Game Theory Without Partitions, and Applications to Speculation and Consensus

John Geanakoplos*

June 18, 2002

Abstract

Decision theory and game theory are extended to allow for information processing errors. This extended theory is then used to reexamine market speculation and consensus, both when all actions (opinions) are common knowledge and when they may not be. Five axioms of information processing are shown to be especially important to speculation and consensus. They are called nondelusion, knowing that you know, nested, balanced, and positively balanced. We show that it is necessary and sufficient that each agent's information processing errors be nondeluded and (1) balanced so that the agents cannot agree to disagree, (2) positively balanced so that it cannot be common knowledge that they are speculating, and (3) KTYK and nested so that agents cannot speculate in equilibrium. Each condition is strictly weaker than the next one, and the last is strictly weaker than partition information.

> "Is there any other point to which you would wish to draw my attention?"
> "To the incident of the dog in the night-time."
> "The dog did nothing in the night-time."
> "That was the curious incident," remarked Sherlock Holmes.
>
> Doyle (1901)

Sherlock Holmes is the perfectly rational Bayesian optimizer that economic models assume all agents are like. Yet most economic actors are probably much more like Dr. Watson, than like Sherlock Holmes. They usually take signals at face value.

*This paper was first presented in July 1988 in Columbus, Ohio. I would like to acknowledge very helpful conversations with Paul Milgrom, Don Brown, Adam Brandenburger, and Eddie Deckel. Don Brown and I are collaborating on a closely related study of common knowledge. Brandenberger, Deckel, and I are writing a sequel to this paper [BDG88] from which I have borrowed the notion of decision-theoretic equivalence used here. Bob Willig prompted the discussion of nested, and John Vickers drew my attention to Sherlock Holmes. A forthcoming paper by Rubinstein and Wolinsky touches on some of the issues raised here.

They often take no notice when something doesn't happen. They occasionally ignore unpleasant information. They forget. And sometimes their opinions won't budge unless confronted by proof.

The aim of this paper is to develop a coherent definition of equilibrium in games that allows for such errors in information processing. My point of view is that behavior may be only boundedly rational, but it need not be any the less goal oriented and purposeful, and therefore any the less predictable.

The bulk of this paper is devoted to analyzing the phenomenon of market speculation, and "agreements to disagree." By now it is well known that neither of these commonplace events can be observed in equilibrium in a model of Bayesian rational agents. When agents are allowed to make errors in information processing, however, these phenomena do eventually emerge in equilibrium. The "curious incident" is the precise stage of irrationality at which they emerge. It is not true, for example, that an agent who always ignores unpleasant information is necessarily a sucker for a bet. It turns out that there is a substantial degree of information processing error (which is characterized here) that can occur, and still agents will not speculate against each other in equilibrium. There is a greater degree of irrationality, which can be specified exactly, that if not exceeded will keep agents from speculating against each other when those actions are common knowledge. And finally there is a still more serious kind of error, which again can be stated precisely, which if not exceeded will never permit agents to agree to disagree.

1 Errors in Information Processing

There are a number of errors that are typically made by decision makers that suggest that we go beyond the orthodox Bayesian paradigm. We list some of them:

1. Agents ignore the subtle information content of some signals, and perceive only their face value. For example, an order to "produce 100 widgets" might convey all kinds of information about the mood of the boss, the profitability of the widget industry, the health of fellow workers and so on, if the receiver of the message has the time and capacity to think about it long enough. Another important example involves prices. It is very easy to compute the cost of a basket of goods at the going prices, but it takes much longer to deduce what the weather must have been like all across the globe to explain those prices. In Bayesian decision making, it is impossible to perform the first calculation without also performing the second.

2. Agents often do not notice when nothing happens. For example, it might be that there are only two states of nature: either the ozone layer is disintegrating or it is not. One can easily imagine a scenario in which a decaying ozone layer would emit gamma rays. Scientists, surprised by the new gamma rays would investigate their cause, and deduce that the ozone was disintegrating. If there were no gamma rays, scientists would not notice their absence, since they might

never have thought to look for them, and so might incorrectly be in doubt as to the condition of the ozone.

3. What one knows is partly a matter of choice. For example, some people are notorious for ignoring unpleasant information. Often there are other psychological blocks to processing information.

4. People often forget.

5. Knowledge derived from proofs is not Bayesian. A proposition might be true or false. If an agent finds a proof for it, he knows it is true. But if he does not find a proof, he does not know it is false.

6. People cannot even imagine some states of the world.

We can model some aspect of all of these non-Bayesian methods of information processing by generalizing the notion of partition from the usual Bayesian analysis. Let Ω, a finite set, represent the set of all possible (physical) states of the world. Let $P : \Omega \to 2^{\Omega} \backslash \phi$ be an arbitrary "possibility correspondence," representing the information processing capacity of an agent. For each $\omega \in \Omega$, $P(\omega)$ is interpreted to mean the collection of states the agent thinks are possible when the true state is ω. Let $\underline{P}$ denote the range of P, so $\underline{P} = \{R \subset \Omega | \exists \omega \in \Omega,\ R = P(\omega)\}$. Given an arbitrary event $A \subset \Omega$, we say that the agent *knows* A *at* ω if $P(\omega) \subset A$, since for any $\omega' \in \Omega$ which he regards as possible at ω, $\omega' \in A$.

Consider, for example, $\Omega = \{a, b\}$ as the state space. Let the possibility correspondence $P : \Omega \to 2^{\Omega}$ take $P(a) = \{a\}$ and $P(b) = \{a, b\}$. We can interpret $\omega = a$ to mean the ozone layer is disintegrating, or a horse is winning, or a proposition is true. Similarly, we can interpret $\omega = b$ to mean the ozone is not disintegrating, or the horse is losing, or the theorem is false.

Since $P(a) = \{a\}$, when $\omega = a$ the agent knows that his horse is winning, or that the ozone is disintegrating, or that the theorem is true. But when $\omega = b$ the agent has no idea whether his horse is winning or losing, or what is happening to the ozone, or whether the theorem is true or false. The reason for the asymmetry in the agent's information processing could be interpreted as any one of the above categories of errors. The agent might take notice of the gamma rays in state a, but not notice that there were no gamma rays in state b. In the horse racing interpretation of the model, the agent might not be able to face the unpleasant news that his favorite horse is damaged. Or he might not remember an event where nothing of interest happened to him.

In Bayesian decision theory, the information possibility correspondence always defines a partition: for every ω, $\omega' \in \Omega$, $\omega \in P(\omega)$ and either $P(\omega) \cap P(\omega') = \phi$ or else $P(\omega) = P(\omega')$. A Bayesian decision-maker could not have the possibility correspondence in the above example. He would reason at $\omega = b$ that since he did not receive the signal $\{a\}$, that in fact the state must be b. In Bayesian decision theory the information at a state ω is always consistent with what could be deduced from knowing the model and the signal:

$$P(\omega) = \{\omega' \in \Omega | P(\omega') = P(\omega)\}.$$

For arbitrary possibility correspondence, such as in the above example, this need not be the case.

Observe that in the generalized possibility approach to knowledge, there need not be any presumption that the agent understands the entire state space Ω. It may well be that the sets of possibilities in $\underline{P}$ are all confined to some small subset of Ω. In that case there would be $\omega \in \Omega$ such that $\omega \not\in P(\omega')$, for any $\omega' \in \Omega$. Such ω are not even imaginable by the agent. Similarly if $\omega \not\in P(\omega)$, then when ω actually occurs the agent does not think of it, although he might at other times.

In the next few sections we shall describe how decision theory and game theory can be extended to generalized partitions. Many phenomenon (such as betting) which cannot be observed in equilibrium when every agent has partition information will now become possible. To give content to our extension, however, it is necessary to categorize precisely the kinds of information processing errors which can occur, and which kinds of errors permit each new phenomenon. Betting, for example, can be an equilibrium even when agents always imagine the truth, provided they make other errors. On the other hand betting is ruled out by a degree of rationality that falls short of partition information.

We shall now describe three limitations on the possibility correspondence. Later we shall introduce two more.

Definition: We say that P is *nondeluded* if $\omega \in P(\omega)$ for all $\omega \in \Omega$. Under this hypothesis the agent who processes information according to P always considers the true state as possible.

Definition (Knowing that you know (KTYK): When Knowledge is Self-Evident): If for all $\omega \in \Omega$, and all $\omega' \in P(\omega)$, we have $P(\omega') \subseteq P(\omega)$, then we say that the agent knows what he knows. If the agent knows some A at ω, and can imagine ω', then he would know A at ω'. Bacharach ([Bac85]), Shin ([Shi87]), and Samet ([Sam87]) have all drawn attention to this property. If the agent can recognize circumstances which confine the possible states of the world to $R \in \underline{P}$, then whenever $\omega \in R$, so that these circumstances do indeed obtain, the agent must realize that.

Definition: The event $E \subset \Omega$ is *self-evident* to the agent who processes information according to P if $P(\omega) \subset E$ whenever $\omega \in E$. A self-evident event can never occur without the agent knowing that it has occurred.

The axiom KTYK implies that every $R \in \underline{P}$ is self-evident to the agent.

Shin [Shi87] has suggested that KTYK and nondelusion are the only properties that need hold true for an agent whose knowledge was derived by logical deductions from a set of axioms.

Definition: We say that P is *nested* if for all ω and ω', either $P(\omega) \cap P(\omega') = \phi$, or else $P(\omega) \subseteq P(\omega')$, or else $P(\omega') \subseteq P(\omega)$.

An example might make the significance of nondelusion, KTYK, and nested clearer. Let there be just two propositions of interest in the universe, and let us suppose that whether each is true or false is regarded as good or bad, respectively. The state space is then $\Omega = \{GG, GB, BG, BB\}$. One type of information processor P might always disregard anything that is bad, but remember anything that is good. Then $P(GG) = \{GG\}$, $P(GB) = \{GG, GB\}$, $P(BG) = \{GG, BG\}$, $P(BB) = \Omega$. (See Diagram 1a.) It is clear that P satisfies nondelusion and KTYK, but does not satisfy nested. Moreover, when the reports are GB the agent chooses to remember only the first, while if they are BG he chooses to remember only the last. Alternatively, consider an agent with possibility correspondence Q who can remember GG and BB because the pattern is simple, and can also remember when he sees GB that the first report was good whereas with BG he remembers nothing. Then $Q(GG) = \{GG\}$, $Q(GB) = \{GB, GG\}$, $Q(BG) = \Omega$, $Q(BB) = \{BB\}$. (See Diagram 1b.) This does satisfy nested, as well as the other two conditions. Nondelusion in these examples means that the agent never mistakes a good report for a bad report, or vice versa. KTYK means that if an agent recalls some collection of reports, then whenever all those reports turn out the same way he must also recall them (and possibly some others as well). Nested means that the reports are ordered in the agent's memory. If he remembers some report, then he must also remember every report that came earlier on the list.

We shall prove in Section 3 that nested can be interpreted as a property of *memory* in this way: Suppose that we think of a set S of fundamental propositions that can be either true or false. A state $\omega \in \Omega$ specifies which of these propositions are true, and which are false. Suppose that knowledge at any ω can be described by a subset $S(\omega) \subseteq S$. The agent knows at ω whether or not each proposition in $S(\omega)$ is true or false. In other words, $P(\omega) = \{\omega' \in \Omega | s \in S(\omega) \Rightarrow [s \text{ is true at } \omega' \text{ iff } s \text{ is true at } \omega]\}$. Finally, let us suppose that the propositions in S can be ordered (say chronologically) and that with respect to this ordering $S(\omega)$ is always an initial set, for any ω. Then P is nested (and nondeluded). Moreover, any nested and nondeluded P can be equivalently described this way. Nested corresponds to memory in the sense that the agent always remembers more or less far down the list of S, perhaps depending on how complicated the pattern of truth valuations is, but always in the same order.

Diagrams 1a

Nondeluded, KTYK, but not nested

Diagram 1b

Nondeluded, nested, and KTYK

Note that nested and KTYK are independent properties. Let $\Omega = \{a, b, c\}$, and let $P(a) = P(c) = \{a, b, c\}$, while $P(b) = \{b, c\}$. Then P is nondeluded and nested, but P does not satisfy KTYK, since $c \in P(b)$ but $P(c) \subseteq P(b)$.

2 Decision Theory Without Partitions

Our purpose in this paper is to analyze decision-making and game theory in environments where information processing is subject to error. Consider the following canonical decision problem:

Let A be a set of possible actions. Let $u : A \times \Omega \to R$. Let π be a measure[1] on Ω. Let $P : \Omega \to 2^{\Omega}$ be a possibility correspondence. We call a decision function $f : \Omega \to A$ *optimal* for the decision problem (A, Ω, P, u, π) iff

Condition (1): $[P(\omega) = P(\omega')] \Rightarrow [f(\omega) = f(\omega')]$.
Condition (2): For all $\omega \in \Omega$ and $a \in A$,

$$\sum_{\omega' \in P(\omega)} u(f(\omega), \omega')\pi(\omega') \geq \sum_{\omega \in P(\omega)} u(a, \omega')\pi(\omega').$$

This definition applies for any possibility correspondence, whether or not it is a partition. Notice both conditions (1) and (2) serve to limit choices to reflect the level of information. Condition (1) requires that the agent's action is a function of what he perceives, and condition (2) requires that the agent optimizes, taking his information at face value.

In the above definition the agent is effectively unaware that he is erring in his information processing. Given the information that the state of nature $\omega' \in R \equiv P(\omega)$, the agent routinely uses Bayes Law to update his beliefs and to optimize.[2] Were he aware of his errors, he would refine the possibility correspondence into a partition by letting $\hat{P}(\omega) = \{\omega' \in \Omega | P(\omega') = P(\omega)\}$.

In the above decision framework the agent does not completely understand the model. He also does not necessarily "know what he is doing." If he knew his optimal plan $f : \Omega \to A$, and knew at each ω what his choice ought to be, then he would further refine his information according to the partition $Q_f(\omega) = \{\omega' \in \Omega | f(\omega') = f(\omega)\}$. It might then turn out that his optimizing behavior would no longer correspond to f.

To illustrate the definition above we shall shortly present three examples which shall be of further use in later sections. At the same time we investigate the precise sense in which an agent who knows more but is boundedly rational may be worse off than if he knew less but was unboundedly rational.

A fundamental consequence of Bayesian decision making, and unbounded rationality, is that knowing more can never be disadvantageous. If $P : \Omega \to 2^{\Omega}$ and $Q : \Omega \to 2^{\Omega}$ then we say that Q *is coarser than* P if $P(\omega) \subseteq Q(\omega)$ for all ω. If g is optimal for (A, Ω, Q, u, π), and Q is coarser than P, and P and Q are partitions, then

Condition (3): $\sum_{\omega \in \Omega} u(g(\omega), \omega)\pi(\omega) \leq \sum_{\omega \in \Omega} u(f(\omega), \omega)\pi(\omega)$.

[1] Typically we shall suppose that π is a probability on Ω, so $\sum_{\omega \in \Omega} = 1$, but it is convenient for technical reasons to allow for the more general situation where π is a measure.

[2] In case $\pi(R) = 0$, the above definition allows the agent's action to be arbitrary. A more sophisticated approach would define conditional probabilities on measure zero events, but we do not consider these extensions here.

It turns out that this property of Bayesian decisions is at the heart of the non-speculation literature. By allowing for less rational information processing it need no longer be the case that more knowledge is better.

In fact, one wonders if there are any general properties at all that can be proved outside the Bayesian framework. We shall show however that there are. Indeed the "more is better property" applies to a more general set of information correspondences than partitions.

Example 2.1: Let $\Omega = \{a, b, c\}$, $P(a) = \{a, b\}$, $P(b) = \{b\}$, $P(c) = \{b, c\}$. (Note that (Ω, P) satisfies nondeluded and KTYK, but not nested.) Let $\pi(a) = \pi(c) = 2/7$ and $\pi(b) = 3/7$. Let the action set be $A = \{B, N\}$, corresponding to bet or not bet. Let the payoffs from not betting be $u(N, a) = u(N, b) = u(N, c) = 0$. Let the payoffs to betting be $u(B, a) = u(B, c) = -1$, while $u(B, b) = 1$. It is easy to calculate that $f(\omega) = B$ for all $\omega \in \Omega$ is optimal for (A, Ω, P, u, π). Yet,

$$\sum_{\omega \in \Omega} u(N, \omega)\pi(\omega) = 0 > (-1/7) = \sum_{\omega \in \Omega} u(B, \omega)\pi(\omega).$$

Of course $g(\omega) = N$ for all $\omega \in \Omega$ is optimal for (A, Ω, Q, u, π) where $Q(\omega) = \Omega$ for all ω, so for this example inequality (3) fails.

Example 2.2: Let $\Omega = \{a, b, c\}$, and let $P(a) = P(c) = \{a, b, c\}$, while $P(b) = \{b, c\}$. Then (Ω, P) is nondeluded and nested, but does not satisfy KTYK. Let $A = \{B, N\}$, let $\pi(\omega) = 1/3$, for all $\omega \in \Omega$, and let $u(N, \omega) = 0$ for all $\omega \in \Omega$, while $u(B, a) = -2$, $u(B, b) = -2$, $u(B, c) = 3$. Then $f = (f(a), f(b), f(c)) = (N, B, N)$ is optimal for (A, Ω, P, u, π), but

$$\sum_{\omega \in \Omega} u(N, \omega)\pi(\omega) = 0 > -\tfrac{2}{3} = \sum_{\omega \in \Omega} u(f(\omega), \omega)\pi(\omega).$$

Once again let the coarse partition be $Q(\omega) = \Omega$ for all ω. Then $g(\omega) = N$ for all $\omega \in \Omega$ is optimal for (A, Ω, Q, u, π), and again (3) is violated.

Example 2.3: Let $\Omega = \{a, b\}$, let $P(a) = \{a\}$, $P(b) = \{a, b\}$. Then (Ω, P) satisfies all these properties nondeluded, nested, and KTYK. Let $A = \{B, N\}$, let $\pi(a) = \pi(b) = 1/2$. Let $u(N, a) \equiv u(N, b) = 0$, let $u(B, a) = 1$, and $u(B, b) = -2$. Then $f(a) = B$, $f(b) = N$ is optimal for (A, Ω, P, u, π), and

$$\sum_{\omega \in \Omega} u(N, \omega)\pi(\omega) = 0 < \tfrac{1}{2} = \sum_{\omega \in \Omega} u(f(\omega), \omega)\pi(\omega).$$

In this case inequality (3) holds. Observe also that if we changed the payoff at (B, b) to $u(B, b) = -1$, then there would be a second optimal decision function $\tilde{f}(a) = B = \tilde{f}(b)$. In that case $\sum_{\omega \in \Omega} u(\tilde{f}(\omega), \omega)\pi(\omega) = 0$ is worse than the payoff arising from f, but still as good as the payoff arising from (A, Ω, Q, u, π) where $Q(\omega) = \Omega$ for all $\omega \in \Omega$.

In Examples 2.1 and 2.2, the agent was (ex ante) worse off knowing more because he did not process information coherently. In Example 2.3 the agent was not worse off, although he also did not process information correctly. In general we have:

Theorem 1 *Let (Ω, P) satisfy nondeluded, nested, and* KTYK. *Let Q be a partition of Ω that is a coarsening of P. Let f, g be optimal for (A, Ω, P, u, π) and (A, Ω, Q, u, π) respectively. Then*

$$\sum_{\omega \in \Omega} u(g(\omega), \omega)\pi(\omega) \leq \sum_{\omega \in \Omega} \leq u(f(\omega), \omega)\pi(\omega).$$

Conversely, suppose that (Ω, P) fails to satisfy one or more of the above hypotheses. Then there is a partition Q of Ω that is a coarsening of P, and an A, u, π such that f, g are optimal for (A, Ω, P, u, π), (A, Ω, Q, u, π), respectively, and yet the above inequality is strictly reversed.

Proof The proof of the theorem proceeds by induction on the cardinality of Ω. Suppose $\#\Omega = 1$. Then $P(\omega) = \Omega = Q(\omega)$ for all $\omega \in \Omega$, and there is nothing to prove. Suppose the theorem is true for $\#\Omega \leq k$. Consider the case where $\#\Omega = k+1$. Let $S = \{\omega \in \Omega : P(\omega) = \Omega\}$. If $S = \Omega$, then again there is nothing to prove. If $S \neq \Omega$, let $\Omega^1 = P(\bar{\omega})$ be any possibility set with the greatest cardinality less than $k+1$. Then by nondelusion $0 < \#\Omega^1 \leq k$, and $\#\Omega \setminus \Omega^1 \leq k$. From knowing that you know, if $\omega \in \Omega^1$, $P(\omega) \subset \Omega^1$ so $S \cap \Omega^1 = \phi$. Let $\Omega^2 = \Omega \setminus [\Omega^1 \cup S]$. From nondeluded and nested, if $\omega \in \Omega^2$, $P(\omega)) \cap \Omega^1 = \phi$. From KTYK, nondeluded, and the definition of S, if $\omega \in \Omega^2$, $P(\omega) \cap S = \phi$. Hence, if $\omega \in \Omega^2$, $P(\omega) \subset \Omega^2$. Let I be the partition of Ω formed by the disjoint sets Ω^1, Ω^2, and S. Consider the partition $Q^* = Q \vee I$, defined by $Q^*(\omega) = Q(\omega) \cap I(\omega)$ for all $\omega \in \Omega$. Let g^* be optimal for (A, Ω, Q^*, u, π). Then because Q and Q^* are partitions, and Q is coarser than Q^* on $\Omega^1 \cup \Omega^2$,

$$\begin{aligned} &\sum_{\omega \in \Omega^1} u(g^*(\omega), \omega)\pi(\omega) + \sum_{\omega \in \Omega^2} u(g^*(\omega), \omega)\pi(\omega) + \sum_{\omega \in S} u(f(\omega), \omega)\pi(\omega) \\ \geq &\sum_{\omega \in \Omega^1} u(g(\omega), \omega)\pi(\omega) + \sum_{\omega \in \Omega^2} u(g(\omega), \omega)\pi(\omega) + \sum_{\omega \in S} u(f(\omega), \omega)\pi(\omega). \end{aligned}$$

But now we can apply the induction hypothesis to (A, Ω^1, P, u, π) and (A, Ω^2, P, u, π), obtaining

$$\begin{aligned} \sum_{\omega \in \Omega^1} u(f(\omega), \omega)\pi(\omega) &\geq \sum_{\omega \in \Omega^1} u(g^*(\omega), \omega)\pi(\omega) \text{ and} \\ \sum_{\omega \in \Omega^2} u(f(\omega), \omega)\pi(\omega) &\geq \sum_{\omega \in \Omega^2} u(g^*(\omega), \omega)\pi(\omega). \end{aligned}$$

Finally, let us observe that if $S = \phi$, we are finished. If there is some $\hat{\omega} \in S$, then Q is the trivial partition and we can assume WLOG that $g(\omega) = f(\hat{\omega})$ for all $\omega \in \Omega$, and in particular $g(\omega) = f(\omega)$ for all $\omega \in S$. This concludes the proof, since the converse follows from Examples 2.1 and 2.2. □

We conclude this section by noting one important extension to decision theory that fits naturally into our framework. Let $\bar{A}$ be a correspondence specifying for each $\omega \in \Omega$ the set of possible actions in some ambient space A, $\bar{A} : \Omega \to 2^A$. Then we would regard u as a function on $A \times \Omega$, and an optimal decision plan for $(\bar{A}, A, \Omega, P, u, \pi)$ would be a function $f : \Omega \to A$ satisfying

Condition (1′): $[P(\omega) = P(\omega')$ and $\bar{A}(\omega) = \bar{A}(\omega')] \Rightarrow [f(\omega) = f(\omega')]$
Condition (2′a): $f(\omega) \in \bar{A}(\omega)$ for all $\omega \in \Omega$ and
Condition (2′b): For all $\omega \in \Omega$ and $a \in \bar{A}(\omega)$

$$\sum_{\omega' \in P(\omega)} u(f(\omega), \omega')\pi(\omega') \geq \sum_{\omega' \in P(\omega)} u(a, \omega')\pi(\omega').$$

This new formulation allows us to model the idea that agents take the face value of a message, and restrict their choices accordingly, without using the subtle content of the message (i.e., without using knowledge of the function $\bar{A}$ to invert the signal and so to deduce more about the state). We shall return to this theme in Section 10.

Corollary 2 *Let $(\bar{A}, A, \Omega, P, u, \pi)$ be a decision problem with variable constraints. Let (Ω, P) satisfy nondeluded, nested, and KTYK. Let $[P(\omega) = P(\omega')] \Rightarrow [\bar{A}(\omega) = \bar{A}(\omega')]$. Let Q be a partition of Ω that is a coarsening of P, and let $\hat{A}(\omega) \subset \bar{A}(\omega)$ for all ω satisfy $[Q(\omega) = Q(\omega')] \Rightarrow [\hat{A}(\omega) = \hat{A}(\omega')]$. If f, g are optimal for $(\bar{A}, A, \Omega, P, u, \pi)$ and $(\hat{A}, A, \Omega, Q, u, \pi)$, respectively, then $\sum_{\omega \in \Omega} u(g(\omega), \omega)\pi(\omega) \leq \sum_{\omega \in \Omega} u(f(\omega), \omega)\pi(\omega)$.*

Proof: The proof is exactly as for Theorem 1. □

3 Equivalent Decision Problems

Evidently the naming of states is somewhat arbitrary. For example, splitting a state into two indistinguishable states, which are physically indistinguishable, should not change the decision problem. By formulating several definitions of equivalent decision problems we can clarify the framework of Section 2.

Definition: We say that the decision problem $(A, \Omega', P', u', \pi')$ is a *renaming* of the decision problem (A, Ω, P, u, π) in the following senses:
Decision-theoretic, if there is a 1–1 and onto map $\delta : \underline{P'} \to \underline{P}$ such that for all $R' \in \underline{P'}$, if $R = \delta(R')$ then $\pi(R) > 0$ if and only if $\pi'(R') > 0$ and if both are positive, then for all $a \in A$,

$$\frac{1}{\pi(R)} \sum_{\omega \in R} u(a, \omega)\pi(\omega) = \sum_{\omega' \in R'} u'(a, \omega')\pi'(\omega');$$

Physical, if there is a function $\varphi : \Omega' \to \Omega$ that is onto and satisfies

$$u(a, \varphi(\omega')) = u'(a, \omega') \text{ for all } \omega' \in \Omega' \text{ and } a \in A$$
$$\pi(\omega) = \sum_{\omega' \in \varphi^{-1}(\omega)} \pi'(\omega') \text{ for all } \omega \in \Omega$$
$$P'(\omega') = \varphi^{-1}(P(\varphi(\omega')) \text{ for all } \omega' \in \Omega'.$$

We say that two decision problems are equivalent (in either of the two senses) if they have a common renaming. An easy argument shows that these are indeed equivalence relations, and that a physical renaming is also a decision-theoretic renaming.

The decision theoretic equivalence notion was formulated by Brandenburger–Deckel–Geanakoplos [BDG88]. If it holds and if agents always optimize according to (1) and (2), then behaviorally the decision problems are equivalent. The following lemma also appears in [BDG88]:

Lemma 3 *For any decision problem* (A, Ω, P, u, π) *there is a decision-theoretic renaming* $(A, \Omega', P', u', \pi')$ *in which* P' *is a partition of* Ω'.

Sketch of Proof: Let $\Omega' = \underline{P} \times \Omega$. Define $u' : A \times \Omega' \to \mathbb{R}$ by $u'(a, (R, \omega)) = u(a, \omega)$, $P'(R, \omega) = \{R\} \times \Omega$, and $\delta(\{R\} \times \Omega) = R$, and let $\pi'(R, \omega) = \left\{ \begin{smallmatrix} 0 & \text{if } \omega \notin R \\ \pi(\omega) & \text{if } \omega \in R \end{smallmatrix} \right.$. $\square$

The upshot of Lemma 5 is that we can understand the information processing and decision problem of Section 2 *as if* the agent is a conventional maximizer, but has got the prior (on Ω') wrong. The "correct" priors would be those that would obtain if the agent knew the function P, namely $\pi^*(R, \omega) = \left\{ \begin{smallmatrix} \pi(\omega) & \text{if } P(\omega) = R \\ 0 & \text{otherwise} \end{smallmatrix} \right.$.

The lemma thus provides us with another interpretation of decision theory with generalized partitions. If the reader wished, he could rewrite all of our results in terms of the consequences of using the wrong priors (or in later sections, of players using different priors). The advantage of the generalized possibility approach is that it explains how the mistaken priors might have arisen. Theorem 1, for example, gives conditions under which the agent does at least as well as he would with the right priors but less information. One perhaps could reformulate the result directly in terms of priors, but nondelusion, KTYK, and nested make clear just what information processing errors are tolerable.

We illustrate Lemma 3 with the ozone example in which $\Omega = \{a, b\}$, $\pi(a) = \pi(b) = 1/2$, $P(a) = \{a\}$, $P(b) = \{a, b\}$. Let $\gamma = \{a\}$ correspond to "gamma rays" and $n = \{a, b\}$ correspond to no gamma rays. Then $\Omega' = \underline{P} \times \Omega = \{\gamma a, \gamma b, na, nb\}$, $P'(\gamma a) = P'(\gamma b) = \{\gamma a, \gamma b\}$, $P'(na) = P'(nb) = \{na, nb\}$, and $\pi'(\gamma a) = 1/2$, $\pi'(\gamma b) = 0$, $\pi'(na) = \pi'(nb) = 1/2$. The correct priors are $\pi^*(\gamma a) = 1/2$, $\pi^*(\gamma b) = \pi^(na) = 0$, and $\pi^*(nb) = 1/2$. Note that the only partition of Ω weaker than P is the trivial partition, which under the transformation of Lemma 3 becomes the trivial partition Q' of Ω'. Theorem 1 can be interpreted to mean that any optimal plan $f : \Omega' \to A$ with respect to the prior π' and the partition P' does at least as well, evaluated according to the correct priors π^*, as any plan feasible (i.e., measurable) with respect to the partition Q'.

Let us now introduce the idea of a particularly simple, concrete description of the set of states of nature. Let us call Ω a *propositional state space* if $\Omega = 2^n$. Each $\omega \in \Omega$ can be interpreted as a truth assignment to each of n ordered propositions, and we can represent ω as an n-tuple of binary digits: $\omega = (\omega(1), ..., \omega(q), ..., \omega(n))$. We call $A \subset 2^n$ a *basic propositional event* if there is some proposition q such that A corresponds to all states assigning the same truth valuation to q. More precisely, there is $\bar{\omega} \in \Omega$ such that $A = \{\omega \in \Omega | \omega(q) = \bar{\omega}(q)\}$. A *propositional event* $A \subset \Omega$ is a nonempty intersection of basic propositional events: there exists $\bar{\omega} \in \Omega$ and $1 \leq q_1 \leq \cdots \leq q_m \leq n$ such that $A = \{\omega \in \Omega | \omega(q_i) = \bar{\omega}(q_i),\ i = 1, ..., m\}$.

Notice that so far the ordering of the propositions has not played an essential role in our definitions. We suggest that a property of memory is that the propositions (or basic propositional events) are arranged in some definite order in the mind, perhaps from most important to least important, or reverse chronologically from most recent to most distant, so that one remembers the outcomes in order. Sometimes one might remember more or less (perhaps depending on the complexity of the outcomes), but always in the same order. More precisely:

Definition: Let $\underline{P}$ be a collection of subsets of a propositional state space $\Omega = 2^n$. We say that $\underline{P}$ has the *memory property* if for any $R \in \underline{P}$, there is $\bar{\omega} \in \Omega$ and $0 \leq k \leq n$ such that $R = \{\omega \in \Omega | \omega(q) = \bar{\omega}(q),\ q = 1, ..., k\}$.

We now show that nested can always be interpreted in terms of memory.

Lemma 4 *Let (A, Ω, P, u, π) be given and let (Ω, P) be nondeluded and nested. Then there exists a physical renaming $(A, \Omega', P', u', \pi')$ which is propositional and has the memory property.*

Proof: The proof proceeds by induction on $\#\Omega$. For $\Omega = \{\omega\}$, let there be one proposition $n = 1$ and let $\Omega' = \{T, F\}$, let $P'(T) = P'(F) = \Omega'$, and let $u'(a, T) = u'(a, F) = u(a, \omega)$ for all $a \in A$. Finally, let $\pi'(T) = \pi'(F) = \frac{1}{2}\pi(\omega)$.

Suppose now that the lemma is true for $\#\Omega \leq k$, and let $\#\Omega = k + 1$. Let $T_0 = \{\omega \in \Omega | \omega \in R \in \underline{P}$ only if $R = \Omega\}$. If $T_0 = \Omega$, there is nothing to prove. So suppose $\#T_0 \leq k$. For each $\omega \in \Omega \backslash T_0$, let $T(\omega)$ be the set in $\underline{P}$ with largest cardinality $\leq k$ containing ω. Then $(T_0, \{T(\omega) | \omega \in \Omega \setminus T_0\})$ is a partition of Ω. By combining sets we may suppose that the partition T consists of two sets, each containing at most k states. For any T_i, and $\omega \in T_i$, either $P(\omega) = \Omega$ or $P(\omega) \subset T_i$. Let $P_i : T_i \to 2^{T_i} \setminus \phi$ be defined by $P_i(\omega) = P(\omega) \cap T_i$. Then each of the decision problems (A, T_i, P_i, u, π) satisfies the induction hypothesis. Hence there are renamings $(A, \Omega'_i, P'_i, u'_i, \pi'_i)$ where each $\Omega'_i = 2^{n_i}$. To every $\omega_i \in \Omega'_i$ there is an integer $0 \leq k_i(\omega_i) \leq n_i$ such that $P'_i(\omega_i) = \{\omega \in \Omega'_i | \omega(q) = \omega_i(q),\ q = 1, ..., k_i(\omega_i)\}$. By adding irrelevant propositions (whose truth is never distinguished by the P'_i) WLOG $n_2 = n_1$ and $\Omega'_2 \approx \Omega'_1 = 2^{1+n_1}$. Let $(A, \Omega', \tilde{P}', \pi', u')$ be defined by $\Omega' = \Omega'_1 \cup \Omega'_2 = 2 \times 2^{n_1}$. $\tilde{P}'$ is defined by $\tilde{P}'(\omega') = P_i(\omega')$ if $\omega' \in \Omega'_i$; π' is defined by $\pi'(\omega') = \pi_i(\omega')$ if $\omega' \in \Omega'_i$; u' is defined by $u'(\omega') = u_i(\omega')$ if $\omega' \in \Omega'_i$. It is clear that $(A, \Omega', \tilde{P}', \pi', u')$ is a physical renaming with the memory property for $(A, \Omega, \tilde{P}, \pi, u)$, where $\tilde{P}(\omega) = P_i(\omega)$ if $\omega \in T_i$. (Each $\tilde{P}'(\omega')$ is characterized by specifying a truth valuations to one proposition distinguishing whether $\omega' \in \Omega'_1$ or $\omega \in \Omega'_2$, to $k_i(\omega')$ of the next n_1 propositions if $\omega' \in \Omega'_i$). Finally, recall that $\tilde{P}(\omega) = P_i(\omega)$ differs from $P(\omega)$ only if $P(\omega) = \Omega$. Define P' on $\Omega' = \Omega'_1 \cup \Omega'_2$, as follows: $P'(\omega') = \begin{cases} \Omega' & \text{if } \omega' \in \Omega'_i \text{ and } P(\varphi_i(\omega')) = \Omega_i \\ \tilde{P}(\omega) & \text{otherwise} \end{cases}$. Clearly, if $P'(\omega') = \Omega'$, then this set can be characterized by specifying the truth valuation of *none* of the $n' + n_1$ propositions. Hence $(A, \Omega', P', u', \pi')$ is a physical renaming of (A, Ω, P, u, π) that has the memory property. $\square$

4 Game Theory with Generalized Partitions

We can extend game theory as well as decision theory to allow for information processing errors. The fundamental notion in game theory is the equilibrium concept provided by John Nash in 1951. Let $G = (I, A_i, \Omega, P_i, u_i, \pi_i)$, $i = 1, ..., I$, be a collection of decision makers. Let u_i have domain $\times_{i=1}^{I} A_i \times \Omega$. Then we say that the functions $(f_i,\ i = 1, ..., I)$ are a *Nash equilibrium for the Game G* iff

1. $[P_i(\omega) = P_i(\omega')] \Rightarrow [f_i(\omega) = f_i(\omega')]$, $\omega,\ \omega' \in \Omega$, $i \in I$.
2. For all $\omega \in \Omega$, $i \in I$, and $a_i \in A_i$,

$$\sum_{\omega' \in P_i(\omega)} u_i(a_i, f_{-i}(\omega'), \omega')\pi_i(\omega') \leq \sum_{\omega' \in P_i(\omega)} u_i(f_i(\omega), f_{-i}(\omega'), \omega')\pi_i(\omega').$$

We can interpret our definition of Nash equilibrium in much the same way as we did the single agent decision-maker. The players themselves do not completely understand the model, and so they are led to make information processing blunders concerning the signals they receive.

One consequence of this point of view is that one of the rationalizations for Nash equilibrium, that each player deduces what he should do as a matter of logical introspection, is no longer tenable. However, many of the other interpretations of equilibrium are still viable. For example, an equilibrium can still be characterized as an agreement from which no agent has an incentive to deviate.

Let us emphasize one limitation of the current model. Agents are permitted to make errors about the significance of their signals. But these errors do not depend on the moves of other agents, even though, for example, the news that some state has occurred might be much more unpleasant depending on what the other players plan to do in that state. If we extended our equilibrium notion to allow for correlated equilibria, as in [BDG (1988)], then it would be natural to allow the errors to depend on the moves of other agents. (E.g. there might be some things that you simply refuse to believe somebody else would ever do.)

We now give three examples illustrating the definition of Nash equilibrium with generalized partitions, which will also serve to introduce the idea of speculation.

Example 4.1: Let $I = \{1, 2\}$. Let $\Omega = \{a, b, c\}$, $P_1(a) = \{a, b\}$, $P_1(b) = \{b\}$, $P_1(c) = \{b, c\}$. Let $P_2(\omega) = \Omega$ for all $\omega \in \Omega$. Let $\pi_1 = \pi_2 = \pi$, with $\pi(a) = \pi(c) = 2/7$, $\pi(b) = 3/7$. Let the action spaces be $A_1 = A_2 = \{B, N\}$. Finally, let the payoffs in the three states be:

$$\begin{array}{c} \\ B \\ N \end{array} \underset{a}{\begin{array}{cc} B & N \\ \hline \end{array}\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\!\begin{array}{c} \quad \\ \left[\begin{array}{cc} -1,1 & 0,\varepsilon \\ \varepsilon,0 & \varepsilon,\varepsilon \end{array}\right]\end{array}} \quad \underset{b}{\begin{array}{c} \begin{array}{cc} B \quad & \quad N \end{array} \\ \left[\begin{array}{cc} 1,-1 & 0,\varepsilon \\ \varepsilon,0 & \varepsilon,\varepsilon \end{array}\right]\end{array}} \quad \underset{c}{\begin{array}{c} \begin{array}{cc} B \quad & \quad N \end{array} \\ \left[\begin{array}{cc} -1,1 & 0,\varepsilon \\ \varepsilon,0 & \varepsilon,\varepsilon \end{array}\right]\end{array}}$$

where $\varepsilon > 0$ is small. It is clear that there are two Nash equilibria. In the first, $f_i(\omega) = N$, $\forall \omega \in \Omega$, $i = 1, 2$. In the second, $f_i(\omega) = B$, $\forall \omega \in \Omega$, $i = 1, 2$. We are most interested in the possibility of the second equilibrium. Here the agents always bet against each other, simply on account of different information. Indeed, although agent 1 always knows strictly more than agent 2, on account of his "irrational" (generalized partition) information processsing, on average he is losing money.

Example 4.2: Let $I = \{1, 2\}$. Let $\Omega = \{a, b, c\}$, $P_1(a) = P_1(c) = \Omega$, $P_1(b) = \{b, c\}$. Let $P_2(\omega) = \Omega$ for all $\omega \in \Omega$. Let $\pi_1(\omega) = \pi_2(\omega) = \pi(\omega) = 1/3$ for all $\omega \in \Omega$. Let the action spaces be $A_1 = A_2 = \{B, N\}$. Finally, let the payoffs in the three states be:

$$\begin{array}{c} \\ B \\ N \\ \\ \end{array} \begin{array}{c} \begin{array}{cc} B & N \end{array} \\ \begin{bmatrix} -2,2 & 0,\varepsilon \\ \varepsilon,0 & \varepsilon,\varepsilon \end{bmatrix} \\ a \end{array} \begin{array}{c} \begin{array}{cc} B & N \end{array} \\ \begin{bmatrix} -2,2 & 0,\varepsilon \\ \varepsilon,0 & \varepsilon,\varepsilon \end{bmatrix} \\ b \end{array} \begin{array}{c} \begin{array}{cc} B & N \end{array} \\ \begin{bmatrix} 3,-3 & 0,\varepsilon \\ \varepsilon,0 & \varepsilon,\varepsilon \end{bmatrix} \\ c \end{array}$$

where $\varepsilon > 0$ is small. Again there are two Nash equilibria. In the trivial one, $f_i(\omega) = N$, $\forall \omega \in \Omega$, $i = 1, 2$. In the second, $f_i(a) = f_1(c) = N$, but $f_1(b) = B$, $f_2(\omega) = B$, $\forall \omega \in N$. Again, in the interesting equilibrium a bet does take place, though not always. Note also that agent 2 is willing to bet (always) because he knows that the only time his bet will be taken up is in state b, where he wins. Agent 1 knows that 2 is always willing to bet, but does not realize that he himself only bets when he is sure to lose. Once again agent 1 loses out to agent 2 despite his superior knowledge at each ω, because he is not perfectly rational.

Example 4.3: Let $I = \{1, 2\}$. Let $\Omega = \{a, b\}$, $P_1(a) = \{a\}$, $P_1(b) = \{a, b\}$, $P_2(a) = P_2(b) = \Omega$. Let $\pi(\omega) = \pi_i(\omega) = 1/2$, $i = 1, 2$, $\omega \in \Omega$. Let the action spaces be $A_1 = A_2 = \{B, N\}$, and the payoffs in the two states be:

$$\begin{array}{c} \\ B \\ N \\ \\ \end{array} \begin{array}{c} \begin{array}{cc} B & N \end{array} \\ \begin{bmatrix} 1,-1 & 0,\varepsilon \\ \varepsilon,0 & \varepsilon,\varepsilon \end{bmatrix} \\ a \end{array} \begin{array}{c} \begin{array}{cc} B & N \end{array} \\ \begin{bmatrix} -2,2 & 0,\varepsilon \\ \varepsilon,0 & \varepsilon,\varepsilon \end{bmatrix} \\ b \end{array}$$

$[1, -1]$where $\varepsilon > 0$ is small. Here there is a unique equilibrium, at which $f_i(\omega) = N$, $i = 1, 2$, $\omega \in \Omega$. If $f_2(\omega) = B$, $\forall \omega \in \Omega$, then agent 1 would choose $f_1(a) = B$, $f_1(b) = N$, in which case agent 2 would no longer be willing to bet.

If we changed the payoff at (B, B, b) from $(-2, 2)$ to $(u_1(B, B, b), u_2(B, B, b)) = (-1/2, 1/2)$, then again there would be a unique equilibrium with $f_i(\omega) = N$, $i = 1, 2$, $\forall \omega \in \Omega$. In this game the information processing error of agent 1 is more serious, since by ignoring the unpleasant information about state b, he is led to make a wrong decision and bet all the time if he thinks agent 2 is always betting. On the other hand, in this game agent 2 is not willing to bet if he thinks agent 1 is always betting.

We might consider a third variant of the game in which the payoffs are as in the second variant, but now $P_2(a) = \{a\}$, $P_2(b) = \{b\}$. Now if agent 1 is always betting, agent 2 can take advantage of the situation, choosing $f_2(a) = N$, $f_2(b) = B$. But in our definition of equilibrium, though agent 1 ignores the exogenous unpleasant news about the state of nature, he does not misunderstand the strategy agent 2 adopts. Hence 1 would choose not to bet, and once again we have a unique equilibrium at which $f_i(\omega) = N$, $i = 1, 2$, $\forall \omega \in \Omega$.

We now show that Nash equilibrium with generalized partitions can always be given a Bayesian interpretation.

Definition: We say that the game $G' = (I, A_i, \Omega', P_i', u_i', \pi_i')$ is a *decision-theoretic renaming of the game* $G = (I, A_i, \Omega, P_i, u_i, \pi_i)$ if (i) there are I 1–1 and onto maps $\delta_i : \underline{P}_i' \to \underline{P}_i$; (ii) moreover, for all $R_i' \in \underline{P}_i'$, and $f_i' : \Omega' \to A_i$ satisfying $f_i'(\tilde{\omega}') = f_i'(\omega')$ whenever $P_i'(\tilde{\omega}') = P_i'(\omega')$, let $R_i = \delta_i(R_i')$, and let $f_i : \Omega \to A_i$ be defined by $f_i(\omega) = f_i'(\omega')$ for ω' with $\delta_i(P_i'(\omega')) = P_i(\omega)$. Then we must have $\pi_i'(R_i') > 0$ if and only if $\pi_i(R_i) > 0$, and then for all $a_i \in A_i$, (iii) $\frac{1}{\pi_i(R_i)} \sum_{\omega \in R_i} u_i(a_i, f_{-i}(\omega), \omega)\pi_i(\omega) = \frac{1}{\pi_i'(R_i')} \sum_{\omega' \in R_i'} u_i'(a_i, f_{-i}'(\omega'), \omega')\pi_i'(\omega')$.

The following lemma appears in [BDG (1988)].

Lemma 5 *Any generalized game $G = (I, A_i, \Omega, P_i, u_i, \pi_i)$ has a decision-theoretic renaming $G' = (I, A_i, \Omega', P_i', u_i', \pi_i')$ in which P_i' is a partition of Ω', for $i = 1, ..., N$.*

Proof: Let $\Omega' = \underline{P}_1 \times \cdots \times \underline{P}_I \times \Omega$. Let $P_i'(R_1, ..., R_I, \omega) = \{R_i\} \times \underline{P}_{-i} \times \Omega$, and let $\pi_i'(R_1, ..., R_I, \omega)) = \begin{cases} \pi_i(\omega) & \text{if } \omega \in R_i \text{ and } R_{-i} = P_{-i}(\omega) \\ 0 & \text{otherwise} \end{cases}$ *for* $i = 1, ..., I$, $R_i \in \underline{P}_i$, $\omega \in \Omega$. Let $u_i'(a, (R_1, ..., R_n, \omega)) = u_i(a, \omega)$ for $i = 1, ..., I$, $a \in X_{i=1}^I A_i$, $R_i \in \underline{P}_i$, $\omega \in \Omega$. Finally, let $\delta(\{R_i\} \times \underline{P}_{-i} \times \Omega) = R_i$. $\square$

As an immediate corollary we have

Theorem 6 *If the action spaces A_i are convex and the u_i concave in A_i (or if the action spaces are discrete but randomization is permitted) then any generalized game has a Nash equilibrium.*

Proof: There is always a decision-theoretic renaming G' of G which is a standard Bayesian game; hence G' has a Nash equilibrium, which induces a Nash equilibrium on G. $\square$

5 Speculation in Equilibrium

Will rational, risk averse agents bet against each other? Can they agree to disagree about the probability of some event? What if they have access to different information? What if some of them make information processing errors?

Aumann (1976) showed that when agents have partition knowledge, it cannot be common knowledge that they disagree. Milgrom–Stokey (1983), and less generally Sebenius–Geanakoplos (1984), showed that when agents have partition knowledge, it cannot be common knowledge that they will speculate or bet against each other. A number of authors, in a long series of papers, have shown that there can be no speculation in a rational expectations equilibrium (e.g., Kreps (1977), Tirole (1982)). The nonspeculation idea could also have been reformulated in terms of a Nash equilibrium. With partition-information all of these theorems are pretty much the same: they can all be instantly derived from one theorem which we shall give below. When knowledge is described by generalized partitions, however, these theorems are distinct; their proofs are different and so are the hypotheses needed for each of them.

Examples 4.1 and 4.2 show that speculation and betting can occur in Nash equilibrium if the possibility correspondences P_i fail to be partitions. The extension of game theory to generalized partitions thus permits us to model a new phenomenon.

On the other hand, perhaps it is not surprising at all that agents can bet against each other even though they have common priors, when their rationality is bounded. After all, one way such generalized partitions can arise is if the agents are faulty in their processing of information for example they may ignore all unfavorable information. Another way to see the same thing, as the proof of Theorem 6 makes clear (see Brandenburger–Deckel–Geanakoplos (1988)) is that a generalized equilibrium is isomorphic to a Bayesian partition equilibrium in which the priors may be different. The agents may have started with common priors, but on account of their faulty information processing, they behave as if their priors were different. It is well-known that gambling can take place between agents with different priors.

In this light the surprise is that any weakening of partition information still retains enough structure to prevent speculation. Recall for instance that nondelusion, knowing that you know, and nested are together still consistent with throwing away unpleasant information at least once. Yet we have:

Theorem 7 *Let $G = (I, A_i, \Omega, P_i, u_i, \pi_i)$ be a generalized game. Suppose each player i has an action z_i such that for all $(f_1, ..., f_I)$, $\sum_{\omega\in\Omega} u_i(z_i, f_{-i}(\omega), \omega)\pi_i(\omega) = \bar{u}_i$. Furthermore, suppose that if for any $(f_1, ..., f_I)$, $\sum_{\omega\in\Omega} u_i(f(\omega), \omega)\pi_i(\omega) = \bar{u}_i$ for all i, then $f_j(\omega) = z_j$ for all $\omega \in \Omega$, $j = 1, ..., I$. Finally, let each P_i satisfy* KTYK, *nondeluded and nested. Then G has a unique equilibrium, in which $f_i(\omega) = z_i$ for all $i = 1, ..., I$, and all $\omega \in \Omega$. Conversely, if any P_i fails to satisfy any of* KTYK, *nondeluded, and nested, then there are A_i, u_i, π_i, $i \in I$, for which the theorem fails.*

Proof: Let $(f_1, ..., f_I)$ be an equilibrium. Fix f_j for all $j \neq i$, and look at the one-person decision problem this induces for player i. Clearly f_i must be an optimal plan for this decision problem. But if i had the trivial partition $Q_i(\omega) = \Omega$ for all $\omega \in \Omega$, he would be able to guarantee himself ex ante utility at least $\bar{u}_i$ by always playing $g_i(\omega) = z_i$. Hence by Theorem 1, $\sum_{\omega\in\Omega} u_i(f(\omega), \omega)\pi_i(\omega) \geq \bar{u}_i$. Since this is true for all i, by hypothesis $f_i(\omega) = z_i$ for each i and $\omega \in \Omega$. $\square$

In Example 4.1 speculation occurs because P_1 does not satisfy nested. In Example 4.2, P_1 does not satisfy KTYK. In Example 4.3, there could be no speculation no matter how we defined the payoffs.

The examples show that KTYK, nondeluded, and nested not only suffice to eliminate speculation in Nash equilibrium, but are necessary as well. If an agent's information P_i fails to satisfy all of the above, we can find other agents with partition information, and properly specified payoffs for all the agents at which there will be some speculation in equilibrium.

Like Theorem 1, Theorem 7 can be extended to allow for variable action spaces.

Corollary 8 *Consider the generalized game $\bar{G} = (I, \bar{A}_i, A_i, \Omega, P_i, u_i\pi_i)$ with variable action spaces. Suppose that $[P_i(\omega) = P_i(\omega')] \Rightarrow [\bar{A}_i(\omega) = \bar{A}_i(\omega')]$, $i = 1, ..., I$. Then if there are $z_i \in \bigcap_{\omega\in\Omega} \bar{A}_i(\omega)$, $i = 1, ..., I$, such that the conditions of Theorem 7 are* *satisfied, then $\bar{G}$ has a unique equilibrium $\bar{f}$ with $\bar{f}_i(\omega) = z_i$ for each $i = 1, ..., I$, $\omega \in \Omega$.*

Proof: Follow the logic of the proof of Theorem 7, and apply Corollary 2 where $\hat{A}_i = \{z_i\}$. $\square$

We can immediately apply the above theorem to show a generalization of the standard no speculation theorem in rational expectations equilibrium.

We define an economy $E = (I, \mathbb{R}^L_+, \Omega, P_i, u_i, \pi_i, e_i)$ by a set of agents I, a commodity space $\mathbb{R}^L_+$, a set Ω of states of nature, endowments $e_i \in \mathbb{R}^{L\Omega}_+$ and utilities $u_i : \mathbb{R}^L_+ \times \Omega \to \mathbb{R}$ for $i = 1, ..., I$, and generalized partitions P_i and measures π_i for each agent $i = 1, ..., I$. We suppose each u_i is strictly monotonic, and strictly concave, and that $[P_i(\omega) = P_i(\omega')] \Rightarrow [e_i(\omega) = e_i(\omega')]$ for all $i = 1, ..., I$.

Definition: A rational expectations equilibrium (REE) (p, I, x_i) for $E = (I, \mathbb{R}^L_+, \Omega, P_i, u_i, \pi_i, e_i)$ is a function p: $\Omega \to \mathbb{R}^L_{++}$ and for each $i \in I$, $x_i \in \mathbb{R}^{L\Omega}_+$ satisfying

(i) $\sum_{i=1}^{I} x_i = \sum e_i$.

(ii) $p(\omega)x_i(\omega) = p(\omega)e_i(\omega)$, for all $i = 1, ..., I$, and all $\omega \in \Omega$.

(iii) $[P_i(\omega) = P_i(\omega') and p(\omega) = p(\omega')] \Rightarrow [x_i(\omega) = x_i(\omega')]$ for $i = 1, ..., I$, and all ω, $\omega' \in \Omega$.

(iv) Let $I(p) = \{\omega : p(\omega) = p\}$. Then $\forall \omega \in \Omega$, and all i, if $y \in \mathbb{R}^L_+$ and $p(\omega)y = p(\omega)e_i(\omega)$, then

$$\sum_{\omega' \in P_i(\omega) \cap I(p(\omega))} u_i(x_i(\omega), \omega')\pi_i(\omega') \geq \sum_{\omega' \in P_i(\omega) \cap I(p(\omega))} u_i(y, \omega')\pi_i(\omega')$$

The reference to rational in REE comes from the fact that agents use the subtle information conveyed by prices in making their decisions. That is, they not only use the prices to calculate their budgets, they also use their knowledge of the function p to learn more about the state of nature. If we modified (iii) and (iv) above to

(iii′) $[P_i(\omega) = P_i(\omega')] \to [x_i(\omega) = x_i(\omega')]$ for $i = 1, ..., I$.

(iv′) $\sum_{\omega' \in P_i(\omega)} u_i(x_i(\omega), \omega')\pi_i(\omega') \geq \sum_{\omega' \in P_i(\omega)} u_i(y, \omega')\pi_i(\omega')$ for all $i = 1, ..., I$, for all $\omega \in \Omega$ and all $y \in \mathbb{R}^L_+$ with $p(\omega)y = p(\omega)e_i(\omega)$.

Then we would have the conventional definition of competitive equilibrium (CE). The following nonspeculation theorem holds for REE, but not for CE. (For proofs when agents have partition information and learn from prices, see Kre (1977), Tir (1981), DGS (1987). For an example with partition information in which agents do not learn from prices, and so speculate, see DGS (1987).) We say that there are only speculative reasons to trade in E if in the absence of asymmetric information there would be no perceived gains to trade. This occurs when the initial endowment allocation is *ex ante* Pareto optimal, that is if $\sum_{i=1}^{I} y_i(\omega) \leq \sum_{i=1}^{I} e_i(\omega)$ for all $\omega \in \Omega$, and if for each $i = 1, ..., I$, $\sum_{\omega\in\Omega} u_i(y_i(\omega), \omega)\pi_i(\omega) \geq \sum_{\omega\in\Omega} u_i(e_i(\omega), \omega)\pi_i(\omega)$, then $y_i = e_i$ for all $i = 1, ..., I$.

Corollary 9 *Let $E = (I, \mathbb{R}_+^L, \Omega, (P_i, u_i, \pi_i, e_i))$ be an economy, and suppose the initial endowment allocation is ex ante Pareto optimal. Let (p, I, x) be a rational expectations equilibrium. Suppose that each P_i is nondeluded, nested, and satisfies knowing that you know. Then $x_i = e_i$ for all $i = 1, ..., I$.*

Proof: The proof follows immediately from Corollary 9. Let $P_i'(\omega) = P_i(\omega) \cap I(p(\omega))$ be the generalized partition for each agent, and let $\bar{A}_i(\omega) = \{y | p(\omega) \cdot y = 0$ and $e_i(\omega) + y \geq 0\}$. Let $z_i = 0 \in \bigcap_{\omega\in\Omega} \bar{A}_i(\omega)$. $\square$

6 Knowing Your Own Action

Consider again the single person decision problem (A, Ω, P, u, π), and suppose that $f : \Omega \to A$ is an optimal plan. What do we mean when we say that the agent "knows what he is doing" at some $\omega \in \Omega$? Simply put, we mean that if the agent regards ω' as possible at ω, then he should take the same action at ω' as at ω: if $\omega' \in P(\omega)$, then $f(\omega') = f(\omega)$.

Recall that in decision theory, the agent begins with a prior π on Ω, then refines his information to $P(\omega)$. If P describes a partition, then the agent behaves as if he has sifted through every possible source of information. On the other hand, if P is a generalized partition, then the agent might forget or ignore information at ω which should have caused him to exclude the possibility of ω'. This is consistent with knowing what he is doing provided that at ω' he would take the same action. Suppose that Watson does not notice whether the dogs did anything in the night time, and chooses some action. If the dogs had barked, Watson would have noticed; he has erred by not deducing from the absence of sound that in fact the dogs did not bark in the night time. Nevertheless we could say that Watson knew what he was doing if the information he would receive from hearing the dogs bark would not change his mind about his best decision.

The agent always knows what he is doing under the action plan f if the action he is supposed to take is always self-evident to him. We shall describe the circumstances in which an agent who "always knows what he is doing" may appear to be perfectly rational even when he is making information processing errors.

Definition: An agent who processes information according to (Ω, P) knows what he is doing at some $\omega \in \Omega$ using the action plan $f : \Omega \to A$ if $P(\omega) \subset Q_f(\omega) = \{\omega' \in \Omega | f(\omega) = f(\omega')\}$.

We shall now show that if an agent always knows what he is doing (i.e., knows for all $\omega \in \Omega$), then better information will make him better off under quite general circumstances. To this end we introduce our fourth and fifth properties of information processing:

Definition: The information processor (Ω, P) is *positively balanced* with respect to some set $E \subset \Omega$ iff there exists a function $\lambda : \underline{P} \to \mathbb{R}_+$, such that (letting χ_A be the characteristic function of any set $A \subset \Omega$)

$$\sum_{C \in \underline{P}} \lambda(C)\chi_C(\omega) = \chi_E(\omega) \text{ for all } \omega \in \Omega.$$

If the same holds true for some λ unrestricted in sign, $\lambda : \underline{P} \to \mathbb{R}$, then we say that (Ω, P) is *balanced* with respect to E. (More generally, for any collection of events, $X \subset 2^\Omega$, and $E \in 2^\Omega$, we say that X is (positively) balanced with respect to E if there is $\lambda : X \to (\mathbb{R}_+)\mathbb{R}$ such that $\sum_{C \in X} \lambda(C)\chi_C = \chi_E$. $\square$

Balancedness gives a condition under which one can say that every element $\omega \in E$ is equally scrutinized by the information correspondence P. Every element $C \in \underline{P}$ has an intensity $\lambda(C)$, and the sum of the intensities with which each $\omega \in E$ is considered possible by P is the same, namely 1. Balancedness is a generalization of partition. If E can be written as a disjoint union of elements of $\underline{P}$, then (Ω, P) is trivially balanced with respect to E.[3]

There is a special class of events for which being balanced is especially important. Recall:

Definition: An event $E \in \Omega$ is *self-evident* to the processor (Ω, P) if $P(\omega) \subset E$ for all $\omega \in E$. The notion of self-evident has been used in Shin [1989], Sam [1987], BG [1988], Gea [1988], and MS [1988]. An event is self-evident if it can never occur without the agent knowing that it has occurred.

Definition: (Ω, P) is balanced (positively balanced) if it is balanced (positively balanced) with respect to every self-evident set.

If (Ω, P) is a partition, then the self-evident sets are (disjoint) unions of elements of $\underline{P}$. Hence (Ω, P) is trivially positively balanced. Note that in Example 2.1 (Ω, P) is balanced with $\lambda(P(a)) = 1 = \lambda(P(c))$, and $\lambda(P(b)) = -1$ but not positively balanced. In Examples 2.2 and 2.3 (Ω, P) is positively balanced since for both of them $\Omega \in P$.

We now use the notion of self-evident events to characterize the relationship between positively balanced and balanced and nondeluded, KTYK, and nested. Positively balanced is a weakening of nested, and balanced is a further weakening that is also a weakening of KTYK.

Lemma 10 *If (Ω, P) is nondeluded and nested, then (Ω, P) is positively balanced.*

Proof: Let E *be self-evident to* (Ω, P). For each $\omega \in E$, let $E(\omega) = \bigcup_{\omega' \in E \quad \omega \in P(\omega')} P(\omega')$. By nondeluded $\omega \in E(\omega)$, and by nested $E(\omega)$ is a partition of E, and each $E(\omega) \in \underline{P}_i$. $\square$

[3]Balancedness is similar to a concept (with the same name) that played an important role in the development of the theory of the core in cooperative game theory.

Lemma 11 *Let (Ω, P) satisfy nondelusion and knowing that you know. Then (Ω, P) is balanced.*

Proof: The proof proceeds by induction. If $\#\Omega = 1$, there is nothing to show. Suppose the truth of the Lemma for $\#\Omega \leq k$. Now let $\#\Omega = k + 1$. Find an element $R \in \underline{P}$ which minimizes $\#R$. From nondeluded, knowing that you know, and the minimality of $\#R$, we deduce that $P(\omega) = R$ for all $\omega \in R$. Moreover, if $P(\omega') \cap R \neq \phi$, then $P(\omega') \supset R$, for otherwise if $\omega'' \in P(\omega') \cap R$, then $\#P(\omega'') < \#R$. Let $\Omega' = \Omega \setminus R$, and let $P' : \Omega' \to 2^{\Omega'}$ be defined by $P'(\omega) = P(\omega) \setminus R$. Then (Ω', P') satisfies KTYK and nondelusion. By the induction hypothesis (Ω', P_i') is balanced. Now, let E be self-evident to (Ω, P). If $E \cap R = \phi$, then E is self-evident to (Ω', P'), and the balancing weights are strictly positive only on $C \in \underline{P} \cap \underline{P}'$. If $E = R$, then the result is obvious. If $R \not\subseteq E$, then $E' = E \setminus R$ is self-evident to (Ω', P'). Use the balancing weights for $\underline{P}'$ and E', and choose the correct weight for R. $\square$

Figure 2

The logical connections between balanced, positively balanced, nested, KTYK, and partition information, assuming nondeluded. In later sections we show that if each agent is nondeluded, then

$$\begin{aligned} \text{all agents balanced} &\Leftrightarrow \text{they cannot agree to disagree} \\ \text{all agents positively balanced} &\Leftrightarrow \text{no common knowledge speculation} \\ \text{all agents nested and KTYK} &\Leftrightarrow \text{no equilibrium speculation} \end{aligned}$$

In our next theorem we show that Theorem 1 can be proved under much weaker hypotheses if we suppose that the agent knows what he is doing. (Needless to say, if an agent is not unboundedly rational, we might not expect him always to know what he is doing.)

Theorem 12 *Let (Ω, P) be nondeluded and positively balanced. Let $f : \Omega \to A$ be optimal for the decision problem (A, Ω, P, u, π). Suppose that the information processor (Ω, P) always knows what he is doing under the action plan f. If Q is a partition of Ω that is a coarsening of P, and if g is optimal for (A, Ω, Q, u, π), then*

$$\sum_{\omega \in \Omega} u(g(\omega), \omega)\pi(\omega) \leq \sum_{\omega \in \Omega} u(f(\omega), \omega)\pi(\omega).$$

If $Q(\omega) \subset Q_f(\omega)$ for all $\omega \in \Omega$, then the above inequality is actually an equality. Conversely, if (Ω, P) fails to be either nondeluded or positively balanced, then there exist decision problems and partitions Q for which the above inequality is strictly reversed.

Proof: Since $Q \vee Q_f$ defined by $(Q \vee Q_f)(\omega) = Q(\omega) \cap Q_f(\omega) \subset Q_f(\omega)$ is a partition which refines Q, and since more information is always better for partitions, it suffices to show that equality holds above whenever $Q(\omega) \subset Q_f(\omega)$ for all $\omega \in \Omega$.

Let $E \in \underline{Q}$. Since Q is a partition of Ω, $Q(\omega) = E$ for all $\omega \in E$, and since Q is a coarsening of P, it follows that E is self-evident to (Ω, P). Write

$$\sum_{C \in \underline{P}} \lambda(C)\chi_C = \chi_E, \text{ where } \lambda \geq 0.$$

By nondelusion, if $\omega \not\in E$, then $\omega \in P(\omega)$ and so $\lambda(P(\omega)) = 0$. Since $E \subset Q_f(\omega)$ for any $\omega \in E$, it follows that $f(\omega) = a$ for all $\omega \in E$, for some $a \in A$. To conclude the proof of the first half of the theorem, it suffices to show that for all $b \in A$,

$$\sum_{\omega \in E} u(f(\omega), \omega)\pi(\omega) = \sum_{\omega \in E} u(a, \omega)\pi(\omega) \geq \sum_{\omega \in E} u(b, \omega)\pi(\omega).$$

But for any $C \in \underline{P}$ with $\lambda(C) > 0$, $\sum_{\omega \in E} u(a, \omega)\pi(\omega) \geq \sum_{\omega \in E} u(b, \omega)\pi(\omega)$, hence $\sum_{C \in \underline{P}} \sum_{\omega \in \Omega} \lambda(C)\chi_C(\omega)u(a, \omega)\pi(\omega) \geq \sum_{C \in \underline{P}} \sum_{\omega \in \Omega} \lambda(C)\chi_C(\omega)u(b, \omega)\pi(\omega)$ hence

$$\sum_{\omega \in \Omega} \chi_E(\omega)u(a, \omega)\pi(\omega) \geq \sum_{\omega \in \Omega} \chi_E(\omega)u(b, \omega)\pi(\omega).$$

To argue in the other direction, note that there are trivial counterexamples if (Ω, P) does not satisfy nondeluded. If it does, then suppose that for all decision problems $(A, \Omega, P, \hat{u}, \hat{\pi})$ and for all $E \subset \Omega$ where E is self-evident to (Ω, P), $[\sum_{\omega \in C}(a, \omega)\hat{\pi}(\omega) = \sum_{\omega \in \Omega} \chi_C(\omega)\hat{u}(a, \omega)\hat{\pi}(\omega) \geq \sum_{\omega \in \Omega} \chi_C(\omega)\hat{u}(b, \omega)\hat{\pi}(\omega)$ for all $C \in \underline{P}$, $C \subset E]$ implies $[\sum_{\omega \in \Omega} \chi_E(\omega)\hat{u}(a, \omega)\hat{\pi}(\omega) \geq \sum_{\omega \in \Omega} \chi_E(\omega)\hat{u}(b, \omega)\hat{\pi}(\omega)]$. Since the $\hat{u}(a, \omega)$, $\hat{u}(b, \omega)$, $\hat{\pi}(\omega)$ are arbitrary, we can apply Farkas' Lemma which assets that there are $\lambda(C) \geq 0$ for $C \subset E$ such that $\sum_{C \in P, C \subset \bar{E}} \lambda(C)\chi_C = \chi_E$. $\square$

The idea behind the proof of Theorem 12 is quite different from that used in the proof of Theorem 1. The following definitions are clarifying.

Definition: A function $\delta : 2^\Omega \rightarrow A$ is said to satisfy the *sure-thing principle* if whenever E, F are disjoint and $\delta(E) = \delta(F)$, then letting $B = E \cup F$, $\delta(B) = \delta(E)$. We say that δ satisfies the *generalized sure-thing principle* if whenever $\delta(E) = d$ for all $E \in X \subset 2^\Omega$, and X is positively balanced with respect to some $B \in 2^\Omega/\phi$, then $\delta(B) = d$.

Behavior which is optimal in the sense we have described, where information processing errors are represented by possibility correspondences or generalized partitions, always satisfies the generalized sure-thing principle.

Theorem 12 has one particularly important consequence. Suppose an agent who is positively balanced optimally chooses some action plan f under which he knows what he is doing. If the agent were to forget everything he knew except what was necessary to implement the decision plan, then he would still choose the same plan if he was nondeluded and positively balanced. To put the matter still differently, the optimal behavior of an agent who is nondeluded and balanced and always knows what he is doing cannot be distinguished from the behavior of an unboundedly rational (partition information processing agent).

7 Common Knowledge of Events vs. Common Knowledge of Actions

Aumann (1976) introduced the idea of common knowledge of events and actions. We investigate what conclusions we can draw (about speculation and consensus) when we add the additional hypothesis that actions are common knowledge. Again we find a nonspeculation theorem, but under weaker conditions than Theorem 12, and with a different proof. We also derive a consensus theorem under still weaker hypotheses, with yet another kind of proof.

We can already get some idea of the importance of the hypothesis that actions are common knowledge from Example 4.2. The agents do bet in equilibrium, but at the moment each commits himself to the bet he does not know whether it will be accepted or not. The bet is not common knowledge. We now formalize this idea.

The possibility correspondence P_i gives rise to a knowledge operator K_i by $K_i(A) = \{\omega : P_i(\omega) \subset A\}$. K_i satisfies the following three properties:

(1) $K_i\Omega = \Omega$;

(2) $A \subset B \Rightarrow K_iA \subset K_iB$;

(3) $K_iA \cap K_iB = K_i(A \cap B)$.

If P_i is nondeluded, then

(4) $K_iA \subset A$ for all A.

Definition (Lewis (1969), Aumann (1976)): We say that an event A is "common knowledge at ω'' if for any sequence of players $i_1, ..., i_n$, $\omega \in K_{i_1}, ..., K_{i_n}A$.

We can give an equivalent definition of common knowledge based on our familiar notion of self-evident event.

Definition: Let agents' knowledge be represented by (Ω, P_i), $i = 1, ..., I$. An event $E \subset \Omega$ is *self-evident to* i if $[\omega \in E] \Rightarrow [P_i(\omega) \subset E]$. Let $\mathcal{E}_i$ be the collection of all self-evident events to i. We call $\mathcal{E} = \bigcap_{i=1}^{I} \mathcal{E}_i$ the collection of *public events.*

The following proposition is taken from Brown–Geanakoplos (1988), or Monderer–Samet (1988), extending Shin (1987), who extended Aumann (1976). The reader interested in a further discussion of common knowledge is referred to any of these works.

Proposition: Let the knowledge of agents $i = 1, ..., I$ be represented by (Ω, P_i), where each P_i is nondeluded. Then an event A is common knowledge at ω if and only if there is a public event $E \in \mathcal{E}$ with $\omega \in E \subset A$.

Not only events, but also actions can be common knowledge.

Definition: Let $f : \Omega \to A$ be a function from Ω to some set A. Given information processors (Ω, P_i), $i = 1, ..., I$, we say that f is common knowledge at some $\omega \in \Omega$ iff the event $Q_f(\omega) = \{\omega' | f(\omega') = f(\omega)\}$ is common knowledge at ω.

Observe that if f is the action plan of some player i, then one consequence of f being common knowledge at ω is that i himself knows what he is doing at ω. Replacing Ω by the smallest public event Ω' containing ω, we see that in fact we can more strongly assert that if f is common knowledge at ω, then i always (in Ω') knows what he is doing.

There would seem to be a wide gulf between the hypothesis that agents know the same things (about *all* events) and the hypothesis that agents know the same things about what they are each planning to do. Indeed it is commonly held that we observe agents interacting (i.e., taking actions) in various ways on account of the fact that they have asymmetric information. The following theorem, however, describes the power of assuming actions are common knowledge. It generalizes a theorem (Geanakoplos (1987)), proved for Nash equilibria, following an idea in Cave (1983). It shows that if in Nash equilibrium the actions are common knowledge at ω, then the information might as well be the same as well. Hence once actions are presumed to be common knowledge, asymmetric information provides no explanation whatsoever of behavior.

Theorem 13 *Let $G = (I, A_i, \Omega, P_i, u_i, \pi_i)$ have an equilibrium $(f_1, ..., f_I)$. Let P_i be nondeluded and positively balanced, for $i = 1, ..., I$. Suppose that it is common knowledge at some ω what moves the players are making. Then we can replace each P_i with $\tilde{P}_i$, creating a new generalized game G, having the same equilibrium $(f_1, ..., f_I)$, and moreover we can choose $\tilde{P}_i(\omega)$ to be independent of i.*

Proof: From the proposition we know that there is a public event E such that $f_i(\omega') = f_i(\omega)$ for all $\omega' \in E$, and all $i = 1, ..., I$. Let $\tilde{P}_i(\omega') = E$ for all $\omega' \in E$, but otherwise leave P_i unchanged. Let $Z_i = f_i(\omega)$. Since Z_i is optimal given the information in any $P_i(\omega')$, for any $\omega' \in E$, by Theorem 12 agent i would choose Z_i at ω if he were only informed of $\tilde{P}_i(\omega)$. Since for $\omega' \notin E$, $\tilde{P}_i(\omega') = P_i(\omega)$, we see that changing the information structure of the game does not affect the equilibrium moves. $\square$

The hypotheses in Theorem 13 are not only sufficient for Theorem 13, but *necessary* as well. If some agent's information processing (Ω, P_i) was not nondeluded or positively balanced, then we could find payoffs, and another player with partition information, such that in the resulting equilibrium actions were common knowledge but the asymmetry of information could not be dispensed with. The same remark applies to Corollaries 14 and 15. They give a necessary and sufficient condition for nonspeculation in equilibrium under common knowledge of actions, which is substantially weaker than the necessary and sufficient condition for nonspeculation in equilibrium.

8 Common Knowledge and Speculation

Corollary 14 *Consider the speculative situation of Theorem 7. Drop the assumption that agents know what they know and are nested. Suppose only that P_i is nondeluded and positively balanced, for each $i = 1, ..., I$. But suppose also that at some $\omega \in \Omega$, it is common knowledge at ω what moves all the players are making. Then at an equilibrium $(f_1, ..., f_I)$, $f_i(\omega) = z_i$ for each $i = 1, ..., I$.*

Proof: From the common knowledge hypothesis, for each i, there is a $B_i \in A_i$ such that each player is choosing $f_i(\omega') = B_i$ for all $\omega' \in E$, where E is a public event. It follows from Theorem 12 that

$$\sum_{\omega' \in E} u_i(B_i, f_{-i}(\omega'), \omega')\pi_i(\omega') \geq \sum_{\omega' \in E} u_i(z_i, f_{-i}(\omega'), \omega')\pi_i(\omega').$$

From this it follows that the ex ante payoffs to each player from the moves

$$g_i(\omega') = \begin{cases} z_i & \text{if} \quad \omega' \notin E \\ B_i & \text{if} \quad \omega' \in E \end{cases}$$

would be at least $\bar{u}_i$. Since by hypothesis that can be achieved only if each player chooses z_i regardless of ω' we have $B_i = z_i$ for each $i = 1, ..., I$. $\square$

Corollary 15 *Let $E = (N, \mathbb{R}^L_+, \Omega, (P_i, u_i, \pi_i, e_i))$ be an economy (not necessarily with strictly convex preferences), and let (p, I, x) be a rational expectations equilibrium. Suppose that each P_i is nondeluded, and positively balanced. If at some ω, for all i, $x_i - e_i$ is common knowledge, then there is a public event F such that the equilibrium would remain the same if we set $P_i(\omega') = F$ for all $i = 1, ..., I$, and all $\omega' \in F$. Furthermore, if in addition the endowments were a Pareto optimal allocation, and no other allocation gave all agents precisely the same utility (as in Corollary 2), then we could conclude that $x_i(\omega) = e_i(\omega)$, for all $i = 1, ..., I$.*

9 Common Knowledge and Consensus

A game of particular interest is the opinion game G^*. Let Ω be a finite set of real numbers. Let A_i be the set of reals for each $i = 1, ..., I$. Let $\pi_i(\omega) = \pi(\omega) > 0$ for all i and ω. Finally, let $u_i((q_1, ..., q_N), \omega) = -(q_i - \omega)^2$. Let P_i be the possibility correspondence. In this game G^*, each player optimizes by giving his conditional expectation of the random variable ω, based on his information $P_i(\omega)$.

Theorem 16 *Suppose that in the generalized opinion game G^*, $(f_1, ..., f_I)$ is an equilibrium. Suppose that it is common knowledge what the moves of the players are at ω. Finally, suppose that each P_i is nondeluded and balanced. Then all the players are taking the same action at ω. Conversely, if P_i is not balanced for some i, then there exists a probability π and renumbering of $\omega \in \Omega$, and partitions for $j \neq i$, such that the players can "agree to disagree" in the equilibrium of G^*.*

Proof: From nondelusion and balancedness, for any public event E we can write $\chi_E = \sum_{C\in\underline{P}_i, C\subset E} \lambda_i(C)\chi_C$. From the common knowledge hypothesis, we may choose E so that for all $C \in \underline{P}_i$, $C \subset E$, $\frac{1}{\pi(C)}\sum_{\omega\in C}\omega\pi(\omega) = k_i$. Hence for all such C, $\sum_{\omega\in C}\omega\pi(\omega) = \sum_{\omega\in C} k_i\pi(\omega)$. Hence

$$\sum_{\substack{C\in\underline{P}_i\\ C\subset E}} \lambda_i(C) \sum_{\omega\in\Omega} \chi_C(\omega)\omega\pi(\omega) = \sum_{\substack{C\in\underline{P}_i\\ C\subset E}} \lambda_i(C) \sum_{\omega\in\Omega} \chi_C(\omega)k_i\pi(\omega).$$

Using nondelusion and balancedness,

$$\sum_{\omega\in\Omega} \chi_E(\omega)\omega\pi(\omega) = \sum_{\omega\in\Omega} \chi_E(\omega)k_i\pi(\omega)$$

or

$$\sum_{\omega\in E} \omega\pi(\omega) = k_i \sum_{\omega\in E} \pi(\omega), \text{ \textit{so} } k_i \text{ is the same for all } i.$$

The converse follows from Farkas' Lemma as in the proof of Theorem 12, except since all the inequalities are equalities here, the $\lambda_i(C)$ can have either sign. $\square$

Aumann (1976) gave the first famous version of this theorem, for partitions. Samet (1987) showed that as long as each P_i satisfies nondeluded and knowing what you know, then if the agents' opinions are common knowledge, they must be the same. By following the logic of Theorem 13, we could have shown that agreement must hold if the P_i are nondeluded and positively balanced, a condition independent of KTYK. For if the opinions are common knowledge, then by Theorem 13 they could have been given with identical information across agents. But in that case they are surely the same. Theorem 16 employs a hypothesis that is weaker than nested and weaker than KTYK, and yields a necessary and sufficient condition for never agreeing to disagree. Note that the proof of Theorem 16 differs from that of Theorem 12 and Theorem 13.

McKelvey–Page (1986) proved the remarkable theorem that if the average (or sum) of different agents' opinions is common knowledge and if each agent began with the same prior, all the opinions must be the same. Here we extend this result to possibility correspondences satisfying only nondeluded, nested, and knowing that you know. The proof can also be regarded as an alternative derivation of McKelvey-Page's average opinion theorem.[4]

Theorem 17 *Suppose that in the generalized opinion game G^*, $(f_1, ..., f_I)$ is an equilibrium. Suppose that the sum $\sum_{i=1}^{I} f_i$ of the opinions is common knowledge at $\bar{\omega}$. Then $f_1(\bar{\omega}) = \cdots = f_I(\bar{\omega})$.*

[4] McKelvey–Page (1986) extends for the case of partitions to situations in which only the average of monotonic transformations is common knowledge. This will hold under the hypothesis that each (Ω, P_i) is nondeluded and balanced, provided that agents also always know their *own* opinions.

Proof: Since $\sum_{i=1}^{I} f_i(\omega)$ is common knowledge at $\bar{\omega}$, there exists $E \in \mathcal{E}$ satisfying $\bar{\omega} \in E$, and $\sum_{i=1}^{I} f_i(\omega) = k$ for all $\omega \in E$. Let $\bar{x} = \frac{1}{\pi(E)}\omega\pi(\omega)$. Then we must have that $\sum_{\omega\in\Omega}(\omega - \bar{x})\sum_{i=1}^{I} f_i(\omega)\pi(\omega) = 0$. To prove the theorem it suffices to show that $\sum_{\omega\in\Omega}(\omega - \bar{x})f_i(\omega)\pi(\omega) > 0$ unless f_i is a constant on all of E, for each $i = 1, ..., I$.

Observe that from Theorem 1 and the fact that $E \in \mathcal{E}_i$,

$$-\sum_{\omega\in E}(\omega - f_i(\omega))^2\pi(\omega) \geq -\sum_{\omega\in E}(\omega - \bar{x})^2\pi(\omega).$$

Multiplying out terms and rearranging yields

$$2\sum_{\omega\in E}\omega f_i(\omega)\pi(\omega) \geq \sum_{\omega\in E}(\bar{x}^2 + f_i^2(\omega))\pi(\omega).$$

Subtracting $2\sum_{\omega\in E} f_i(\omega)\pi(\omega)$ from both sides yields

$$\begin{aligned} 2\sum_{\omega\in E}(\omega - \bar{x})f_i(\omega)\pi(\omega) \;&\geq\; \sum_{\omega\in E}(\bar{x}^2 + f_i^2(\omega))\pi(\omega) - 2\sum_{\omega\in E} f_i(\omega)\pi(\omega) \\ &= \sum_{\omega\in E}(\bar{x} - f_i(\omega))^2\pi(\omega) > 0 \end{aligned}$$

unless $f_i(\omega) = \bar{x}$ for all $\omega \in E$. $\square$

10 Generalized Games in Extensive Form

Let us now apply our theory to games in extensive form. In this theory players act over time, and in particular the same player may move many times. The inconsistency that occurs when information is not given by partitions and total recall can now reveal itself in time inconsistency of behavior.

Consider a *tree*, consisting of a finite set of nodes $\omega \in \Omega$ with a partial order. A tree has a first element, called the *root*; maximal or *terminal* nodes are associated with payoffs for each of the I players. To every node we associate either "nature" or one player "who has the move." Let $\#\omega$ be the number of immediate successor nodes of ω, and let player i be on the move at ω. Then the convex feasible move set for i at ω is $\bar{A}_i(\omega) \subset S^{\#\omega-1}$, where $S^{\#\omega-1}$ is the simplex of dimension $\#\omega - 1$. (Note: implicitly we have numbered the immediate successor nodes $1, ..., \#\omega$.) For convenience we set $A_j(\omega)$ equal to a single point if $j \neq i$ is not on the move at ω. If nature is on the move at ω, then there is also given an element in $S^{\#\omega-1}$, which for simplicity we take as objectively given.

To every node ω we associate a possibility set $P_i(\omega)$ for each player i, in such a way that if $\omega' \in P_i(\omega)$, then $\bar{A}_i(\omega) \subset S^{\#\omega'-1}$. Furthermore, if $\omega' \in P_i(\omega)$, then ω' is not comparable with ω in the tree ordering (that is, it is on a different branch of the tree). We shall maintain $\omega \in P_i(\omega)$, although this is not actually necessary. Note that these assumptions do not imply perfect recall. This completes our description of a generalized game G in extensive form.

A strategy for a player i is an association with each $\omega \in \Omega$ of an element $f_i(\omega) \in \bar{A}_i(\omega)$ in such a way that if $(P_i(\omega), \bar{A}_i(\omega)) = (P_i(\omega'), \bar{A}_i(\omega'))$, then $f_i(\omega) = f_i(\omega')$. Given strategies for each of the players, one can calculate "correctly" the expected payoffs to each player, and then define Nash equilibrium, or perfect Nash equilibrium, as in conventional game theory. We shall describe a different, generalized notion of equilibrium.

Given the strategies $(f_1, ..., f_N) = f$ of the N players, we can always calculate the probability $\pi_i(\omega, f) = \pi(\omega, f)$ that the node ω will be reached. Note that it is possible that $\pi(\omega, f) = 0$, and also in general $\sum_{\omega \in \Omega} \pi(\omega, f) > 1$. Note however that $\pi(\omega, f)$ depends only on $f_i(\omega')$ for ω' that precede ω. In particular, if $\omega' \in P_i(\omega)$, then $\pi(\omega, f)$ does not depend on $f_i(\omega')$.

To each $\omega \in \Omega$ and $i = 1, ..., I$ let us also associate the payoff $u_i(a_i, a_{-i}, \omega, f)$, which is the conditional payoff to player i given that ω has been reached, calculated using the strategies in f, except that $f_j(\omega)$ is replaced by a_j, for $j = 1, ..., N$. This is calculated in the conventional manner, and is obviously well-defined. Note that if i is on the move at ω and $j \neq i$ then a_j does not affect $u_k(a, \omega, f)$ for any $k = 1, ..., N$. Note furthermore that $u_i(a_i, a_{-i}, \omega, f)$, is continuous in (a_i, a_{-i}, f), and concave, in fact linear, in a_i.

We define an equilibrium, for the generalized game G as a tuple of strategies $(f_1, ..., f_N)$ satisfying:

$$f_i(\omega) \in \operatorname*{ArgMax}_{a_i \in A_i(\omega)} \sum_{\omega' \in P_i(\omega)} u_i(a_i, f_{-i}(\omega'), \omega', f)\pi_i(\omega', f).$$

Theorem 18 *Every generalized game in extensive form has an equilibrium.*

Here is not the place to go into details, but one can also define various refinements of equilibrium for generalized games. A perfect equilibrium for a generalized game in extensive form can be defined analogously to a perfect equilibrium for a conventional game in extensive form, since in the latter case one must appear also to the agent normal form.

One could give many examples of generalized games in extensive form, but let us concentrate on the story of Odysseus and the Sirens, and the problem of time consistency.

Recall that one day Odysseus was told by Circe that his boat was sailing near the island home of the famous Sirens, whose beautiful singing lured many sailors to crash against the terrible rocks that studded the shore. Anticipating that once they heard the music, he and his men would not be able to resist the seductive temptation to sail nearer to the shore to better hear the songs, he ordered his men to sail clear of the shore, and he put wax in their ears to make sure that they could hear neither the Sirens nor himself. He also had himself tied to the mast, so that he could hear the Sirens but could do nothing to change the course of the boat. When the boat finally came within earshot of the Sirens, Odysseus struggled violently to free himself from his bonds and to exhort his men toward shore. Fortunately for him, both the wax and the bonds held firm, and his boat sailed safely past.

Odysseus' decision problem is one of the best known in history, precisely because Odysseus is so famous for his cunning, which indeed this story seems to confirm, and yet his behavior before and after hearing the Sirens is apparently inconsistent. A celebrated explanation, given by Strotz, suggests that Odysseus was a man who greatly discounted the future, but did not discount the distant future much more than the near future. According to this theory, when Odysseus first realized where he was, he weighed the near future of sailing near shore to better hear the sirens against the distant future of crashing on the rocks, and decided to avoid the bargain. But when the near future became the present, so that the trade-off was between hearing better now and crashing later, he wanted the bargain.

Although this impatience explanation is quite striking, it is not clear that it is faithful to the story, nor that it is the most interesting explanation. In the first place, one of the salient characteristics of Odysseus' personality is that, unlike most of the other Greeks, he was always planning for the future. These plans and preparations often involved initial sacrifices for future rewards. (Meticulously arranging the wax for his men is a perfect example.) It does not seem credible that wily Odysseus would trade his life for a song just because the song came first. (Indeed if anything he was willing to give up his life to the Trojan expedition in order to be part of the immortal song of the poet, because that song is last, and everlasting.) There is a further technical problem with Strotz's explanation, namely that it suggests discontinuous preferences. Until the moment Odysseus hears the song, he keeps his wits and tries to avoid the rocks. It is only at the instant that he hears the music that he forgets himself. This behavioral change cannot be accounted for by continuous time preferences, that ignore the information content of the song.

The theory of extensive form games proposed here is designed to model behavior that is purposeful and cunning, but based on information processing that is not perfect. My interpretation is that when Odysseus hears the Sirens and "forgets himself," he literally forgets what he knew before, namely that the Sirens are dangerous. His behavior after he hears the Sirens is not less purposeful or less skillful than before. The difference is that it is constrained and it is based on different information. Typically such a beautiful song deserves a better hearing, and having forgotten the warning of Circe, Odysseus struggles to land his boat closer to shore. The subtlest part of this information explanation of Odysseus, and one that requires all of the apparatus of the model, is that Odysseus recognizes full well that he is ensnared in the ropes and cannot get the attention of his men. But he never asks himself how he got in that situation. If he did he might have inferred his predicament. Thus the information explanation, which I believe expresses the paradox of the inconsistent but cunning planner in a way which impatience cannot, rests on the two ideas which are the basis of our extension of game theory. Knowledge is not necessarily describable by partitions, and even the most clever men do not necessarily make inferences from the constraints they face.

The Odysseus game can be formally modeled in our framework, as the following diagram makes clear. Nature moves first and chooses to blow Odysseus' boat near the island of the Sirens, or near some other harmless island (on which there is also singing, but perhaps less good). Odysseus then has the choice of binding himself and putting wax in the ears of his men, or else leaving them all free. Finally he hears the music, and must decide whether to give the order to stay clear of the shore, or to move toward the shore. Note that if he has put wax in his mens' ears, then he is constrained not to give the order to head closer to shore, although the payoff if he were able is still defined.

The novelty about this game is the information Odysseus has. Odysseus hears very well the advice of Circe, and so knows whether the boat will be sailing past the island of the Sirens or a harmless island. If he hears songs from the harmless island, he remembers well whether he put wax in the ears of his men. A best strategy from that point is to sail close to shore. But if Odysseus is near the Sirens' island, and hears their singing, then he forgets completely Circe's advice. Moreover, although he recognizes whether he is bound or not, he does not infer anything from this.

Figure 3

There is a unique equilibrium to this generalized game in extensive form. Odysseus applies the wax if Circe advises him that he will pass the Sirens, and otherwise, he does not. After hearing either song, he always tries to head for shore, but when he is constrained from giving such an order he does the only thing he can, and permits the boat to stay clear. The reason this last move of Odysseus is optimal is because Odysseus computes that 99.9% of the time he is called upon to make a decision about whether to better hear a beautiful song, it is worth doing. Knowing that he will decide this way later, Odysseus earlier on has no choice but to put the wax in his mens' ears and tie himself to the mast, even though he would be much better off sailing clear of the island unfettered.

The time inconsistency of Odysseus' behavior is mirrored in a host of similar examples usually having to do with temptation. Typically the optimal response to a pleasant sensation is to increase it. Life's experiences strongly encourage such priors. There are some pleasant experiences, like some drugs or cigarette smoking that some people recognize to be harmful for them. However, when under their influence, or sometimes just in their presence, they forget the particular, and reason only from the general principle that pleasure is desirable. One occasionally meets modern day Odysseuses who deliberately leave their money home so they will not be tempted by anything fattening, or who join clubs like alcoholics anonymous so that their drinking will be punished by shame as well as hangovers.

References

[1] Aumann, R., "Agreeing to Disagree," *The Annals of Statistics* (1976) 4:1236–1239.

[2] Aumann, R., "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica* (1987), 55:1–18.

[3] Binmore, K. and A. Brandenburger, "Common Knowledge and Game Theory," London School of Economics, 1988.

[4] Brandenburger, A. and J. Geanakoplos, "Common Knowledge of Summary Statistics," unpublished, Cowles Foundation, Yale University, 1986.

[5] Brandenburger, A., E. Deckel, and J. Geanakoplos, "Correlated Equilibrium with Generalized Information Structures," Yale University, 1988.

[6] Brown, D. and J. Geanakoplos, "Common Knowledge without Partitions," unpublished, 1988.

[7] Cave, J., "Learning to Agree," *Economic Letters* (1983), 12:147–152.

[8] Doyle, Arthur Conan, "The Adventure of Silver Blaze," *The Strand Magazine*, 1901.

[9] Dubey, P., J. Geanakoplos, and M. Shubik, "The Revelation of Information in Strategic Market Games: A Critique of Rational Expectations Equilibrium," *Journal of Mathematical Economics* (1987), 16(2):105–138.

[10] Geanakoplos, J. "Common Knowledge of Actions Negates Asymmetric Information," mimeo, Yale University, 1987.

[11] Geanakoplos, J., "Common Knowledge, Bayesian Learning, and Market Speculation with Bounded Rationality," mimeo, Yale University, 1988.

[12] Geanakoplos, J. and P. Milgrom, "A Theory of Hierarchies Based on Limited Managerial Attention," CFDP No. 775, Cowles Foundation, Yale University, October 1985.

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

[14] Kreps, D., "A Note on Fulfilled Expectations Equilibrium," *Journal of Economic Theory* (1977), 32–43.

[15] Lewis, D., *Conventions: A Philosophical Study.* Cambridge: Harvard University Press, 1969.

[16] McKelvey, R. and T. Page, "Common Knowledge, Consensus, and Aggregate Information," *Econometrica* (1986), 54:109–127.

[17] Milgrom, P., "An Axiomatic Characterization of Common Knowledge," *Econometrica* (1981), 49:219–222.

[18] Milgrom, P. and N. Stokey, "Information, Trade, and Common Knowledge," *Journal of Economic Theory* (1982), 26:17–27.

[19] Monderer, D. and D. Samet, "Approximating Common Knowledge with Common Beliefs," mimeo, Northwestern, 1988.

[20] Samet, D., "Ignoring Ignorance and Agreeing to Disagree," MEDS Discussion Paper, 1987, Northwestern University.

[21] Rubinstein, A. and A. Wolinsky, "A Comment on the Logic of 'Agreeing to Disagree' Type Results," mimeo.

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

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

[24] Shin, H., "Logical Structure of Common Knowledge, I and II," unpublished, Nuffield College, Oxford, 1987.

[25] Strotz, R. H., "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies* (1955–6), 23:165–180.

[26] Tirole, J. "On the Possibility of Speculation under Rational Expectations," *Econometrica* (1982), 1163–1181.