintaining the same qualitative results.

Thus if signals were not independent and individuals did not agree on the posterior probabilities of the signals, the results of

Section 2 would still go through if individuals had a *common ordinal interpretation of signals* (Morris, 1992) – i.e. there exists an ordering of signals of each individual's possible signals such that conditional on any observation of any individual, each individual agrees that a 'higher' signal makes the event they are betting about more likely. This assumption guarantees the crucial strategic complementarities, and thus monotonic best response functions, in the betting game.

Common knowledge of the other individual's prior and posterior distribution could be relaxed to allow each individual to have a probability distribution over the possible priors and posteriors of the other individual. Of course, such uncertainty about the other individual's type would play a role exactly analogous to the existing two sided private information in the model. The bid ask spread properties demonstrated in this paper would be maintained as long as this new private information did not lead to a violation of the common ordinal interpretation of signals property.

NOTES

1. Ramsey (1927), pages 34/35.
2. de Finetti (1937) in footnote (a) on page 62. The footnote was added in 1964.
3. Ellsberg (1961).
4. Knight (1921).
5. Bewley (1986, 1987, 1988), Gilboa and Schmeidler (1989), Wakker (1991). Work with non-additive probability distributions is also closely related in terms of both motivation and mathematical structure, see e.g. Schmeidler (1989).
6. Bewley (1989).
7. Dow and Werlang (1992), Epstein and Wang (1994).
8. Dow and Werlang (1994), Klibanoff (1993), Mukerji (1994b), Lo (1994).
9. Bewley (1986, pages 1/2) explicitly rejects this argument: "One may try to explain the lack of betting (between individuals with different prior beliefs) by mutual suspicion that the other decision maker has secretly acquired superior information. I find this hard to reconcile with the observation that people usually seem very fond of their own decisions".
10. This is essentially the static betting game studied (under the common prior assumption) in Sebenius and Geanakoplos (1983).
11. There is also always a Nash equilibrium where both individuals always reject the bet, and possibly other Nash equilibria where one agent always rejects the bet. The argument that follows does not apply to those equilibria which are formally ruled out consideration when other (non-trivial) equilibria exist in the next section.

12. To be precise, ρ_A is undefined at $\bar{s} = 1$; but the limit as $\bar{s}$ tends to 1 is 0 and it is convenient and natural here and in the rest of paper to assume the reaction function takes on the value of its limit at boundary points 0 and 1.
13. The analysis could easily be generalized to the extended real line, i.e. $\mathcal{R} \cup \{\infty, -\infty\}$, if limits of distribution functions on the real line existed.
14. r_F^{-1} is well-defined on the interval specified because r_F is strictly increasing.
15. Notice that there is a perfection argument implicitly built into the reaction functions. If $\underline{s} = 0$, then with probability 1, A never accepts the bet, so F's reaction ought to be indeterminate. We are assuming that F's reaction function is single-valued and continuous. It is straightforward to justify this by a perfection argument. In an analogous setting in Morris (1992), an *ad hoc* pessimistic conjecture is used to support this outcome. Restricting attention to proper equilibria would also suffice.
16. The terminology 'increasing' and 'decreasing' are used in the lemma and throughout the paper in the weak sense i.e. non-decreasing and non-increasing.
17. Thus for an appropriate simplification of the game and ordering of signals, the game is one with *strategic complementarities* (Bulow *et al*., 1985; Milgrom and Roberts, 1990). Specifically, suppose one round of dominated strategies are deleted, so that all remaining strategies are of the form 'accept if $s_F \geq \bar{s}'$ for F and 'accept if $s \leq \underline{s}$' for A, and strategies are ordered by the critical acceptance signal, then ordinal strategic complementarity conditions (Milgrom and Shannon, 1994) are satisfied. Morris (1992) studies these and other issues concerning strategic complementarities in more general Bayesian 'acceptance games'. The results of that paper could thus be used to show in more detail the relation between Nash equilibria in this paper.
18. Thus it is equivalent to the requirement that both $r_i(s)/R_i(s)$ and $r_i(s)/\underline{R}_i(s)$ are increasing functions of s.
19. Where $r^1(s) = g^1(s|E)/g^1(s|{\sim} E)$, $R^1(s) = G^1(s|E)/G^1(s|{\sim} E)$ and $\underline{R}^1(s) = \underline{G}^1(s|E)/\underline{G}^1(s|{\sim} E)$.
20. Notice that by construction $r_i^\lambda(s) \geq 1$ if and only if $r_i^1(s) \geq 1$.
21. Borel (1924), page 58.
22. Hartigan (1983), page 7.
23. Borel (1924), page 50.
24. It would be interesting to test how sensitive Ellsberg-paradox-type phenomena are to varying emphasis in the experimental design on the experimenter's incentives.
25. Thaler (1991).
26. Dow and Werlang (1992) and Epstein and Wang (1994).

REFERENCES

Bewley, T.: 1986, 'Knightian Decision Theory: Part I', Cowles Foundation Discussion Paper No. 807.

Bewley, T.: 1987, 'Knightian Decision Theory: Part II', Cowles Foundation Discussion Paper No. 835.

Bewley, T.: 1988, 'Knightian Decision Theory and Econometric Inference', Cowles Foundation Discussion Paper No. 868.
Bewley, T.: 1989, 'Market Innovation and Entrepreneurship: A Knightian View', Cowles Foundation Discussion Paper No. 905.
Blackwell, D.: 1951, 'The Comparison of Experiments', in *Proceedings, Second Berkeley Symposium on Mathematical Statistics and Probability*, pp. 93–102, University of California Press.
Borel, E.: 1924, 'Apropos of a Treatise on Probability', *Revue Philosophique*, translated and reprinted *Studies in Subjective Probability*, Kyburg and Smokler eds. (1st edition). John Wiley and Sons (1964).
Bulow, J., Geanakoplos, J. and Klemperer, P.: 1985, 'Multimarket Oligopoly: Strategic Substitutes and Complements', *Journal of Political Economy* **93**, 488–511.
de Finetti, B.: 1937, 'La Prevision: Ses Lois Logiques, Ses Sources Subjectives', *Annales de l'Institut Henri Poincaré* **7**, 1–68, translated in English as 'Foresight: its logical laws, its subjective sources' and reprinted in Kyburg and Smokler (1980), pp. 53–118.
Dempster, A.: 1967, 'Upper and Lower Probabilities Induced by a Multivalued Mapping', *Annals of Mathematical Statistics* **38**, 325–339.
Dempster, A.: 1968, 'A Generalization of Bayesian Inference', *Journal of the Royal Statistical Society, Series B* **30**, 205–247.
Dow, J. and Werlang, S.: 1992, 'Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio', *Econometrica* **60**, 197–204.
Dow, J., and Werlang, S.: 1994, 'Nash Equilibrium Under Knightian Uncertainty: Breaking Down Backward Induction', *Journal of Economic Theory* **64**, 305–324.
Ellsberg, D.: 1961, 'Risk, Ambiguity and the Savage Axioms', *Quarterly Journal of Economics* **75**, 643–669.
Epstein, L. and Wang, T.: 1994, 'Intertemporal Asset Pricing under Knightian Uncertainty', *Econometrica* **62**, 283–322.
Gilboa, I., and Schmeidler, D.: 1989, 'Maxmin Expected Utility with a Non-Unique Prior', *Journal of Mathematical Economics* **18**, 141–153.
Gilboa, I. and Schmeidler, D.: 1994, 'Additive Representations of Non-Additive measures and the Choquet Integral', *Annals of Operations Research* **52**, 43–65.
Glosten, L. and Milgrom, P.: 1985, 'Bid, Ask, and Transactions Prices in a Specialist Market with Heterogeneously Informed Traders', *Journal of Financial Economics* **14**, 71–100.
Hartigan, J.: 1983, *Bayes Theory.* New York: Springer-Verlag.
Kahneman, D. and Tversky, A.: 1974, 'Judgement under Uncertainty: Heuristics and Biases', *Science* **185**, 1124–1131.
Kelsey, D. and Milne, F.: 1993, 'Induced Preferences and Decision Making under Risk and Uncertainty', University of Birmingham Discussion Paper.
Keynes, J.: 1921, *A Treatise on Probability.* London and New York: Macmillan & Co.
Klibanoff, P.: 1993, 'Uncertainty, Decision and Normal Form Games'.
Knight, F.: 1921, *Risk, Uncertainty and Profit.* Boston and New York: Houghton-Miller
Kreps, D. and Wilson, R.: 1982, 'Reputation and Imperfect Information', *Journal of Economic Theory* **27**, 253–279.

Kyburg, H. and Smokler, H.: 1980, *Studies in Subjective Probability* (2nd edition). New York: Robert E. Krieger Publishing Co.
Leamer, E.: 1986, 'Bid Ask Spreads for Subjective Probabilities', in *Bayesian Inference and Decision Techniques*, P. Goel and A. Zellner, Eds., Amsterdam: Elsevier Science Publishers.
LeRoy, S. and Singell, L.: 1987, 'Knight on Risk and Uncertainty', *Journal of Political Economy* **95**, 394–406.
Lo, K.: 1995, 'Equilibrium in Beliefs under Uncertainty', University of Toronto.
Milgrom, P. and Roberts, J.: 1982, 'Predation, Reputation and Entry Deterrence', *Journal of Economic Theory* **27**, 280–312.
Milgrom, P. and Roberts, J.: 1990, 'Rationalizability, Learning and Equilibrium in Games with Strategic Complementarities', *Econometrica* **58**, 1255–1278.
Milgrom, P. and Shannon, C.: 1994, 'Monotone Comparative Statics', *Econometrica* **62**, 157–180.
Milgrom, P. and Stokey, N.: 1982, 'Information, Trade and Common Knowledge', *Journal of Economic Theory* **26**, 17–27.
Morris, S.: 1992, 'Acceptance Games and Protocols', CARESS Working Paper No. 92–22, University of Pennsylvania.
Morris, S.: 1993, 'Bid Ask Spreads with Two Sided Private Information', CARESS Working Paper No. 93–09, University of Pennsylvania.
Morris, S.: 1994, 'Trade with heterogeneous prior beliefs and asymmetric information', *Econometrica* **62**, 1327–1347.
Mukerji, S.: 1994a, 'Foundations of Decision Making under Subjective Uncertainty and Ambiguity: An Epistemic Approach', Yale University, forthcoming in *Economic Theory*.
Mukerji, S.: 1994b, 'A Theory of Play for Games in Strategic Form When Rationality is not Common Knowledge', Yale University.
Ramsey, F.: 1926, 'Truth and Probability', reprinted in Kyburg and Smokler (1980), pp. 23–52.
Savage, L.: 1954, *The Foundations of Statistics.* John Wiley and Sons.
Schmeidler, D.: 1989, 'Subjective Probability and Expected Utility Without Additivity', *Econometrica* **57**, 571–587.
Sebenius, J. and Geanakoplos, J.: 1983, 'Don't Bet On It: Contingent Agreements with Asymmetric Information', *Journal of the American Statistical Association* **78**, 424–426.
Shin, H.: 1991, 'Optimal Betting Odds Against Insider Traders', *Economic Journal* **101**, 1179–1185.
Shin, H.: 1992, 'Prices of State Contingent Claims with Insider Traders, and the Favourite-Longshot Bias', *Economic Journal* **102**, 426–435.
Shin, H.: 1993, 'Measuring the Incidence of Insider Trading in a Market for State-Contingent Claims', *Economic Journal* **103**, 1141–1153.
Thaler, R.: 1991, *QuasiRational Economics.* New York: Russell Sage Foundation.
Walley, P.: 1991, *Statistical Reasoning with Imprecise Probabilities.* New York: Chapman and Hall.

Department of Economics,
University of Pennsylvania,
3718 Locust Walk, Philadelphia PA 19104-6297, U.S.A.