ROBERT F. NAU

THE INCOHERENCE OF AGREEING TO DISAGREE

ABSTRACT. The agreeing-to-disagree theorem of Aumann and the no-expected-gain-from-trade theorem of Milgrom and Stokey are reformulated under an operational definition of Bayesian rationality. Common knowledge of beliefs and preferences is achieved through transactions in a contingent claims market, and mutual expectations of Bayesian rationality are defined by the condition of joint coherence, *i.e.*, the collective avoidance of arbitrage opportunities. The existence of a common prior distribution and the impossibility of agreeing to disagree follow from the joint coherence requirement, but the prior must be interpreted as a 'risk-neutral' distribution: a product of probabilities and marginal utilities for money. The failure of heterogenous information to create disagreements or incentives to trade is shown to be an artifact of overlooking the potential role of trade in constructing the initial state of common knowledge.

KEY WORDS: Arbitrage, joint coherence, common knowledge, subjective probability, revising probabilities, consensus, rational expectations.

1. INTRODUCTION

If two agents hold the same prior distribution and subsequently receive information via commonly-known partitions, and if they expect each other to rationally update their beliefs according to Bayes' theorem, then they cannot 'agree to disagree' (Aumann, 1976): their posterior distributions must be identical if they, too, are common knowledge. Similarly, risk-neutral agents with common priors cannot devise contingent gambling agreements yielding positive conditional expectations to both sides (Sebenius and Geanakoplos, 1983); and risk-averse traders with concordant beliefs and rational expectations cannot agree to non-null state-contingent trades (Milgrom, 1981; Milgrom and Stokey, 1982). These results are perceived to be "a problem for the theory of speculative markets: asymmetric information alone cannot be responsible for the existence of large stock exchanges. A very important research project in the finance literature is to find where Milgrom–Stokey's model departs from reality. It is a point which is crucial for the understanding of the very complex speculative markets we see nowadays" (Werlang, 1987).

The setting for such results is as follows. Let Ω denote a finite set of states of nature, about which uncertainty will be resolved at some future date. Consider a group of Bayesian rational agents who expect to obtain private information partially resolving their uncertainty about the element $\omega \in \Omega$ which will occur, and who hold commonly known conditional beliefs given such information. In Aumann's (1976) formulation, it is assumed that the agents have a common prior distribution on Ω and that their information partitions are common knowledge.[1] These assumptions suffice to determine each agent's conditional probabilities given her possible states of information and guarantee that the conditional distributions of different agents are consistent. The conditional probabilities are interpreted as the posterior probabilities that agents will actually hold upon receipt of their private information.[2]

The key step in Aumann's argument is the formal definition of common knowledge: the occurrence of an event A is common knowledge at state ω if A includes the member of the *meet* (finest common coarsening) of the agents' information partitions which includes ω. The resulting theorem asserts that whenever the values of the agents' posterior probabilities are common knowledge, they must be identical. Thus, it appears that differential information can never produce a divergence of beliefs under conditions of common knowledge. Geanakoplos and Polemarchis (1982), Sebenius and Geanakoplos (1983), and McKelvey and Page (1986) build explicit models of this disclosure process. Extending this type of result to a market setting, Milgrom and Stokey (1982) show that if agents' prior wealth endowments are Pareto optimal and their beliefs are concordant with respect to payoff-irrelevant events, then the receipt of information cannot create commonly-known incentives to trade: prior endowments remain Pareto optimal *a posteriori*, although prices may change. Bacharach (1985) embeds Aumann's and Milgrom and Stokey's results in a more general framework of epistemic models; Samet (1990) generalizes the agreeing-to-disagree result to non-partitional information structures; Monderer and Samet (1989) derive an approximate version of it by substituting common belief for common knowledge; and Rubinstein and Wolinsky (1990) comment on distinctions between the logical structure of results based on equality judgments (assertions of conditional probability) and

those based on inequality judgments (assertions of conditional preference).

The agreeing-to-disagree and no-expected-gain-from-trade results typically assume that the rationality of the agents and the numerical values of their prior probabilities and utilities are also common knowledge, without specifying how such information comes to be disseminated or verified. This paper will reformulate those results under an operational view of Bayesian rationality, in which information about agents' beliefs and preferences is assumed to be revealed by their behavior in a public market for contingent claims. The focus of the paper is not the definition of common knowledge *per se*: market activity is taken to be common knowledge in the familiar sense of the term. Instead, we inquire into the nature of the common knowledge which can plausibly be sustained in a market context – *e.g.*, to what extent can agents be expected to know each others' beliefs and preferences, and to what extent will their beliefs be consistent?

The assumption that there exists a contingent-claims market through which beliefs and preferences are revealed is not as stylized or restrictive as it might first appear. In the seminal work on subjective probability and expected utility by de Finetti (1937) and Savage (1954), a pervasive theme is that probabilities and utilities should be defined and measured by their manifestations in material behavior – *i.e.*, through some kind of overt gambling or choices among lotteries. For example, one agent may elicit another's probability for an event by asking at what price she is willing to buy or sell lottery tickets contingent on that event, with the understanding that transactions may actually take place. If measurements of this kind are carried out reciprocally by two or more agents with respect to the same events, the result is naturally a market for lottery tickets or other kinds of contingent claims.

Mutual expectations of Bayesian rationality are formalized here by the requirement of *joint coherence*, namely that the public elicitation of beliefs and preferences should create no arbitrage opportunities (Nau and McCardle, 1990; Nau, 1992a). The existence of a common prior distribution and the agreeing-to-disagree and no-expected-gain-from-trade results are shown to follow from the joint coherence requirement. However, all is not what it seems: the agents' marginal utilities for money are generally confounded with the appar-

ent probabilities revealed by their behavior in the market (Kadane and Winkler, 1988; Nau and McCardle, 1991). The common prior therefore must be interpreted as a product of possibly-heterogeneous probabilities and marginal utilities: a 'risk-neutral probability distribution' in the terminology of financial markets.

The contribution of this paper is to show that Aumann's and Milgrom–Stokey's results lose their paradoxical quality when the process which generates common knowledge is explicitly modeled in this way. In a market setting, trade is the precursor of commonly known risk-neutral probabilities, hence it is implausible to assert that agents will ever know (let alone embrace) each other's true probabilities, and tautological to assert that common knowledge removes the incentive to trade.

2. MARKET PRICES AS COMMON KNOWLEDGE

Assume the existence of a market in which an agent can publicly announce 'bid' or 'ask' prices for contingent claims, where a contingent claim is a vector of monetary payoffs indexed by $\omega \in \Omega$ and conditioned on an event in Ω. The symbols E and F will be used interchangeably as names for contingent claims and events and also as the corresponding payoff or indicator vectors. Thus, $E(\omega) \in \Re$ is the payoff of claim E in state ω, $F(\omega) = 1(0)$ if the event F is true (false) at ω, and $\neg F$ is synonymous with $1 - F$. In the market, an agent may *bid* a conditional price of p for a claim E given an event F (written $E|F$), meaning that any other agent is free to announce a non-negative number α of such claims that she will sell to the first agent at this price,[3] and a contract is then enforced in which the net payment to the first agent from the second in state ω will be $\alpha(E(\omega) - p)F(\omega)$. In this case the vector $(E - p)F$ constitutes an *acceptable trade* for the first agent. Similarly, an agent may *ask* a conditional price of q for $E|F$, meaning that another agent may announce a number β of claims she will buy from the first agent at this price, and the state-contingent payoff to the first agent from the second will be $\beta(q - E(\omega))F(\omega)$. In this case, the vector $(q - E)F$ is an acceptable trade for the first agent. Trades are assumed to be additive as well as non-negatively scalable, so that if an agent is currently bidding p_1 and p_2 for $E_1|F_1$ and $E_2|F_2$, then she is assumed willing

to accept a net trade of $\alpha_1(E_1 - p_1)F_1 + \alpha_2(E_2 - p_2)F_2$ for small, non-negative α_1, α_2. Trading takes place in continuous time, so that an arbitrarily short interval may elapse between the announcement of an acceptable trade and the response to it by another agent, and prices may be adjusted continuously.

Let T denote the time index and suppose that common knowledge of conditional bid/ask prices is achieved at $T = 0$, while the state of nature will be fully revealed at $T = 1$. (At some $T < 1$, the agents may receive private information concerning which of their conditioning events has occurred, but the timing of this receipt of information will not be critical here.) Let n denote the number of agents and let J_i denote a set of index numbers for conditional claims on which bid/ask prices are posted by agent i. Thus, E_{ij} is a claim on which agent i posts bid and ask prices conditional on event F_{ij} for all $j \in \{1, \ldots, J_i\}, i \in \{1, \ldots, n\}$. The conditional claims $\{E_{ij}|F_{ij}\}$ need not be distinct.

To characterize the initial common-knowledge state, suppose the market is opened at some $T < 0$ and, perhaps after an initial flurry of trading and price adjustments, a cleared market is achieved at $T = 0$ in which prices are stable and trading has ceased, after which the market is closed. In particular, suppose that at $T = 0$ agent i is bidding p_{ij} and asking q_{ij} for claim E_{ij} conditional on F_{ij}, for all $j \in \{1, \ldots, J_i\}, i \in \{1, \ldots, n\}$, but no agent finds any takers at these prices. Then in a practical sense it is common knowledge at $T = 0$ that agent i's conditional valuation of $E_{ij}|F_{ij}$ is between p_{ij} and q_{ij}. This is the common knowledge of 'infinite specularity' (Dupuy, 1989) in which agents face each other in continuous time and see the same persistent truth reflected in each other's eyes. Here, the possibility of actual trade at advertised prices is what gives rise to specularity: the canonical language of communication consists of price quotations $(\{p_{ij}\}, \{q_{ij}\})$ and numbers of claims offered for purchase or sale $(\{\alpha_{ij}\}, \{\beta_{ij}\})$ at the quoted prices. At $T = 0$, agent i bids p_{ij} for $E_{ij}|F_{ij}$ ('this is my price'), the other agents decline by announcing $\alpha_{ij} = 0$ ('we're not interested'), agent i reiterates her bid in light of this knowledge ('I hear you but this is still my price'), they reiterate their declination in light of this knowledge ('we hear you but we're still not interested'), and so on. This yields common knowledge in the formal sense if infinitely many such messages

are imagined to flash back and forth in finite time. If the sequence is finitely truncated, neither party can be sure that the other is not on the verge of changing her response, as in the 'electronic mail game' (Rubinstein, 1989). Nonetheless, the degree of reciprocal knowledge achieved in a public market is arguably the inspiration for, and the closest real approximation to, the mathematical ideal of common knowledge. As Monderer and Samet (1989) observe: "Clearly not even *homo rationalis* checks the validity of infinitely many statements one by one. Still, everybody understands public announcements to [be] common knowledge."

3. COHERENCE AS BAYESIAN RATIONALITY

In the tradition of de Finetti (1937, 1974) and Smith (1961), we consider agent i to be Bayesian rational if her announced bid and ask prices for contingent claims are *coherent – i.e.,* do not expose her to arbitrage. In this case, as is well known, there exists a probability distribution on Ω under which all of her acceptable trades have non-negative expected value. Thus, an agent is Bayesian rational if and only if her acceptance of trades appears to be predicated on calculations of expected value with respect to some probability distribution on states. More precisely, in view of the additivity and scalability assumptions, agent i's announcement of bid/ask prices $\{p_{ij}, q_{ij}\}$ for $\{E_{ij}|F_{ij}\}$ means that her *set of acceptable trades* is the set of all payoff vectors of the form

$$(1) \qquad \sum_{j=1}^{J_i} \alpha_{ij}(E_{ij}(\omega) - p_{ij})F_{ij}(\omega) + \beta_{ij}(q_{ij} - E_{ij}(\omega))F_{ij}(\omega)$$

for (small) non-negative $\{\alpha_{ij}\}$, $\{\beta_{ij}\}$. An arbitrage opportunity (alias a 'Dutch book') against agent i is an acceptable trade whose payoff to her is negative in all states of nature. By a separating-hyperplane theorem (Gale, 1960, Theorem 2.10), there is no arbitrage opportunity against agent i if and only there exists a distribution π on Ω such that

$$(2) \qquad \sum_{\omega \in \Omega} \pi(\omega)(E_{ij}(\omega) - p_{ij})F_{ij}(\omega) \geqslant 0 \quad \text{and}$$

$$\sum_{\omega \in \Omega} \pi(\omega)(q_{ij} - E_{ij}(\omega))F_{ij}(\omega) \geqslant 0$$

for all $j \in \{1, \ldots, J_i\}$. Let the 'π' notation be stretched to define $\pi(F)$ and $\pi(E|F)$ as the probability of F and conditional expectation of $E|F$ induced by the distribution π:

$$\pi(F) \equiv \sum_{\omega \in \Omega} \pi(\omega) F(\omega), \quad \text{and}$$

$$\pi(E|F) \equiv \pi(EF)/\pi(F) \quad \text{if} \quad \pi(F) > 0.$$

The inequalities (2) are then equivalent to the condition that, for every $j \in \{1, \ldots, J_i\}$, either $p_{ij} \leqslant \pi(E_{ij}|F_{ij}) \leqslant q_{ij}$ or else $\pi(F_{ij}) = 0$. On this basis, we define agent i's bid and ask prices $\{p_{ij}\}$ and $\{q_{ij}\}$ to be her *lower and upper conditional expectations* for the claims $\{E_{ij}|F_{ij}\}$, and the effect of market clearing at $T = 0$ is to render these expectations common knowledge. This definition of lower and upper conditional expectations can be extended to claims other than $\{E_{ij}|F_{ij}\}$ as follows:

DEFINITION. Agent i's conditional expectation for $E|F$ is at least p [not more than q] if her set of acceptable trades includes a vector which is equal to or weakly dominated by $\alpha(E - p)F[\alpha(q - E)F]$ for some $\alpha > 0$.

In the special case where E is (like F) the indicator vector of an event, p and q are *lower and upper conditional probabilities* for $E|F$ as defined by Smith (1961).

4. JOINT COHERENCE AS MUTUAL EXPECTATIONS OF BAYESIAN RATIONALITY; THE NATURE OF THE COMMON PRIOR

Assuming that trades are additive between agents (*i.e.*, that they are in a common currency), the *aggregate set of acceptable trades* at $T = 0$ is the set of all vectors of the form:

$$(3) \quad \sum_{i=1}^{n} \sum_{j=1}^{J_i} \alpha_{ij}(E_{ij}(\omega) - p_{ij})F_{ij}(\omega) + \beta_{ij}(q_{ij} - E_{ij}(\omega))F_{ij}(\omega)$$

for (small) nonnegative $\{\alpha_{ij}\}, \{\beta_{ij}\}$. Bayesian rationality of all agents, as defined in the preceding section, requires that each agent's conditional expectations should 'cohere' in the sense of not admitting arbitrage opportunities. Mutual expectations of Bayesian rationality

require something more, namely that the conditional expectations of all agents should cohere 'jointly'. This is formalized in the following

DEFINITION. The agents' conditional expectations are *jointly coherent* if the aggregate set of acceptable trades contains no arbitrage opportunities (strictly negative vectors).

This requirement has been shown by Nau and McCardle (1990) to capture the intuitive idea of mutually expected Bayesian rationality in the setting of a noncooperative game, where uncertainty exists with respect to strategies of agents as well as states of nature. In the present setting, it captures the idea that not only is it irrational to create arbitrage opportunities against oneself, but it is also irrational to overlook arbitrage opportunities created by the collective behavior of the group.[4] For, common knowledge of the agents' valuations of contingent claims cannot be said to exist until the market has cleared and trading has ceased, and agents who know each other to be rational will not expect trading to cease in a state where unexploited arbitrage opportunities remain. (If an arbitrage opportunity existed, then any agent could arrange a transaction in which she would receive a strictly positive additional amount of money in every state of nature, possibly in conjunction with a trade acceptable to herself. A minimal standard of mutually expected rationality is that the agents should not observe each other passing up such free lunches.) The separating hyperplane argument now yields:

THEOREM 1. *The agents' conditional expectations are jointly coherent if and only if there exists a distribution π such that for every i and j, either $p_{ij} \leqslant \pi(E_{ij}|F_{ij}) \leqslant q_{ij}$ or else $\pi(F_{ij}) = 0$.*

The distribution π, which need not be unique, plays the role of the 'common prior' whose existence is taken as a primitive assumption by Harsanyi (1967) and Aumann (1976, 1987). Here, its existence is a consequence of the operational definition of common knowledge of Bayesian rationality: probabilities and expectations which are common knowledge in a market setting must also be common property. For example, the fact that agent i's ultimate bid of p_{ij} for $E_{ij}|F_{ij}$ is not acted upon by the other agents evidently means that no other agent's valuation of $E_{ij}|F_{ij}$ is strictly less than p_{ij}.

The apparent common prior need not represent the 'true' probabilities of any agent. As pointed out by Kadane and Winkler (1988), probabilities elicited via monetary transactions will be confounded with marginal utilities for wealth. When making comparisons among gambles, an agent will weigh the monetary payoff in each state not only in proportion to its probability but also in proportion to her marginal utility for additional wealth in that state. Hence, the apparent probabilities revealed through her acceptance of small gambles or choices among small lotteries will equal the renormalized product of her true probabilities and her marginal utilities for money in different states of nature. In the literature of finance, the probability distribution on states of nature whose existence is necessary and sufficient for the absence of arbitrage opportunities in a contingent claims market is known as a 'risk neutral' probability distribution (Cox and Ross, 1976), because it is as if the representative investor is risk neutral and sets prices equal to expected returns using this distribution, even though actual investors are presumably risk averse.[5]

This interpretation of the common distribution helps to explain the convergence to market-clearing prices: if the agents' true beliefs are heterogeneous and they are risk averse, then it will be mutually profitable for them to engage in trade up to the point at which the renormalized products of their probabilities and marginal utilities are equalized. Part of this convergence may be due to updating of beliefs as information is communicated back and forth, but the remainder will be due to shifts in marginal utility due to transfers of wealth. For example, an agent who initially attaches exceptionally high probability to a given state will wish to undertake trades increasing her wealth in that state, other things being equal, thereby eventually offsetting her high probability with decreases in marginal utility. Conversely, an agent whose initial wealth in a given state is exceptionally low will wish to undertake trades increasing her wealth in that state, other things being equal, because this will have a relatively high marginal utility. In general, it will be difficult for one agent to tell whether another's enthusiasm for increasing her wealth in a particular state is due to high initial probability or low initial wealth or both, hence communication-through-trade is unlikely to terminate in a reconciliation of 'true' beliefs.

To illustrate this process, suppose that there is a single conditional claim $E|F$ on which the agents bid directly, so that the ij subscript may be suppressed, and suppose that E is simply the indicator of an event (*i.e.*, a claim with payoffs of 0 or 1). The minimal relevant state space is then $\Omega = \{\omega_1, \omega_2, \omega_3\}$, with $\omega_1 = E \cap F, \omega_2 = \neg E \cap F$, and $\omega_3 = \neg F$. Let π_i denote agent i's (unobservable) true subjective probability distribution on Ω, and let v_i denote agent i's (unobservable) state-dependent marginal utility for money, assumed to be strictly positive. Then agent i should trade as though her probability distribution were proportional to $\pi_i v_i$. In particular, letting $\pi_i(E|F) \equiv \pi_i(\omega_1)/(\pi_i(\omega_1) + \pi_i(\omega_2))$ denote agent i's true conditional probability for event E given event F, and letting $v_i(E|F) \equiv v_i(\omega_1)/(v_i(\omega_1) + v_i(\omega_2))$ denote her relative marginal utility for money in event E given event F, her bid and ask prices for $E|F$ should satisfy:

$$p_i \leqslant \frac{\pi_i(E|F)v_i(E|F)}{\pi_i(E|F)v_i(E|F) + (1 - \pi_i(E|F))(1 - v_i(E|F))} \leqslant q_i,$$

where the quantity in the denominator is merely the renormalizing constant. (Note that $1 - \pi_i(E|F)$ and $1 - v_i(E|F)$ are, respectively, agent i's conditional probability for $\neg E$ given F and her relative marginal utility in the event $\neg E$ given F.) Assume that neither agent maintains a bid–ask spread, so that these relations hold with equality. Then, for example, if $\pi_1(E|F) = 0.85$ and $v_1(E|F) = 0.4$, we obtain $p_1 = q_1 \approx 0.71$, whereas if $\pi_2(E|F) = 0.15$ and $v_2(E|F) = 0.6$ we obtain $p_2 = q_2 \approx 0.29$. If these are the values which prevail at some $T < 0$ when the market opens, then because $p_1 > q_2$, agent 1 will buy claims to $E|F$ while agent 2 sells. If both agents are risk averse, this will produce a decrease [increase] in $v_1(E|F)[v_2(E|F)]$, and there may also be a downward [upward] revision in $\pi_1(E|F)[\pi_2(E|F)]$ due to the informativeness of the trades. Consequently, p_1 will fall and q_2 will rise until they meet (and trade ceases) at time $T = 0$, which occurs when $\pi_1(E|F) = v_2(E|F)$ and $\pi_2(E|F) = v_1(E|F)$ as illustrated in Figure 1. (There, the argument '$(E|F)$' has been suppressed: π_1 stands for $\pi_1(E|F)$ etc.) Suppose that equality is obtained at $\pi_1(E|F) = v_2(E|F) = 0.75$ and $\pi_2(E|F) = v_1(E|F) = 0.25$, yielding (by symmetry) $p_1 = q_2 = 0.5$. This is the apparent common prior probability, $\pi(E|F)$, to which the players have converged, while their true probabilities
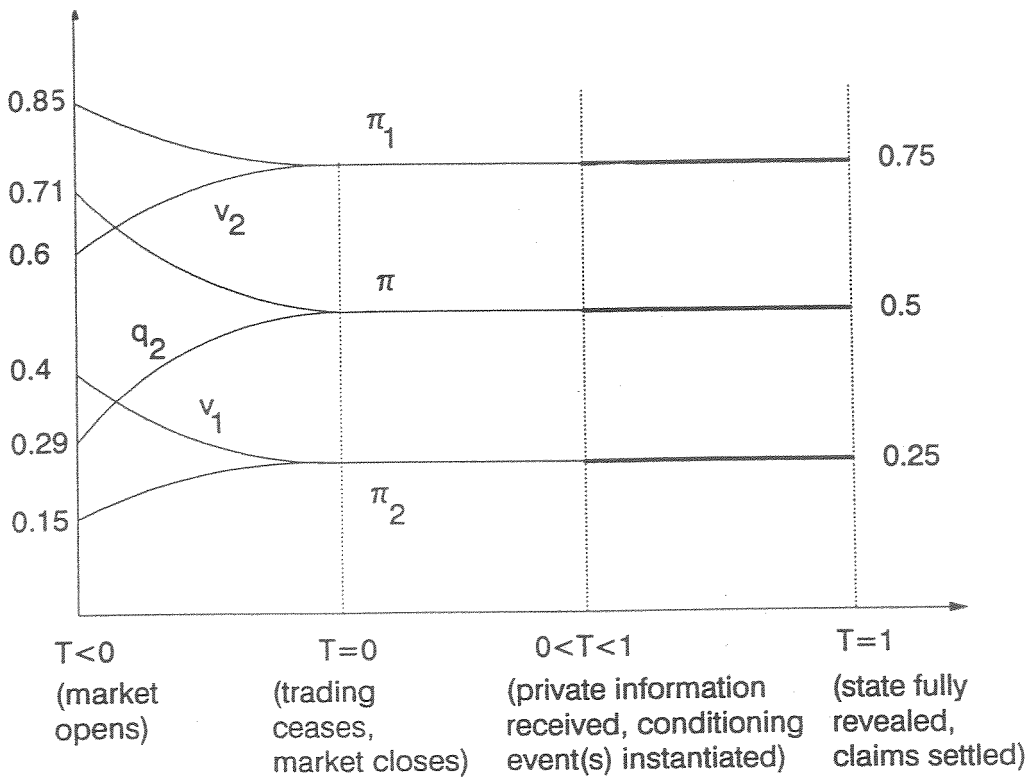
Fig. 1. Market dynamics in the case of a single conditional event $E|F$. π_i =agent i's true probability for $E|F$; v_i =agent i's relative marginal utility in event E given f; $p_i[q_i]$ =agent i's buying [selling] price for a claim on $E|F$; π =common risk neutral probability for $E|F$ after $T = 0$.

remain stuck apart at 0.75 and 0.25, respectively. To complete the scenario, if the conditioning event F is instantiated at some $0 < T < 1$, then (assuming the market is closed after $T = 0$, so that no other information is received) the conditional probabilities become 'posterior' probabilities at that point, and the E-claims are in force. Finally, the remainder of the state specification (E or $\neg E$) is revealed at $T = 1$, and the claims are settled. It is significant that all the interesting activity occurs prior to $T = 0$ as the agents approach – or rather *construct* – the initial state of common knowledge.

5. INFORMATION PARTITIONS

The role of information partitions in agreeing-to-disagree models is to allow commonly-known conditional ('posterior') probabilities to be inferred from common prior probabilities. Here, in contrast, con-

ditional probabilities are directly observed in market prices, and they are *ipso facto* common knowledge; the existence of a common prior is inferred from the absence of arbitrage. However, if the agents have information partitions, these can be revealed through the structure of the contingent claims priced in the market. For example, if agent i has partition $\mathcal{P}_i$, then it is prudent for her to bid only on contingent claims whose conditioning events are measurable with respect to $\mathcal{P}_i$. By so doing, she is able to offer the narrowest possible bid-ask spreads without accepting trades she might regret after receiving her information, or without disclosing that information if she has already received it. By observing each other's conditioning events for contingent claims, the agents can thus infer each other's partitions.

For example, consider a situation in which there are two agents and the state space comprises four events $\{A, B, C, D\}$. Suppose agent 1 asserts that 3/5 is her lower probability for $B|A \cup B$ and 4/5 is her lower probability for $C|C \cup D$; while agent 2 asserts that 1/3 is his lower probability for $A|A \cup C$ and 1/4 is his lower probability for $D|B \cup D$. Then these four lower conditional probabilities determine a unique common prior distribution $\{0.2, 0.3, 0.4, 0.1\}$, even though neither agent has revealed a complete distribution of his or her own. These assertions also suggest that the first agent's information partition is $\{A \cup B, C \cup D\}$ while the second agent's partition is $\{A \cup C, B \cup D\}$, and that the agents have not yet received (or least have not disclosed) their private information.

6. AGREEING TO DISAGREE

To replicate the main results of Aumann (1976) and Sebenius and Geanakoplos (1983) in the present framework, it is not necessary to start with the assumption of fully-specified information partitions. It suffices instead to assume that there is some claim for which the lower conditional expectation of one agent and the upper conditional expectation of another can both be inferred. The following will also be needed:

DEFINITION. An event F is *non-null* if the aggregate set of acceptable trades includes a trade which is equal to or weakly dominated by $\alpha(F - p')$ for some $p' > 0$ and $\alpha > 0$.

Equivalently, F is non-null if $\pi(F) \geqslant p' > 0$ for every π which satisfies (2) for all i and j. In other words, a non-null event is one to which the agents have collectively assigned a strictly positive subjective probability.

THEOREM 2. *If* (i) *the agents' conditional expectations are jointly coherent, and* (ii) F *is non-null, and* (iii) *agent 1's conditional expectation for $E|F$ is at least p and agent 2's conditional expectation for $E|F$ is no more than q, then $p \leqslant q$.*

Proof. By assumptions (ii) and (iii), the trade $p' - F$ is acceptable in the aggregate for some $p' > 0$, and $(E - p)F$ and $(q - E)F$ are acceptable to agents 1 and 2, respectively. By additivity, $\alpha(F - p') + (E - p)F + (q - E)F = (\alpha + q - p)F - \alpha p'$ is also acceptable in the aggregate for $\alpha > 0$, and this is strictly negative for any α between 0 and $p - q$ if $p > q$. Consequently, (i) holds along with (ii) and (iii) only if $p \leqslant q$.[6] ∎

Hence, the agents cannot hold discordant expectations for a claim which is conditioned on an event to which they assign positive probability. This is the analog of Sebenius and Geanakoplos' Proposition 2 in our framework, and it implies Aumann's theorem if E is the indicator of an event. To make these connections clearer, the assumption of fully-specified information partitions can now be introduced. Suppose that agent 1's and agent 2's conditioning events are all measurable with respect to partitions $\mathcal{P}_1$ and $\mathcal{P}_2$, respectively, with the interpretation that these partitions describe the information they expect to hold at some $T < 1$. Suppose that all events in Ω are non-null, and that there is some event E for which both agents have revealed their lower and upper probabilities conditional on all possible states of their information; and let these lower and upper probabilities be denoted by $p_i(E|F')$ and $q_i(E|F')$, for all $F' \in \mathcal{P}_i, i \in \{1, 2\}$. Now let F denote an event which is an element of the meet of their partitions, and define

$$p_1 \equiv \min_{F' \in \mathcal{P}_1, F' \subseteq F} p_1(E|F') \quad \text{and}$$

$$q_2 \equiv \max_{F' \in \mathcal{P}_2, F' \subseteq F} q_2(E|F')$$

Then it is common knowledge in Aumann's sense, as well as ours, that agent 1's [2's] conditional probability for $E|F$ is at least p_1 [not

more than q_2], and F is non-null, whence Theorem 2 requires $p_1 \leqslant q_2$. Similarly, we can infer $p_2 \leqslant q_1$ under corresponding definitions of p_2 and q_1, and in the special case where $p_1 = q_1$ and $p_2 = q_2$ it follows that $p_1 = p_2 = q_1 = q_2$: if their posterior probabilities are not only common knowledge but 'sharp', they must be identical.

7. NO EXPECTED GAIN FROM TRADE

The no-expected-gain-from-trade results of Milgrom (1981) and Milgrom and Stokey (1982) follow from the joint coherence requirement in a similar fashion. Unlike Aumann, Milgrom and Stokey allow the agents to hold different subjective probabilities for 'payoff-relevant' events, but they implicitly assume a common prior risk-neutral distribution, as will be seen. The state space in their model is partitioned as $\Omega = S \times X$, where S is the set of payoff-relevant events and X is an additional set of purely informational events, and agents hold state-contingent allocations of l different commodities. Let $e_i : \Omega \mapsto \Re_+^l$ denote the state-contingent prior commodity endowment of agent i, and let $t_i : \Omega \mapsto \Re^l$ be defined as an *acceptable trade* for agent i if she weakly prefers $e_i + t_i$ to her current holdings e_i. e_i and t_i will be manipulated as vectors with generic elements $e_i(\omega, c)$ and $t_i(\omega, c)$, where $\omega \in \Omega$ is the index for states of nature and $c \in \{1, \ldots, l\}$ is the index for commodities. Although endowments are written as functions of $\omega = (s, x)$, they are assumed to depend only on s, the payoff-relevant component of the state; and a trade which depends only on s will be called an 's-trade'. If $F \subset \Omega$ is an event, t_i is defined to be *conditionally acceptable* to agent i given the occurrence of F if the trade Ft_i is acceptable to her.

DEFINITION. $t = (t_1, \ldots, t_n)$ is a *feasible trade* among the agents if:

$$\sum_{i=1}^{i=n} t_i(\omega, c) \leqslant 0 \quad \forall \omega, c \text{ (it requires no net inputs)}, \quad \text{and}$$

$$e_i(\omega, c) + t_i(\omega, c) \geqslant 0 \quad \forall i, \omega, c$$
   (it yields no short positions).

Implicit in this definition of feasibility is the assumption that all commodities have non-negative intrinsic worth: they can be freely disposed of but not freely created. Milgrom and Stokey assume that: (i) all events are non-null; (ii) agents have information partitions, whose structure is common knowledge, according to which their beliefs will be updated; (iii) their conditional probabilities for payoff-irrelevant events given payoff-relevant events are concordant (*i.e.,* identical); (iv) their preferences among endowments are described by strictly concave (risk-averse) von Neumann–Morgenstern utility functions which depend only on s; and (v) their initial endowments are *ex ante* Pareto optimal relative to s-trades. They consider the situation in which it is common knowledge in Aumann's sense that an s-trade is *ex post* acceptable to all agents, and show that this can only be the null trade. (Under weak risk aversion, the proposed trade need not be null, but the agents must all be indifferent between it and the null trade.)

To recast these results in our framework, first observe that Milgrom–Stokey's assumptions (iii) and (v) imply a common risk-neutral prior distribution on Ω if one of the commodities plays the role of money. To see this, note that *ex ante* Pareto optimality of the initial endowments relative to s-trades means precisely that the agents already hold common risk-neutral probabilities for all payoff-relevant events. The assumption that conditional beliefs about payoff-irrelevant events are concordant then serves to extend this common risk-neutral distribution to all of Ω. [7] Once this fact has been observed, it is intuitively obvious that the agents have no incentives to make s-trades conditioned on the occurrence of commonly-known events – the existing endowments are effectively Pareto optimal on $S \times X$, not merely S, by virtue of the common risk-neutral distribution – but let us continue anyway.

Assume a market in which, as before, agents may publicly announce trades they are willing to accept, and other agents are free to enforce such trades. Suppose that a cleared market is reached at $T = 0$ in which $\mathcal{T}_i$ is the set of trades asserted to be acceptable by trader i, and the market is then closed. Thus, the preferences of all agents for state-contingent trades are rendered common knowledge at $T = 0$, and they evolve subsequently, if at all, only through the receipt of private information. Assume that commodity 1 is money, of which

more is strictly preferred to less by all agents, and let m denote the vector representing one unit of money. (That is, $m(\omega, 1) = 1$ and $m(\omega, c) = 0$ if $c \neq 1$, for all ω.)

DEFINITIONS. The preferences of agent i are *strictly convex* if $t_i \in \mathcal{T}_i$ and $t_i \neq 0$ imply that for any $\alpha \in (0, 1)$ there exists $\epsilon > 0$ such that $\alpha t_i - \epsilon m \in \mathcal{T}_i$. An *arbitrage opportunity* is a feasible trade which is acceptable to all agents and which has a strictly negative aggregate input of money in every state of nature. The preferences of all agents are *jointly coherent* if there are no arbitrage opportunities.

Strict convexity of preferences replaces the assumptions of intra-agent scalability and additivity of trades which were used earlier. This condition is implied by, and hence is weaker than, Milgrom and Stokey's assumptions that all events are non-null and every agent has strictly concave von Neumann–Morgenstern utility. Joint coherence embodies mutual expectations of Bayesian rationality and also guarantees Pareto optimality of the prior wealth allocation: violation of joint coherence would permit money to be pumped out of the economy while leaving every agent in an at-least-as-preferred position. Redistribution of this money to one or more agents would then make these agents strictly better off. (For a more general duality theorem for exchange economies with multiple commodities, see Nau and McCardle, 1991). Under these assumptions and definitions, we have:

THEOREM 3. *If the agents' preferences are strictly convex and jointly coherent, then a trade is feasible and acceptable to all agents only if it is the zero trade.*

Proof. Suppose that t is feasible and acceptable, $t \neq 0$, and preferences are strictly convex. Then for any $\alpha \in (0, 1)$ and every i there exists ϵ_i such that $\alpha t_i - \epsilon_i m$ is also acceptable to agent i. It follows that $(\alpha t_1 - \epsilon_1 m, \ldots, \alpha t_n - \epsilon_n m)$ is an arbitrage opportunity, violating the condition of joint coherence. ■

In particular, for any event F and any s-trade t, the conditional trade Ft can be feasible and acceptable without giving rise to arbitrage opportunities *ex ante* only if $Ft = 0$, which means that t can be

acceptable given F only if it is the zero trade *ex post*. This is theorem 2 of Milgrom (1981) and Theorem 1 of Milgrom and Stokey (1982) for the case of strict risk aversion. As in the preceding section, the correspondence can be made more explicit by introducing information partitions and taking F to be an element of the meet, but this additional structure is inessential to the main result. The role of information partitions in Milgrom–Stokey's model is to identify a set of potential conditioning events (namely, the elements of the meet) with respect to which the conditional preferences of all the agents are commonly known *ex post* – given that their prior probabilities and utility functions are somehow already known. But, if it is common knowledge *ex post* at ω that an s-trade t is acceptable, and if F is the element of the need that includes ω, then it must also be common knowledge *ex ante* that the conditional trade Ft is acceptable. The same effect is obtained if the agents simply announce in the marketplace the conditional trades they find acceptable and discover that, for some event F and s-trade t, Ft is feasible and acceptable to all.

8. DISCUSSION

The objectively given prior probabilities in Aumann's model have been replaced here by subjective probabilities which are revealed by the behavior of agents in a public market. In this setting, trade is the precursor of the state of common prior beliefs which is assumed to exist at $T = 0$, and the revealed probabilities are risk neutral probabilities – *i.e.* products of true probabilities and relative marginal utilities for money. The assumption that true prior probabilities are common property as well as common knowledge, which is standard in much of game theory and financial economics, therefore appears incompatible with a realistic view of communication and trade in markets. In Milgrom and Stokey's model, the failure of new information to generate trade under conditions of common knowledge has been traced to the fact that the allocation of state-contingent wealth is effectively assumed to be already Pareto optimal on the entire state space: the role of trade is merely obscured because it has been relegated to an earlier point in time. Taken together, these results reinforce rather than contradict the intuitive conviction that

agents with heterogeneous information will maintain heterogeneous beliefs and engage in trade.

These results also have implications concerning the characterization of *correlated equilibria* in noncooperative games (Nau, 1992b). 'Objective' correlated equilibria are Bayesian game solutions predicated on a common prior distribution, while 'subjective' correlated equilibria admit different priors (Aumann, 1974, 1987). The latter concept appears more harmonious with a personalistic view of probability, but it is unreasonably coarse, since it places virtually no restrictions on the players' beliefs. However, it can be shown that if the players' beliefs and preferences are articulated via gambles in a public market, then joint coherence requires the play of the game to be consistent with a subjective correlated equilibrium in which the *risk neutral* probability distributions of the players are commonly held, though their true probabilities may differ. This may be considered a 'refinement' of subjective correlated equilibrium, insofar as it imposes a strong consistency condition on otherwise-heterogeneous beliefs.

The models presented here have been interpreted from the viewpoint of subjective-expected-utility theory, but the full axiomatic structure of SEU theory has not been employed. Instead, the key axiom which has been used to define Bayesian rationality is that of joint coherence – *i.e.*, no arbitrage. This requires agents to behave individually in ways which are consistent with SEU-maximization and to behave collectively in a way which is Pareto efficient, but it does not require them to reveal complete preference orderings nor to uniquely separate probabilities from utilities. This approach is consistent with a view that rationality is fundamentally a property of markets, which are composed of individuals with bounded cognitive and computational abilities.

ACKNOWLEDGMENTS

supported by the Business Associates Fund at the Fuqua School of Business.

NOTES

[1] The assumption that information partitions are common knowledge is to some extent self-referential, although the self-references can be removed by embedding the original state space in a larger one which allows the construction of an infinite regress (Tan and Werlang 1988).

[2] The distinction between 'conditional' and 'posterior' probabilities is itself controversial – e.g., see Hacking (1967) and Goldstein (1983, 1985). This distinction is germane to understanding the dynamic behavior of agents in markets, but beyond the scope of the present paper.

[3] It is assumed that the number of claims for sale at a given price is finite but 'small'. Without loss of generality, explicit bounds could be imposed. Claims are assumed to be infinitely divisible; thus, α may be fractional.

[4] The definition of joint coherence given in Nau and McCardle (1990) and Nau (1992a) is slightly stronger than that used here, inasmuch as it rules out arbitrage opportunities *ex post* as well as *ex ante*. The *ex post* version is appropriate if the outcome space includes decisions of agents as well as states of nature, or if the possibility exists that agents already possess their private information at $T = 0$. Under the *ex post* version of joint coherence, the players are considered irrational if the outcome which occurs is one to which they have collectively assigned a subjective probability of zero.

[5] For a more thorough discussion of risk neutral probabilities, see Nau and McCardle (1991). The interpretation of the probability distribution derived from equilibrium contingent-claim prices as a product of probabilities and relative marginal utilities is discussed in Drèze's classic paper (1970) on market allocation under uncertainty (based on a 1965 lecture). The same notion appears, although somewhat less transparently, in Wilson's (1968) paper on Pareto optimal sharing rules and Beja's (1967) thesis on capital markets, and it implicitly underlies much of the arbitrage literature on financial asset pricing of the 1970's and 1980's.

[6] The proof is trivial once the framework has been set in place, but the same is true of Aumann's (1976) result, as Aumann himself points out. Of course, if p and q are directly asserted bid/ask prices for $E|F$ and F is the certain event, then $p \leqslant q$ must hold or else a trivial arbitrage opportunity exists and trade should not have ceased. The content of the theorem is merely that the same is true even if p and q are indirectly determined by application of the additivity and scalability assumptions and/or F is not certain.

[7] Agent i's risk neutral probability for $\omega = (s, x) \in \Omega$ is proportional to $\pi_i(s, x) v_i(s, x)$, where $\pi_i(s, x)$ is her true probability and $v_i(s, x)$ is her marginal utility for money in state (s, x) under her initial endowment. By definition, x is payoff-irrelevant, so $v_i(s, x)$ must be independent of x, hence agent i's risk neutral probability for (s, x) is proportional to $\pi_i(s, x) v_i(s)$, which can be factored as $\pi_i(x|s)(\pi_i(s) v_i(s))$. But, *ex ante* Pareto optimality of initial endowments relative to s-trades implies that the factor $\pi_i(s) v_i(s)$ is the same (up to a multiplicative

constant) for all agents, and concordance implies that $\pi_i(x|s)$ is also the same for all agents, so their product is the same for all agents upon renormalization.

REFERENCES

Aumann, R.: 1974, 'Subjectivity and Correlation in Randomized Games', *J. Math. Econ.* **1**, 67–96.

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

Aumann, R.: 1987, 'Correlated Equilibrium as an Expression of Bayesian Rationality', *Econometrica* **55**, 1–18.

Bacharach, M.: 1985 'Some extensions of a Claim of Aumann in an Axiomatic Model of Knowledge', *J. Econ. Theory* **37**, 167–190.

Beja, Avraham: 1967, 'Capital Markets with Delayed "Learning".' Ph.D. Thesis, Stanford University.

Cox, J. and Ross S.: 1976, 'The Valuation of Options for Alternative Stochastic Process', *Financial Econ.* **3**(1/2), 145–166.

De Finetti, B.: 1937, 'Foresight: Its Logical Laws, Its Subjective Sources', Translation reprinted in H.E. Kyburg and H.E. Smokler (Eds.): 1980, *Studies in Subjective Probability*, Robert Krieger, New York, pp. 53–118.

De Finetti, B.: 1974, *Theory of Probability, Vol. 1* Wiley, New York.

Drèze, J.: 1970, 'Market Allocation Under Uncertainty', *European Econ. Rev.* **2**, 133–165.

Dupuy, J.P.: 1989, 'Common Knowledge, Common Sense', *Theory and Decision* **27**, 37–62.

Gale, D.: 1960, *The Theory of Linear Economic Models*, McGraw-Hill, New York.

Geanakoplos, J. and Polemarchakis, H.: 1982, 'We Can't Disagree Forever', *J. Econ. Theory* **28**(1), 192–200.

Goldstein, M.: 1983, 'The Prevision of a Prevision', *J. Amer. Statist. Assoc.* **78**, 817–819.

Goldstein, M.: 1985, 'Temporal Coherence (with discussion)', in Bernardo, J., DeGroot, M., Lindley, D. and Smith, A. (eds.): *Bayesian Statistics 2: Proceedings of the Second Valencia International Meeting* North-Holland, New York, pp. 231–248.

Hacking, I.: 1967, 'Slightly More Realistic Personal Probability', *Philosophy of Science* **34**, 311–325.

Harsanyi, J.: 1967, 'Games with Incomplete Information Played by 'Bayesian' Players: I–III', *Management Science* **14**, 159–182, 320–334, 486–502.

Kadane, J. and Winkler, R.: 1988, 'Separating Probability Elicitation from Utilities', *J. Amer. Statist. Assoc.* **83**(402), 357–363.

McKelvey, R. and Page, T.: 1986, 'Common Knowledge, Consensus, and Aggregate Information', *Econometrica* **54**(1), 109–127.

Milgrom, P.: 1981, 'An Axiomatic Characterization of Common Knowledge', *Econometrica* **49**(1), 219–222.

Milgrom, P. and Stokey, N.: 1982, 'Information, Trade, and Common Knowledge', *J. Econ. Theory* **26**, 17–27.

Monderer, D. and Samet, D.: 1989, 'Approximating Common Knowledge with Common Beliefs', *Games and Economic Behavior* **1**(2), 170–190.

Nau, R. and McCardle, K.: 1990, 'Coherent Behavior in Noncooperative Games', *J. Econ. Theory* **50**(2), 424–444.

Nau, R. and McCardle, K.: 1991, 'Arbitrage, Rationality, and Equilibrium', *Theory and Decision* **31**, 199–240.

Nau, R.: 1992a, 'Joint Coherence in Games of Incomplete Information', *Management Science* **38**(3), 374–387.

Nau, R.: 1992b, 'The Relativity of Subjective and Objective Correlated Equilibria', Fuqua School of Business Working Paper 9102, December 1992.

Rubinstein, A.: 1989, 'The Electronic Mail Game: Strategic Behavior Under Almost Common Knowledge', *Amer. Econ. Rev.* **79**(3), 385–391.

Rubinstein, A. and Wolinsky, A.: 1990, 'On the Logic of "Agreeing to Disagree" Type Results', *J. Econ. Theory* **51**, 184–193.

Samet, D.: 1990, 'Ignoring Ignorance and Agreeing to Disagree', *J. Econ. Theory* **52**, 190–207.

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

Sebenius, J. and Geanakoplos, J.: 1983, 'Don't Bet On It: Contingent Agreements with Asymmetric Information', *J. Amer. Statist. Assoc.* **78**(382), 424–426.

Smith, C.A.B.: 1961, 'Consistency in Statistical Inference and Decision', *J. Roy. Stat. Soc. B* **23**, 1–25.

Tan, T.C. and da C. Werlang, S.R.: 1988, 'On Aumann's Notion of Common Knowledge – An Alternative Approach', CARESS Working Paper No. 88–09.

Werlang, S.R. da C.: 1987, 'Common Knowledge', in Eatwell, J., Milgate, J. and Newman, P. (Eds.), *Game Theory: The New Palgrave.*, W.W. Norton, New York, pp. 74–85

Wilson, R.: 1968, 'The Theory of Syndicates', *Econometrica* **36**(1), 119–132.

Robert F. Nau
The Fuqua School of Business,
Duke University,
Durham, North Carolina 27708-0120,
U.S.A.