Decentralized Betting under Heterogeneous State-Contingent Preferences

Abstract

We revisit the classic no-bet results of Geanakoplos and Sebenius (1983) (etc.) and consider the implications of introducing some heterogeneity in state-contingent preferences. In particular, we find that the no-bet result can withstand a moderate degree of preference heterogeneity, and we fully characterize the situations for which betting does obtain, in the special case that agents are risk-neutral in each state. Furthermore, we identify two ways in which the problem is badly behaved. We find, first, that even arbitrarily small changes to the distribution of signals and preferences can produce discrete jumps in the volume of betting that takes place; and second, that even arbitrarily small transaction costs may do the same.

Introduction

Since the inception of information economics over forty years ago, by the seminal work of Robert Aumann, David Blackwell, John Harsanyi, and others, economic theorists have formally explored the nature and implications of disagreement. Perhaps the most interesting fruits of this exploration have been the relationships uncovered between information and "speculative trade"—that is, trade for which some kind of disagreement (about the expected future value of the state-contingent asset being traded) is necessary. In particular, it often happens that trade that initially appears feasible on account of two parties' information differences will not take place: each party can update his own beliefs in light of the other's willingness to trade, ultimately causing the difference in expectations to vanish.

Central to all of this early work concerning the infeasibility of trade was the assumption that, for a trade to take place, the parties involved had to have common knowledge of their willingness to trade. Accordingly, when the trade is predicated entirely on a difference in expectations—such as a bet among risk-neutral parties—it can only go through when the parties have common knowledge of the fact that their expectations differ. But Aumann (1976) shows that, so long as the parties have a common prior and common knowledge of their rationality, they cannot have common knowledge of a difference in expectations. Indeed, as Geanakoplos and Polemarchakis (1982) spell out, the sequence of tentative offers and counteroffers that presumably precedes the state-contingent contract will cause the parties' expectations to converge (at least so long as their partitions are finite; and if countably infinite, the sequence brings their expectations arbitrarily close). It immediately follows that, under the above assumptions, no trading predicated on a difference in expectations can take place. This argument is made explicit for simple bets in Geanakoplos and Sebenius (1983) (and more famously for general commodity spaces in Milgrom and Stokey (1982), so long as the endowments are ex-ante efficient and the agents have convex preferences).

Since betting and informationally-motivated trade are commonly observed, researchers went on to consider various ways of relaxing the assumptions behind these early theorems. The most obvious place to start is to relax the common knowledge requirement to one of "*k*-mutual knowledge." That is, instead of stipulating that both parties know of each other's rationality (or prior, or willingness to trade), know of that knowledge, and so on *ad infinitum*, we may stipulate that they only have that knowledge, know of that knowledge, and so on, for k steps. (The notion of k-mutual knowledge is most famously illustrated and explored, though under a different name, in the "electronic mail game" scenario of Rubenstein (1989).) It is easy to observe—and, in a limited way, was noted even within Geanakoplos and Sebenius (1983)—that when we replace common with k-mutual knowledge, all the results outlined above, with respect to agreement, betting, and trade, can fail spectacularly, so long as the partitions of the agents involved are sufficiently fine.[1]

There is also, it turns out, a more promising way to relax the common knowledge assumption. Two agents have "common p-belief" in E, as defined by Monderer and Samet (1989), if they both assign probability at least p to E, both assign at least probability p to the idea that they both assign probability at least p to E, and so on, *ad infinitum*. In general, if two parties share common p-belief in E, their beliefs can differ by up to $2(1 - p)$. As a result, if these people can attain only common p-belief in each other's willingness to trade—a situation that could arise, for instance, if their decision to trade is to some extent "noisy," in that the button to accept or reject an offer misfires $1 - p$ of the time—some trade may yet go on between them.[2] Finally, a broad

[1] A more thorough investigation of the bounds imposed by k-mutual knowledge of rationality can be found in Takamiya and Tanaka (2006).

[2] The relevant connections to betting and trade are drawn in Neeman (1993) and Sonsino (1995).

literature considers other ways to approximate common knowledge, and the relevant implications for speculation.

There are a handful of less mainstream ways to relax the assumptions built into the original framework, of which at least two are worth mentioning. First, we may consider "non-partitional" signal structures, in which each person's interim-period signal takes the form of a family of (not necessarily disjoint) subsets of the state space. Though such structures, and their implications for speculation, have been analyzed,[3] most research in information economics continues to rely, implicitly or explicitly, on partitions, since they are the only signal structures that satisfy certain basic rationality assumptions.[4] Second, we may permit our agents to have heterogeneous priors. (One motivation for doing so is Aumann's argument[5] that, strictly speaking, we ought to incorporate even the agents' beliefs into our "states." When we do so, heterogeneous priors over the state space are necessary if we are to relax any part of the seemingly strong condition that both the priors and the partitions are common knowledge.) The implications of heterogeneous priors have been studied thoroughly as well,[6] but this approach too is open to philosophical criticism[7]: heterogeneous priors can be introduced only at the cost of violating the "Harsanyi doctrine"—that is, at the cost of making the strange claim that one's beliefs are, at least to some extent, truly arbitrary,

---

[3] Originally in Geanakoplos (1989).

[4] As explained in Bacharach (1985).

[5] Made explicit in Aumann (1999), etc.

[6] Seminally in Morris (1994); see also Sethi and Yildiz (2012), etc.

[7] See Morris (1995) for a discussion of the debate. (He ultimately argues on behalf of the legitimacy of heterogeneous priors.)

and not functions of the information one has received. In the end, therefore, common priors and partitions continue to be the standard tools used in analyzing the economic implications of differences in information.

In this theoretical tradition, however, little has been written exploring the consequences of relaxing the assumption of ex-ante efficiency—that is, the consequences of diverse beliefs in the presence of the kinds of heterogeneous preferences that allow for ordinary trade as well. Accordingly, we will introduce such preference heterogeneity as simply as possible: the good in question will be some payment in the event of E, its price being some obligation in the event of $\neg E$ (so that the trade amounts to a bet), and the relevant circumstances will be distributions of linear state-contingent preferences for money. This could be understood as a simple case of the situation that arises when, for instance, an agent sells a rarely traded, risky bond over the counter. In the absence of an exchange (or recent price data), neither party can deduce anything about the true state from a single, public price; each agent updates his beliefs only on the evidence of the other's willingness to trade. Furthermore, though the buyer and the seller privately assign probabilities to the event of default, they might find a price at which to trade even if they had the same beliefs, since the expectation of default may be correlated or anticorrelated with each agent's other positions.

Presented here is a more precise statement of the problem and a demonstration of the key findings.

Framework

The model of information used here can be summarized as follows: time is divided into three periods. Ex-ante, everyone has a common prior probability distribution over an event space, along with a unique partition over it. Both the prior and the partitions are common knowledge. At an intermediate stage, each person learns which element of his own partition contains the true state of the world, allowing him to rule out the others and update his beliefs accordingly. (At this point, people have the opportunity to trade goods whose value depends on the true state, each presumably hoping to profit from his own private information.) At an ex-post stage, the true state is revealed.

Here, we will consider a population of common-knowledge-rational agents who receive independent signals and then have the opportunity to bet over E, an event with respect to which they have a common prior but heterogeneous preferences for money. (That is, interpreting this bet as the trade of a simple state-contingent good, we will reject the ex-ante efficiency assumption of Milgrom and Stokey (1982) but maintain the others.) The agents are then randomly matched and given the opportunity to bet, each without knowing the belief- or preference-type of the other.

Note that the classic no-bet results can be interpreted as resting on an infinite sequence of conditioning and counter-conditioning whose ultimate result leaves both agents with the same beliefs. The framework presented here hopes to shed more light on this process by studying the kind of *partial* conditioning in which agents engage when they do not know to what extent their partner is acting on different beliefs and to what extent he is acting on different preferences.[8]

For instance, the analysis of betting from Geanakoplos and Sebenius, etc. was surprising, at the time, because it showed that under certain standard assumptions, a difference in information is never enough to produce betting; no matter how much people's signals differ, common knowledge of a willingness to bet must cause their beliefs about the event to converge. What we will observe below is this: not only do the induced differences in information fail to give rise to betting, they are generally an active force in killing off betting that would have taken place before the signals were received. That is, betting generally fails not *despite* differences in information, but somewhat *because of* them, and the greater the information differences, the rarer betting becomes. (This observation is reminiscent of the insurance-market-destroying effects of asymmetric information discussed at length in Arrow (1963) and the subsequent literature.)

Statement of the Problem

For simplicity, we will only consider the case in which each person's utility for money is linear in each state:

$$u_E(x) = ax + b \quad \text{and} \quad u_{\neg E}(x) = cx + d$$

[8] Note also the way in which our approach differs from the approach to trade under uncertainty taken in Arrow and Debreu (1954) and the subsequent literature on general equilibrium under uncertainty. Since betting is decentralized, we are not solving for general equilibrium; bets at different odds may be taking place simultaneously, between different pairs of agents. Indeed, the very presence of a public, equilibrating price would usually render information trade impossible, as per Nielsen, Brandenburger, Geanakoplos, McKelvey, and Page (1990).

Accordingly, we can assign each person a number $p = a/(a + c) \in [0,1]$, which represents, essentially, the extent to which he would prefer a marginal penny under E to a marginal penny under $\neg E$.

Our agents have a common prior q for E; their different beliefs about E result from different, independent signals. A signal is just a number $b \in [0,1]$ assigning some probability to E. Upon receiving his signal, each agent can be characterized by a pair $(p,b) \in [0,1]^2$.

Let $k \in (0, \infty)$ denote the odds: the size of the payment if $\neg E$ per dollar paid if E. After receiving their signals, agents can be characterized in terms of their "naïve willingness to bet on E."

DEFINITION: Agent A's naïve willingness to bet n is the lowest (highest) odds at which he is willing to betting on E ($\neg E$)—that is, the odds at which he is indifferent to betting either way—*without conditioning* on the other person's willingness to bet at that odds. More formally, it is the n satisfying $b_A p_A - (1 - b_A)(1 - p_A)n = 0$.

Rearranging the above to get b as a function of p, we have the hyperbolic function

$$b = \frac{n-np}{p+n-np}$$

FIGURE 1.

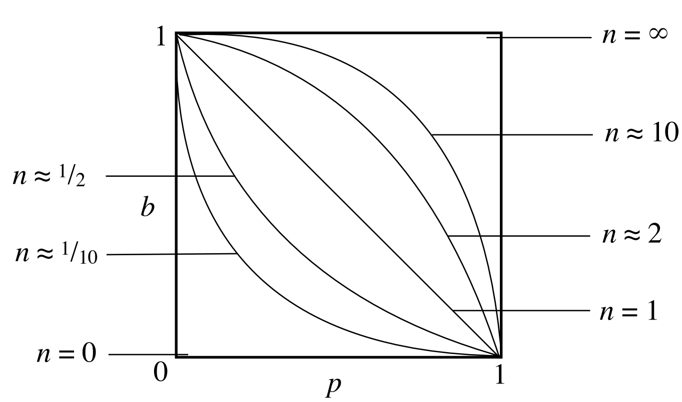

PROPOSITION 1: If an agent is naïvely willing to bet at odds k, paying if E ($\neg E$), he is naïvely willing to bet at odds $k' \leq k$ ($k' \geq k$).

Since our agents are not naïve, however, they condition on the fact that the other agent belongs to the opposite set, and they raise or lower their beliefs accordingly. In general, these new beliefs render it unprofitable for those on the margin to go ahead with the bet; the frontiers of the willingness-to-bet sets retreat toward the lower left and upper right corners. Each party then reconditions on the new, even more distant set to which he knows the other belongs; the frontiers retreat further; and so on. There will be some betting if and only if this process reaches Bayesian equilibrium before our sets "hit the corners."

More formally: a bet of odds k can go through between two agents, *1* and *2*, if and only if there is a pair of subsets $F_1 \ni (p_1,b_1)$ and $F_2 \ni (p_2,b_2)$ of the unit square such that each set consists precisely of the agents who would be willing to bet at odds k, even after incorporating into his beliefs the (common) knowledge that the other agent's type is in the other set.

Let *1* denote the person paying if E and receiving if $\neg E$. It is clear that if *1* is willing to make a bet, conditional on the knowledge that $2 \in F_2$, then so is any agent A with $p_A \leq p_1$ and $b_A \leq b_1$. Furthermore, if *1* and some A are both indifferent to taking the bet, and $b_A < b_1$, then $p_A > p_1$. Likewise, if *2* is willing to make a bet, conditional on the knowledge that $1 \in F_1$, then so is any A with $p_A \geq p_2$ and $b_A \geq b_2$; and if *2* and A are both indifferent to taking the bet, and $b_A > b_2$, then $p_A > p_2$. It follows that we can characterize our sets F_i by monotonically decreasing functions $f_i(p)$, defined on [0,1]. These tell us the signals that someone of preference-type p would have to have received in order to be indifferent to offering, or taking, a given bet.

FIGURE 2.

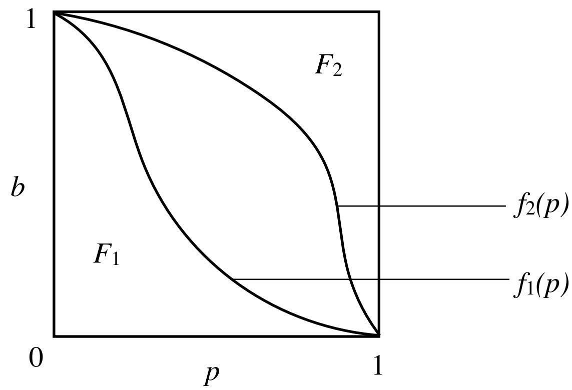

If E, our agents receive one distribution of signals, given by probability density function (or, if the distribution is discrete, mass function) $g_E(p,b)$. If $\neg E$, they receive a different distribution, given by $g_{\neg E}(p,b)$. For now, we will assume that the distributions of preference- and signal-types are integrable and continuous. If either or both is discrete, one may simply replace the relevant integral with a summation.

The only restrictions placed on this pair of distributions are those two which follow directly from the above setup:

1) $\frac{g_{\mathrm{E}}(p,b)}{g_{\neg\mathrm{E}}(p,b)} = b$ for all b

This guarantees that the probability an agent assigns to E, upon receiving signal b, is in fact equal to b.

2) $q\int_0^1 g_{\mathrm{E}}(p,b)db + (1-q)\int_0^1 g_{\neg\mathrm{E}}(p,b)db = q$, for all p

This guarantees that each person's expected belief, upon receiving his signal, is equal to the prior.

Bayes' Rule tells us that the probability that *1* assigns to E, after learning both his private signal and the fact that $2 \in F_2$, is

$$\Pr(E|F_2,b_1) = \Pr(E|b_1) \times \Pr(F_2|E,b_1) \,/\, \Pr(F_2|b_1) = b_1\Pr(F_2|E) \,/\, \Pr(F_2|b_1).$$

If $b_1 = f_1(p)$, this equals

$$\frac{f_1(p) \int_0^1 \int_{f_2(x)}^1 g_E(x,b)\, db dx}{\int_0^1 \int_{f_2(x)}^1 f_1(p) g_E(x,b) + \left(1\text{-}f_1(p)\right) g_{\neg E}(x,b)\, db dx}$$

Of course, *2*'s subjective probability of E after receiving both his private signal and the fact that $1 \in F_1$ is the same, but with $f_2(p)$ in place of $f_1(p)$, and integrating from 0 to $f_1(x)$ instead of from $f_2(x)$ to 1. Let us denote these posterior beliefs b_1' and b_2'.

Finally, to simplify notation, let

$$X_1 = \int_0^1 \int_{f_2(x)}^1 g_E(x,b)\, db dx \qquad \text{and} \qquad Y_1 = \int_0^1 \int_{f_2(x)}^1 g_{\neg E}(x,b)\, db dx$$

and let X_2 and Y_2 represent the corresponding terms in b_2'. Now we have

$$b_i' = f_i(p)X_i \,/(f_i(p)X_i + (1 - f_i(p))Y_i).$$

To be indifferent to the marginal bet at odds k, each agent's (posterior) expected utility gain from it must be 0. That is, we must have, for both the *1* and the *2*,

$$b_i'p - (1 - b_i')(1 - p)k = 0, \forall p \in [0,1].$$

Plugging in for b_i', and rearranging algebraically, we see that each "indifference function" takes the form

$$f_i(p) = \frac{c_i - c_i p}{p + c_i - c_i p}$$

where c_i denotes kY_i / X_i. As we can see, this curve is hyperbolic, just like the curve defining the naïve willingness-to-bet types. Furthermore, since it is plain that $X_1 \geq Y_1$ and $Y_2 \geq X_2$ for any distribution and any pair of sets defined by $f_1(p)$ and $f_2(p)$, we have:

PROPOSITION 2: If two agents A and A' are indifferent to betting at odds k, paying if E ($\neg E$), there is some odds $k' \leq k$ ($k' \geq k$) at which A and A' are naïvely indifferent to betting.

Visually, this tells us that our sets F_1 and F_2—just like the sets of naïve willingness-to-bet types—must take the following form for some pair of nonnegative c_1, c_2 ($c_1 < c_2$):

FIGURE 3.

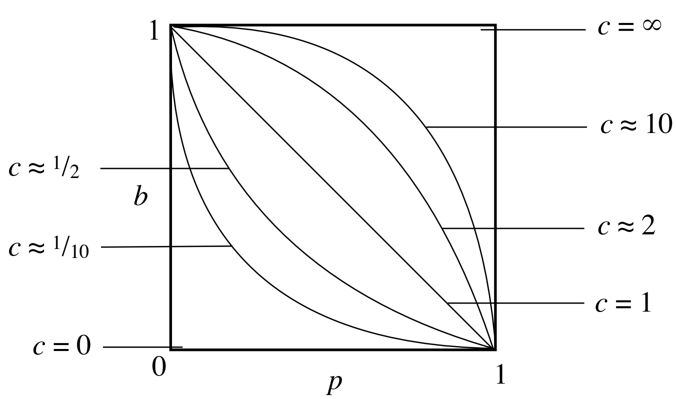

COROLLARY 1: If an agent is willing to bet at odds k, he is naïvely willing to bet at odds k.

Proof: By Proposition 2, if an agent is willing to bet at odds k (paying if E), he is naïvely willing to bet at some odds $k' \geq k$. Then by Proposition 1, he is naïvely willing to bet at odds k.

Now, to determine the odds at which our (p,b)-distribution will permit betting, it is clear that everything relevant about our distribution can be captured by a pair of one-variable cumulative distribution functions $N_E(c)$ and $N_{\neg E}(c)$: the fraction of the population with naïve willingness to bet no greater than c under E and $\neg E$, respectively. (If two distributions of preferences and signals give rise to the same distributions of naïve willingness to bet, they ultimately exhibit the same betting characteristics.) So the proposition above can be restated thus: our agents can reach Bayesian equilibrium in considering a bet at odds k if there exist c_1 and c_2 such that

$$c_1 = k\,(1 - N_{\neg E}(c_2))/(1 - N_E(c_2)) \quad \text{and} \quad c_2 = k\,N_{\neg E}(c_1) \,/\, N_E(c_1).$$

As expected, this has a trivial solution, for all k, of $c_1 = 0$ and $c_2 = \infty$. The question is, are there any other solutions? That is—plugging c_2 into the equation for c_1—are there any zeros of the function

$$\frac{1 - N_{\neg E}\left(\frac{N_{\neg E}(c)}{N_E(c)}\right)}{1 - N_E\left(\frac{N_{\neg E}(c)}{N_E(c)}\right)} k - c$$

besides that which we know exists (in the limit) at $c = 0$?

A Simpler Characterization

Consider the functions $c_1(c_2)$ and $c_2(c_1)$ defined above. To say that there will be no betting at some odds k, say $k = 1$, is to say that the graphs of these functions (where $k = 1$) never intersect:

FIGURE 4.

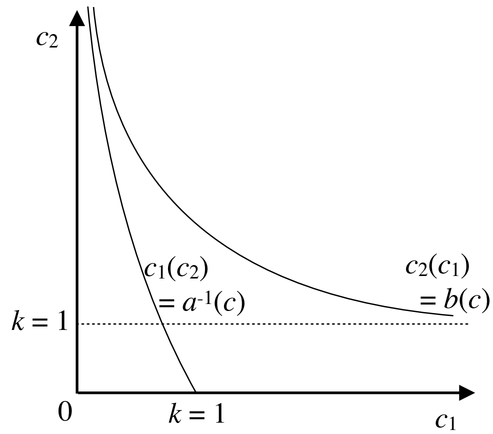

For now, let us think of these graphs as both representing c_2 in terms of c_1, and let us denote them $a(c)$ and $b(c)$. (Note that $a(c)$ may be a relation, rather than a function, since we have no reason to believe that $c_1(c_2)$ is monotonically decreasing in c_2. Furthermore, $c_1(c_2)$ and $c_2(c_1)$ may be discontinuous, in the event that our distribution g contains point masses; if so, let us connect our graphs with the relevant vertical or horizontal line segments, and have $c_1(c_2)$ and $c_2(c_1)$ refer henceforth to the corresponding relations. For the purpose of the diagrams that follow, however, we will generally assume that g is continuous and supported everywhere.) There will be no betting at our chosen odds ($k = 1$) if and only if the graphs never intersect—that is, if $a(c) < b(c)$ for all $x > 0$.

Furthermore, as we know from the definitions of $c_1(c_2)$ and $c_2(c_1)$, choosing another odds k will produce a pair of graphs like those above but with $a(c)$ shifted right, and $b(c)$ shifted up, by a factor of k. Therefore, to say that betting won't take place for *any* k is to say that $a(x/k) < kb(x)$ for all x, $k > 0$. It follows from this that there is a function $c(x)$ satisfying the following two conditions:

1) $a(x) < c(x) \leq b(x)$, for all $x > 0$, and

2) $c(x/k) = kc(x)$, for all x, $k > 0$.

Note that (2) is equivalent to the condition that $c(x) = d/x$, for some constant d.
Therefore, it is also equivalent to the condition that $c(1/k)/k = c(1)$, for all $k > 0$.

To see this, observe that (2) essentially consists of an infinite set of requirements on $c(1)$: two requirements for every positive real number k, taking the form

$$a(1/k)/k < c(1) \leq b(1/k)/k\ .$$

So long as there is some $c(1)$ that satisfies all these requirements, we can define $c(x)$ in general by $c(1/k) = kc(1)$, and both conditions will be satisfied.

Naturally, there is such a potential $c(1)$ if, for all pairs k_1, k_2, we have

$$a(1/k_1)/k_1 < b(1/k_2)/k_2\ .$$

In fact this inequality does hold: it can be generated by substituting $k = k_1/k_2$ and $x = 1/k_2$ into the inequality stipulated above, $a(x/k) < kb(x)$, and then dividing both sides by k_1.

So there will be no betting if, for some constant d, we have $a(x) < d/x \leq b(x)$, for all $x > 0$. Since d/x is symmetric about $y = x$, we may replace $a(x)$ with $c_1(c_2)$ from above (which, by construction, is simply $a(x)$ but flipped about $y = x$). It follows that there will be betting if, for some d and some pair c_1, c_2 ($c_1 < c_2$), $c_1(c_2)c_2 = d = c_2(c_1)c_1$. Or—since the choice of d is now irrelevant, and since $c_1(c) < c_2(c)$—

PROPOSITION 3: There will be betting if and only if

1) $\inf\limits_{c} \dfrac{N_{\neg E}(c)}{N_E(c)}c \;<\; \sup\limits_{c} \dfrac{1 - N_{\neg E}(c)}{1 - N_E(c)}c$, $c \in (0, \infty)$, or

2) $\min\limits_{c} \dfrac{N_{\neg E}(c)}{N_E(c)}c \;=\; \max\limits_{c} \dfrac{1 - N_{\neg E}(c)}{1 - N_E(c)}c$, $c \in (0, \infty)$.

That is, there will betting if and only if a horizontal line can be drawn that intersects the graphs of $c_1(c)c$ and $c_2(c)c$ at values other than 0 or ∞.

FIGURE 5.

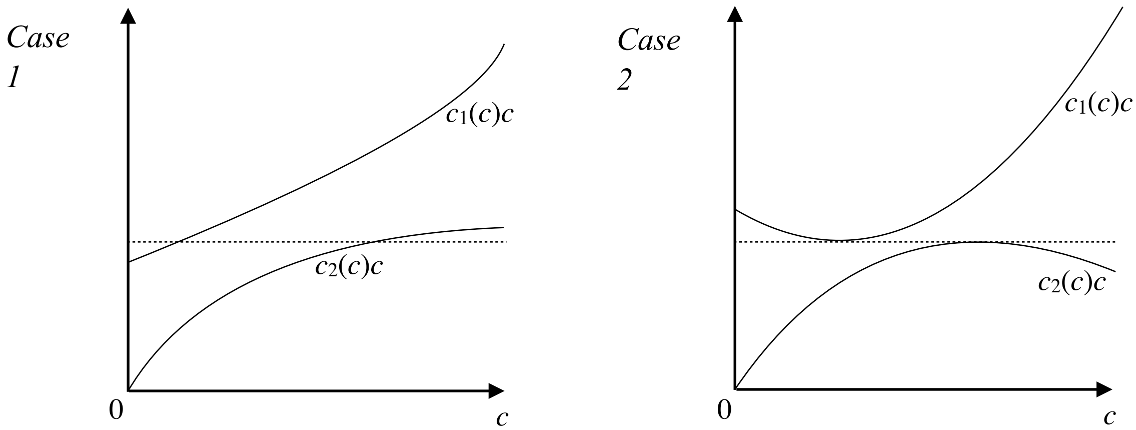

This visualization helps to clarify the way in which an increase in information quality decreases the prevalence of betting.

FIGURE 6.

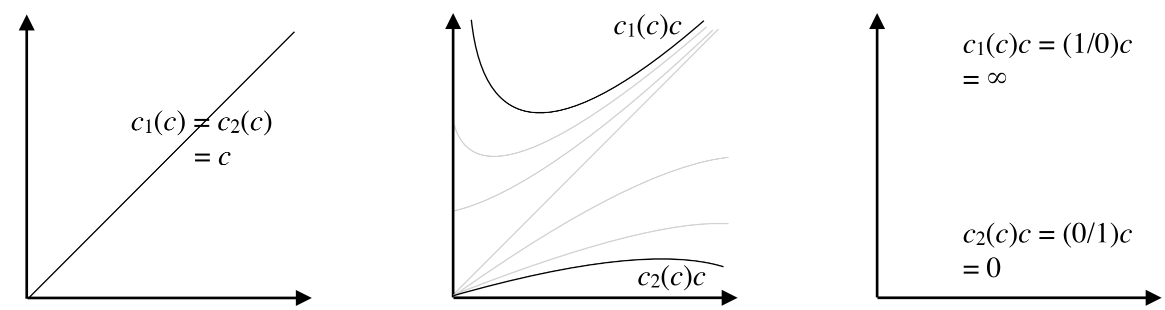

- *Left*: When everyone has the same beliefs—that is, when everyone's signal is completely uninformative—our distribution $g_E = g_{\neg E}$ takes the form of a horizontal line across the unit square, at $b = q$; so $N_{\neg E}(c) = N_E(c)$ for all c. It follows that

$$c_1(c)c = ((1 - N_{\neg E}(c))/(1 - N_E(c)))c = c \ \forall c, \text{ and}$$

$$c_2(c)c = (N_{\neg E}(c)/N_E(c))c = c \ \forall c.$$

Accordingly, any horizontal line drawn across the graph above—any choice of d—intersects both graphs; and at the corresponding odds, a bet will go through for some pair of sets denoted by $c_1 = c_2$.

- *Middle*: A strict improvement in the quality of information received by any individual increases $N_{\neg E}(c)$ and decreases $N_E(c)$; accordingly, it weakly increases $c_1(c)c$ and decreases $c_2(c)c$. A geometrically obvious consequence of this is that it weakly increases the minimum wedge between c_1 and c_2, and so the wedge between F_1 and F_2 (minimized over the choice of d, and so over the choice of k).
- *Right*: If everyone receives perfect information, then $N_E(c) = 0$ and $N_{\neg E}(c) = 1$, $\forall c \in (0, \infty)$; so $c_1(c)c = \infty$ and $c_2(c)c = 0$, $\forall c \in (0, \infty)$.

Further Examples

The characterization presented above gives us enough of a handle on the problem to determine whether betting will take place under various classes of distributions. Two examples:

- Suppose that the expected distribution of signals is uniform on [0,1] for each preference-type (that is, $g_E(p,b) = 2bg(p)$). Then there will not be betting if $p_{max} < 4p_{min}/(1 + 2p_{min})$, where p_{min} and p_{max} are respectively the lowest and highest preference-types to be found in the population.

  Proof: We have not specified the distribution of preference-types, but we know that $N_{\neg E}(c)/N_E(c)$ increases, for all c, when preference-types increase. Therefore $c_2(c)c$ is minimized for all c, within our constraints, when the entire population possesses preferences $p = p_{min}$. In this case, simple algebra reveals $c_2(c)c = 2p_{min}/(1 - p_{min}) + c$; so $c_2(c)c$ attains a minimum of $2p_{min}/(1 - p_{min})$ at $c = 0$. Likewise, $(1 - N_{\neg E}(c))/(1 - N_E(c))c = c_1(c)c$ is maximized for all c when the entire population possesses preferences $p = p_{max}$, and in this case $c_1(c)c = cp_{max}/(p_{max} + 2c - 2p_{max})$, which approaches a supremum (in the infinite limit) of $p_{max}/(2 - 2p_{max})$. From Proposition 3, then, there will not be betting if $p_{max}/(2 - 2p_{max}) < 2p_{min}/(1 - p_{min})$; rearranging, we get the result above.

This example illustrates the principle that the introduction of even a moderately informative signal will eliminate betting, if preferences are insufficiently diverse.

- Suppose that the distribution of preference-types is uniform on [0,1], and that each person, independently of preference, receives one of three signals: 0, 1, or q. That is, one either learns the truth-value of E, with probability r, or learns nothing, with probability $1 - r$. Then, regardless of q, betting obtains if and only if $r < 1/2$.

Proof: The proof proceeds along the same lines as above. Routine calculations tell us that $c_2(c)c = rq/((1 - r)(1 - q)) + c/(1 - r)$, which has a clear minimum at $c = 0$ of $rq/((1 - r)(1 - q))$. And $c_1(c)c = c(q - rq)/(q + c(r - rq))$, which approaches a supremum (in the infinite limit) of $(q - rq)/(r - rq)$. The inequality $rq/((1 - r)(1 - q)) < (q - rq)/(r - rq)$ then reduces to $r < 1/2$.

This example again illustrates that an increase in the quality of information decreases the incidence of betting. (In particular, it illustrates an interesting pattern that repeatedly appears, upon closer investigation, though is difficult to formalize in its own right: that, in general, we only have to move "halfway" from the ex-ante state of uncertainty to a state of complete information in order to eliminate all betting.)

Extreme Sensitivity in the Distribution

Having determined the betting behavior possible under a given distribution, one may wonder whether, in some sense, similar betting behavior must occur under similar distributions. We will see that the answer is resoundingly *no*. The problem may be formalized thus:

DEFINITION: A distribution G supports a betting volume of v if there exists odds k such that individuals matched at random would bet with ex-ante probability at least v.

DEFINITION: A distribution G is "sensitive" if it supports a betting volume of $v > 0$ and if there exists a distribution G', centered around the same prior, such that the distribution $G^\varepsilon = (1 - \varepsilon)G + \varepsilon G'$ supports no betting, $\forall \varepsilon > 0$.

That is, a distribution G is sensitive if it supports some positive quantity of betting, but slight deviations from G (in the direction of some G') support no betting.

PROPOSITION 4. A distribution is sensitive if and only if it satisfies the condition specified in Case 2 of Proposition 3 above.

Proof: Suppose G is such that $min_c\, c_2(c)c = max_c\, c_1(c)c$. Consider the distribution G' defined by $G'_E(p,b) = 0$, for $b < 1$, and $G'_E(p,b) = p$, for $b = 1$: that is, the distribution in which preference-types are uniformly distributed on [0,1] and everyone learns the true state. Then, for any $\varepsilon > 0$, $N^\varepsilon_{\neg E}(c) > N_{\neg E}(c)$ and $N^\varepsilon_E(c) < N_E(c)$ $\forall c > 0$. It follows that $c^\varepsilon_2(c)c > c_2(c)c$ and $c^\varepsilon_1(c)c < c_1(c)c$ $\forall c > 0$; accordingly, $min_c\, c^\varepsilon_2(c)c > max_c\, c^\varepsilon_1(c)c$. By Proposition 3, G^ε supports no betting.

Conversely, suppose that there is a (non-singleton) range of values d attained by both $c_1(c)c$ and $c_2(c)c$ on $c \in (0, \infty)$, as in Case 1. In particular, let us consider values d_1 and d_2 where $0 < d_1 < d_2$ and $d_2 - d_1 = \delta$. We will find an ε such that, given any distribution G', the corresponding functions $c^{\varepsilon}{}_1(c)c$ and $c^{\varepsilon}{}_2(c)c$ attain the value $(d_1 + d_2)/2$.

Let C_2 denote the value c such that $c_2(c)c = d_1$, and let C_1 denote the value c such that $c_1(c)c = d_2$. The distribution G' that maximizes the term $c^{\varepsilon}{}_2(C_2)C_2$ will be a distribution such that $N'_{\neg E}(C_2) = 1$ and $N'_E(C_2) = 0$; in this case, the difference between the two terms, which we may denote δ_2, equals $c^{\varepsilon}{}_2(C_2)C_2 - c_2(C_2)C_2$

$$= ((1-\varepsilon)N_{\neg E}(C_2) + \varepsilon) / ((1-\varepsilon)N_E(C_2))C_2 - (N_{\neg E}(C_2) / N_E(C_2))C_2$$

$$= \varepsilon C_2 / ((1-\varepsilon)N_E(C_2)).$$

To ensure that this difference not exceed $\delta/2$, choose

$$\varepsilon \leq A \equiv \delta N_E(C_2) / (2C_2 - \delta N_E(C_2)).$$

Likewise, by considering the distribution G' that minimizes the term $c^{\varepsilon}{}_1(C_1)C_1$, we find that we can ensure that $\delta_1 = c^{\varepsilon}{}_1(C_1)C_1 - c_1(C_1)C_1 \leq \delta/2$ by choosing

$$\varepsilon \leq B \equiv \delta(2N_E(C_1) - N_E(C_1)^2 - 1) / (\delta(N_E(C_1) - N_E(C_1)^2) + 2N_{\neg E}(C_1)C_1 - 2C_1).$$

Given the positive, finite values d_1 and d_2 attained, we know that $N_E(C_1)$, $N_{\neg E}(C_1)$, $N_E(C_2)$, and $N_{\neg E}(C_2)$ must all lie strictly within the interval (0,1). One can then see that A and B must be positive and finite. Letting $\varepsilon = min(A, B) > 0$, we have that, regardless of the choice of G', $c^{\varepsilon}{}_2(C_2)C_2 = d_1 + \delta_2 \leq d_1 + \delta/2 = (d_1 + d_2)/2$, and $c^{\varepsilon}{}_1(C_1)C_1 = d_2 - \delta_1 \geq d_2 - \delta/2 = (d_1 + d_2)/2$. From the connectedness of the graphs, the fact that $c^{\varepsilon}{}_1(0)0 = 0$, and the fact that $c^{\varepsilon}{}_2(c)c$ has no upper bound, it follows that $c^{\varepsilon}{}_1(c)c$ and $c^{\varepsilon}{}_2(c)c$ both attain $(d_1 + d_2)/2$. Therefore G^{ε} supports betting.

PROPOSITION 5. There are sensitive distributions.

Example: Consider the distribution defined by

$$G^{\alpha}_{E}(p,b) = \begin{cases} 0, & b < 1/2 \\ \alpha 2^{(p-1)/p}, & 1/2 \leq b < 1,\ p < 1/2 \\ \alpha(1 - 2^{p/(p-1)}), & 1/2 \leq b < 1,\ p \geq 1/2 \\ (1-\alpha)p + \alpha 2^{(p-1)/p}, & b = 1,\ p < 1/2 \\ (1-\alpha)p + \alpha(1 - 2^{p/(p-1)}), & b = 1,\ p \geq 1/2. \end{cases}$$

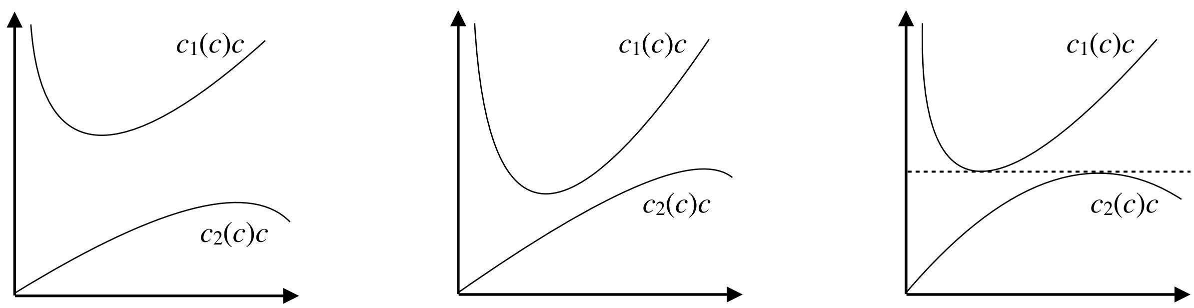

$\alpha = 0.5$ $\alpha = 0.9$ $\alpha \approx 0.790$

Note that G^{α} is the α-average of

(i) a distribution G defined by

$$\begin{aligned} G_E(p,b) &= 0,\ b < 1/2\,; \\ &= 2^{(p-1)/p},\ b \geq 1/2,\ p < 1/2\,; \\ &= 1 - 2^{p/(p-1)},\ b \geq 1/2,\ p \geq 1/2 \end{aligned}$$

(ii) the perfect-information distribution G' (described in the first half of the proof of Proposition 4 above).

By numerical approximation, we achieve the equality of Case 2 at $\alpha^* \approx 0.790$. Under this distribution G^{α^*}, betting obtains at exactly one odds ($k \approx 6.00$), with a betting volume of $v \approx .08$. Even an infinitesimal decrease in α, however—that is, even an infinitesimal averaging of G^{α^*} with G'—renders all betting impossible.

Transaction Costs

Given this sensitivity, another question arises: if a distribution G supports a betting volume of $v > 0$, when the transaction costs of betting are 0 (as implicitly assumed so far), will betting be possible if there are arbitrarily small transaction costs?

In the event of transaction costs, *1* does not receive every dollar that *2* pays him if $\neg E$. That is, whenever *2* effectively faces odds k_2, *1* effectively faces odds $k_1 < k_2$. Agent i's indifference function, therefore, must satisfy $b_i'p - (1-b_i')(1-p)k_i = 0$, $\forall p \in [0,1]$. Now, remembering Figure 2 above: upon multiplying $a(x)$ by a factor of k_2, we only shift $b(x)$ right by a factor of $k_1 < k_2$. It follows that if there exists only a single value d such that, for some (c_1, c_2), $a(c_1) = d/c_1$ and $b(c_2) = d/c_2$, then introducing transaction costs will not allow for any such d; but if there is a range of acceptable d, by the same reasoning as used to prove Proposition 5 above, small transaction costs may be introduced without entirely eliminating betting. In other words,

PROPOSITION 6: "Sensitivity in the distribution" as defined above is equivalent to sensitivity to transaction costs.

Conclusion

Our main result is an extension of the no-bet result to a class of situations involving preference heterogeneity. In particular, the situations considered are those in which individuals receive independent signals and hold linear preferences in each of two states of the world.

Of course, the latter assumption is unrealistic; it entails, for example, that our randomly matched agents will (upon finding mutually agreeable odds) make their bets infinitely large. An original motivation for this simplification, however, was the observation that, so long as people's utility functions are smooth, they will be "locally linear" (around whatever their endowments in each state happen to be). As a result (the thinking went), if we pretend that they have those linear utility functions—those tangent to their actual utility functions at their ex-ante endowments—and then determine the odds at which betting can take place, we will have determined the odds at which, given their *actual* utility functions, "marginal bets" (bets of infinitesimal size) at a given odds will take place. Misguided intuition and an erroneous sketch of a proof then suggested that if we could find the distributions under which no marginal bets would take place, we would more generally have found the distributions under which no bets at all would take place, since increasing the size of the bet could only (it appeared) leave fewer people willing to bet.

This reasoning *does hold* when the distributions of signals and preferences are independent. Under such circumstances, the offer of a large bet cannot convey any more or less information to the other party about the probability of the event than would the offer of a small bet at the same odds. However, this "no small bet implies no large bet" reasoning does not hold in general. If highly risk-averse members of our population tend to receive high-quality signals, the offer of a large bet may go through when the offer of a small bet would not. The acceptance of the large bet makes it common knowledge between the agents that they are less risk-averse, and therefore that their signals are relatively weak; so an agent will not alter his beliefs as much, in light of the other's acceptance, as he would upon observing the other's acceptance of a small bet. This phenomenon is explored in Easley and O'Hara (1987) and related literature.

There is hope, however, that the "linear-preferences-based" analysis presented above is useful even in cases where the distributions of signals and preferences do exhibit some mutual information. If the bets being analyzed here are interpreted to be trades of an uncommon sort of financial product—a particular risky bond, for instance, as proposed in the introduction—then large bets may be infeasible. Since the seller's holdings of this state-contingent good may be negligible in comparison to either party's endowment in each state, their behavior could be expected to approximate the behavior that would be exhibited under linear preferences.

Furthermore, independence is not a necessary condition for more general applications of above analysis. There is in fact a broader class of distributions under which, as desired, the impossibility of small bets entails the impossibility of large ones. Instead of taking the time to characterize these distributions, however, perhaps a more promising project would be to extend this analysis by positing that individuals have a single, concave utility function but different endowments in each state, and by allowing bets to differ in both odds and size. For now, then, though the analysis begun here only partially characterizes the distributions of preferences and signals under which a no-bet result holds, it may serve to shed some light on the broader forces at play in decentralized betting under heterogeneous state-contingent preferences, and to illustrate some of the complications that can arise.

References

Arrow, K. and G. Debreu. 1954. "Existence of an Equilibrium for a Competitive Economy." *Econometrica*, 22(3):265–290

Arrow, K. 1963. "Uncertainty and the Welfare Economics of Medical Care." *American Economic Review*, 53(5):941–973

Aumann, R. 1976. "Agreeing to Disagree." *Annals of Statistics*, 4:1236–1239

Aumann, R. 1999. "Interactive Epistemology I: Knowledge." *International Journal of Game Theory*, 28:263–300

Bacharach, M. 1985. "Some Extensions of a Claim of Aumann in an Axiomatic Model of Knowledge." *Journal of Economic Theory*, 37:167–190

Easley, D. and M. O'Hara. 1987. "Price, Trade Size, and Information in Securities Markets." *Journal of Financial Economics*, 19:69–90

Geanakoplos, J. 1989. "Game Theory without Partitions, and Applications to Speculation and Consensus." *Cowles Foundation Discussion Papers*, No. 914

Geanakoplos, J. and H. Polemarchakis. 1982. "We Can't Disagree Forever." *Journal of Economic Theory*, 28:192–200

Geanakoplos, J. and J. Sebenius. 1983. "Don't Bet on It: Contingent Agreements with Asymmetric Information." *Journal of the American Statistical Association*, 78:424–426

Kim, Jinbeom and T. Leung. 2014. "Impact of Risk Aversion and Belief Heterogeneity on Trading of Defaultable Claims." *Annals of Operation Research*, forthcoming.

Milgrom, P. and N. Stokey. 1982. “Information, Trade and Common Knowledge.” *Journal of Economic Theory*, 26(1):17–27

Monderer, D. and D. Samet. 1989. “Approximating Common Knowledge with Common Beliefs.” *Games and Economic Behavior* 1:170–190

Morris, S. 1994. “Trade with Heterogeneous Prior Beliefs and Asymmetric Information.” *Econometrica*, 62(6):1327–1347

Morris, S. 1995. “The Common Prior Assumption in Economic Theory.” *Economics and Philosophy*, 11:227–253

Morris, S. 1999. “Approximate Common Knowledge Revisited.” *International Journal of Game Theory*, 28:385–408

Neeman, Z. 1993. “Common Beliefs and the Existence of Speculative Trade.” *Games and Economic Behavior* 16(1):77–96

Nielsen, L.T., A. Brandenburger, J. Geanakoplos, R. McKelvey, and T. Page. 1990. “Common Knowledge of an Aggregate of Expectations.” *Econometrica*, 58(5): 1235–1239

Rubenstein, A. 1989. “The Electronic Mail Game: Strategic Behavior under ‘Almost Common Knowledge’.” *American Economic Review*, 79(3):385–391

Sethi, R. and M. Yildiz. 2012. “Public Disagreement.” *American Economic Journal: Microeconomics*, 4(3):57–95

Sonsino, D. 1995. “‘Impossibility of Speculation’ Theorems with Noisy Information.” *Games and Economic Behavior* 8:406–423

Takamiya, K. and A. Tanaka. 2006. "Mutual Knowledge of Rationality in the Electronic Mail Game." *Osaka University Institute of Social and Economic Research Discussion Papers*, No. 650