

Strategies as states*

Forthcoming *Journal of Economic Theory*

Michael Mandler
Department of Economics
Royal Holloway College, University of London
Egham, Surrey TW20 0EX
United Kingdom
m.mandler@rhul.ac.uk
44-1784-443985

December 2002

Revised: May 2006

Running title: Strategies as states

* I thank Robert Aumann, Ken Binmore, Stephen Morris, Abhinay Muthoo, and Marciano Siniscalchi for helpful advice and discussions. A referee suggested several important improvements.

Abstract:

We define rationality and equilibrium when states specify agents' actions and agents have arbitrary partitions over these states. Although some suggest that this natural modeling step leads to paradox, we show that Bayesian equilibrium is well-defined and puzzles can be circumvented. The main problem arises when player j 's partition informs j of i 's move and i knows j 's strategy. Then i 's inference about j 's move will vary with i 's own move, and i may consequently play a dominated action. Plausible conditions on partitions rule out these scenarios. Equilibria exist under the same conditions, and more generally ε equilibria usually exist.

JEL Numbers: C70, D81, D82

Keywords: Bayesian decision theory, interactive epistemology, Newcomb's paradox, causal decision theory

1. Introduction

Agents in decision theory take as given a set of states that specify every feature of the world that affects them and the probabilities of those states. When placed in a game, therefore, an agent i should take as given the probabilities of the actions that any other agent j plays. But although the probabilities of j 's actions are given from i 's perspective, they cannot be given from j 's own perspective. To deal with this complexity, one can let states specify profiles of the agents' actions. Every agent then faces the same state space and all can simultaneously maximize utility, with a move for an agent i leading i to update probabilities accordingly. But a difficulty arises: if j 's partition informs j of i 's move and i knows j 's strategy (i.e., j 's move as a function of the state), then i 's own move can reveal to i what move j is taking. Agents can thus gather substantive information from their own moves and the customary independence between actions and knowledge breaks down. Even worse, in this scenario agent i can rationally play a dominated action. Some (Binmore [5], Shin [16]) have recognized these possibilities and suggested that allowing agents to draw Bayesian inferences from their own actions leads too far astray from orthodox game theory. We will show that these concerns are exaggerated; with the right conditions on knowledge in place, Bayesian rationality can be applied to states that specify actions and yet remain consistent with traditional thinking about games.

The keys that open the door to the play of dominated actions are first, states that specify actions, and second, agents who receive partitional information about other players' moves. The less distinctive equilibrium assumption that agents know other players' strategies is also important. While Aumann [3] and the literature in its wake give agents partitions over states that specify actions, we will argue that the Aumann model evades the consequences of these assumptions.¹ Curiously it is the traditional dictates of Bayesian rationality – states that specify

¹ Much of the remainder of the literature on Bayesian games employs a base state space that does not specify actions (as in the Harsanyi model) or, when actions states do list actions, assumes that agents do not have partitions over the state space. For a sampling of the literature, see Armbruster and Böge [1], Aumann and Brandenburger [4], Böge and Eisele [7], Brandenburger and Dekel [8], Dekel and Gul [9], Mertens and Zamir [14], Tan and Werlang [18]. Brandenburger and Dekel [8] also considers Newcomb's paradox, which we discuss below.

every relevant contingency, and partitions – that can lead to anomaly.

The classic illustration of the feedback from action to knowledge when states specify actions, and the resulting potential for the play of dominated actions, is the famous Newcomb problem. An agent c (for chooser) is presented with two boxes, one of which is opaque and may or may not contain a million dollars and one of which is transparent and visibly contains a thousand dollars: c may either follow the one-box action c_1 , and take only the opaque box or the two-box action c_2 , and take both boxes. At an earlier point in time, the other agent d (for demon) either places a million dollars in the opaque box, action d_2 , in which case both boxes contain money, or leaves the opaque box empty, d_1 , in which case only the transparent box has money. Agent c 's utility is given by his money payoff, while d wants justice to prevail: d prefers the outcomes where c modestly chooses only the opaque box and it has money or c greedily chooses both boxes but the opaque box is empty over the outcomes where c chooses the single opaque box but finds it empty or c chooses both boxes and discovers the opaque box is full. The payoffs are given below, where $x > y$.

	d_1	d_2
c_1	0, y	10^6 , x
c_2	1000, x	$10^6 + 1000$, y

In standard game theory, c will play c_2 , which dominates c_1 , and d will therefore select d_1 . The supposed paradox lies in the long history of past plays of the game. Some c 's have chosen c_1 – perhaps they have different payoffs or think that sometimes it can be rational to play a dominated action. And it turns out that whenever c selects c_1 , d had earlier chosen d_2 . Agent d is somehow able to predict with perfect accuracy if he is dealing with a c who chooses c_1 or a c who chooses c_2 . Some conclude that there is then a logic to choosing c_1 . If d can indeed predict with perfect accuracy the sort of c he faces, then c will benefit from deciding to be the sort of c who chooses c_1 . The following story sometimes bolsters this argument: if d has extensive knowledge of the physiology of c 's decision-making and sufficient prior information,

then he or she should be able to predict c 's actions flawlessly. This story is implausible, but it underscores the need for a minimal and plausible formal condition that will rule it out.

The above case for playing c_1 is noticeably vague about the nature of d 's knowledge of c 's actions. If we let the state space specify agents' actions and model d 's knowledge by a partition, we can put the argument for choosing c_1 in Bayesian terms. Suppose the state space Ω specifies just c 's and d 's actions: $\Omega = \{(c_1, d_1), (c_1, d_2), (c_2, d_1), (c_2, d_2)\}$. Agent d 's flawless knowledge of c 's move may then be modeled with the following partition of Ω :

$$\{\{(c_1, d_1), (c_1, d_2)\}, \{(c_2, d_1), (c_2, d_2)\}\}.$$

Assuming c has no comparable partitioned knowledge of d 's move, c 's partition is just $\{\Omega\}$.

When a player i moves after having originally been informed of the cell P_i , the player updates using the additional information implied by his move: the original cell P_i is replaced by the

subset of P_i that consists of those states that report the move that i takes. For example, if d is originally informed of $\{(c_1, d_1), (c_1, d_2)\}$ and moves d_1 then d knows that $\{(c_1, d_1)\}$ obtains,

while if c , who is always informed of just Ω , moves c_1 then c knows that $\{(c_1, d_1), (c_1, d_2)\}$

obtains. By itself, d 's information about c 's move will not lead c to play the dominated action

c_1 : if the probabilities that d plays d_1 or d_2 are fixed, c will prefer to take the move c_2 . What

leads to trouble is if d 's action depends on d 's information and c knows this. So suppose that

each player knows the equilibrium strategy of the other player and that d plays the utility-

maximizing strategy of choosing d_2 when facing $\{(c_1, d_1), (c_1, d_2)\}$ and d_1 when facing

$\{(c_2, d_1), (c_2, d_2)\}$. Then c knows that only the states (c_1, d_2) and (c_2, d_1) are possible and so

c , by playing c_1 , can lead the more preferred of the two, (c_1, d_2) , to obtain. If c plays c_1 , c 's

Bayesian inference is that d must have had information that led to the play of d_2 , while if c plays

c_2 then c would infer that d had information that led to the play of d_1 .

Notice that c 's partition is completely uninformative; it is d 's information combined with

the equilibrium assumption that players know each other's strategy that leads c to play c_1 . More

generally, the fact that an agent i receives partitioned information about another player j 's move

will not lead i to play a dominated action – it just helps i play a better response. For a dominated

action to be rational, it is one's opponents' information that matters. It is also important that d 's partition carries information about c 's move, not just about c 's type (payoffs).

The present example of playing a dominated action is the simplest of the cases we will consider. Here d 's partition gives d perfect knowledge of c 's move, but all that is necessary is that d has some partitional information. In fact, a player i can still play a dominated action even when some other agent j has partitional knowledge of i 's move that is so weak that j always judges each of i 's possible actions to have positive probability.

It is the converse to these examples that is most important: if no agent has any partitional information about other agents' moves then no agent will play a dominated action. Indeed in the absence of partitional information about others' moves (and with a plausible restriction on how to update on 0 probability events), we will arrive at the same predictions as Aumann's model of correlated equilibrium. An absence of partitional knowledge of others' moves is a natural and plausible way to get rid of the paradoxes that can accompany states that specify actions. Regarding Newcomb's paradox, these results allow us to identify the conditions under which c_1 is rational. Agent d must have some partitional knowledge of c 's move, it is not enough for d merely to be an accurate predictor of c (see sections 5.ii and 9).

Our model weakens Savage's independence between knowledge and actions but not so drastically that we end up on the Jeffrey [12] side of the divide in decision theory, where agents' actions can change the probabilities of states in arbitrary ways. Instead the updating that takes place as an agent moves occurs only via the refinement of the agent's prespecified partition. As a consequence, an agent i 's act of moving does not *directly* inform i of another player j 's move; it is the impact of i 's move on j 's information and hence on j 's equilibrium action that gives i information and leads i to play a dominated action.

Traditional game and decision theory blocks the rational play of dominated actions either by eliminating actions from the description of states or by not letting agents update as they move. Agents then cannot infer anything about the state space from their own moves. But these crude steps both go too far and fail to pin down where the problem lies. Agents should be able

to draw at least some conclusions about the world from their moves – they are subject to physical laws, after all, and therefore ought to be able to deduce some physical facts from their moves. If we abide by the Savage dictum that states omit no relevant detail about the world, such inferences become unavoidable. But with appropriate and plausible auxiliary assumptions in place, they do not lead to trouble.

In addition to analyzing dominated move pathologies, we also aim to show that games and information updating can be analyzed coherently when states specify actions. Rational play and equilibrium are readily definable even when agents do receive partitional information about others' moves. The existence of equilibrium involves some complications not present in ordinary game theory, but equilibria or ε equilibria exist in the important cases.

2. Strategies as states: the problem

Aumann's [3] theory of Bayesian decision-makers playing a game was the first to endow agents with partitions over states that specify actions, and it illustrates the characteristic problem: in order to preserve the independence of actions and knowledge, utility maximization is defined so as to ignore the effect of an agent's action on the state.

Let $I = \{1, \dots, n\}$ be the set of players with each player i having a set of actions or moves S_i and define $S = S_1 \times \dots \times S_n$. Each i has the utility function $u_i: S \rightarrow R$. Uncertainty is described by a set of states Ω , where each $\omega \in \Omega$ specifies (among other things) the players' actions. The agents share a common prior π on Ω while each i 's private information is modeled by a partition \mathcal{P}_i of Ω such that each cell $P_i \in \mathcal{P}_i$ has $\pi(P_i) > 0$. Let $P_i(\omega)$ denote the cell of \mathcal{P}_i that contains ω . Agent i chooses a strategy, $\sigma_i: \Omega \rightarrow S_i$, that specifies i 's preferred action in S_i as a function of the state. The function σ_i is required to be *measurable* with respect to \mathcal{P}_i . That is, for all $\omega, \omega' \in \Omega$,

$$P_i(\omega) = P_i(\omega') \Rightarrow \sigma_i(\omega) = \sigma_i(\omega').$$

It is not clear if the Aumann model pertains to a point in time before or after actions are taken. If before, then measurability means that i cannot vary his or her move as a function of states in

the same cell. If after, and if we assume informally that i knows σ_i , then measurability means that i knows his or her own move.

Let the strategy functions $(\sigma_i)_{i \in I}$ be *Aumann rational* if and only if, for all $\omega \in \Omega$, $i \in I$, and all σ_i' that are measurable with respect to \mathcal{P}_i ,

$$E(u_i(\sigma_i, \sigma_{-i})|P_i(\omega)) \geq E(u_i(\sigma_i', \sigma_{-i})|P_i(\omega)),$$

where $\sigma_{-i} = (\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n)$ and $E(u_i(\sigma_i, \sigma_{-i})|A)$ is the expectation of $u_i(\sigma_i, \sigma_{-i})$ conditional on the event A .

Aumann rationality has several drawbacks (see Binmore [5], Binmore and Brandenburger [6], Shin [16]). The left hand side of the above inequality gives i 's expected utility in the event $P_i(\omega)$. By measurability, i takes the same action, say $s_i = \sigma_i(\omega)$, at each $\hat{\omega} \in P_i(\omega)$; any $\hat{\omega} \in P_i(\omega)$ should therefore describe a state at which s_i occurs. But if i were instead to play the strategy σ_i' where $\sigma_i'(\omega) = s_i' \neq s_i$ for all $\omega \in P_i$, the right hand side of the inequality says that i 's payoff would be $E(u_i(\sigma_i', \sigma_{-i})|P_i(\omega))$, the expected utility of playing s_i' given that the event $P_i(\omega)$ obtains. But $P_i(\omega)$ contains only states at which i plays s_i , not s_i' . Agent i would thus have to consider his payoff in a state that i deems to be impossible – a state not in the cell $P_i(\omega)$.

Second, how should we interpret the σ_i ? The function reports the action s_i chosen as a function of the state ω . But if a state specifies what actions are taken, then σ_i evaluated at ω could report only the action for i that is given by ω (σ_i would simply project ω onto the coordinate that gives i 's action). Hence Ω normally cannot include every possible profile of actions $s \in S$ if Aumann rationality is to obtain. When for example $\omega \in \Omega$ specifies a strictly dominated action $s_i \in S_i$, and therefore $\sigma_i(\omega) = s_i$, then i cannot be Aumann rational at this ω .

If Ω has states that specify every possible profile of actions, then an agent i 's action must reveal information to i , and thus i should condition on different events as i considers different possible moves. The rest of this paper considers the consequences of this conditioning, in particular that as the events on which i conditions change, i may be able to infer the moves of other players. If for instance i were to deviate to the action s_i' from the action $\sigma_i(\omega)$, then some

$j \neq i$ might be informed that a different cell of \mathcal{P}_j obtains and thus i might be able to infer, from knowledge of σ_j , that j will play a distribution of actions that differs from $\sigma_j|P_i(\omega)$.

3. Partitional rationality

Again the set of players is $\mathbb{I} = \{1, \dots, n\}$ where each $i \in \mathbb{I}$ has the set of actions S_i with typical element s_i , and S again denotes $S_1 \times \dots \times S_n$. To make players' actions explicit in the description of a state, the state space Ω will now be a subset of $S \times \Gamma$, with typical element $\omega = (s, \gamma)$, where Γ indicates all relevant features of the world besides players' moves. We assume throughout that each S_i is a compact subset of some Euclidean space, e.g., all mixtures of a finite set of pure actions. But no structure is imposed on Γ , and so γ could, for example, specify an infinite hierarchy of each player i 's beliefs, beliefs about j 's beliefs, etc. Since we sometimes need to consider convex action sets, we allow Ω to be uncountable and endow Ω with a σ -algebra of measurable subsets.

Notation: Let $S_i(A)$, where $A \subset \Omega$, denote the projection of A onto S_i , $S(A)$ denote the projection of A onto S , etc. We use $s_i(\omega)$, $i \in \mathbb{I}$, and $\gamma(\omega)$ to denote the coordinates of ω . When $S_i(A)$ is a singleton, we also use $s_i(A)$ to denote $s_i \in S_i(A)$. As usual, $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$, $S_{-i} = S_1 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_n$, etc.

We assume throughout that $S(\Omega) = S$. In every important example, Ω in fact equals the product $S \times \Gamma$ (one exception arises in this section and two others in section 5).

Each agent i is described by a utility $u_i: \Omega \rightarrow R$ and a *premove partition* \mathcal{P}_i of Ω consisting of measurable cells that indicate i 's information prior to moving. We assume for simplicity that i cannot exclude ex ante the possibility of any of his or her moves: for each agent i , each $s_i \in S_i$, and each $P_i \in \mathcal{P}_i$, there exists a $\omega \in P_i$ such that $s_i(\omega) = s_i$. Let P_{-i} denote a $(P_1, \dots, P_{i-1}, P_{i+1}, \dots, P_n) \in \mathcal{P}_1 \times \dots \times \mathcal{P}_{i-1} \times \mathcal{P}_{i+1} \times \dots \times \mathcal{P}_n$.

The act of moving refines agent i 's partition, leading to a *postmove partition* \mathcal{P}_i^* of measurable cells of Ω , where \mathcal{P}_i^* refines \mathcal{P}_i . We interpret $P_i^* \in \mathcal{P}_i^*$ both as what i knows after having taken the move $s_i(P_i^*)$ and as what i anticipates knowing if i were to take the action

$s_i(P_i^*)$. The latter vantage point is the relevant one when we consider what moves are rational for i . The coincidence of the two interpretations amounts to an assumption of rational expectations.

We impose the following requirements on the \mathcal{P}_i^* . Given any $P_i, P_i^* \subset P_i, P_i^{*'} \subset P_i, \omega \in P_i^*,$ and $\omega' \in P_i^{*'},$ then

KYOM (know your own move) $P_i^* = P_i^{*'} \Rightarrow s_i(\omega) = s_i(\omega'),$

NEI (no extra information) $P_i^* = P_i^{*'} \Leftarrow s_i(\omega) = s_i(\omega').$

KYOM says that if two actions are in the same cell of the postmove partition, then they are the same action: after moving, i knows his or her own action. This has the same meaning as the ‘after the action’ interpretation of measurability in the previous section: an agent knows and remembers his move. Notice that when S_i consists of an infinite number of actions, then KYOM implies that the postmove partition \mathcal{P}_i^* contains an infinite number of cells.

NEI says that if two cells of i 's postmove partition report the same move for i (and they originate from the same cell of i 's premove partition) then they are the same cells: when i cannot distinguish ex ante between two states at which i makes the same move, the act of moving does not by itself distinguish the states. Assuming that nothing else besides i 's move occurs when i chooses a $P_i^*,$ we view NEI as a rationality requirement. To see this, suppose \mathcal{P}_i^* refines \mathcal{P}_i and satisfies KYOM but violates NEI: then there exist distinct \tilde{P}_i^* and $P_i^{*'} in some P_i and a \hat{s}_i such that $s_i(\omega) = \hat{s}_i$ for all $\omega \in \tilde{P}_i^* \cup P_i^{*}'$. If taking the move \hat{s}_i necessarily leads to one of the postmove cells, say $P_i^{*}',$ then i ought to be able to deduce premove that the states in P_i^* cannot obtain. Moreover some $P_i^* \in \mathcal{P}_i^*$ should necessarily obtain when i selects \hat{s}_i since we think of i as choosing a P_i^* in P_i . If both \tilde{P}_i^* and P_i^{*}' could obtain, and again assuming that nothing else occurs when i moves, then how does i select, say, \tilde{P}_i^* rather than $P_i^{*}'?$$

It is readily confirmed that given a partition $\mathcal{P}_i,$ there is a unique partition \mathcal{P}_i^* that both refines \mathcal{P}_i and satisfies KYOM and NEI.² Henceforth all postmove partitions will be generated

² Let \mathcal{P}_i^* and \mathcal{P}_i^{*}' denote partitions that refine \mathcal{P}_i and satisfy KYOM and NEI. For any $P_i^* \in \mathcal{P}_i^*,$ there exists $P_i^{*'} \in \mathcal{P}_i^{*}'$ such that $P_i^* \cap P_i^{*}' \neq \emptyset$ and hence a $\omega \in P_i^* \cap P_i^{*}'$. Then, for any $\omega, \omega' \in P_i^* \cup P_i^{*}',$

from premove partitions by KYOM and NEI.

KYOM and NEI do not imply that agents learn nothing from the act of moving above and beyond their own move. For instance, if \mathcal{P}_i consists of a single cell $\{(s_i, s_{-i}, \gamma), (s'_i, s_{-i}, \gamma')\}$ then \mathcal{P}_i^* is $\{\{(s_i, s_{-i}, \gamma)\}, \{(s'_i, s_{-i}, \gamma')\}\}$. So if i moves s_i then i may infer that γ obtains, while if i moves s'_i then i knows that γ' obtains. The variables γ and γ' might denote different arrays of information about the physical world at an earlier date, and so i might know that the move s_i is perfectly correlated with γ and that s'_i is perfectly correlated with γ' . The following assumption excludes this possibility.

NI (no information): for any P_i , $\omega \in P_i$, and $s_i \in S_i$, there exists $P_i^* \subset P_i$ such that

$$(s_i, s_{-i}(\omega), \gamma(\omega)) \in P_i^*.$$

NI says that every (s_{-i}, γ) that arises in some cell of i 's postmove partition arises in every other postmove cell that originates from the same premove cell: each P_i is the product of S_i and some $Q_{P_i} \subset S_{-i} \times \Gamma$. (And so, given KYOM and NEI, each $P_i^* \subset P_i$ has the form $\{s_i\} \times Q_{P_i}$ for some $s_i \in S_i$.)

NI, although implausibly strong, is implicit in most models, such as Aumann's, of interacting Bayesian agents: agents know nothing after moving that they do not know before moving. NI nevertheless does not eliminate the rational play of dominated actions: the simple Newcomb example in the introduction satisfies NI as will virtually every other example we consider where agents rationally play dominated actions (there is one minor exception in section 5). We will see in section 6 that NI also need not obtain when it is not rational to play dominated actions. Despite its strength, NI is neither necessary nor sufficient for the traditional view of rational play.

In addition to the state space Ω and the *partition profile* $(\mathcal{P}_i)_{i \in \mathcal{I}}$, the final primitives of the model are the agent utilities $u_i: \Omega \rightarrow \mathbb{R}$, $i \in \mathcal{I}$. We assume that each u_i is integrable with respect to any probability measure on Ω . Each u_i may be constant as a function of γ , in which

KYOM implies $s_i(\hat{\omega}) = s_i(\omega)$ and $s_i(\hat{\omega}) = s_i(\omega')$. Hence $s_i(\omega) = s_i(\omega')$ and so, by NEI, $P_i^* = P_i^{*'}$.

case u_i has the specification of the previous section. To avoid distraction, we use the same probability measure π on Ω to calculate expected utilities for each u_i . Although agents' choice of actions affect the probabilities of states and hence one cannot view π as wholly exogenous, we could allow the elements of the meet (finest common coarsening) of the partitions $\mathcal{P}_i, i \in \mathbb{I}$, to have exogenously given probability.

The rationality of agents turns on how they select elements of the postmove partition \mathcal{P}_i^* . The *strategy* of agent i is a function $h_i: \mathcal{P}_i \rightarrow \mathcal{P}_i^*$ such that, for each $P_i, h_i(P_i) \subset P_i$. A profile of strategies is denoted $h = (h_1, \dots, h_n)$.

We assume that an agent i who uses the strategy h_i against opponents playing $h_{-i} = (h_1, \dots, h_{i-1}, h_{i+1}, \dots, h_n)$ knows that the event $h_i(P_i)$ obtains when informed initially of the cell P_i . But if i knows the functions h_{-i} , as one usually assumes in equilibrium analysis, i can infer more: since each $j \neq i$ selects only states that lie in some $h_j(P_j)$, any state that obtains should be in some $h_j(P_j)$ for $j \neq i$ as well as in $h_i(P_i)$, assuming such states exists. To state the restriction that only these states have positive probability, let $\text{Range } h_i$ (in a slight abuse of notation) denote $\{\omega \in \Omega: \omega \in h_i(P_i) \text{ for some } P_i \in \mathcal{P}_i\}$, and define $\text{Range}_{\mathbb{I}} h = \bigcap_{j \in \mathbb{I}} \text{Range } h_j$ and $\text{Range}_{-i} h = \bigcap_{j \in \mathbb{I} \setminus \{i\}} \text{Range } h_j$.

Definition 1. The probability π is *accurate* with respect to the strategies h if and only if $\pi(\text{Range}_{\mathbb{I}} h) = 1$.

In order to optimize, an agent i must take into account that as i varies the P_i^* in P_i that he selects, the $h_j(P_j), j \neq i$, that intersect P_i^* may change. Each agent i thus anticipates what he would know if he were to undertake the various P_i^* in P_i . Agent i therefore may be able to infer information about the moves of $j \neq i$ move from his own move, and this information can influence which move is optimal for i . If i is optimizing, i should choose an action whose expected utility, conditioning on what i knows given that i takes that action, is at least as great as the expected utility of any alternative action, conditioning on what i would know if i were indeed to take that alternative. So i when facing the premove cell P_i should choose a $h_i(P_i)$ such

that

$$E(u_i | h_i(P_i) \cap \text{Range}_{\neg i} h) \geq E(u_i | P_i^* \cap \text{Range}_{\neg i} h)$$

for each $P_i^* \subset P_i$.³ Unfortunately, the conditioning events above can be empty, in which case expectations are not well-defined. The empty event can arise, for instance, if \mathcal{P}_i informs i of some j 's move since then some P_i can imply moves for j that j does not take with h_j . This possibility is important. For example, in the Newcomb example in the introduction where c 's premove partition consists of the entire state space Ω , let h_c specify the move c_2 (formally $h_c(\Omega) = \{(c_2, d_1), (c_2, d_2)\}$). And suppose again that d faces the premove cells $\{(c_1, d_1), (c_1, d_2)\}$ and $\{(c_2, d_1), (c_2, d_2)\}$. In order to specify rational play for d in all eventualities, we must define d 's expected utility when c plays c_1 and d therefore faces $\{(c_1, d_1), (c_1, d_2)\}$. Whatever move d plays at this cell, $P_d^* \cap \text{Range } h_c$ will equal the empty set since h_c specifies c_2 but P_d^* specifies c_1 , and so d 's expected utility is not well-defined. Moreover, we cannot leave d 's expected utility undefined in the event that c plays c_1 , since, for c to optimize, c must know how d would act if c hypothetically were to play c_1 .

To deal with the empty conditioning event, we assume that when i contemplates a cell P_i^* that does not intersect $\text{Range}_{\neg i} h$, then i takes P_i^* itself to be the set of possible states. Let $K(P_i^*)$ denote $P_i^* \cap \text{Range}_{\neg i} h$ if $P_i^* \cap \text{Range}_{\neg i} h \neq \emptyset$ and P_i^* otherwise. So our assumption is that i knows $K(P_i^*)$ when taking the move P_i^* . Also, π will denote both a probability measure $\pi(\cdot)$ on Ω and, for any measurable $P \subset \Omega$, a conditional probability measure $\pi(\cdot | P)$ on Ω . Given $\pi(\cdot | P)$ and an integrable u , $E(u | P)$ (the conditional expectation of u given P) will denote $\int u(\omega) d\pi(\omega | P)$ rather than some arbitrary version of the conditional expected value.

Definition 2. The strategies and probability (h, π) form a *partitionally rational equilibrium* if and only if π is accurate with respect to h and for all $i \in \mathbb{I}$, P_i , and $P_i^* \subset P_i$,

$$E(u_i | K(h_i(P_i))) \geq E(u_i | K(P_i^*)).$$

To see how Definition 2 is applied when agents face the empty event, see the matching

³ Shin [16], cited in Binmore and Brandenburger [6], appears to propose a similar definition of rationality.

pennies example in section 4 or section 5. Even when the empty event does not arise, many action profiles are not taken when the agents play a given h . Consequently zero-probability events appear routinely as each i considers various P_i^* . Conditional probabilities are then not uniquely determined. Definition 2 requires only that the h_i are optimal for some set of conditional probabilities.

4. Existence and nonexistence of partitionally rational equilibria

We begin by using matching pennies to illustrate partitionally rational equilibria and clarify the role played by accuracy, and then turn to an example that satisfies standard convexity and continuity conditions but that has no partitionally rational equilibria. Nonexistence stems from the fact that the cells of \mathcal{P}_i typically cannot all be compact. But we show that models with the more important types of partition profiles always possess equilibria and that ε equilibria exist in two-player models if non-move uncertainty (information about γ) is symmetric. This section is self-contained except for our later use of Definition 4; the reader may wish just to peruse some of the examples.

Keep in mind that, unlike the Aumann [2] model of correlated equilibrium, the state space and agent partitions are among the primitives $(\Omega, (\mathcal{P}_i, u_i)_{i \in \mathbb{I}})$ rather than part of the definition of equilibrium. So for example a statement that “no equilibrium exists” does not mean that no equilibrium exists for a different specification of partitions.

Matching pennies. There are two players a and b . Let $\Omega = \{(H, T), (H, H), (T, T), (T, H)\}$, where the first coordinate of each ω is a 's move and the second is b 's. There is no non-move uncertainty: γ is constant across states and suppressed in the notation. Utilities are given by

$$\begin{aligned} u_a(H, T) &= u_a(T, H) = u_b(H, H) = u_b(T, T) = 1, \\ u_a(H, H) &= u_a(T, T) = u_b(H, T) = u_b(T, H) = 0. \end{aligned}$$

Assume first that before moving each player knows the other's move:

$$\mathcal{P}_a = \{\{(H, T), (T, T)\}, \{(H, H), (T, H)\}\}, \mathcal{P}_b = \{\{(H, T), (H, H)\}, \{(T, T), (T, H)\}\}.$$

Then, given KYOM and NEI,

$$\mathcal{P}_a^* = \mathcal{P}_b^* = \{\{(H, T)\}, \{(H, H)\}, \{(T, H)\}, \{(T, T)\}\}.$$

One might think that partitional rationality would be possible since a for example will reason that if he were to move H , b will know that and move T , and if he were to move T , b would know that too and move H . So a should be indifferent between his two moves. The difficulty is that *both* players know the other's move. When, for instance, a knows that b plays H and therefore plays T – so $h_a(\{(H, H), (T, H)\}) = \{(T, H)\}$ – the state must be (T, H) . But if h_b assigns $\{(T, T)\}$ to $\{(T, T), (T, H)\}$ then state (T, H) is not in Range h_b . And similarly if a knows b plays T , we have $h_a(\{(H, T), (T, T)\}) = \{(H, T)\}$ and $h_b(\{(H, T), (H, H)\}) = \{(H, H)\}$, and again h_a and h_b do not intersect. Hence there cannot be an accurate π .

If instead just one player, say b , knows the other's move, then a partitionally rational equilibrium does exist. Suppose b knows a 's move but not vice versa:

$$\mathcal{P}_a = \{\Omega\}, \mathcal{P}_b = \{\{(H, T), (H, H)\}, \{(T, T), (T, H)\}\}.$$

If a were to play H , then b will play H , and if a were to play T , then b will play T . Two pairs of strategies are therefore possible in partitionally rational equilibrium. In one, a plays H (i.e., $h_a(\Omega) = \{(H, H), (H, T)\}$) and in the other, a plays T , while in both,

$$h_b(\{(H, T), (H, H)\}) = \{(H, H)\}, \text{ and } h_b(\{(T, T), (T, H)\}) = \{(T, T)\}.$$

Accuracy is achieved in the first case if

$$\pi((H, H)) = 1, \pi((H, T)) = \pi((T, H)) = \pi((T, T)) = 0,$$

and in the second if

$$\pi((T, T)) = 1, \pi((H, T)) = \pi((T, H)) = \pi((H, H)) = 0.$$

Notice that although a partitionally rational equilibrium now exists, one of the agents, b , must specify play in the face of the empty event. For example, in the equilibrium where a plays H , b must still specify a move when informed of the cell $\{(T, T), (T, H)\}$. The dictates of partitional rationality are entirely noncontroversial, however: b plays (T, T) when $\{(T, T), (T, H)\}$ occurs.

The nonexistence that occurs above when both players know the other's move disappears

if mixed actions are allowed. If the action sets contain all mixtures of H and T – so $S_i = [0, 1]$ rather than $S_i = \{H, T\}$ – and payoffs are extended accordingly, then equilibrium for matching pennies will always exist regardless of the partition profile. But this repair works only because matching pennies is zero-sum: if both agents play their secure actions, any deviation for i to give j different information and thereby induce a change in j 's move will not raise i 's payoff. In non-zero-sum games, such deviations can be profitable and can lead to nonexistence of equilibrium even when all mixed actions are available, as the following example shows. ■

Nonexistence with mixed actions. The payoffs to pure actions are given by the following matrix.

	a	b		L	R
	T		0, 2	6, 1	
	B		2, 1	4, 2	

Let p denote the probability that a plays T and q denote the probability that b plays L . Payoffs to mixed actions are given by the expectations of pure payoffs. There is no non-move uncertainty, and so $\Omega = [0, 1] \times [0, 1]$, with typical state $(p, q) \in \Omega$. Premove partitions are given by $\mathcal{P}_a = \{[0, 1] \times [0, 1]\}$, and $\mathcal{P}_b = \{[0, .4] \times [0, 1], [.4, 1] \times [0, 1]\}$. Thus, a receives no partitional information about b 's move, whereas b knows that a is either playing $p \in [0, .4)$ or $p \in [.4, 1]$.

The sole Nash equilibrium of the standard model (that is, where $\mathcal{P}_a = \mathcal{P}_b = \{[0, 1] \times [0, 1]\}$) occurs at $p = q = .5$. We can use this fact to exclude several equilibrium possibilities. If h_a selects any $p \in [.4, .5)$, b will play R ($q = 0$) when informed that $p \in [.4, 1]$; but if h_b selects R at the cell where $p \in [.4, 1]$, h_a must select T ($p = 1$). Similarly, if h_a selects any $p \in (.5, 1]$, b will play L when informed that $p \in [.4, 1]$; but if h_b selects L at the cell where $p \in [.4, 1]$, h_a must select B . Equilibrium also cannot occur at the standard Nash equilibrium $p = q = .5$. Although $q = .5$ is optimal for b when informed that $p \in [.4, 1]$ and h_a selects $p = .5$, $p = .5$ is not optimal for a . To see this, note that h_b must select R at the cell where $p \in [0, .4)$. So, by choosing p less than but near to $.4$, a can achieve expected utility arbitrarily near to $.4 \times 6 + .6 \times 4 = 4.8$, while, by choosing $p = .5$, a receives expected utility equal to 3 if h_b selects $q = .5$

at the cell where $p \in [.4, 1]$. We may (again invoking standard reasoning) also exclude the possibility of an equilibrium in which h_a selects $p = .5$ and h_b selects $q < .5$ or $q > .5$ at the cell where $p \in [.4, 1]$. Summing up, we conclude there can be no equilibrium in which h_a selects $p \geq .4$. But an h_a that selects $p < .4$ is also impossible since $p + \varepsilon$ for small $\varepsilon > 0$ gives a a higher expected utility (nearer to 4.8) than p does.

The example is robust in that for a sufficiently small open set of payoffs for the pure actions and boundaries for the partition cells, nonexistence will persist. Equilibria will exist however if we “perturb” \mathcal{P}_b by changing the cell $[0, .4) \times [0, 1]$ to $[0, .4] \times [0, 1]$ (and hence changing b ’s other cell from $[.4, 1] \times [0, 1]$ to $(.4, 1] \times [0, 1]$). ■

Noncompact partition cells, which are key to the example, unfortunately can be unavoidable. If one of two agents, say b , has a set of actions formed from at least two pure actions and their probability mixtures, and if the number of cells in \mathcal{P}_a is finite and greater than 1, then at least one cell of \mathcal{P}_a cannot be closed. And so it becomes possible for b to have a sequence of actions each of which lies in one P_a and along which b ’s expected utility is increasing but whose limit action switches a to another cell P_a' ; if a ’s action changes as a result and discretely lowers b ’s utility, equilibrium may not exist.

There are two types of remedies for the nonexistence problem. First, we can impose restrictions on information that imply that convergent sequences of actions cannot lead in the limit to a discrete fall in utility. The second remedy, perhaps the more promising path, is to consider ε equilibria.

Several prominent classes of partition profiles qualify under the first cure. Suppose in a two-agent model that one agent a knows *ex ante* as much as b , which we define to mean that \mathcal{P}_a refines \mathcal{P}_b^* . Agent a would then know b ’s action *ex ante* but b would not know a ’s action even *ex post* (unless b can infer it from $\text{Range } h_a$). The Newcomb example in the introduction (or see section 5.i below) provides an example, letting $a = d$. When a knows *ex ante* as much as b , then any change in b ’s action will *always* shift a to a new cell of \mathcal{P}_a . But with appropriate continuity

assumptions on u_a , a 's best response will move continuously as a function of b 's action, thus eliminating the discrete changes in utility that can lead to existence trouble. The existence result in this case is given as Theorem 1 below and is proved by backward induction. Even though a and b move simultaneously, if a knows as much ex ante as b , then it is 'as if' b moves first: b in effect acts as a Stackelberg leader and a as a Stackelberg follower. The appearance of an 'as if' temporal order of play is characteristic when some agents have partitional knowledge of other agents' moves.

Definition 3. *Partitional continuity* is satisfied if and only if, for any $i \in \mathbb{I}$ and $\mathcal{P}_i' \subset \mathcal{P}_i$, there exists a function $\gamma_{\mathcal{P}_i'}: S_i \times S_{-i}(\bigcup_{P_i \in \mathcal{P}_i'} P_i) \rightarrow \Gamma$ such that (1) for each $s \in S_i \times S_{-i}(\bigcup_{P_i \in \mathcal{P}_i'} P_i)$, $(s, \gamma_{\mathcal{P}_i'}(s)) \in \bigcup_{P_i \in \mathcal{P}_i'} P_i$, and (2) for all $j \in \mathbb{I}$, $v_j: S_i \times S_{-i}(\bigcup_{P_i \in \mathcal{P}_i'} P_i) \rightarrow R$ defined by $v_j(s) = u_j(s, \gamma_{\mathcal{P}_i'}(s))$ is continuous.

One simple way to satisfy partitional continuity is for γ to be payoff-irrelevant (see below) and for each agent's utility to be a continuous function of actions.

Theorem 1. If for $(\Omega, (\mathcal{P}_i, u_i)_{i=a,b})$ agent a knows ex ante as much as b , and partitional continuity is satisfied, then a partitionally rational equilibrium exists.

Proofs are in the appendix. Existence of equilibrium also obtains in the n -agent case where, for $i \in \{1, \dots, n-1\}$, i knows ex ante as much as $i+1$.

Next, suppose each i 's partition informs i of the moves of $j \neq i$. Subject to standard technical caveats, each i can then best respond to the actions of $j \neq i$ reported by P_i and the existence of a standard Nash equilibrium will ensure that $\text{Range}_{\mathbb{I}} h$ is nonempty. A sequence of actions for i now cannot in the limit lead to a discrete change in i 's utility since the premove cells of $j \neq i$ already inform j of s_i ; so s_{-i} is fixed along any sequence of postmove cells.

Definition 4. Agent i *partitionally knows* j 's move if and only if, for each P_i ,

$$\omega, \omega' \in P_i \Rightarrow s_j(\omega) = s_j(\omega').$$

Definition 5. The variable γ is *payoff-irrelevant* if and only if, for all $i \in \mathbb{I}$ and $\omega, \omega' \in \Omega$,

$$s_j(\omega) = s_j(\omega') \text{ for } j \in \mathbb{I} \Rightarrow u_i(\omega) = u_i(\omega').$$

We say that *convexity* is satisfied if and only if, for all $i \in \mathbb{I}$, S_i is convex. *Simple continuity* is satisfied if and only if for all $i \in \mathbb{I}$ and for all functions $\hat{\gamma}: S \rightarrow R$ with $(s, \hat{\gamma}(s)) \in \Omega$ for $s \in S$: (1) for any $s_{-i} \in S_{-i}$, $s_i \mapsto u_i(s_i, s_{-i}, \hat{\gamma}(s_i, s_{-i}))$ is quasiconcave, and (2) $s \mapsto u_i(s, \hat{\gamma}(s))$ is continuous. For later reference, *γ -simple continuity* is satisfied if and only if (1) and (2) are imposed only on functions $\hat{\gamma}$ such that there is a $\bar{\gamma} \in \Gamma$ with $\hat{\gamma}(s) = \bar{\gamma}$ for all $s \in S$ (as well as $(s, \hat{\gamma}(s)) \in \Omega$ for $s \in S$).

Theorem 2. If for $(\Omega, (\mathcal{P}_i, u_i)_{i \in \mathbb{I}})$ each i partitionally knows the move of each $j \neq i$, γ is payoff-irrelevant, and convexity and simple continuity are satisfied, then a partitionally rational equilibrium exists.

We omit the proof, which recapitulates the standard Nash argument in the current notation (the only extra step is to set a state where a standard Nash equilibrium occurs to have probability 1).

As we make clear in section 7, existence is also assured if we replace each i partitionally knowing the move of each $j \neq i$ in Theorem 2 with the “opposite” assumption, namely that as some i changes his move the other agents remain partitionally uninformed of this fact (see Definition 10 below). Assuming γ is payoff-irrelevant, we then return to the Aumann [2, 3] model of correlated equilibrium. Existence problems do not arise since a sequence of actions for i transmit no information to $j \neq i$ and hence cannot induce in the limit a discrete change in i ’s utility.

We turn to ε equilibria, where for any $\varepsilon > 0$ there exist strategies such that each agent achieves utility within ε of the supremum of the utility levels attainable given the other agent’s strategy. We do not settle the existence question; given strategies h_j for $j \neq i$, an agent i ’s best responses when facing the cell P_i need not form a convex set, and so standard existence proofs do not apply. But with two agents, and when agents’ uncertainty about the non-move variable γ is symmetric, then ε -equilibria do always exist.

Definition 6. *Non-move uncertainty is symmetric* if and only if there exists a partition \mathcal{P}_Γ of Γ

such that for $i \in \mathbb{I}$: $P_\Gamma \in \mathcal{P}_\Gamma \Leftrightarrow (S_i \times S_{-i}(P_i) \times P_\Gamma) \in \mathcal{P}_i$.

In words, each agent faces the same partition of Γ and each cell of this partition can arise whatever i 's information about the other agents' moves. While symmetric non-move uncertainty is restrictive, it permits a wide range of possibilities, including the Newcomb example in the introduction, all of the matching pennies examples, and the nonexistence example.

Definition 7. An ε partitionally rational equilibrium exists if and only if for all $\varepsilon > 0$ there exists (h, π) such that π is accurate with respect to h and, for all $i \in \mathbb{I}$, P_i , and $P_i^* \subset P_i$,

$$E(u_i | K(h_i(P_i))) + \varepsilon \geq E(u_i | K(P_i^*)).$$

Theorem 3. If for $(\Omega, (\mathcal{P}_i, u_i)_{i=a,b})$ non-move uncertainty is symmetric, convexity and γ -simple continuity are satisfied, and $S_{-i}(P_i)$ is convex for each i and P_i , then an ε partitionally rational equilibrium exists.

The idea of the proof is that, conditional on γ , standard arguments show that, for any pair (P_a, P_b) , the closure of $P_a \cap P_b$ has a ‘‘constrained Nash equilibrium’’ in which each i is required to choose a s_i in the closure of $S_i(P_{-i})$, and so there is a ε constrained Nash equilibrium in $P_a \cap P_b$ itself. We then use these ε constrained equilibria to determine which states have conditional probability 1, ensuring that when i faces P_i and contemplates choosing a P_i^* that intersects, say, $\hat{P}_{-i} \in \mathcal{P}_{-i}$, he or she anticipates that the ε constrained equilibrium of $P_i \cap \hat{P}_{-i}$ will obtain. Each i achieves ε rationality by selecting a P_i^* whose corresponding ε constrained equilibrium is near the upper bound of the set of utilities that can be reached by the ε constrained equilibria. With a careful adjustment of the actions in the $P_a \cap P_b$ that contains the standard (unconstrained) Nash equilibrium and of the pertinent conditional probabilities, accuracy is assured as well.

5. The rational play of dominated actions

Definition 8. The strategy h_i plays a *dominated action* if and only if there exist P_i and $P_i^{*'} \subset P_i$ such that, for all $s_{-i} \in S_{-i}$, if $(s_i, s_{-i}, \gamma) \in h_i(P_i)$ and $(s_i', s_{-i}, \gamma') \in P_i^{*'}$ then

$$u_i((s_i, s_{-i}, \gamma)) < u_i((s_i', s_{-i}, \gamma')).^4$$

An agent i can play a dominated action in partitionally rational equilibrium either because switching to a dominated action can lead P_i^* to intersect different subsets of the range of the other agents' strategies or because switching actions can change the conditional probabilities of the other agents' moves or some non-move feature of the world. We illustrate both possibilities with various Newcomb examples. Until warning to the contrary, the payoffs are those given in the introduction.

(i) To save on notation, let an action denote the states where the action is taken, e.g., when d faces P_d , d_1 will indicate the states in P_d that have a s_d coordinate equal to d_1 .

The simplest case where it is rational for c to play the dominated action c_1 , sketched in the introduction, occurs where d *partitionally knows* c 's move, that is, where every cell of d 's partition consists only of states whose s_c coordinates agree (see Definition 4 in section 4):

$$\Omega = \{(c_1, d_1), (c_1, d_2), (c_2, d_1), (c_2, d_2)\},$$

$$\mathcal{P}_d = \{\{(c_1, d_1), (c_1, d_2)\}\{(c_2, d_1), (c_2, d_2)\}\}, \mathcal{P}_c = \{\Omega\}.$$

With these \mathcal{P}_i , the only partitionally rational equilibrium occurs where $h_c(\Omega) = c_1$, $h_d(\{(c_1, d_1), (c_1, d_2)\}) = d_2$, $h_d(\{(c_2, d_1), (c_2, d_2)\}) = d_1$.

Neither agent i in this example learns anything directly from P_i^* that he/she did not know before moving (besides s_i itself). But indirectly, via $P_i^* \cap \text{Range } h_{-i}$, $i = c$ acquires information about d 's move from c 's own move, and this leads c to play a dominated action. The absence of direct informational content to the P_i^* will be retained through sections 5.i and 5.ii. The principal reason for this lack of informational content is that the agent partitions satisfy the NI axiom of section 3, which shows that NI does not eliminate the play of dominated actions. In fact, except for the next paragraph, every Newcomb example we consider satisfies NI.

Partitional knowledge does not require d to observe c 's move directly. Instead, d may

⁴ We use this “ex post” definition of domination, rather than a less demanding “interim” definition (where h_i only has to take actions whose expected utility condition on $h_i(P_i)$ is less than the expected utility of some $P_i^* \subset P_i$, conditional on P_i^*) since we wish dominated actions to be as plainly suboptimal as possible.

observe γ at an earlier date, and the different γ 's inform d of c 's move. For example, if we set

$$\mathcal{P}_d = \{\{(c_1, d_1, \alpha), (c_1, d_2, \alpha)\}, \{(c_2, d_1, \beta), (c_2, d_2, \beta)\}\}, \mathcal{P}_c = \{\Omega\},$$

then d 's partitional knowledge of c 's move can be ascribed to d 's observation of α or β . Think of α as infallible physical evidence that c will play c_1 and β as infallible evidence that c will play c_2 .

When d partitionally know c 's move, every cell of \mathcal{P}_d informs d of c 's action. But c can still always play c_1 even if d partitionally knows c 's move only at some of d 's premove cells.

Definition 9. Agent j occasionally partitionally knows i 's move if and only if, for some P_j ,

$$\omega, \omega' \in P_j \Rightarrow s_i(\omega) = s_i(\omega').$$

For an example where d only occasionally partitionally knows c 's move but where c nevertheless always chooses c_1 , suppose c does not know when d is informed about c 's action. Let α_{c_i} denote $\{(c_i, d_1, \alpha), (c_i, d_2, \alpha)\}$, $i = 1, 2$. Also, when convenient let a nonmove coordinate represent the four states that have that coordinate, e.g., $\beta = \{(c_1, d_1, \beta), (c_1, d_2, \beta), (c_2, d_1, \beta), (c_2, d_2, \beta)\}$. With this notation, set $\mathcal{P}_c = \{\Omega\}$, $\mathcal{P}_d = \{\alpha_{c_1}, \alpha_{c_2}, \beta\}$. Then the strategies

$$h_c(\Omega) = c_1, h_d(\alpha_{c_1}) = d_2, h_d(\alpha_{c_2}) = d_1, h_d(\beta) = d_2$$

are utility-maximizing if $\pi((c_1, d_2, \alpha))$ is sufficiently large. For accuracy to obtain, set $\pi((c_1, d_2, \alpha)) = 1 - \delta$ and $\pi((c_1, d_2, \beta)) = \delta$ for nonnegative δ sufficiently near 0.

Not surprisingly, if we let c distinguish between the β states and the α states, c will still play c_1 at the α states. Replace $\mathcal{P}_c = \{\Omega\}$ with $\{\alpha, \beta\}$. Then the strategies

$$h_c(\alpha) = c_1, h_c(\beta) = c_2, h_d(\alpha_{c_1}) = d_2, h_d(\alpha_{c_2}) = d_1, h_d(\beta) = d_1$$

form a partitionally rational equilibrium when $\pi((c_1, d_2, \alpha)) = 1 - \delta$ and $\pi((c_2, d_1, \beta)) = \delta$, $\delta \in [0, 1]$.

(ii) It may seem that if an agent a is to play a dominated action, then b must at least in some cells of \mathcal{P}_b be certain of a 's move (i.e., there must be a P_b that contains only states that specify just one move for a). This is not the case. To build a Newcomb example where no agent partitionally knows the other's move even occasionally, we employ states with four values

for the non-move coordinate, α, β, η , and τ . Suppose when either $\gamma = \beta$ or $\gamma = \eta$ that $\min_{i=1,2} u_c(c_2, d_i, \gamma) > \max_{i=1,2} u_c(c_1, d_i, \gamma)$ thus ensuring that c always chooses c_2 at the β or η states. Similarly, suppose that at the α states payoffs are such that c always chooses c_1 . At the τ states, c has the Newcomb game payoffs as does d in all states. Set $\Omega = S_c \times S_d \times \{\alpha, \beta, \eta, \tau\}$ and as before let each of α, β, η , and τ denote the four states with the corresponding value of γ . For example, $\alpha = \{(c_1, d_1, \alpha), (c_1, d_2, \alpha), (c_2, d_1, \alpha), (c_2, d_2, \alpha)\}$. Finally, for $i = 1, 2$, let $\tau_i \subset \tau$ denote $\{(c_i, d_1, \tau), (c_i, d_2, \tau)\}$. Now set

$$\mathcal{P}_c = \{\alpha, \beta \cup \eta, \tau\},$$

$$\mathcal{P}_d = \{\alpha \cup \eta \cup \tau_2, \beta \cup \tau_1\},$$

$$h_c(\alpha) = c_1, h_c(\beta \cup \eta) = c_2, h_c(\tau) = c_1, h_d(\alpha \cup \eta \cup \tau_2) = d_1, h_d(\beta \cup \tau_1) = d_2,$$

$$\pi((c_1, d_2, \tau)) = 1 - 2\varepsilon - \delta, \pi((c_1, d_1, \alpha)) = \delta, \pi((c_2, d_2, \beta)) = \varepsilon, \pi((c_2, d_1, \eta)) = \varepsilon.$$

View d 's first cell as a noisy signal that c has moved c_2 when facing the τ cell and the second as a noisy signal that c has moved c_1 at the τ cell. We specify the needed conditional probabilities momentarily.

Agent c 's preferences ensure that $h_c(\alpha) = c_1$ and $h_c(\beta \cup \eta) = c_2$ are partitionally rational (i.e., best responses in the sense of Definition 2). As for $h_c(\tau)$, if c plays c_1 then (given the specified h_d) the state (c_1, d_2, τ) must obtain and c 's payoff is 10^6 , whereas if c plays c_2 then (c_2, d_1, τ) must obtain and c 's payoff is 1000. Hence $h_c(\tau) = c_1$ is partitionally rational.

Turning to h_d , as $\frac{\varepsilon}{\delta} \rightarrow \infty$, the payoff to d of playing d_1 when facing $\alpha \cup \eta \cup \tau_2$ converges to x . The payoff to d of playing d_2 when facing $\alpha \cup \eta \cup \tau_2$ depends on the probabilities conditional on 0-probability events. Let $\pi(s_i | s_{-i}, A)$, where $A \subset \Gamma$, serve as shorthand for

$$\pi(\{\omega: s_i(\omega) = s_i\} | \{\omega: s_{-i}(\omega) = s_{-i}, \gamma \in A\} \cap \text{Range } h_i).$$

Assuming $\delta + \varepsilon > 0$, we have

$$\pi(c_1 | d_1, \{\alpha, \eta, \tau\}) = \frac{\delta}{\delta + \varepsilon}, \pi(c_2 | d_1, \{\alpha, \eta, \tau\}) = \frac{\varepsilon}{\delta + \varepsilon}.$$

And we set

$$\pi(c_1 | d_2, \{\alpha, \eta, \tau\}) = \frac{\delta}{\delta + \varepsilon}, \pi(c_2 | d_2, \{\alpha, \eta, \tau\}) = \frac{\varepsilon}{\delta + \varepsilon}.$$

So, as $\frac{\varepsilon}{\delta} \rightarrow \infty$, $\pi(c_2 | d_2, \{\alpha, \eta, \tau\}) \rightarrow 1$ and the payoff to d of playing d_2 converges to y . Hence $h_d(\alpha \cup \eta \cup \tau_2) = d_1$ is partitionally rational for large $\frac{\varepsilon}{\delta}$. Finally, consider d 's payoffs when

facing $\beta \cup \tau_1$. If $\delta < 1$ and ε becomes small, then $\pi(c_1|d_2, \{\beta, \tau\})$ approaches 1; hence the payoff to playing d_2 converges to x . If we set $\pi(c_1|d_1, \{\beta, \tau\}) = \pi(c_1|d_2, \{\beta, \tau\})$, the payoff to playing d_1 approaches y .

Thus, if we set $\varepsilon > 0$ small and $\frac{\varepsilon}{\delta}$ large, the specified strategies and probabilities are partitionally rational and the specified probabilities are accurate. We conclude that c rationally plays the dominated action c_1 when facing the cell τ .

Agent d 's lack of information about c 's move in this example goes beyond not having occasional partitional knowledge of c 's move: both c_1 and c_2 are played with positive probability in each of d 's premove cells. But despite d never being sure of c 's move, c still rationally chooses c_1 at some cells. We could tweak the example so that in addition both d_1 and d_2 are played with positive probability in each of c 's premove cells.⁵

The source of trouble in the above cases is that there are P_c and P_d such that $P_c \cap P_d \neq \emptyset$ but where, for some $P_c^* \subset P_c$, $P_c^* \cap P_d = \emptyset$. Thus d sometimes gets a (possibly noisy) signal of c 's move, namely that c has not moved $s_c(P_c^*)$, not just a signal about c 's type. The following requirement bars an agent j from receiving such a signal, not just at each of j 's premove cells but given any move j might take.

Definition 10. Agent j is *partitionally ignorant* of i 's move if and only if, for all $P_j^*, P_i, P_i^* \subset P_i, P_i^{*'} \subset P_i$,

$$P_i^* \cap P_j^* \neq \emptyset \Rightarrow P_i^{*'} \cap P_j^* \neq \emptyset.$$

In words, j is partitionally ignorant of i 's move if whenever j believes it possible that i could take some action s_i when facing P_i then j also believes it possible that i could take any other action s_i' when facing P_i . The set of *conceivable* P_i^* therefore do not change as P_j^* varies with a single P_j . So, for instance, if j is partitionally ignorant of i 's move then j does not occasionally know i 's move. But more strongly, partitional ignorance implies that j never has better

⁵ For instance, add a small-but-positive probability set of states σ at which d always moves d_2 , let μ denote another small-but-positive probability state at which c always moves c_1 , let the premove partitions be $\mathcal{P}_c = \{\alpha \cup \mu, \beta \cup \eta, \sigma \cup \tau\}$, $\mathcal{P}_d = \{\alpha \cup \eta \cup \tau_2, \sigma \cup \mu \cup \beta \cup \tau_1\}$, and let each player i 's move at his j th cell coincide with the move taken at the j th cell in the previous example (where cells are numbered in the order we have written them).

information about i 's move than i has ex ante. Partitional ignorance nevertheless allows j to have considerable knowledge of i , e.g., \mathcal{P}_i could refine \mathcal{P}_j .

Notice that even if each agent is partitionally ignorant of the other's move the NI axiom of section 3 need not be satisfied. For instance let

$$\begin{aligned}\mathcal{P}_c = \mathcal{P}_d &= \{(c_1, d_1, \alpha), (c_1, d_2, \alpha), (c_2, d_1, \beta), (c_2, d_2, \beta)\}, \\ \mathcal{P}_c^* &= \{(c_1, d_1, \alpha), (c_1, d_2, \alpha)\}, \{(c_2, d_1, \beta), (c_2, d_2, \beta)\}, \\ \mathcal{P}_d^* &= \{(c_1, d_1, \alpha), (c_2, d_1, \beta)\}, \{(c_1, d_2, \alpha), (c_2, d_2, \beta)\}.\end{aligned}$$

Evidently NI is violated but each agent is partitionally ignorant of the other's move. As in section 5.i, think of α (resp. β) as unerring physical evidence that c will move c_1 (resp. c_2) except that here d does not receive this evidence in advance. After moving c_1 , say, c will know that α must obtain. And this possibility indicates again why NI is so implausible: agents' own moves can and do inform them of things other than the moves themselves.

(iii) But partitional ignorance is not enough. With no restrictions on how agents form conditional probabilities, agents may still select dominated actions. To that end, we construct an example where c has the original Newcomb preferences throughout some cell and takes the dominated action c_1 in that cell, but where each agent is partitionally ignorant of the other's move. We use states with the four non-move coordinates $\alpha, \beta, \eta,$ and τ . Assume as in 5.ii that at the β and η states, each possible payoff for c when playing c_2 is strictly larger than each possible payoff when playing c_1 . In all other states, let c have the Newcomb game preferences as will d in all states. Using the notational conventions of 5.i and 5.ii, set

$$\begin{aligned}\mathcal{P}_c &= \{\beta \cup \eta, \alpha \cup \tau\}, \mathcal{P}_d = \{\alpha \cup \beta, \tau \cup \eta\}, \\ h_c(\beta \cup \eta) &= c_2, h_c(\alpha \cup \tau) = c_1, h_d(\alpha \cup \beta) = d_1, h_d(\tau \cup \eta) = d_2, \\ \pi((c_1, d_2, \tau)) &= \pi((c_2, d_1, \beta)) = .5.\end{aligned}$$

The partitional rationality of $h_c(\beta \cup \eta) = c_2$ is assured, given c 's payoffs in the β and η states. If d when facing the cell $\alpha \cup \beta$, plays d_1 , the state (c_2, d_1, β) obtains with probability 1 and d receives the payoff x , while if d plays d_2 , either (c_1, d_2, α) or (c_2, d_2, β) obtains and hence d receives a weighted average of x and y : hence $h_d(\alpha \cup \beta) = d_1$ is partitionally rational. When d faces $\tau \cup \eta$ and plays d_1 , either (c_1, d_1, τ) or (c_2, d_1, η) obtains, and so d 's payoff is a

weighted average of x and y , while if d plays d_2 , (c_1, d_2, τ) obtains with probability 1 and so d 's payoff is x . So $h_d(\tau \cup \eta) = d_2$ is also consistent with partitional rationality. It remains to check the most important case, c 's move when facing the cell $\{\alpha, \tau\}$. If c plays c_1 , (c_1, d_2, τ) obtains with probability 1 and so c 's payoff is 10^6 , while if c plays c_2 , either (c_2, d_1, α) or (c_2, d_2, τ) obtains and so c 's payoff is a weighted average of 1000 and $10^6 + 1000$. Both (c_2, d_1, α) and (c_2, d_2, τ) have probability 0. But if the conditional probability $\pi(d_1 | c_2, \alpha)$ is sufficiently large, then the move c_1 will be partitionally rational.

A twist in this case is that it is the updating induced by c 's moves that drives the paradox. Agent d is partitionally ignorant of c 's move and the updating on 0-probability events that occurs when d moves does not have to occur in a skewed or suspicious way.

6. Eliminating the play of dominated actions

One step to guaranteeing that a partitionally rational agent i will not play a dominated action is that the remaining agents are partitionally ignorant of i 's move. The second is to prohibit the suspicious updating of probabilities that occurred in section 5.iii. The NI axiom plays no role. When there are n agents, the first step requires an appropriate extension of Definition 10.

Definition 11. Agents $j \neq i$ are *mutually partitionally ignorant* of i 's move if and only if, for all $P_i, P_i^* \subset P_i, P_i^{*'} \subset P_i$, and P_{-i}^* ,

$$P_i^* \cap \left(\bigcap_{j \neq i} P_j^* \right) \neq \emptyset \Rightarrow P_i^{*'} \cap \left(\bigcap_{j \neq i} P_j^* \right) \neq \emptyset.$$

Definition 12. Agent i is *update independent* at (h, π) if and only if for all $P_i, P_i^* \subset P_i, P_i^{*'} \subset P_i$, and P_{-i}^* ,

$$\left[\left(\bigcap_{j \neq i} P_j^* \right) \cap K(P_i^*) \neq \emptyset \Leftrightarrow \left(\bigcap_{j \neq i} P_j^* \right) \cap K(P_i^{*'}) \neq \emptyset \right] \Rightarrow \pi \left(\bigcap_{j \neq i} P_j^* \mid K(P_i^*) \right) = \pi \left(\bigcap_{j \neq i} P_j^* \mid K(P_i^{*'}) \right).$$

Behind the notation, Definition 12 says something simple: if i contemplates two different moves, P_i^* and $P_i^{*'}$, at P_i and the set of possible profiles of moves for the other players does not change, then i regards each profile of moves for the other players to have the same posterior probability whether i moves P_i^* or $P_i^{*'}$. Since update independence applies only when the act

of moving has no informational content, it is akin to and as plausible as NEI.

If π is accurate, then for any P_i at most one of the family of conditioning events $\{K(P_i^*)\}_{P_i^* \subset P_i}$ can have positive probability and therefore conditional probabilities may always be set so as to satisfy update independence. Up until section 5.iii, every example we considered satisfied update independence – or, when some conditional probabilities went unspecified, they could be set consistently with update independence. And the equilibria in the proofs of the existence theorems in section 4 could be supplemented to satisfy update independence.

Theorem 4. If $j \neq i$ are mutually partitionally ignorant of i 's move and i is update independent at the partitionally rational equilibrium (h, π) , then h_i does not play a dominated action.

7. Aumann redux

We make our model compatible with Aumann [2, 3] by assuming that γ is *payoff-irrelevant* (Definition 5).

Definition 13. A partitionally rational equilibrium (h, π) leads to a correlated distribution of actions if and only if there exist a probability space (Ψ, \mathcal{F}, μ) and, for $i \in \mathbb{I}$, a partition \mathcal{Q}_i of Ψ and a strategy function $g_i: \Psi \rightarrow S_i$ measurable with respect to \mathcal{Q}_i such that (1) $Eu_i(g_i, g_{-i}) \geq Eu_i(f_i, g_{-i})$ for any $f_i: \Psi \rightarrow S_i$ measurable with respect to \mathcal{Q}_i , and (2) for all $A \subset S$, $\Omega_A = \{\omega \in \Omega: s(\omega) \in A\}$ is measurable $\Leftrightarrow \Psi_A = \mu(\{\psi \in \Psi: g(\psi) \in A\})$ is measurable, and $\pi(\Omega_A) = \mu(\Psi_A)$.

Theorem 5. If (h, π) is a partitionally rational equilibrium, γ is payoff-irrelevant, and, for each $i \in \mathbb{I}, j \neq i$ are mutually partitionally ignorant of i 's move and i is update independent, then (h, π) leads to a correlated distribution of actions.

If we did not assume a partitionally rational equilibrium, then the remaining assumptions in Theorem 5 would imply that a partitionally rational equilibrium exists as long as a standard Nash equilibrium exists for the game given by $(u_i)_{i \in \mathbb{I}}$: simply let each i play his or her Nash action at every P_i and assign probability 1 to one of the Nash equilibria. In terms of a converse

to Theorem 5, it is not hard to adapt the proof of Theorem 5 to show that if (1) $(\Psi, \mathcal{F}, \mu, (\mathcal{Q}_i, g_i)_{i \in \mathbb{I}})$ is a correlated equilibrium, (2) Ψ is isomorphic to the join (coursest common refinement) of $\mathcal{P}_i, i \in \mathbb{I}$, with bijection φ , and (3) $Q_i \in \mathcal{Q}_i \Leftrightarrow$ there exists $P_i \in \mathcal{P}_i$ with $\bigcup_{\psi \in Q_i} \varphi(\psi) = P_i$, then a partitionally rational equilibrium exists that leads to the same distribution of actions as $(\Psi, \mathcal{F}, \mu, (\mathcal{Q}_i, g_i)_{i \in \mathbb{I}})$.

9. Discussion

(i) The possibility that one agent a might rationally take a dominated action hinges on the nature of some other agent b 's knowledge of a 's actions. There is a world of difference between a b who can make accurate or even flawless predictions of a 's actions and a b who would know if a were to make a move that a will not in fact make. As we have seen, the second partitional type of knowledge can be of various strengths. An agent b might partitionally know another agent a 's moves, or might occasionally partitionally know a 's moves, or might simply not be partitionally ignorant of a 's moves. In all cases, it becomes possible for a rational agent a to take a dominated action.

We can characterize flawless prediction, say on b 's part, by assuming that for each P_b only those states in P_b with the same s_a coordinate have nonzero probability. So let us say that b knows a 's move with probability 1 if and only if, for all P_b there exists $A \subset P_b$ and $s_a \in S_a$ such that $S_a(A) = \{s_a\}$ and $\pi(A|P_b) = 1$. We interpret π as indicating objective rather than subjective probability.

Evidently, b can know a 's move with probability 1 even when partitionally ignorant of a 's move. For example, if \mathcal{P}_a and \mathcal{P}_b each consists of a single cell (neither agent receives any partitional information) and agents play in partitionally rational equilibrium, then by accuracy $\pi(h_a(P_a)) = \pi(h_b(P_b)) = 1$. Hence each agent knows the other agent's move with probability 1. Moreover it is only a failure of partitional ignorance that can allow a to make the counterfactual inferences that can justify the play of a dominated action, where, e.g., a can reason "if I were to play the 0-probability action s_a then I can infer that b must be selecting the 0-probability action s_b ." Knowledge with probability 1 in contrast does not entail this type of reasoning. Theorem 4

accordingly reports that if agents are partitionally ignorant of others' moves and update independence holds, dominated actions cannot be rational.

Other players' knowledge with probability 1 of an agent a 's move is not sufficient for a to play a dominated action, no matter what payoffs agents have. It is not necessary either. In section 5.ii, we saw a Newcomb example where d does not know c 's action with probability 1 (at each of d 's cells, both of c 's actions have positive probability). But due to the failure of partitional ignorance, it is rational for c to play c_1 .

A violation of partitional ignorance means that some agent b at some cell partitionally knows more about another agent a 's move than a does prior to moving. One way to eliminate the rational play of dominated actions is therefore to impose the following rule (on top of update independence): if a state space Ω specifies some agent a 's moves, then another agent b cannot not have a partition over Ω that violates partitional ignorance until the point in time at which a actually moves. Leaving aside stories like Newcomb that are designed to scrutinize the foundations of decision theory, it is hard to think of a decision problem where violating this rule would be justified.

(ii) The distinction between premove and postmove knowledge clarifies what it means for a profile of agent actions to be specified by a state. Different choices of actions for an agent i refine i 's premove partition in different ways, and so each i evidently has the leeway or "freedom" to choose whatever action he likes. Agents thus would not see their choices as somehow fixed in advance (Aumann [3] and Aumann and Brandenburger [4] assert the same conclusion, but as argued in section 2, their models belie their point). Notice that i 's leeway to choose continues to hold even when some other agent partitionally knows i 's action.

The premove-postmove distinction supplies a formalism by which agents can regard their moves as knowable events. With the right supplementary conditions, such as partitional ignorance, in place no paradoxes need result (see, e.g., Gilboa [11] for the view that there is a paradox). An agent can regard each of his possible moves in turn as a separate event that through updating induces distinct consequences.

(iii) Although the Newcomb paradox has served primarily to illustrate the difficulties that

arise when states specify actions, partitional knowledge also sheds light on the substantial literature on Newcomb. The literature sometimes argues that different approaches to decision theory back different conclusions about which actions are rational for c to take. Nozick's [15] original presentation of the Newcomb problem argued that the "principle of dominance" supports playing c_2 , while the "principle of expected utility maximization" supports playing c_1 . Gibbard and Harper [10], although c_2 partisans, came to a similar judgment: "causal" utility theory endorses c_2 as rational while a Jeffrey-like utility theory can support playing c_1 . By distinguishing between probabilistic and partitional knowledge, we can formalize within the confines of Bayesian decision theory both the view c_1 is rational and the view that c_2 is rational. As we have seen, the rationality of c_1 vs. c_2 then hinges on specific assumptions on information and partitions in particular. Playing c_1 is not ruled out a priori, but depends either on d having a highly refined information partition – indeed, so refined as to be implausible – or on c using a skewed rule to update probabilities.

The Newcomb literature introduces a red herring when it suggests that the rationality of c_1 rests on d being a flawless predictor of c 's move. As we have seen, d knowing c 's move with probability 1 does not imply that c_1 is a rational move. And even when d never assigns probability 1 to any of c 's actions it can be rational for c to play c_1 . So the rationality of c_1 and d 's complete accuracy as a predictor are separate issues. What matters for the rationality of c_1 is the *type* of knowledge that d has about c 's move.

(iv) The possibility of Newcomb-style paradoxes have led some to argue for a new species of decision analysis – causal decision theory – that classifies how an agent i 's action s_i causes changes to the world; this classification in turn presupposes a metric that judges which "possible world" is nearest to one in which s_i obtains (see Joyce and Gibbard [13] and the references therein and Shin [17]). The rule suggested in (i) above – that an agent remains partitionally ignorant of any other agent's move until the latter moves – also draws upon temporal and causal information. But that information is used only to determine partitions and the timing of their refinement; no broader overhaul of expected utility theory or metric on states is necessary.

Appendix: Proofs of Theorems 1, 3, 4 and 5.

Proof of Theorem 1. For each P_b^* , select an arbitrary $P_a \in \mathcal{P}_a$ such that $P_a \subset P_b^*$ and label it $P_a(P_b^*)$. Define, for each $P_b \in \mathcal{P}_b$, $Q(P_b) = \bigcup_{P_b^* \subset P_b} P_a(P_b^*)$. By partitional continuity, for any P_b there is a function $\gamma_{Q(P_b)}: S \rightarrow \Gamma$ such that (i) $(s, \gamma_{Q(P_b)}(s)) \in Q(P_b)$ for all $s \in S$, and (ii) for $i = a, b$, $s \mapsto u_i(s, \gamma_{Q(P_b)}(s))$ is continuous. Hence, for any $s_b \in S_b$, $\operatorname{argmax}_{s_a \in S_a} u_a(s_a, s_b, \gamma_{Q(P_b)}(s_a, s_b)) \neq \emptyset$ and, when seen as a correspondence of s_b , has a closed and hence compact graph. The problem

$$\begin{aligned} & \max_{s_a, s_b} u_b(s_a, s_b, \gamma_{Q(P_b)}(s_a, s_b)) \\ \text{s.t.} \quad & s_a \in \operatorname{argmax}_{s_a \in S_a} u_a(s_a, s_b, \gamma_{Q(P_b)}(s_a, s_b)), s_b \in S_b. \end{aligned} \quad (1.1)$$

therefore has a solution, which we label $(s_a[P_b], s_b[P_b])$.

For each P_b , set $h_b(P_b) \subset P_b$ so that

$$s_b(h_b(P_b)) = s_b[P_b]. \quad (1.2)$$

When $P_a \in \{P_a(P_b^*): P_b^* \in \mathcal{P}_b^*\}$, set $h_a(P_a) \subset P_a$ so that

$$s_a(h_a(P_a)) \in \operatorname{argmax}_{s_a \in S_a} u_a((s_a, s_b(P_a), \gamma_{Q(P_b)}(s_a, s_b(P_a)))), \quad (1.3)$$

where $Q(P_b) \supset P_a$, but if in addition $P_a \subset h_b(P_b)$ for some P_b , use the particular element of the argmax , $s_a[P_b]$, i.e., set $h_a(P_a)$ so that

$$s_a(h_a(P_a)) = s_a[P_b]. \quad (1.4)$$

When $P_a \notin \{P_a(P_b^*): P_b^* \in \mathcal{P}_b^*\}$, let $\gamma_{\{P_a\}}$ be the function defined on $S_a \times S_b(P_a)$ given by partitional continuity and set $h_a(P_a) \subset P_a$ so that

$$s_a(h_a(P_a)) \in \operatorname{argmax}_{s_a \in S_a} u_a((s_a, s_b(P_a), \gamma_{\{P_a\}}(s_a, s_b(P_a)))). \quad (1.5)$$

As for the probabilities, set $\pi((s_a[P_b], s_b[P_b], \gamma_{Q(P_b)}(s_a[P_b], s_b[P_b]))) = 1$ for some P_b and $\pi(\omega) = 0$ for all other $\omega \in \Omega$. When $P_a^* \subset P_a \in \{P_a(P_b^*): P_b^* \in \mathcal{P}_b^*\}$, set

$$\pi((s_a(P_a^*), s_b(P_a^*), \gamma_{Q(P_b)}(s_a(P_a^*), s_b(P_a^*))) | P_a^*) = 1, \quad (1.6)$$

where $P_b \supset P_a$. When $P_a^* \subset P_a \notin \{P_a(P_b^*): P_b^* \in \mathcal{P}_b^*\}$, set

$$\pi((s_a(P_a^*), s_b(P_a^*), \gamma_{\{P_a\}}(s_a(P_a^*), s_b(P_a^*))) | P_a^*) = 1. \quad (1.7)$$

For all P_b^* , set

$$\pi((s_a(h_a(P_a(P_b^*))), s_b(P_b^*), \gamma_{Q(P_b)}(s_a(h_a(P_a(P_b^*))), s_b(P_b^*))) | P_b^* \cap \operatorname{Range} h_a) = 1, \quad (1.8)$$

where $P_b \supset P_b^*$. All remaining conditional probabilities may be set arbitrarily.

To see that h_a is partitionally rational, note first that, since \mathcal{P}_a refines \mathcal{P}_b^* , if $P_a \cap \text{Range } h_b \neq \emptyset$ then, for all $P_a^* \subset P_a$, $P_a^* \cap \text{Range } h_b = P_a^*$, while if $P_a \cap \text{Range } h_b = \emptyset$ then, for all $P_a^* \subset P_a$, $P_a^* \cap \text{Range } h_b = \emptyset$. Hence, for all P_a^* ,

$$K(P_a^*) = P_a^*. \quad (1.9)$$

So, for $P_a^* \subset P_a \in \{P_a(P_b^*): P_b^* \in \mathcal{P}_b^*\}$, (1.6) and (1.9) imply

$$E(u_a | K(P_a^*)) = u_a((s_a(P_a^*), s_b(P_a^*), \gamma_{Q(P_b^*)}(s_a(P_a^*), s_b(P_a^*))))$$

where $P_b \supset P_a$. Then (1.3) implies the partitional rationality of h_a . For $P_a^* \subset P_a \notin \{P_a(P_b^*): P_b^* \in \mathcal{P}_b^*\}$, (1.7) and (1.9) imply

$$E(u_a | K(P_a^*)) = u_a((s_a(P_a^*), s_b(P_a^*), \gamma_{\{P_a\}}(s_a(P_a^*), s_b(P_a^*))).$$

Then (1.5) implies the partitional rationality of h_a .

As for h_b , since \mathcal{P}_a refines \mathcal{P}_b^* , we have $K(P_b^*) = P_b^* \cap \text{Range } h_a$ for all P_b^* . Hence, for all P_b^* , (1.8) implies

$$E(u_b | K(P_b^*)) = u_b((s_a(h_a(P_a(P_b^*))), s_b(P_b^*), \gamma_{Q(P_b^*)}(s_a(h_a(P_a(P_b^*))), s_b(P_b^*))),$$

where $P_b \supset P_b^*$, and so (1.1), (1.2), (1.3), and (1.4) imply the partitional rationality of h_a . ■

Proof of Theorem 3. Fix arbitrary $P_\Gamma \in \mathcal{P}_\Gamma$ and $\gamma \in P_\Gamma$. Until the end of the proof, we suppress γ from most of the notation. For example, (s_a, s_b) will refer to (s_a, s_b, γ) . Also, fix $\varepsilon > 0$.

For each $P_i \in \mathcal{P}_i$, $i = a, b$, let \bar{P}_i denote the closure of P_i . Let $f_{\bar{P}_a, \bar{P}_b}: \bar{P}_a \cap \bar{P}_b \Rightarrow \bar{P}_a \cap \bar{P}_b$ denote the best-response correspondence defined by

$$f_{\bar{P}_a, \bar{P}_b}(\omega) = \{ \hat{\omega} \in \bar{P}_a \cap \bar{P}_b: (s_i', s_{-i}(\hat{\omega})) \in \bar{P}_{-i} \Rightarrow u_i(\hat{\omega}) \geq u_i((s_i', s_{-i}(\hat{\omega}))), i = a, b \}.$$

Let $k_{\bar{P}_a, \bar{P}_b}$ denote the fixed points of $f_{\bar{P}_a, \bar{P}_b}: \omega \in k_{\bar{P}_a, \bar{P}_b} \Leftrightarrow \omega \in f_{\bar{P}_a, \bar{P}_b}(\omega)$. Given convexity and γ -simple continuity, and the convexity of the partition cells, Kakutani's theorem implies that $k_{\bar{P}_a, \bar{P}_b}$ is nonempty. For any $(P_a, P_b) \in (\mathcal{P}_a, \mathcal{P}_b)$, γ -simple continuity implies for $i = a, b$ that $\max_{s_i \in S_i(\bar{P}_{-i})} u_i((s_i, s_{-i}))$ is a continuous function of s_{-i} . Hence there exists a $\omega_{P_a, P_b} \in P_a \cap P_b$ near enough to some point of $k_{\bar{P}_a, \bar{P}_b}$ to satisfy the inequality

$$\max_{s_i \in S_i(\bar{P}_{-i})} u_i((s_i, s_{-i}(\omega_{P_a, P_b}))) - u_i(\omega_{P_a, P_b}) < \frac{\varepsilon}{2} \quad (3.i)$$

for both $i = a$ and $i = b$. For at least one (P_a, P_b) , say (P_a^N, P_b^N) , $P_a^N \cap P_b^N$, contains a Nash equilibrium (s_a^N, s_b^N) of the conventionally defined game where each i has the action set S_i and utility u_i . In conformity with (3.i), set $\omega_{P_a^N, P_b^N} = (s_a^N, s_b^N)$.

Let $P_{-i}(P_i^*)$ denote the unique cell of \mathcal{P}_{-i} such that $P_{-i}(P_i^*) \cap P_i^* \neq \emptyset$, and let $g_i(P_i)$ denote some $\hat{P}_i^* \subset P_i$ such that

$$\sup_{P_i^* \subset P_i} \{u_i(\omega_{P_i, P_{-i}(P_i^*)})\} - u_i(\omega_{P_i, P_{-i}(\hat{P}_i^*)}) < \frac{\varepsilon}{2} \quad \text{and} \quad s_i(\hat{P}_i^*) = s_i((\omega_{P_i, P_{-i}(\hat{P}_i^*)})). \quad (3.ii)$$

If $\text{Range } g_a \cap \text{Range } g_b \neq \emptyset$ (call this case I), then for $i = a, b$ set $h_i(P_i) = g_i(P_i)$ for all P_i . If $\text{Range } g_a \cap \text{Range } g_b = \emptyset$ and for some i there exists some $\tilde{P}_i^* \subset P_i^N$ such that $\tilde{P}_i^* \cap \text{Range } g_{-i} \neq \emptyset$ and $u_i(\omega_{P_i^N, P_{-i}(\tilde{P}_i^*)}) > u_i(\omega_{P_a^N, P_b^N})$ (case II), then label one such i as \hat{i} and set $h_{-\hat{i}}(P_{-\hat{i}}) = g_{-\hat{i}}(P_{-\hat{i}})$ for all $P_{-\hat{i}}$, $h_{\hat{i}}(P_{\hat{i}}) = g_{\hat{i}}(P_{\hat{i}})$ for all $P_{\hat{i}} \neq P_{\hat{i}}^N$ and let $h_{\hat{i}}(P_{\hat{i}}^N)$ equal some $\hat{P}_{\hat{i}}^* \subset P_{\hat{i}}^N$ such that

$$\begin{aligned} \hat{P}_{\hat{i}}^* \cap \text{Range } g_{-\hat{i}} &\neq \emptyset, \\ \sup_{P_{\hat{i}}^* \subset P_{\hat{i}}^N: P_{\hat{i}}^* \cap \text{Range } g_{-\hat{i}} \neq \emptyset} \{u_{\hat{i}}(\omega_{P_{\hat{i}}^N, P_{-\hat{i}}(P_{\hat{i}}^*)})\} - u_{\hat{i}}(\omega_{P_{\hat{i}}^N, P_{-\hat{i}}(\hat{P}_{\hat{i}}^*)}) &< \frac{\varepsilon}{2}, \\ s_{\hat{i}}(\hat{P}_{\hat{i}}^*) &= s_{\hat{i}}((\omega_{P_{\hat{i}}^N, P_{-\hat{i}}(\hat{P}_{\hat{i}}^*)})). \end{aligned} \quad (3.iii)$$

Finally, in the remaining possibility, case III – where $\text{Range } g_a \cap \text{Range } g_b = \emptyset$ and there is no $\tilde{P}_i^* \subset P_i^N$ such that $\tilde{P}_i^* \cap \text{Range } g_{-i} \neq \emptyset$ and $u_i(\omega_{P_i^N, P_{-i}(\tilde{P}_i^*)}) > u_i(\omega_{P_a^N, P_b^N})$ for either $i = a$ or $i = b$ – then, for $i = a, b$, set $h_i(P_i) = g_i(P_i)$ for all $P_i \neq P_i^N$ and set $h_i(P_i^N)$ equal to the $P_i^* \subset P_i^N$ such that $s_i(P_i^*) = s_i^N$. Notice that in all cases $\text{Range } h_a \cap \text{Range } h_b \neq \emptyset$.

Let us say that P_i^* for $i = a$ or $i = b$ *requires Nash adjustment* if case II obtains, $i = \hat{i}$, $P_i^* \subset P_i^N$, and $P_i^* \cap \text{Range } h_{-i} = \emptyset$; or if case III obtains, $P_i^* \subset P_i^N$, and $P_i^* \cap \text{Range } h_{-i} = \emptyset$. Otherwise we say P_i^* *does not require Nash adjustment*. When P_i^* for $i = a$ or $i = b$ requires Nash adjustment, set

$$\pi((s_i(P_i^*), s_{-i}^N) | K(P_i^*)) = 1.$$

When P_i^* for $i = a$ or $i = b$ does not require Nash adjustment, set

$$\pi((s_i(P_i^*), s_{-i}(\omega_{P_i, P_{-i}(P_i^*)})) | K(P_i^*)) = 1,$$

where $P_i \supset P_i^*$. The remaining conditional probabilities can be set arbitrarily. As for the unconditional probabilities, set $\pi(\omega') = 1$ for some $\omega' \in \text{Range } h_a \cap \text{Range } h_b$ and $\pi(\omega) = 0$ for $\omega \neq \omega'$. Thus accuracy obtains.

To show that the h_i satisfy the ε rationality condition in Definition 7, conditional on ω being in the fixed $P_\Gamma \in \mathcal{P}_\Gamma$, we first calculate the expected utility of arbitrary P_i^* actions. Suppose for the remainder of this paragraph that P_i^* does not require Nash adjustment. Our

specification of the conditional probabilities then implies

$$E(u_i | K(P_i^*)) = u_i((s_i(P_i^*), s_{-i}(\omega_{P_i, P_{-i}(P_i^*)}))),$$

where $P_i \supset P_i^*$. Since furthermore

$$u_i((s_i(P_i^*), s_{-i}(\omega_{P_i, P_{-i}(P_i^*)}))) \leq \max_{s_i \in S_i(\text{cl } P_{-i}(P_i^*))} u_i((s_i, s_{-i}(\omega_{P_i, P_{-i}(P_i^*)}))),$$

(3.i) implies

$$E(u_i | K(P_i^*)) \leq u_i(\omega_{P_i, P_{-i}(P_i^*)}) + \frac{\varepsilon}{2}.$$

Suppose that it is *not* the case that simultaneously (1) case II obtains, (2) $i = \hat{i}$, (3) $P_i^* \subset P_{\hat{i}}^N$, and (4) $P_i^* \cap \text{Range } h_{-i} \neq \emptyset$. Then the fact that $u_i(\omega_{P_i, P_{-i}(P_i^*)}) \leq \sup_{\tilde{P}_i^* \subset P_i} \{u_i(\omega_{P_i, P_{-i}(\tilde{P}_i^*)})\}$ and (3.ii) imply

$$u_i(\omega_{P_i, P_{-i}(P_i^*)}) < u_i(\omega_{P_i, P_{-i}(h_i(P_i))}) + \frac{\varepsilon}{2}. \quad (3.iv)$$

(In case III and $P_i = P_i^N$, (3.iv) follows from the fact that (s_a^N, s_b^N) is a Nash equilibrium of the conventionally defined game with action set S_i and utility u_i .) When (1) through (4) do obtain, the fact that

$$u_i(\omega_{P_i, P_{-i}(P_i^*)}) \leq \sup_{\tilde{P}_i^* \subset P_i^N: \tilde{P}_i^* \cap \text{Range } g_{-i} \neq \emptyset} \{u_i(\omega_{P_i^N, P_{-i}(\tilde{P}_i^*)})\}$$

and (3.iii) imply that (3.iv) still holds. Hence, whenever P_i^* does not require Nash adjustment,

$$E(u_i | K(P_i^*)) < u_i(\omega_{P_i, P_{-i}(h_i(P_i))}) + \varepsilon. \quad (3.v)$$

When P_i^* requires Nash adjustment, then, given our specification of the conditional probabilities,

$$E(u_i | K(P_i^*)) = u_i((s_i(P_i^*), s_{-i}^N)),$$

and, since the Nash action s_i^N is utility maximizing against s_{-i}^N ,

$$E(u_i | K(P_i^*)) \leq u_i((s_i^N, s_{-i}^N)) = u_i(\omega_{P_a^N, P_b^N}).$$

Hence,

$$E(u_i | K(P_i^*)) < u_i(\omega_{P_a^N, P_b^N}) + \varepsilon. \quad (3.vi)$$

To conclude the demonstration of ε rationality, we calculate the expected utility of the $h_i(P_i)$ actions and compare them to the expected utilities of the arbitrary P_i^* actions. When case I obtains, or when case II obtains but $P_i \neq P_i^N$, or when case III obtains and $P_i \neq P_i^N$ for either $i = a$ or $i = b$, then

$$E(u_i | K(h_i(P_i))) = u_i(\omega_{P_i, P_{-i}(h_i(P_i))}) \quad (3.vii)$$

and so (3.v) establishes the ε rationality of h_i . When case II obtains and $P_i = P_i^N$, we again have (3.vii), and if additionally $P_i^* \subset P_i^N$ does not require Nash adjustment, then (3.v) again establishes the ε rationality of h_i . When $P_i^* \subset P_i^N$ does require Nash adjustment, (3.vi) and the fact that $u_i(\omega_{P_i^N, P_{-i}(h_i(P_i^N))}) > u_i(\omega_{P_a^N, P_b^N})$ (see the definition of case II) establish the ε rationality of h_i . When finally case III obtains, then, for either $i = a$ or $i = b$,

$$E(u_i | K(h_i(P_i^N))) = u_i(\omega_{P_a^N, P_b^N}). \quad (3.viii)$$

Hence, if $P_i^* \subset P_i^N$ requires Nash adjustment, (3.vi) and (3.viii) establish the ε rationality of h_i . If $P_i^* \subset P_i^N$ does not require Nash adjustment, then $P_i^* \cap \text{Range } h_{-i} \neq \emptyset$ and hence $u_i(\omega_{P_a^N, P_b^N}) \geq u_i(\omega_{P_i^N, P_{-i}(P_i^*)})$ (see the definition of case III). So $u_i(\omega_{P_i^N, P_{-i}(P_i^*)}) \geq u_i((s_i(P_i^*), s_{-i}(\omega_{P_i, P_{-i}(P_i^*)}))) = E(u_i | K(P_i^*))$ implies the ε rationality of h_i .

So far we have fixed P_Γ and $\gamma \in P_\Gamma$. For any $P_{\Gamma'} \neq P_\Gamma$, again select an arbitrary $\gamma' \in P_{\Gamma'}$, define the g_i as before, and set $h_i(P_i) = g_i(P_i)$ for each $P_i, i = a, b$, implying that the h_i are ε rational on the entirety of their domain. ■

Proof of Theorem 4. For any P_i^* , let $\mathcal{P}_{P_i^*}$ be the partition of $K(P_i^*)$ with typical element $P_{P_i^*}(s_{-i})$ defined, for all s_{-i} , by

$$\omega \in P_{P_i^*}(s_{-i}) \Leftrightarrow (s_{-i} = s_{-i}(\omega) \text{ and } \omega \in K(P_i^*)).$$

Then $E(u_i | K(P_i^*)) = \int_{S_{-i}} E(u_i | P_{P_i^*}(s_{-i})) d\pi(s_{-i} | K(P_i^*))$.

Suppose $P_i^* \cap \text{Range}_{-i} h \neq \emptyset$. Given some s_{-i} , consider an arbitrary P_{-i}^* such that $s_{-i} = s_{-i}(\bigcap_{j \neq i} P_j^*)$ and $(\bigcap_{j \neq i} P_j^*) \cap (P_i^* \cap \text{Range}_{-i} h) \neq \emptyset$. Then $P_j^* \subset \text{Range } h_j, j \neq i$, and of course $(\bigcap_{j \neq i} P_j^*) \cap P_i^* \neq \emptyset$. For $P_i^{*'} \subset P_i^*$, where $P_i \supset P_i^*$, the mutual partitional ignorance of i 's move implies $P_i^{*'} \cap (\bigcap_{j \neq i} P_j^*) \neq \emptyset$ and so $(\bigcap_{j \neq i} P_j^*) \cap (P_i^{*' } \cap \text{Range}_{-i} h) \neq \emptyset$. Hence i 's update independence implies $\pi(\bigcap_{j \neq i} P_j^* | P_i^{*' } \cap \text{Range}_{-i} h) = \pi(\bigcap_{j \neq i} P_j^* | P_i^* \cap \text{Range}_{-i} h)$ and therefore

$$\pi(s_{-i} | K(P_i^{*' })) = \pi(s_{-i} | K(P_i^*)). \quad (4.1)$$

Now suppose $P_i^* \cap \text{Range}_{-i} h = \emptyset$. Given some s_{-i} , let P_{-i}^* satisfy $s_{-i}(\bigcap_{j \neq i} P_j^*) = s_{-i}$ and $P_i^* \cap (\bigcap_{j \neq i} P_j^*) \neq \emptyset$. For $P_i^{*' } \subset P_i^*$, the mutual partitional ignorance of i 's move implies $P_i^{*' } \cap (\bigcap_{j \neq i} P_j^*) \neq \emptyset$. Since $P_i^* \cap \text{Range}_{-i} h = \emptyset$, mutual partitional ignorance also implies $P_i^{*' } \cap \text{Range}_{-i} h = \emptyset$. Hence $K(P_i^*) = P_i^*$ and $K(P_i^{*' }) = P_i^{*' }$. Update independence therefore again implies (4.1).

Thus for any $P_i^{*'} \subset P_i$

$$E(u_i|K(P_i^{*'})) = \int_{S_{-i}} E(u_i|P_{P_i^{*'}(s_{-i})})d\pi(s_{-i}|K(P_i^{*'})). \quad (4.2)$$

If h_i plays a dominated action, there exist $P_i \in \mathcal{P}_i$ and $P_i^{*'} \subset P_i$ such that, for each s_{-i} , $u_i((s_i', s_{-i}, \gamma')) > u_i((s_i, s_{-i}, \gamma))$ whenever $(s_i', s_{-i}, \gamma') \in P_i^{*'}$ and $(s_i, s_{-i}, \gamma) \in h_i(P_i)$. Hence, for each s_{-i} , $E(u_i|P_{P_i^{*'}(s_{-i})}) > E(u_i|P_{h_i(P_i)(s_{-i})})$ and so (4.2) implies $E(u_i|K(P_i^{*'})) > E(u_i|K(h_i(P_i)))$, violating partitional rationality. ■

Proof of Theorem 5. Let Ψ be isomorphic to the join of \mathcal{P}_i , $i \in \mathbb{I}$, denoted $\bigvee_{i \in \mathbb{I}} \mathcal{P}_i$, with accompanying bijection $\varphi: \Psi \rightarrow \bigvee_{i \in \mathbb{I}} \mathcal{P}_i$. Define \mathcal{F} by $A \in \mathcal{F}$ if and only if $\varphi(A)$ is a measurable subset of Ω , and \mathcal{Q}_i by $Q_i \in \mathcal{Q}_i$ if and only if there exists $P_i \in \mathcal{P}_i$ such that $\bigcup_{\psi \in Q_i} \varphi(\psi) = P_i$. Let $Q_i(\psi)$ denote the $Q_i \in \mathcal{Q}_i$ such that $\psi \in Q_i$. Set $g_i(\psi) = s_i(h_i(\bigcup_{\psi' \in Q_i(\psi)} \varphi(\psi')))$ for all $\psi \in \Psi$, and $\mu(A) = \pi(\varphi(A))$ for all $A \in \mathcal{F}$. Then (2) of Definition 13 obtains.

To show that (1) of Definition 13 obtains, we first observe that, for all P_i^* , $K(P_i^*) = P_i^* \cap \text{Range}_{-i}h$. Without loss of generality set $i = 1$. Given some P_1^* , let P_{-1}^* be such that $P_1^* \cap (\bigcap_{j \neq 1} P_j^*) \neq \emptyset$. Then since the agents $j \neq 2$ are partitionally ignorant of 2's move, $P_1^* \cap h_2(P_2) \cap (\bigcap_{j \neq 1, 2} P_j^*) \neq \emptyset$ where $P_2 \supset P_2^*$. Proceeding by induction, we conclude that $P_1^* \cap (\bigcap_{j \neq 1} h(P_j)) \neq \emptyset$. Hence $P_1^* \cap \text{Range}_{-1}h \neq \emptyset$.

Therefore $E(u_i|K(P_i^*))$

$$\begin{aligned} &= \int_{\{h_{-i}(P_{-i}) \cap P_i^*: P_{-i} \in \mathcal{P}_{-i}\}} E(u_i|P_i^* \cap h_{-i}(P_{-i}))d\pi(h_{-i}(P_{-i})|P_i^* \cap \text{Range}_{-i}h), \\ &= \int_{\{h_{-i}(P_{-i}) \cap P_i^*: P_{-i} \in \mathcal{P}_{-i}\}} u_i(s_i(P_i^*), s_{-i}(h_{-i}(P_{-i})))d\pi(h_{-i}(P_{-i})|P_i^* \cap \text{Range}_{-i}h). \end{aligned}$$

As in the proof of Theorem 4, partitional ignorance and update independence imply

$$\pi(\bigcap_{j \neq i} P_j^* | P_i^{*'}) \cap \text{Range}_{-i}h = \pi(\bigcap_{j \neq i} P_j^* | P_i^* \cap \text{Range}_{-i}h)$$

for all $P_i^{*' \subset P_i$ and in particular for $P_i^{*' = h_i(P_i)$. So

$$\begin{aligned} &\int_{\{h_{-i}(P_{-i}) \cap P_i^*: P_{-i} \in \mathcal{P}_{-i}\}} u_i(s_i(P_i^*), s_{-i}(h_{-i}(P_{-i})))d\pi(h_{-i}(P_{-i})|P_i^* \cap \text{Range}_{-i}h) = \\ &\int_{\{h_{-i}(P_{-i}) \cap h_i(P_i): P_{-i} \in \mathcal{P}_{-i}\}} u_i(s_i(P_i^*), s_{-i}(h_{-i}(P_{-i})))d\pi(h_{-i}(P_{-i})|h_i(P_i) \cap \text{Range}_{-i}h). \end{aligned}$$

Therefore, for any P_i and $P_i^* \subset P_i$,

$$\begin{aligned} &\int_{\{h_{-i}(P_{-i}) \cap h_i(P_i): P_{-i} \in \mathcal{P}_{-i}\}} u_i(s_i(h_i(P_i)), s_{-i}(h_{-i}(P_{-i})))d\pi(h_{-i}(P_{-i})|h_i(P_i) \cap \text{Range}_{-i}h) \geq \\ &\int_{\{h_{-i}(P_{-i}) \cap h_i(P_i): P_{-i} \in \mathcal{P}_{-i}\}} u_i(s_i(P_i^*), s_{-i}(h_{-i}(P_{-i})))d\pi(h_{-i}(P_{-i})|h_i(P_i) \cap \text{Range}_{-i}h). \end{aligned}$$

Taking expectations over $\{h_i(P_i): P_i \in \mathcal{P}_i\}$ yields

$$\int_{\{h_i(P_i) \cap h_{-i}(P_{-i}): P_i \in \mathcal{P}_i, P_{-i} \in \mathcal{P}_{-i}\}} u_i(s_i(h_i(P_i)), s_{-i}(h_{-i}(P_{-i}))) d\pi(h_i(P_i) \cap h_{-i}(P_{-i})) \geq \int_{\{h_i(P_i) \cap h_{-i}(P_{-i}): P_i \in \mathcal{P}_i, P_{-i} \in \mathcal{P}_{-i}\}} u_i(s_i(P_i^*), s_{-i}(h_{-i}(P_{-i}))) d\pi(h_i(P_i) \cap h_{-i}(P_{-i})).$$

Since accuracy implies that $\pi(h_i(P_i) \cap h_{-i}(P_{-i}): P_i \in \mathcal{P}_i, P_{-i} \in \mathcal{P}_{-i}) = 1$, we conclude that

$Eu_i(g_i, g_{-i}) \geq Eu_i(f_i, g_{-i})$ for any f_i measurable with respect to \mathcal{Q}_i . ■

References

- [1] W. Armbruster, W. Böge, Bayesian game theory, in: O. Moeschlin, D. Pallaschke (Eds.), *Game Theory and Related Topics*, North Holland, Amsterdam, 1978, pp. 17-28.
- [2] R. Aumann, Subjectivity and correlation in randomized strategies, *J. Math. Econ.* 1 (1974), 67-96.
- [3] R. Aumann, Correlated equilibrium as an expression of Bayesian rationality, *Econometrica* 55 (1987), 1-18.
- [4] R. Aumann, A. Brandenburger, Epistemic conditions for Nash equilibrium, *Econometrica* 63 (1995), 1161-80.
- [5] K. Binmore, Foundations of game theory, in: J.-J. Laffont (Ed.), *Advances in Economic Theory*, Sixth World Congress Vol. 1, Cambridge University Press, Cambridge, 1992, pp. 1-31.
- [6] K. Binmore, A. Brandenburger, Common knowledge and game theory, in: K. Binmore, *Essays on the Foundations of Game Theory*, Blackwell, Oxford, 1990, pp. 105-150.
- [7] W. Böge, T. Eisele, On solutions of Bayesian games, *Int. J. Game Theory* 8 (1979), 193-215.
- [8] A. Brandenburger, E. Dekel, Rationalizability and correlated equilibria, *Econometrica* 55 (1987), 1391-1402.
- [9] E. Dekel, F. Gul, Rationality and knowledge in game theory, in: D. Kreps, K. Wallis (Eds.), *Advances in Economics and Econometrics: Theory and Applications*, Vol. I, Cambridge University Press, Cambridge, 1997, pp. 87-172.
- [10] A. Gibbard, W. Harper, Counterfactuals and two kinds of expected utility, in: C. Hooker, J. Leach, and E. McClennen (Eds.), *Foundations and Applications of Decision Theory*, Reidel, Dordrecht, 1978, pp. 125-162.
- [11] I. Gilboa, Can free choice be known?, in: C. Bicchieri, R. Jeffrey, B. Skyrms (Eds.), *The Logic of Strategy*, Oxford University Press, Oxford, 1994, pp. 163-174.
- [12] R. Jeffrey, *The Logic of Decision*, Chicago University Press, Chicago, 1983.
- [13] J. Joyce, A. Gibbard, Causal decision theory, in: S. Barbera, P. Hammond, C. Seidl (Eds.), *Handbook of Utility Theory*, Kluwer Academic Press, New York, 1998. pp. 627-666.

- [14] J.-F. Mertens, S. Zamir, Formulation of Bayesian analysis for games with incomplete information, *Int. J. Game Theory* 14 (1985), 1-29.
- [15] R. Nozick, Newcomb's problem and two principles of choice, in N. Rescher (Ed.), *Essays in Honor of Carl G. Hempel*, Reidel, Dordrecht, 1969, pp. 114-146.
- [16] H. Shin, A comment on Aumann's definition of Bayes-rationality, mimeo, Nuffield College, Oxford, 1988.
- [17] H. Shin, A reconstruction of Jeffrey's notion of ratifiability in terms of counterfactual beliefs, *Theory and Decision* 31 (1991), 21-47.
- [18] T. Tan, S. Werlang, The Bayesian foundations of solution concepts of games, *J. Econ. Theory* 45 (1988), 370-391.