

CAUSAL ASSESSMENT IN FINITE-LENGTH EXTENSIVE-FORM GAMES*

Jose Penalva-Zuasti
Universitat Pompeu Fabra

Michael D. Ryall[†]
University of Rochester, and
University of California, Los Angeles

December 30, 2002

*We thank Adam Brandenburger, Colin Camerer, David M. Chickering, Eddie Dekel, Bryan Ellickson, Ehud Kalai, David Levine, Steven Lippman, Glenn MacDonald, Leslie Marx, Joseph Ostroy, Scott Page, Judea Pearl, Larry Samuelson, Eduardo Zambrano, Bill Zame, participants in the MEDS and Georgetown theory workshops, the Stanford SITE and Stonybrook theory conferences, two anonymous referees and an anonymous coeditor, for thoughtful comments. Ryall thanks the John M. Olin foundation for financial support.

[†]Correspondence should be sent to: M. D. Ryall, Economics Dept., 8283 Bunche Hall, Box 951477, Los Angeles, CA. 90095; 310-825-3925; ryall@simon.rochester.edu.

Abstract

We analyze the extent to which the information structure in an extensive form game can be inferred from the distribution on action profiles generated by player strategies. Two games are said to be *empirically compatible* when the distribution on action profiles generated by every behavior strategy in one can also be generated by an appropriately chosen behavior strategy in the other. Our central idea is to relate a game's information structure to the conditional independencies in the empirical distributions it generates. We present a new analytical device, the *influence opportunity diagram of a game*, and demonstrate that it provides, for a large class of economically interesting games, a complete summary of the information needed to test empirical compatibility. A new equilibrium concept, *causal Nash equilibrium*, is presented and compared to several other well-known alternatives. Cases in which causal Nash equilibrium seems especially well-suited are explored.

Keywords: causal inference, information structure, extensive form, empirical compatibility, Bayesian network

1 Introduction

This paper focuses on the question of what can be said about situations in which the information structure of a finite-length extensive-form game is not known. There are two cases in which the answer to this question matters. The first is when an individual outside the game, say a researcher, would like to infer something about the information structure of the game from the observed behavior of its participants. The second is when individuals in a game are unsure of the information upon which their opponents condition their decisions. In this latter case, beliefs about causal structure – who influences whom – should play an important role in determining equilibrium behavior.

We explore both cases and, specifically, analyze what can be inferred about a game’s information structure solely from the probability distribution on action profiles induced by actual player strategies (which we refer to as the “empirical distribution of play”). The main idea is to connect a game’s information structure, which identifies the individual histories upon which players condition their behavior, to the corresponding set of conditional independencies that must be observed in all of its empirical distributions. When two games with different information structures imply different sets of such independencies, then knowledge of the empirical distribution provides a basis upon which to distinguish one from the other.

The first part of the paper considers information structure assessment from the perspective of an outsider who only observes player behavior (actions, not strategies). To this end, we introduce the notion of *empirical compatibility* between games. One game is said to be empirically compatible with another when the empirical distribution induced by any behavior strategy profile in the first game can also be induced by an appropriately chosen behavior strategy profile in the second. Thus, even under infinite repetition, it is impossible to distinguish between empirically compatible games based solely upon the observed behavior of players.

Our analysis is facilitated by the introduction of a graphical device, the *influ-*

ence opportunity diagram (hereafter, IOD). We define the IOD constructively for a broad class of finite-length extensive-form games, including those with infinite action sets. A basic result is that such diagrams summarize information about the conditional independencies that must be observed in all empirical distributions arising from play of the underlying game. Thus, a necessary condition for two games to be empirically compatible is that their IODs imply a consistent set of conditional independencies. This condition is not, in general, sufficient because differences in the specific information upon which players condition their behavior may imply empirical incompatibility. However, we do identify a broad class of games, termed games of *perfect observation*, for which this condition is also sufficient. For games of this type, we apply a well-known result from the artificial intelligence literature (see Appendix A) to show that empirical compatibility can be determined by simple visual comparison of their respective IODs.

The second part of the paper shifts the focus to influence assessment from the inside – that is, to games in which the players themselves are uncertain about the information structure governing their play. If equilibrium is interpreted as the outcome of some generic learning process (as is typical in the literature on learning in noncooperative games), then a player’s equilibrium beliefs regarding the underlying influence relationships should be consistent with reality. This idea leads to a new equilibrium notion, that of a *causal Nash equilibrium*, which imposes such consistency on player beliefs. We demonstrate the relationship between causal Nash equilibrium and other well-known equilibrium ideas.

An obvious question is whether this new equilibrium concept holds useful implications for situations of genuine economic interest. We can think of at least two cases in which it does. The first and, perhaps most obvious, is when payoffs are systematically related to information structure. In such situations, refining beliefs with respect to the true information structure may well lead to a better assessment of the payoffs faced both by oneself and one’s opponents. The second case, which to our knowledge

has previously received no explicit attention in the economics literature, is when a player (or players) must choose an appropriate ‘intervention’ in the activities of one or more of their opponents. A player is said to have intervention ability when his or her choice of action determines, non-trivially, the feasible actions available to others.¹ Here, an accurate assessment of the game’s influence relationships may be crucial to the success of the interventionist. We term these *intervention games* and present an example of causal Nash equilibrium applied to such a game.

The remainder of the paper is organized as follows. The next section presents several simple examples designed to illustrate the notion of empirical compatibility. Section 3 lays out the definition of a finite-length extensive-form game (which differs in some ways from the usual setup) and defines empirical compatibility. In Section 4, we present our main results regarding the analysis of empirical compatibility from the outside perspective. Section 4.1 shows how to construct an IOD from an extensive-form game. Section 4.2 connects information structure to empirical compatibility through the IOD. Section 4.3 shows how to test empirical compatibility between two games by visual inspection of their respective IODs. Section 5 shifts the focus to influence uncertainty within the game. First, we give a motivating example in which uncertainty about who takes the role of Stackleberg leader may cause potential entrants to stay out of a market. Section 5.2 introduces our definition of causal Nash equilibrium and makes formal comparisons to several well-known equilibrium concepts. Section 5.3 presents an extended example of causal Nash equilibrium applied to an intervention game. We conclude in Section 6 with a more thorough discussion of related research and potential extensions.

2 Examples and Intuition

Consider the game trees presented in Figures 1 through 3. The first, Γ^A , has the familiar structure of a standard “signalling” game. The other three are variations

involving the same players who have the same feasible actions at the time of their moves. We wish to show that Γ^A and Γ^B are in an equivalence class in the sense that any distribution on action profiles generated by (behavioral) strategies in one can also be generated by an appropriate choice of strategies in the other. Γ^C , on the other hand, is not a member of this class.

Let $\mathbf{A} \equiv \{(u, L, U), \dots, (d, R, D)\}$ be the set of possible action profiles in each of the three games (up to a permutation of the components). We refer to a single profile $\mathbf{a} \in \mathbf{A}$ as an “outcome” of play. Let $\boldsymbol{\theta}^k \equiv (\theta_N^k, \theta_I^k, \theta_{II}^k)$ denote a behavior strategy profile in Γ^k where θ_i^k is the strategy chosen by player i in game k . Every behavior strategy in each of the three games implies a probability distribution $m_{\boldsymbol{\theta}^k}$ on \mathbf{A} constructed as follows, for all $\mathbf{a} \in \mathbf{A}$,

$$\begin{aligned} m_{\boldsymbol{\theta}^A}(\mathbf{a}) &\equiv \theta_{II}^A(a_{II}|a_I) \theta_I^A(a_I|a_N) \theta_N^A(a_N), \\ m_{\boldsymbol{\theta}^B}(\mathbf{a}) &\equiv \theta_N^B(a_N|a_I) \theta_I^B(a_I|a_{II}) \theta_{II}^B(a_{II}), \\ m_{\boldsymbol{\theta}^C}(\mathbf{a}) &\equiv \theta_{II}^C(a_{II}) \theta_I^C(a_I) \theta_N^C(a_N). \end{aligned}$$

Now, suppose Γ^A is repeated a large number of times under a fixed strategy profile $\boldsymbol{\theta}^A$. Assume the outcomes are recorded and reported to an outside observer who knows that one of Γ^A , Γ^B , or Γ^C is the game responsible for generating the data (but not which). The question we wish to answer is whether there are any strategies $\boldsymbol{\theta}^A$ that would allow the outsider to correctly identify Γ^A as the underlying game.

First, note that the construction of $m_{\boldsymbol{\theta}^A}$ immediately implies that, for all $\mathbf{a} \in \mathbf{A}$, the following factorization holds

$$m_{\boldsymbol{\theta}^A}(\mathbf{a}) = m_{\boldsymbol{\theta}^A}(a_{II}|a_I) m_{\boldsymbol{\theta}^A}(a_I|a_N) m_{\boldsymbol{\theta}^A}(a_N).$$

Of course, $m_{\boldsymbol{\theta}^B}$ and $m_{\boldsymbol{\theta}^C}$ can be factored analogously. By the definition of conditional

probability, for every θ^A ,

$$\begin{aligned}
m_{\theta^A}(\mathbf{a}) &= m_{\theta^A}(a_{II}|a_I) m_{\theta^A}(a_I|a_N) m_{\theta^A}(a_N) \\
&= \frac{m_{\theta^A}(a_{II}, a_I)}{m_{\theta^A}(a_I)} \frac{m_{\theta^A}(a_I, a_N)}{m_{\theta^A}(a_N)} m_{\theta^A}(a_N) \\
&= \frac{m_{\theta^A}(a_N, a_I)}{m_{\theta^A}(a_I)} \frac{m_{\theta^A}(a_I, a_{II})}{m_{\theta^A}(a_{II})} m_{\theta^A}(a_{II}) \\
&= m_{\theta^A}(a_N|a_I) m_{\theta^A}(a_I|a_{II}) m_{\theta^A}(a_{II}).
\end{aligned}$$

This is significant because it implies that for every behavior strategy in Γ^A , one can find a corresponding strategy in Γ^B that generates exactly the same distribution on \mathbf{A} ; given θ^A , simply construct θ^B such that, for all $\mathbf{a} \in \mathbf{A}$, $\theta_N^B(a_N|a_I) \equiv m_{\theta^A}(a_N|a_I)$, and so on. Then, $m_{\theta^A} = m_{\theta^B}$. Therefore, the outside observer – even with very exact information about the true distribution on outcomes implied by some behavior strategy – can never distinguish between Γ^A and Γ^B . Since the converse is also true, we say that Γ^A and Γ^B are empirically equivalent.

On the other hand, it should be clear that Γ^C is not a member of the empirical equivalence class containing Γ^A and Γ^B . Barring correlated strategies without an explicit correlating device, there are many strategy profiles in Γ^A (and, therefore, in Γ^B as well) that generate distributions over action profiles that could not possibly correspond to any strategy profile in Γ^C . In particular any θ^A in which either player I 's behavior varies with Nature's play or player II 's behavior varies with the play of player I results in a m_{θ^A} that cannot be arranged by an appropriate choice of strategy in Γ^C .

3 The Model

Wherever possible, capital letters (X, Z) denote sets, small letters (a, w) either elements of sets or functions, and script letters, (\mathcal{A}, \mathcal{F}) collections of sets. Sets with product structure are indicated by bold capitals (\mathbf{A}, \mathbf{E}) with small bold (\mathbf{a}, \mathbf{e}) denoting typical elements (ordered profiles) in such sets. Graphs and probability spaces

play a large role in the following analysis. Standard notation and definitions are adopted wherever possible.

3.1 Extensive-Form Games

We begin with a finite-length, extensive-form game of perfect recall. The game Γ has a game tree (X, \mathbf{E}) with nodes X and edges \mathbf{E} . Players are indexed by $N \equiv \{1, \dots, n\}$ with $n < \infty$. The terminal nodes are $Z \subset X$ with typical element z . Payoffs are given by $u : Z \rightarrow \mathbb{R}^n$. Attention is restricted to games in which influence opportunities between players are fixed.² Specifically, assume that all paths are of length $t < \infty$ and that the player-move order is summarized by an onto function $o : T \rightarrow N$ where $T \equiv \{1, \dots, t\}$ and $i = o(r)$ means that i is the player who (always) has the r^{th} move.³

Every $(x_r, x_{r+1}) \in \mathbf{E}$ corresponds to an action available at x_r . For all $r \in T$, let A_r be the union of the actions available at the nodes in X_r . Edges are labeled in such a way that every history is unique. In particular, every $z \in Z$ corresponds to a unique *action profile* $\mathbf{a}_z = (a_{1_z}, \dots, a_{t_z})$. The set of all action profiles is $\mathbf{A} \equiv \cup_{z \in Z} \mathbf{a}_z$. Each A_r comes equipped with a σ -algebra \mathcal{A}_r . The σ -algebra for \mathbf{A} is $\mathcal{A} \equiv \sigma(\{\mathbf{F} \in \times_{r \in T} \mathcal{A}_r \mid \mathbf{F} \subset \mathbf{A}\})$. Assume all measure spaces are standard.⁴ We call $(\mathbf{A}, \mathcal{A})$ the *outcome space*. This, coupled with an appropriate probability measure, is the focal object of our analysis. Let $\mathbf{a}_z \mapsto v(\mathbf{a}_z) \equiv u(z)$ translate payoffs on Z to payoffs on \mathbf{A} .

For $r \in T$, the *history at r* is an \mathcal{A} -measurable function $\mathbf{a} \mapsto \tilde{h}_r(\mathbf{a}) \equiv (a_1, \dots, a_{r-1})$. We use \mathbf{h}_r to denote a typical element of $\tilde{h}_r(\mathbf{A})$ and define \tilde{h}_1 to be a constant equal to the null history \mathbf{h}_0 . For every move r , there is a bijective relationship between $\tilde{h}_r(\mathbf{A})$, the set of all $(r-1)$ -length action profiles, and X_r . In general, players do not know the full profile of actions leading up to their move. To reflect this, X_r is partitioned into a collection of subsets called the *move- r information partition* and whose elements are called *move- r information sets*. Given the bijective relationship between Z and \mathbf{A} (and the fact that every path in the tree contains exactly one

node in X_r), the move- r information partition implies a corresponding partition of \mathbf{A} whose elements, we assume, are \mathcal{A} -measurable. Define the *information at move r* , \mathcal{I}_r , to be the sub- σ -algebra of \mathcal{A} generated by this partition; note that $\mathcal{I}_r \subseteq \sigma(\tilde{h}_r)$.

Typically, not all of A_r is available to player $o(r)$ given a particular history \mathbf{h}_r . The feasible actions at r are given by the \mathcal{I}_r -measurable *move- r feasible action constraint* $\tilde{c}_r : \mathbf{A} \rightarrow \mathcal{A}_r$. The measurability condition that is implied by the standard assumption that feasible action sets are equal for all histories in the same information set. This allows us to write $\tilde{c}_r(\tilde{h}_r(\mathbf{a}))$ or $\tilde{c}_r(\mathbf{h}_r)$ without ambiguity.

Let $\Delta(X, \mathcal{X})$ denote the set of probability measures on a measure space (X, \mathcal{X}) ; when X is countable, we simply write $\Delta(X)$ where it is to be understood that $\mathcal{X} = 2^X$. Traditionally, a behavior strategy at a move is a function from the information sets at that move to probability measures on the player's feasible actions. Equivalently, we implement this idea by defining a *behavior strategy at move- r* to be an \mathcal{I}_r -measurable function $\theta_r : \mathbf{A} \rightarrow \Delta(A_r, \mathcal{A}_r)$ where $\theta_r(F|\tilde{h}_r(\mathbf{a}))$ is the probability that player $o(r)$ takes an action in $F \in \mathcal{A}_r$ given her arrival at the node corresponding to the partial history $\tilde{h}_r(\mathbf{a})$. The measurability requirement achieves the effect of making θ_r constant on all histories in the same information set. Naturally, $\theta_r(\cdot|\tilde{h}_r(\mathbf{a}))$ is restricted to assign positive probability only to measurable subsets of $\tilde{c}_r(\tilde{h}_r(\mathbf{a}))$. Player i 's *behavior strategy* is defined as the profile $\boldsymbol{\theta}_i \equiv (\theta_r)_{r \in o^{-1}(i)}$. Σ_i is the set of all behavior strategies available to i . A *strategy profile* is an element $\boldsymbol{\theta} \in \Sigma \equiv \times_{i \in N} \Sigma_i$. When convenient, we use the familiar shorthand $\boldsymbol{\theta} = (\boldsymbol{\theta}_i, \boldsymbol{\theta}_{-i})$.

3.2 Empirical Distribution

Given a game meeting the conditions of the previous section, every behavior strategy profile $\boldsymbol{\theta}$ induces a probability space, denoted $(\mathbf{A}, \mathcal{A}, m_\theta)$. The measure m_θ can be constructed directly from $\boldsymbol{\theta}$ as follows: for all $\mathbf{F} \in \mathcal{A}$,

$$m_\theta(\mathbf{F}) \equiv \int_{A_1} \dots \int_{A_t} I_{\mathbf{F}}(\mathbf{a}) \theta_t(da_t|a_1, \dots, a_{t-1}) \cdots \theta_2(da_2|a_1) \theta_1(da_1), \quad (1)$$

where \int indicates Lebesgue integration and $I_{\mathbf{F}}$ is the indicator function for \mathbf{F} . We call m_{θ} the *empirical distribution induced by θ* . For all $r \in T$, define $\tilde{a}_r : \mathbf{A} \rightarrow A_r$ so that $\tilde{a}_r(\mathbf{a})$ is the projection of \mathbf{a} into its r^{th} dimension. Then, for all $\mathbf{a} \in \mathbf{A}$, $F_r \in \mathcal{A}_r$,

$$m_{\theta} \left(F_r | \tilde{h}_r \right) (\mathbf{a}) = \theta_r \left(F_r | \tilde{h}_r (\mathbf{a}) \right) \quad (2)$$

where $m_{\theta} \left(F_r | \tilde{h}_r \right)$ denotes the conditional probability of $\tilde{a}_r^{-1}(F_r)$ given $\sigma \left(\tilde{h}_r \right)$.⁵

We use $m_{\theta} \left(\tilde{a}_r | \tilde{h}_r \right)$ to denote the m_{θ} -conditional distribution of \tilde{a}_r given $\sigma \left(\tilde{h}_r \right)$. Since θ_r is \mathcal{I}_r -measurable, $m_{\theta} \left(\tilde{a}_r | \tilde{h}_r \right)$ is equal to $m_{\theta} \left(\tilde{a}_r | \mathcal{I}_r \right)$. This, combined with (1) and (2), implies that, for all $\theta \in \Sigma$,

$$m_{\theta} = \prod_{r \in T} m_{\theta} \left(\tilde{a}_r | \mathcal{I}_r \right), \quad (3)$$

in the sense that, for all $\mathbf{F} \in \mathcal{A}$, $m_{\theta}(\mathbf{F}) = \int_{\mathbf{F}} m_{\theta}(\mathbf{a}) d\mathbf{a} = \int_{\mathbf{F}} \left(\prod_{r \in T} m_{\theta}(\tilde{a}_r | \mathcal{I}_r)(\mathbf{a}) \right) d\mathbf{a}$.

Equation (3) says that the information structure of an extensive-form game implies certain conditional independencies in *every* empirical distribution that could arise as a result of play. Alternatively, given an arbitrary m_{θ} , is it possible to use the relationship in (3) to deduce the information structure of the underlying game? The answer is: yes, up to an equivalence class of games as described in the next section.

3.3 Empirical Compatibility and Equivalence

A game Γ' is said to be empirically compatible with Γ when the empirical distribution induced by any strategy profile in Γ can also be induced by an appropriately chosen strategy profile in Γ' . Consider a situation in which the data generated by a game is cross-sectional; i.e., a listing of the specific actions taken by each player without reference to the timing of the moves. Then, an individual observing outcomes generated by repeated play of θ in Γ , eventually, develops a fairly precise estimate of m_{θ} . However, when Γ' is empirically compatible with Γ , then there is no collection of Γ -generated data capable of ruling out Γ' as the true underlying game.

An obvious necessary condition for empirical compatibility (in the sense described above) is that the games have consistent player sets and outcome profiles. Given a permutation $f : T \rightarrow T$, let $f(\mathbf{a})$ denote the permuted profile $(a_{f(r)})_{r \in T}$ and, for $\mathbf{F} \subseteq \mathbf{A}$, let $f(\mathbf{F})$ be the set whose elements are the permuted elements of \mathbf{F} . Then, Γ' is *outcome compatible* with Γ if and only if: (1) $N = N'$; (2) there exists a permutation f such that $f(\mathbf{A}) = \mathbf{A}'$; (3) for all $r \in T$, $o(r) = o'(f(r))$; and, (4) for all $r \in T$, $\mathcal{A}_r = \mathcal{A}'_{f(r)}$. Let \mathbb{O}_Γ denote the class of games that are outcome compatible with Γ . If $\Gamma' \in \mathbb{O}_\Gamma$, then there may exist a $\theta' \in \Sigma'$ that induces an empirical distribution on $(\mathbf{A}, \mathcal{A})$; i.e., constructed as in (1) but using the appropriate permutation. When this is the case, we write $m_{\theta'}$ without ambiguity.

Definition 1 *A game Γ' is said to be empirically compatible with Γ , denoted $\Gamma \preceq \Gamma'$, if $\Gamma' \in \mathbb{O}_\Gamma$ and there exists a function $g : \Sigma \rightarrow \Sigma'$ such that $\forall \theta \in \Sigma$, $m_\theta = m_{g(\theta)}$.*

If both $\Gamma \preceq \Gamma'$ and $\Gamma' \preceq \Gamma$, then Γ and Γ' are said to be *empirically equivalent*, denoted $\Gamma \sim \Gamma'$. The interpretation is that when Γ' and Γ are empirically equivalent, any behavior observed under Γ (“observed” in the sense of knowing m_θ) could also be observed under Γ' and visa versa. When $\Gamma \sim \Gamma'$, Γ differs from Γ' in terms of its information and, possibly, payoff structures. Note that empirical compatibility is strong in the sense that the condition must hold for *all* $\theta \in \Sigma$. Alternatively, for example, one might be interested in a notion of empirical compatibility defined only for specific (e.g., equilibrium) profiles.

Lemma 1 *Empirical equivalence is an equivalence relation on the space of finite-length extensive form games. Moreover, if $\Gamma \sim \Gamma'$, then $(f(\mathbf{A}'), f(\mathcal{A}')) = (\mathbf{A}, \mathcal{A})$ for some permutation f and there exists an onto correspondence $g : \Sigma \rightarrow \Sigma'$ such that $\forall \theta \in \Sigma, \theta' \in g(\theta)$, $m_\theta = m_{\theta'}$.*

To help fix ideas, let us revisit the examples in Section 2. Starting with Γ^A , for all $\theta^A \in \Sigma^A$, the empirical distribution m_{θ^A} is constructed by: for all $\mathbf{a} \in \mathbf{A}$,

$$m_{\theta^A}(\mathbf{a}) = \theta_{II}^A(a_{II}|a_I) \theta_I^A(a_I|a_N) \theta_N^A(a_N).$$

Let $\mathbf{A}_{a_i} \subset \mathbf{A}$ be the event in \mathbf{A} corresponding to player i playing action a_i ; e.g., $\mathbf{A}_U = \{\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_5, \mathbf{a}_6\}$. Then, it is easy to check that, for all $\boldsymbol{\theta}^A \in \Sigma^A$,

$$m_{\boldsymbol{\theta}^A} = m_{\boldsymbol{\theta}^A}(\tilde{a}_{II}|\mathcal{I}_{II}^A) m_{\boldsymbol{\theta}^A}(\tilde{a}_I|\mathcal{I}_I^A) m_{\boldsymbol{\theta}^A}(\tilde{a}_N|\mathcal{I}_N^A),$$

where, $\mathcal{I}_N^A = \{\emptyset, \mathbf{A}\}$ (i.e., θ_N^A is a constant), $\mathcal{I}_I^A = \{\emptyset, \mathbf{A}_U, \mathbf{A}_D, \mathbf{A}\}$ and $\mathcal{I}_{II}^A = \{\emptyset, \mathbf{A}_L, \mathbf{A}_R, \mathbf{A}\}$.

Clearly, $\Gamma^B \in \mathcal{O}_{\Gamma^A}$. Moreover, as we saw in the example, for any $\boldsymbol{\theta}^A \in \Sigma^A$, there corresponds a $\boldsymbol{\theta}^B \in \Sigma^B$ such that, for all $\mathbf{a} \in \mathbf{A}$,

$$m_{\boldsymbol{\theta}^A}(\mathbf{a}) = \theta_N^B(a_N|a_I) \theta_I^B(a_I|a_{II}) \theta_{II}^B(a_{II}).$$

Therefore, $\Gamma^A \preceq \Gamma^B$. Since this works in both directions, it is also true that $\Gamma^B \preceq \Gamma^A$, thereby implying $\Gamma^A \sim \Gamma^B$.

4 Assessing Empirical Compatibility

In this section we analyze empirical compatibility from the perspective of an outside observer who, we suppose, observes a large number of outcomes generated by repetition of a game with unknown information structure. To what extent does such data illuminate the game's underlying information structure? Given a candidate game, empirical compatibility can be checked with the same “brute-force” approach used in the motivating examples. In simple cases, the analysis is relatively straightforward. On the other hand, consider the game in Figure 4. Here, 5 players interact under a relatively complex information structure. The implications of this structure for the empirical distributions on actions arising from player strategies are not obvious. We now develop results by which these implications are neatly analyzed.

4.1 Influence Opportunity Diagrams

Loosely, player $o(r)$ is said to have the opportunity to influence play at move s if he has a choice of feasible actions available under some conceivable play of the game

that permits player $o(s)$ to alter her behavior regardless of what her other opponents do (i.e., opponents other than $o(r)$). The following definition formalizes this idea.

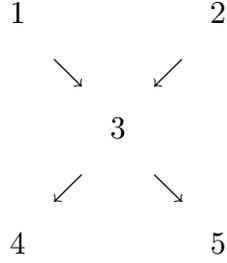
Definition 2 *The influence opportunity diagram of Γ is a graph (T, \rightarrow) such that $r \rightarrow s$ if and only if $r < s$, and there exist $\mathbf{a}, \mathbf{a}' \in \mathbf{A}$ satisfying each of the following conditions: (1) $\tilde{h}_r(\mathbf{a}) = \tilde{h}_r(\mathbf{a}')$; (2) $\exists \mathbf{F} \in \mathcal{I}_s$ such that $\tilde{h}_{r+1}^{-1}(a_1, \dots, a_r) \subset \mathbf{F}$ and $\mathbf{a}' \in \mathbf{F}^c$; (3) $a'_s \neq a_s$; and, (4) $\tilde{a}_j(\mathbf{a}')_{j \in \{k | k > r, k \rightarrow s\}} = \tilde{a}_j(\mathbf{a})_{j \in \{k | k > r, k \rightarrow s\}}$.*

The meat of the definition is that $r \rightarrow s$ when there is some move- r history (item 1) at which player $o(r)$ has a choice of actions that cause play at s to be at different information sets (item 2) and to which player $o(s)$ can respond differently (item 3). Note that item 2 implies that there are at least two distinct actions available at r , one that guarantees the occurrence of \mathbf{F} and another that is necessary for the occurrence of \mathbf{F}^c (but may not guarantee it). Influence is only an “opportunity” since this condition is neither necessary nor sufficient for move r actions to have an actual effect on move s behavior. For example, the player at move s may choose to ignore the action taken at move r (e.g., when θ_s is constant on \mathbf{A}). Alternatively, the player at move r may influence play at move s indirectly through other players (e.g., when $r \rightarrow q \rightarrow s$ even though $r \not\rightarrow s$). Item 4 is a technical condition that rules out spurious influence due to feasible action restrictions that force the move at s to be independent of actions taken at r given actions taken at some subset of moves following r . Although spurious influence due to game structure is a technical possibility, it does not arise in any games of economic interest with which we are familiar.

Return to game Γ^A in Figure 1. Here, player I observes player N and player II observes player I , which suggests the IOD should be $N \rightarrow I$ and $I \rightarrow II$. To see that this is correct, first check $N \rightarrow I$. In this case, $\mathcal{I}_I = \{\emptyset, \mathbf{F}_U, \mathbf{F}_D, \mathbf{A}\}$ where $\mathbf{F}_U \equiv \{\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_5, \mathbf{a}_6\}$. Then, $(\mathbf{a}_5, \mathbf{a}_4)$ establish the result: (1) $\tilde{h}_N(\mathbf{a}_5) = \tilde{h}_N(\mathbf{a}_8) = \mathbf{h}_0$, (2) $\tilde{h}_I^{-1}(\mathbf{h}_0, U) = \{\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_5, \mathbf{a}_6\} = \mathbf{F}_U$ and $\mathbf{a}_4 \in \mathbf{F}^c = \mathbf{F}_D$, and (3) $\tilde{a}_I(\mathbf{a}_5) = (R) \neq \tilde{a}_3(\mathbf{a}_4) = (L)$. Item (4) is automatically satisfied since there are not moves between

N and II . Similarly, $I \rightarrow II$ is established by $(\mathbf{a}_1, \mathbf{a}_6)$. However, $N \not\rightarrow II$ since the smallest \mathcal{I}_{II} -measurable event containing either $\tilde{h}_N^{-1}(U)$ or $\tilde{h}_N^{-1}(D)$ is \mathbf{A} .⁶

By identical reasoning, the IODs for the game in Figure 3 is $N \leftarrow I \leftarrow II$. The IOD for Figure 2 is a graph with three nodes and no edges. The IOD for the Gatekeeper game (Figure 4) is simply:



Player 3 is the “gatekeeper” of information flowing from players 1 and 2 to players 4 and 5.

To understand item (4) of the definition, consider Game I in Figure 5. Notice that this game has the unusual feature that player 2’s feasible action sets are different at every information set. Without item (4), the IOD would be $1 \rightarrow 2$, $2 \rightarrow 3$ and $1 \rightarrow 3$. However, by condition (4), $1 \rightarrow 3$ is removed. Intuitively, the game’s structure implies that knowing the action chosen by 1 is *always* irrelevant in assessing 3’s behavior when the action taken by 2 is already known. If the feasible actions at information set $2b$ are $\{U, D\}$, as in Game II, then the IOD is $1 \rightarrow 2$, $2 \rightarrow 3$ and $1 \rightarrow 3$.

Lemma 2 *Let (T, \rightarrow) be an IOD for some game Γ . If $r, s \in T$ such that $r < s$ and $r \not\rightarrow s$, then $\mathcal{I}_s \subseteq \sigma(\tilde{h}_{s \setminus r})$ where $\tilde{h}_{s \setminus r}(\mathbf{a}) \equiv (a_1, \dots, a_{r-1}, a_{r+1}, \dots, a_{s-1})$.*

The preceding lemma is helpful in establishing the conditional independencies implied by game structure (Proposition 1, below). The following lemma follows immediately from the requirement that $r \rightarrow s$ only if $r < s$. This is needed for the results in Section 4.3.

Lemma 3 *Every influence opportunity diagram is a directed, acyclic graph.*

We now demonstrate that the IOD summarizes certain conditional independencies that must arise in every empirical distribution of play.⁷ Given (T, \rightarrow) , the set of moves at which players may exert a direct influence upon player $o(r)$ at move r is $\{s \in T | s \rightarrow r\}$. Define $\tilde{\pi}_r(\mathbf{a}) \equiv (\tilde{a}_k)_{k \in \{s \in T | s \rightarrow r\}}$ to be the $\sigma(\tilde{h}_r)$ -measurable projection of \mathbf{a} into the dimensions indexed by $\{s \in T | s \rightarrow r\}$. If $\{s \in T | s \rightarrow r\} = \emptyset$, let $\tilde{\pi}_r$ be an arbitrary constant (in which case, $\sigma(\tilde{\pi}_r) = \{\emptyset, \mathbf{A}\}$).

Proposition 1 *Given a game Γ with influence opportunity diagram (T, \rightarrow) ,*

$$\forall \theta \in \Sigma, m_\theta = \prod_{r \in T} m_\theta(\tilde{a}_r | \tilde{\pi}_r). \quad (4)$$

4.2 Compatibility and Independence in Probability

If $\Gamma \preceq \Gamma'$, then there does not exist a strategy profile in Γ that, sufficiently repeated, generates data that distinguishes it from Γ' . In particular, equation (1) and the measurability restriction on behavior strategies imply that $\Gamma \preceq \Gamma'$ if and only if $\Gamma' \in \mathbb{O}_\Gamma$ and

$$\forall \theta \in \Sigma, m_\theta = \prod_{r \in T} m_\theta(\tilde{a}_r | \mathcal{I}'_r), \quad (5)$$

where $\mathcal{I}'_r \subset \mathcal{A}$ corresponds to the information at move $f(r)$ in Γ' .⁸ In other words, testing whether Γ' is empirically compatible with Γ is equivalent to checking for outcome compatibility and then checking whether every empirical distribution induced by a strategy in Γ can be factored according to the information algebras implied by Γ' .⁹ Since $\mathcal{I}'_r \subseteq \sigma(\tilde{\pi}'_r)$, we have the following corollary to Proposition 1.

Corollary 1 *If $\Gamma \preceq \Gamma'$ then*

$$\forall \theta \in \Sigma, m_\theta = \prod_{r \in T} m_\theta(\tilde{a}_r | \tilde{\pi}'_r). \quad (6)$$

To see why condition (6) is only necessary (i.e., as opposed to necessary and sufficient when $\Gamma' \in \mathbb{O}_\Gamma$), consider the two games in Figure 6. Both have the same IOD: $I \rightarrow II$. Moreover $\Gamma_2 \in \mathbb{O}_{\Gamma_1}$ and $\Gamma_1 \in \mathbb{O}_{\Gamma_2}$. Clearly, however, there are empirical

distributions that can arise in Γ_1 but not in Γ_2 . Indeed, $\Gamma_2 \preceq \Gamma_1$ but $\Gamma_1 \not\preceq \Gamma_2$. The issue is the measurability distinction between conditions (5) and (6). If \mathcal{I}_{II}^2 is the information algebra for II in Γ_2 , then there are behavior strategies for II in Γ_1 that are not \mathcal{I}_{II}^2 -measurable even though every such strategy is $\sigma(\tilde{\pi}_{II}^2)$ -measurable.

For certain types of games, Corollary 1 can be strengthened. Γ is said to be a game of *perfect observation* if, for all $r \in T$, $\mathcal{I}_r = \sigma(\tilde{\pi}_r)$; a game of perfect observation is one in which a player observes the moves of those preceding her either perfectly or not at all. This class contains many extensive-form games of economic interest: all of the games in Section 2 meet this requirement as do many standard market games such as Cournot, Stackleberg, etc.

Proposition 2 *Let Γ' be a game of perfect observation. Then, $\Gamma \preceq \Gamma'$ if and only if $\Gamma' \in \mathbb{O}_\Gamma$ and condition (6) hold.*

Suppose one is interested in determining whether two games, say Γ and Γ' , happen to be empirically compatible. Then, Corollary 1 implies a sufficient condition by which to reject empirical compatibility.¹⁰

Proposition 3 *Suppose $\Gamma' \in \mathbb{O}_\Gamma$, then $\Gamma \not\preceq \Gamma'$ if $\exists \theta \in \Sigma$, $r \in T$ such that $m_\theta(a_r | \tilde{\pi}_r)$ is not $\sigma(\tilde{\pi}_r')$ -measurable.*

Proposition 3 provides a sufficient test for *rejecting* empirical compatibility. What about a sufficient test to establish it? Corollary 1 seems impractical as a basis for this since it appears to require checking an infinite number of strategies. As we now show, however, it is not necessary to test every empirical distribution. Rather, very generally it is possible to select one distribution that illuminates all of the conditional dependencies that may arise and all the information available to each agent.

A strategy profile $\theta \in \Sigma$ is said to be *maximally revealing* if and only if there does not exist $S \subset \{s \in T | s \rightarrow r\}$ such that $m_\theta(\tilde{a}_r | \tilde{\pi}_S) = m_\theta(\tilde{a}_r | \tilde{\pi}_r)$, where $\tilde{\pi}_S()$ is the projection of \mathbf{a} to the dimensions indexed by S . If θ is maximally revealing,

then the IOD is minimal in the sense that removal of any edge causes equation (4) to fail. Game II in Figure 5 illustrates a maximally revealing strategy (where θ is identified by the edge weights shown). To see this, recall that the IOD is given by $1 \rightarrow 2$, $1 \rightarrow 3$ and $2 \rightarrow 3$. Note that $m_\theta(U|L) = .50$ but $m_\theta(U) = .41$. In addition, $m_\theta(x|L, D) = 1$ while $m_\theta(x|D) = .54$, $m_\theta(x|L) = .50$ and $m_\theta(x) = .41$. So, $m_\theta(\mathbf{a}_3) = m_\theta(L)m_\theta(D|L)m_\theta(x|L, D)$ exactly per the IOD and this equality fails by the removal of any conditioning variables.

Proposition 4 *For every game there exists a maximally revealing strategy profile.*

If θ is maximally revealing, then an observer who knows the move order o and empirical distribution m_θ can accurately infer the game's IOD. A strategy profile θ' that is not maximally revealing implies that $m_{\theta'}$ contains additional conditional independencies relative to m_θ (due, i.e., to players who choose to ignore certain information or who play certain actions with zero probability). Note, however, that such a θ' can never introduce *new* conditional dependencies relative to m_θ .

Thus, a test for empirical compatibility between Γ and Γ' , where $\Gamma \in \mathbb{O}_{\Gamma'}$, is: 1) take any maximally revealing strategy profile θ in Γ , and 2) check whether equation (4) holds. If not, the games are not empirically compatible. If Γ and Γ' are games of perfect observation then the ability to factor m_θ according to (4) also provides a sufficient test of empirical compatibility. We formalize this below.

Proposition 5 *Suppose Γ' is a game of perfect observation and that $\Gamma \in \mathbb{O}_{\Gamma'}$. If there exists a maximally revealing $\theta \in \Sigma$ such that $m_\theta = \prod_{r \in T} m_\theta(a_r | \tilde{\pi}'_r)$, then $\Gamma \preceq \Gamma'$.*

Corollary 2 *Suppose Γ and Γ' are games of perfect observation, $\Gamma' \in \mathbb{O}_\Gamma$ and $\Gamma \in \mathbb{O}_{\Gamma'}$. If there exist maximally revealing $\theta \in \Sigma$ and $\theta' \in \Sigma'$ such that $m_\theta = \prod_{r \in T} m_\theta(\tilde{a}_r | \tilde{\pi}'_r)$ and $m_{\theta'} = \prod_{r \in T} m_{\theta'}(a'_{f'(r)} | \tilde{\pi}_{f'(r)})$, then $\Gamma \sim \Gamma'$.*

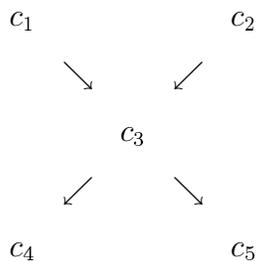
4.3 A Simple Visual Test for Empirical Compatibility

The number of possible conditional independencies implied by the IOD of any reasonably complex game will be quite large. This section presents a result that allows the testing of empirical equivalence by direct comparison of IODs without reference to specific probability parameters. For the following proposition, given a directed graph (T, \rightarrow) , let \mathcal{E} be the set of edges without reference to direction; i.e., $\{i, j\} \in \mathcal{E}$ if and only if $(i \rightarrow j)$ or $(j \rightarrow i)$. Let $\mathbf{S} \subset T^3$ be the set of all ordered triples such that $(i, j, k) \in \mathbf{S}$ if and only if $(i \rightarrow j), (k \rightarrow j)$ and $\{i, k\} \notin \mathcal{E}$.

Proposition 6 *Assume $\Gamma \in \mathbb{O}_{\Gamma'}$. If $\Gamma \preceq \Gamma'$ then (T, \rightarrow) and (T, \rightarrow') are such that $\mathcal{E} \subseteq \mathcal{E}'$ and $\mathbf{S} \subseteq \mathbf{S}'$. If Γ' is also a game of perfect observation, then these conditions are also sufficient.*

Corollary 3 *Assume Γ and Γ' are games of perfect observation and $\Gamma \in \mathbb{O}_{\Gamma'}$ and $\Gamma' \in \mathbb{O}_{\Gamma}$. Then, $\Gamma \sim \Gamma'$ if only if (T, \rightarrow) and (T, \rightarrow') are such that $\mathcal{E} = \mathcal{E}'$ and $\mathbf{S} = \mathbf{S}'$.*

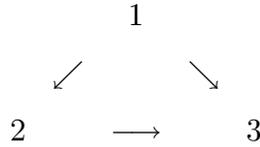
This is an especially nice result in games of perfect observation because it allows us to test empirical compatibility between games by direct visual inspection of their IODs. To see this, consider once again games Γ^A and Γ^B in Section 2. These are both games of perfect observation. The IODs are $N \rightarrow I \rightarrow II$ and $N \leftarrow I \leftarrow II$, respectively. By Corollary 3, we know almost immediately that these games are empirically equivalent: $\mathcal{E}^A = \mathcal{E}^B = \{\{N, I\}, \{I, II\}\}$ and $\mathbf{S}^A = \mathbf{S}^B = \emptyset$. Recall that the IOD for the Gatekeeper game is:



Since this is a game of perfect observation, Corollary 3 tells us that there are no other games with which this game is empirically equivalent. Any empirically compatible

IOD would have the same set of edges, some with different directions. However, reversing any arrow above either breaks a converging pair of arrows or creates a new one.¹¹

Finally, the requirement that elements of \mathbf{S} not include convergent arrows with linked tails (i.e., $(i, j, k) \in \mathbf{S} \Rightarrow \{i, k\} \notin \mathcal{E}$) has an implication for three-move games that should be kept in mind when reviewing the examples in Sections 5.1 and 5.3. Namely, all outcome compatible games whose IOD is a variation of the fully connected graph



are empirically compatible.

5 Influence from the Player’s Perspective

In this section, we consider the implications of empirical compatibility from the perspective of the players inside a game. There are at least two cases in which uncertainty about a game’s information structure may have equilibrium implications. The first is when payoffs are correlated with game structure. That is, when knowledge about the structure of the game may allow players to infer something about their own type. We begin this section with a motivating example of this kind. We then present a new equilibrium notion, causal Nash equilibrium, for games in which uncertainty about who influences whom is an important factor. Finally, we close with an example of the second case, games in which the central interest is in the ability of one player to intervene in the activities of others. In such situations, the interventionist’s beliefs about his influence relationships may have important behavioral implications.

5.1 Causal Uncertainty as a Barrier to Entry

Consider a situation in which a firm must decide whether or not to enter an industry. Assume the potential entrant is a short-term player (i.e., will play for one period only) which, upon entry, challenges a long-term incumbent in a market game of quantity competition. Suppose the challenger is uncertain both about the information structure and its own marginal cost. Imagine the challenger has in its possession cross-sectional quantity data from a long sequence of interactions in which entry occurred (by other short-term competitors). Assume that the data indicates a noisy process with a strong negative correlation between the quantity choices of the incumbent and those of its competitors. Demand parameters are known, but actual cost information is not publicly available. Entrants share a common cost.

What should the challenger do? The correlated quantity choices suggest that someone, either the entrant or the incumbent, takes the role of Stackleberg leader. To make things concrete, suppose the market game is parameterized as follows. The market leader and follower have constant marginal costs of $c_l = 2$ and $c_f = 1$, respectively. Inverse demand is given by $P \equiv 7 - q_l - q_f$ where $(q_l, q_f) \in \mathbb{R}_+^2$ are the quantities chosen by the two firms. Firm production processes are prone to stochastic shocks with actual output for firm k given by $q_k^a \equiv q_k + \varepsilon_k$ where $\varepsilon_k \sim (0, \sigma_k^2)$ is an *i.i.d.* random noise term. The Nash equilibrium expected output is $\bar{q}_l = 2$ and $\bar{q}_f = 2$. The expected profit for the leader is $\bar{v}_l = 2$ and for the follower is $\bar{v}_f = 4$. Actual observations (i.e., the data available to the challenger) are generated by $q_l^a = 2 + \varepsilon_l$ and $q_f^a = 2 - \frac{1}{2}\varepsilon_l + \varepsilon_f$. This implies that $Cov(q_l^a, q_f^a) < 0$. The challenger knows these parameters, but not the role to which it will be assigned upon entry.

Let Γ_1 be the game in which the incumbent is the leader and Γ_2 be the one in which entrants lead. Assume once and for all that entrants are always Stackleberg followers; i.e., the true game is Γ_1 . If the challenger enters, it pays a one-time entry fee of 3. If it stays out, it receives a payoff of zero. In this situation, the Nash equilibrium of the game is for our challenger to enter with a net expected payoff of 1.

The incumbent is assumed to know the truth and to play optimally in every period (which is simply to play his part of the static Nash equilibrium in the market stage game).

The problem with applying Nash here is highlighted by Corollary 3. The true stage game has three moves and a fully connected IOD $\{(E_1 \rightarrow I), (I \rightarrow E_2), (E_1 \rightarrow E_2)\}$ where E_1 is the entrant's decision to enter or not, and I and E_2 are the incumbent's and the entrant's quantity choices, respectively. As we discuss on page 19, since Stackleberg is a game of perfect observation, this is empirically equivalent to the subgame in which the entrant is the leader. No quantity of data (of the type specified above) can identify which is the true game. Specifically, suppose the challenger has initial prior $\mu \in [0, 1]$ that Γ_1 is the true stage game. If $\mu \leq \frac{1}{2}$, the subjectively rational challenger enters, otherwise it does not. Notice that, if entry occurs, the challenger learns the game is, indeed, Γ_1 and, upon learning this, has no regrets about its decision. On the other hand, if the firm stays out, it receives a payoff of zero (as expected) and no sequence of additional entry data generated by future challengers will ever reveal its mistake.

One well-known solution concept that may seem appropriate in this situation is Bayesian Nash equilibrium (hereafter, BNE). However, BNE requires players to have common and correct priors which, in this context, implies either that all challengers enter or all challengers stay out. In particular, assuming the challenger has a wealth of data from previous entries but decides to stay out is inconsistent with BNE. Apparently, some solution concept other than Nash or BNE is required. This is the subject to which our analysis now turns.

5.2 Causal Nash Equilibrium

In the spirit of the literature on game theoretic learning, we wish to develop an equilibrium concept whose interpretation is consistent with situations like the one described above. In equilibrium, players begin with priors that are consistent with

some (unmodelled) pre-play learning process that (implicitly) generates information about the game's influence relationships, they choose strategies that are optimal with respect to these priors and, as play unfolds, nothing observed with positive probability refutes their initial priors. Specifically, suppose players in some game Γ are uncertain about the game's information structure and payoffs; that is, everyone knows they are playing some game in \mathbb{O}_Γ . Let $\hat{\mu}_i$ denote player i 's initial prior regarding which of the games in \mathbb{O}_Γ is the one actually being played. For simplicity, assume that $\Lambda_i \equiv \text{support}(\hat{\mu}_i)$ is finite. We do not require players to have common priors, but we do impose a minimal amount of consistency with the underlying game: for all $i \in N$, $\Gamma \in \Lambda_i$.

In this context, each player needs to know what she will do at any information set that could be reached in any of the games she believes she might be playing. Recall that the information sets at a move in Γ correspond to a partition of \mathbf{A} . Therefore, for all $\Gamma^k \in \Lambda_i$, let $\mathcal{C}_r^k \subset \mathcal{A}$ denote the partition of \mathbf{A} that corresponds to player i 's move- r information sets in Γ^k . Define $\mathcal{C}_i^k \equiv \cup_{r \in \mathcal{O}^{-1}(i)} \mathcal{C}_r^k$. For each $\mathbf{C} \in \mathcal{C}_i^k$, there is a corresponding set of feasible actions for player i , which we now denote $A_{\mathbf{C}}^k$ (the definition of \mathbb{O}_Γ ensures measure consistency across games, so we suppress reference to the associated σ -algebras). Thus, reaching an information set in game Γ^k is equivalent to being told $(\mathbf{C}, A_{\mathbf{C}}^k)$. Define $\Omega_{\Lambda_i} \equiv \cup_{\Gamma^k \in \Lambda_i} \cup_{\mathbf{C} \in \mathcal{C}_i^k} (\mathbf{C}, A_{\mathbf{C}}^k)$; that is, Ω_{Λ_i} is the set of all $(\mathbf{C}, A_{\mathbf{C}}^k)$ upon which i may condition her behavior given her belief that the true game is one contained in Λ_i . When Λ_i is clear from the context, we simply write Ω_i . Let ω denote a typical element of Ω_i and A_ω its feasible action component.

To illustrate, suppose player II from the examples in Section 2 places positive weight on Γ^A (Figure 1) and Γ^B (Figure 2); so, $\Lambda_{II} = \{\Gamma^A, \Gamma^B\}$. As we know, $\Gamma^A, \Gamma^B \in \mathbb{O}_\Gamma$. Player II has one move. If the true game is Γ^A , then at the time of her move, she knows either $\mathbf{C}_L \equiv \{\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3, \mathbf{a}_4\}$ or $\mathbf{C}_R \equiv \{\mathbf{a}_5, \mathbf{a}_6, \mathbf{a}_7, \mathbf{a}_8\}$ and that she is to choose one of $\{u, d\}$. If, on the other hand, Γ^B is the game, her knowledge at the time of her move is completely unrefined; that is, she knows $\mathbf{C}_\emptyset \equiv \mathbf{A}$ and

that her feasible actions are $\{u, d\}$. Therefore, $\mathcal{C}_{II}^A = \{\mathbf{C}_L, \mathbf{C}_R\}$ and $\mathcal{C}_{II}^B = \{\mathbf{A}\}$. Given her uncertainty, player II must develop an action plan that allows for any of $\Omega_{II} = \{(\mathbf{C}_L, \{u, d\}), (\mathbf{C}_R, \{u, d\}), (\mathbf{C}_\emptyset, \{u, d\}), \}$.

A *subjective behavior strategy for player i given μ_i* is a function ϕ_i such that, for all $\omega \in \Omega_i$, $\phi_i(\omega) \in \Delta(A_\omega, \mathcal{A}_\omega)$. Given μ_i , let Φ_i be the set of all subjective behavior strategies for i . It is easy to see that ϕ_i restricted to the ω implied by Γ^k , for example, corresponds to a unique behavior strategy for i in Γ^k , written $\phi_i|_{\Gamma^k} \in \Sigma_i^k$. Thus, given a game Γ , a profile of subjective strategies $\phi = (\phi_1, \dots, \phi_n)$ implies a probability space (A, \mathcal{A}, m_ϕ) where m_ϕ is the measure induced by $(\phi_1|_\Gamma, \dots, \phi_n|_\Gamma) \in \Sigma$.

For all $\Gamma^k \in \Lambda_i$, player i makes some assessment, denoted θ_{-i}^k , of the strategies adopted by the other players when the true game is Γ^k . Let $\hat{\Theta}_i \equiv (\theta_{-i}^k)_{\Gamma^k \in \Lambda_i}$ be the profile summarizing i 's assessment of opponent behavior in each of these games. Given a subjective behavior strategy ϕ_i and beliefs $(\hat{\mu}_i, \hat{\Theta}_i)$, we can define the expected payoff

$$E_v(\phi_i | \hat{\mu}_i, \hat{\Theta}_i) \equiv \sum_{\Gamma^k \in \Lambda_i} \hat{\mu}_i(\Gamma^k) \int_{\mathbf{A}} v_i^k(\mathbf{a}) m_{(\phi_i|_{\Gamma^k}, \theta_{-i}^k)}(d\mathbf{a}).$$

For beliefs $(\hat{\mu}_i, \hat{\Theta}_i)$, the best reply correspondence is

$$BR(\hat{\mu}_i, \hat{\Theta}_i) \equiv \left\{ \phi_i \in \Phi_i \mid \forall \phi'_i \in \Phi_i, E_v(\phi_i | \hat{\mu}_i, \hat{\Theta}_i) \geq E_v(\phi'_i | \hat{\mu}_i, \hat{\Theta}_i) \right\}.$$

Let $\hat{\mu} \equiv (\hat{\mu}_1, \dots, \hat{\mu}_n)$ and $\hat{\Theta} \equiv (\hat{\Theta}_1, \dots, \hat{\Theta}_n)$ denote profiles of player beliefs regarding the underlying game and opponent behavior, respectively.

In a causal Nash equilibrium (hereafter, CNE), we desire agents to maintain beliefs that are consistent both with the information they observe as well as with the best assessment of influence relations possible from (hypothetical) data on pre-equilibrium play. We have not yet identified what events i thinks he might observe at the conclusion of play if, say, Γ^k is the true game. What we assume is that i knows exactly as much as is implied by his payoff v_i^k . Formally, \mathcal{V}_i^k is the partition of \mathbf{A} implied by v_i^k . Lastly, for the upcoming influence-consistency condition, let Ξ_Γ denote the set of games that are empirically compatible with Γ .

Definition 3 A profile ϕ is a causal Nash equilibrium if there exist beliefs $(\hat{\mu}, \hat{\Theta})$ such that, for all $i \in N$: (1) Subjective optimization: $\phi_i \in BR(\hat{\mu}_i, \hat{\Theta}_i)$; (2) Uncontradicted beliefs: (i) for all $\Gamma^k \in \Lambda_i$, $\mathbf{B} \in \mathcal{C}_i^k \cup \mathcal{V}_i^k$, $m_{(\phi_i|_{\Gamma^k}, \theta_{-i}^k)}(\mathbf{B}) = m_\phi(\mathbf{B})$, (ii) $v_i^k = v_i$ m_ϕ -a.s.; and, (3) Learned structure: $\Lambda_i \subseteq \Xi_\Gamma$.

The first condition says that players play best responses to their beliefs. The second imposes consistency between a player's expectations and the true distribution induced on their own observables by ϕ . That is, a player's expectations are correct with respect to information sets arrived at with positive probability during the game (\mathcal{C}_i^k) as well as with whatever information is reported at its conclusion (\mathcal{V}_i^k). Moreover, conditional expectations over own outcomes upon arriving at a particular information set are also correct. The last requirement limits the set of games under consideration to those meeting the "pre-equilibrium learning" consistency condition. The interpretation of this is that, as players grope their way toward equilibrium during the (unmodelled) learning phase, they discover the influence relationships implied by the structure of their game. Finally, although a subjective strategy must provide for the possibility that player i observes the same histories with two distinct action sets (i.e., $(\mathbf{C}, A_{\mathbf{C}}^k)$ or $(\mathbf{C}, A_{\mathbf{C}}^l)$ with $A_{\mathbf{C}}^k \neq A_{\mathbf{C}}^l$), items (2) and (3) combined with the assumption of perfect recall imply that this never occurs in equilibrium (e.g., Γ^1 and Γ^2 in the preceding illustration are not empirically compatible).

Returning to the entry example, let ϕ be given by: (i) $q_E = 0$, and (ii) $q_I = \frac{5}{2}$ if $q_E = 0$ and $q_I = 2$ otherwise. Assume the challenger's beliefs about which game is being played is given by $\mu_E > \frac{1}{2}$. Regarding the incumbent's strategy, the challenger correctly believes the incumbent produces $\frac{5}{2}$ when there is no entry and 2 otherwise. The incumbent knows the game and assumes the challenger produces 2 if it enters. Equilibrium payoffs are as expected. These strategies and beliefs constitute a causal Nash equilibrium.

CNE places no explicit restrictions on players' beliefs about the rationality or payoffs of their opponents. Of course, $(\hat{\mu}_i, \hat{\Theta}_i)$ may explicitly include such additional

restrictions. For example, a self-confirming equilibrium (SCE) in Γ is a CNE such that, for all $i \in N$, $\hat{\mu}_i(\Gamma) = 1$. So, $SCE_\Gamma \subseteq CNE_\Gamma$ where SCE_Γ is the set of self-confirming equilibria associated with Γ , etc. A Nash equilibrium (NE) is an SCE such that, for all $i \in N$, $\hat{\Theta}_i = \theta_{-i}$. Therefore, $NE_\Gamma \subseteq SCE_\Gamma$. A Bayesian Nash equilibrium (BNE) is an SCE in which, for all $i, j \in N$, $\hat{\mu}_i = \hat{\mu}_j$ and $\hat{\Theta}_i = \theta_{-i}$. So, $BNE_\Gamma \subseteq SCE_\Gamma$. Summing up:

Proposition 7 *For any finite-length extensive-form game Γ , $NE_\Gamma \subseteq SCE_\Gamma \subseteq CNE_\Gamma$ and $BNE_\Gamma \subseteq CNE_\Gamma$.*

Kalai and Lehrer (1995) present the notion of a “subjective game” and a corresponding definition of subjective Nash equilibrium (SNE). Their novelty is to demonstrate that a player need not know the entire game or his co-players’ strategies in order to compute a best response. Rather, it suffices for a player to know his own “environment response function,” a mapping from his available actions probability distributions on the consequences he might experience as a result of those actions. We wish to show that, in the context of the games studied by Kalai and Lehrer, CNE is a refinement of SNE.

In order to make the comparison formal we must introduce some new concepts and the corresponding notation (much of the latter adopted directly from Kalai and Lehrer). For consistency, we restrict attention to finite subjective games, which are simultaneous-move games played in stages. In each stage, player i has a countable set of actions, denoted A_i . The outcome of any stage is an element in $\mathbf{A}_{\text{stg}} \equiv \times_{i \in N} A_i$. Let T index repetitions of the stage, so the outcome space for the game as a whole is the Cartesian product $\mathbf{A} \equiv (\mathbf{A}_{\text{stg}})^t$. (Much of the analysis in Kalai and Lehrer focuses on games with infinitely repeated stages, but introducing a dynamic version of CNE is beyond the scope of this paper).

As before, when game Γ and player i are understood from the context, \mathcal{C}_r denotes the partition of \mathbf{A} that summarizes what player i knows at the start of stage r ; in

this setup, e.g., for all players, $\mathcal{C}_1 = \{\mathbf{A}\}$. \mathcal{V}_i is defined with respect to v_i as before. In a subjective game, $\mathcal{C}_i \equiv \mathcal{V}_i \cup_{r \in \mathcal{O}^{-1}(i)} \mathcal{C}_r$ is called i 's set of *consequences*. Since \mathbf{A}_{stg} is countable and $t < \infty$, \mathcal{C}_i is also countable. The idea is that, in stage r , player i recalls the consequence reported at the end of stage $(r-1)$, $\mathbf{C}_r \in \mathcal{C}_i$, and chooses an action in A_i . He is then informed of a new consequence $\mathbf{C}_{r+1} \in \mathcal{C}_i$. Under perfect recall, consequences do double-duty – they imply both the history known to i at the start of a stage as well as the event reported to i at its conclusion.

The *environment response function* for a player summarizes his individual decision problem. Formally, $e_i|_{\mathbf{C}_{r-1}a_{i,r}}(\mathbf{C}_r)$ denotes the probability of player i 's stage r consequence being \mathbf{C}_r given that his last consequence was \mathbf{C}_{r-1} and that he took action $a_{i,r}$. Given a profile of opponent strategies $\boldsymbol{\theta}_{-i}$, the computation of i 's environment response function is straightforward: for all $r \in T$, $a_{i,r} \in A_i$ and $\mathbf{C}_{r-1}, \mathbf{C}_r \in \mathcal{C}_i$,

$$e_i|_{\mathbf{C}_{r-1}a_{i,r}}(\mathbf{C}_r) = m_{\boldsymbol{\theta}}(\mathbf{C}_r|\mathbf{C}_{r-1})$$

where $\boldsymbol{\theta}_i$ is chosen such that $\mathbf{C}_{r-1} \cap \{\mathbf{a}' \in \mathbf{A} | a'_{i,r} = a_{i,r}\}$ occurs with positive probability. If $\boldsymbol{\theta}_{-i}$ is such that \mathbf{C}_{r-1} is impossible no matter what strategy i chooses, then $e_i|_{\mathbf{C}_{r-1}a_{i,r}}$ can be defined arbitrarily (since this situation never comes up).

Thus, e_i summarizes all the stochastic information i needs to calculate an optimal strategy. Let $(\mathbf{A}, \sigma(\mathcal{C}_i), m_{\boldsymbol{\theta}_i, e_i})$ be the probability space in which $m_{\boldsymbol{\theta}_i, e_i}$ is the measure on i 's individually observable events $\sigma(\mathcal{C}_i) \subset \mathcal{A}$, induced by $m_{\boldsymbol{\theta}}$. Define

$$E_v(\boldsymbol{\theta}_i | e_i) \equiv \sum_{\mathbf{C} \in \mathcal{V}_i} v_i(\mathbf{C}) m_{\boldsymbol{\theta}_i, e_i}(\mathbf{C}).$$

Given an e_i , the best-response correspondence is

$$BR(e_i) \equiv \{\boldsymbol{\theta}_i \in \Sigma_i | \forall \boldsymbol{\theta}'_i \in \Sigma_i, E_v(\boldsymbol{\theta}_i | e_i) \geq E_v(\boldsymbol{\theta}'_i | e_i)\}.$$

The last piece of the analysis is to assume that players do not know their true environment response function. Instead, player i assesses e_i by a *subjective* environment response function \hat{e}_i . That is, $\hat{e}_i|_{\mathbf{C}_{r-1}a_{i,r}}(\mathbf{C}_r)$ is i 's subjective assessment that

C_r occurs after having been told C_{r-1} and having taken action $a_{i,r}$. Then, m_{θ_i, \hat{e}_i} represents i 's belief on observable events given his choice of θ_i and his assessment \hat{e}_i .

Definition 4 *The pair (θ, \hat{e}) , $\hat{e} \equiv (\hat{e}_1, \dots, \hat{e}_n)$, is a subjective Nash equilibrium (SNE) if, for all $i \in N$: (1) subjective optimization: $\theta_i \in BR(\hat{e}_i)$; and, (2) uncontradicted beliefs: $m_{\theta_i, e_i} = m_{\theta_i, \hat{e}_i}$.*

To prove that CNE refines SNE we need only demonstrate that beliefs $(\hat{\mu}_i, \hat{\Theta}_i)$, as defined above, imply a subjective environment response function in the finite subjective game. Then, items (1) and (2) in Definition 4 are implied by items (1) and (2) in Definition 3 (players are assumed to know their payoffs in a subjective game). So, the only difference is that CNE has the learned influence requirement, item (3), that is not imposed in SNE. It is assumed that players know they are playing a finite subjective game.

Proposition 8 *Given a finite subjective game Γ , $CNE_\Gamma \subseteq SNE_\Gamma$.*

5.3 Intervention Games

We now turn to a class of games in which the distinctions of Definition 3 are meaningful. Define an *intervention game* as one in which some player must choose an appropriate intervention, meaning taking an action that changes the feasible actions available to some other player or players. Consider the following extended example of such a game.

A manager, denoted M , is responsible for the output of two departments, denoted A and B . The firm's profits, which the manager wishes to maximize, depend upon coordination between the departments. The options available to M are: 1) pursue a decentralized strategy and permit the two departments to engage in activities as they see fit, or 2) implement an intervention strategy to improve coordination by setting departmental actions (e.g., by monitoring and policing that department's

behavior). Assume, perhaps due to resource constraints, that M can only intervene in the activities of one department or the other.

Referring to Figure 7, suppose the actual departmental subgame is Γ_1 : A moves, then B attempts to coordinate. We suppress payoffs to A and B and assume they play fixed strategies. The order of moves is: 1) M chooses from the set of actions $A_M \equiv \{a_M^0, a_M^L, a_M^R, a_M^l, a_M^r\}$ where

$$\begin{aligned} a_M^0 &= \{L, R, l, r\} && \text{(do nothing),} \\ a_M^L &= \{L, l, r\} && \text{(make } A \text{ play } L), \\ a_M^R &= \{R, l, r\} && \text{(make } A \text{ play } R), \\ a_M^l &= \{L, R, l\} && \text{(make } B \text{ play } l), \\ a_M^r &= \{L, R, r\} && \text{(make } B \text{ play } r), \end{aligned}$$

2) A moves by choosing $a_A \in \{L, R\} \cap a_M$, and 3) B moves by choosing $a_B \in \{l, r\} \cap a_M$. M receives $v_M = 1$ if A and B coordinate (i.e., $\{L, l\}$ or $\{R, r\}$) and 0 otherwise. Any choice other than the “do nothing” option by M is an *intervention*. The idea is that M can either sit by and let the game run its natural course, or (imperfectly) influence the joint behavior of A and B .

Suppose, however, that M does not know the structure of the interaction between departments. For simplicity, assume that the departments play according to: A operates independently with $\theta_A(L, R) = (.4, .6)$ and B attempts to coordinate with the following probabilities

A action	$\theta_B(l, r a_A)$
L	(.8, .2)
R	(.1, .9)

After a sufficient history of unmanaged departmental interaction, M observes the

following outcome frequencies

Empirical Distribution m	
Dept. Activity Profiles	Observed Frequency
(L, l)	.48
(L, r)	.12
(R, l)	.04
(R, r)	.36

It is clear that A and B already do a reasonable job of coordinating. Left to their own devices, coordinate 84% of the time. Thus, the expected payoff of the decentralized (do nothing) approach .84.

From the history of interaction described in the preceding table, it is clear that either A or B plays a leadership role with the counterpart attempting to coordinate (with mixed success). Clearly, the simultaneous-move subgame, Γ_3 , can be ruled out. Γ_1 and Γ_2 on the other hand, are empirically equivalent. This is easily seen since the respective IODs are $\{(M \rightarrow A), (M \rightarrow B), (B \rightarrow A)\}$ and $\{(M \rightarrow A), (M \rightarrow B), (A \rightarrow B)\}$, both of which conform to the conditions in Corollary 3. Decomposing the empirical distribution into departmental strategies consistent with Γ_2 , we have the following: B operates independently with $\theta_B(l, r) = (.52, .48)$ and A attempts to coordinate with these probabilities

B action	$\theta_A(L, R a_B)$
l	(.92, .08)
r	(.25, .75)

Can M do better with an intervention? Since Γ_3 can be ruled out, suppose $\mu_M(\Gamma_1) = \mu_M(\Gamma_2) = 0.5$. Then, the expected payoffs associated with the available interventions are:

Action	Payoff in Γ_i		Expected Payoff
	Γ_1	Γ_2	
Do nothing	.84	.84	.84
Fix L	.80	.52	.66
Fix R	.90	.48	.69
Fix l	.40	.92	.66
Fix r	.60	.75	.68

These beliefs and doing nothing constitute a CNE. Objectively, of course, M should intervene and fix $a_A = R$, thereby increasing the expected payoff from .84 to .90. Thus, doing nothing is not a NE. Since positive weight is placed by M on Γ_2 , neither is it a SCE.

Suppose $\mu_M(\Gamma_3) = 1$ with $\tilde{\theta}_A(L, R) = (.4, .6)$ and $\tilde{\theta}_B(l, r) = (.5, .5)$. The subjective expected intervention-contingent payoffs are

Do nothing	.5
Fix $a_A = L$.5
Fix $a_A = R$.5
Fix $a_B = l$.4
Fix $a_B = r$.6

The subjectively rational manager sets $a_B = r$, observes $m_\theta(L) = (.4)$ and $m_\theta(R) = (.6)$ as expected and receives the expected payoff of .6. This is an SNE, but *not* a CNE since condition (3) of the CNE definition fails.

6 Discussion

Although our definition of empirically compatible games is new, the idea of empirically equivalent *strategies* is introduced at least as early as Kuhn (1953). Two strategies, behavior or mixed, are equivalent if they lead to the same probability distribution

over outcomes for all strategies of one's opponents. Kuhn demonstrates that, in games of perfect recall, every mixed strategy is equivalent to the unique behavior strategy it generates and each behavior strategy is equivalent to every mixed strategy that generates it (see Aumann, 1964, for an extension to infinite games). It follows immediately that every extensive-form game of perfect recall is empirically compatible with its reduced normal form and *visa versa*. Our results demonstrate that, generally, the set of empirically compatible games is much larger.

We have interpreted the results in Section 4 as consistent with the inferences that would be made by an outside observer with sufficiently informative empirical data. One question that immediately comes to mind is whether these ideas can be extended to construct an econometric test for game structure given cross sectional data on player actions. For example, the maximum likelihood estimate of the information structure for an industry could be useful in refining cost estimation in I/O empirical work (as suggested by the example in Section 5.1). This is the subject of on-going research.

The literature contains two primary approaches to analyzing situations in which players do not know the structure of the game. The first, and closest to ours in spirit, is Kalai and Lehrer's (1993, 1995) work on subjective games and their notion of subjective equilibrium. Kalai and Lehrer show that, provided beliefs are sufficiently close to the truth, play converges to a SNE. Moving to an infinitely repeated version of CNE and exploring the convergence properties of noisy learning processes strikes us as a worthwhile extension of this paper; we conjecture that results along the lines of Kalai and Lehrer also hold in our setting. The second approach, is to encode a player's uncertainty regarding the information structure of the game into his or her type (*à la* Harsanyi, 1967 – 68). When players have correct (and, therefore, common) priors, there is nothing in our methodology that is inconsistent with the Harsanyi approach.

Several authors have proposed other equilibrium definitions whose interpretations

are consistent with the idea that equilibria arise as the result of learning.¹² The structure shared by these definitions is: 1) players have prior beliefs about certain unknowns (i.e., competitor strategies and/or various elements of game structure), and 2) choose strategies that are best-replies to these beliefs, which then, 3) generate “observables” that do not refute the priors upon which the strategy choices were based. CNE has the novelty that beliefs are restricted to the set of empirically compatible games rather than, say, the set of games empirically consistent with the specific equilibrium strategy profile (typically, a much larger set). The stronger condition is appropriate if players observe a wide range of behavior prior to settling down into equilibrium.

A game’s IOD is a graph that summarizes information about its empirical distributions. This idea (i.e., using graphs to encode probabilistic information) is not new outside economics. In particular, there is a burgeoning artificial intelligence literature on the use of graphs to simultaneously model causal hypotheses and to encode the conditional independence relations implied by these hypotheses. Such graphs are called *probabilistic networks*.¹³ An important distinction in our work is that the IOD is derived from the primitives of a game and not from the properties of a single, arbitrarily-specified probability distribution. Thus, the information encoded in an IOD holds for all empirical distributions arising from play in the underlying game. Moreover, many of the results in the first part of the paper rely on the special structure implied by distributions of this kind and, as mentioned earlier, may not hold in a non-game-theoretic context.

Until recently, work on probabilistic networks in the artificial intelligence community focused upon the decision problem of a single individual. Thus, another aspect separates our work from the existing AI literature is its use of these objects in the solution of game-theoretic (i.e., interactive) decision problems. Two other papers, one by Koller and Milch (2002) and another by La Mura (2002) also use probabilistic networks to derive results of interest to game theorists along a different

line. Both of these papers develop alternative representations for interactive decision problems based upon probabilistic networks (i.e., as opposed to a game's strategic or normal form) and argue that these representations are not only computationally advantageous but also provide qualitative insight into the structural interdependencies between player decisions.

A AI Results

The next proposition is due to Verma and Pearl (1990; see also Geiger et al. 1990). It demonstrates that a directed, acyclic graph can be used to summarize the conditional independencies implied by a joint distribution on a finite set of random variables. Proposition 9 was originally proven for finite probability spaces. For the extension to the infinite case, see Cowell et al. (1999, p. 63).

Proposition 9 *Suppose $V = \{v_1, \dots, v_n\}$ is a finite collection of random variables on a probability space (Ω, \mathcal{F}, m) . Assume (V, \rightarrow) and $(V, \rightarrow)'$ are minimal directed, acyclic graphs such that $m = \prod_{j=1}^n m(v_j | pa_j)$. Then, $m = \prod_{j=1}^n m(v_j | pa'_j)$ if and only if $\mathcal{E} = \mathcal{E}'$ and $\mathbf{S} = \mathbf{S}'$.*

Corollary 4 *Suppose $V = \{v_1, \dots, v_n\}$ is a finite collection of random variables on a probability space (Ω, \mathcal{F}, m) . Assume (V, \rightarrow) and $(V, \rightarrow)'$ are directed, acyclic graphs such that $m = \prod_{j=1}^n m(v_j | pa_j)$. If (V, \rightarrow) is minimal, then $m = \prod_{j=1}^n m(v_j | pa'_j)$ if and only if $\mathcal{E} \subseteq \mathcal{E}'$ and $\mathbf{S} \subseteq \mathbf{S}'$.*

B Proofs of the propositions

B.1 Lemma 1

Part I (Equivalence relation) Reflexivity: Given Γ and the identity mappings $f(r) = r$ and $g(\theta) = \theta$ implies $\Gamma \preceq \Gamma$. Transitivity: Assume $\Gamma \sim \hat{\Gamma}$ and $\hat{\Gamma} \sim \Gamma'$. Suppose $\Gamma \preceq \hat{\Gamma}$ with permutation f and strategy mapping g , and $\hat{\Gamma} \preceq \Gamma'$ with permutation \hat{f} and mapping \hat{g} . Then, $\Gamma \preceq \Gamma'$ under $\tilde{f} \equiv f \circ \hat{f}$. and $\tilde{g} \equiv g \circ \hat{g}$. By similar reasoning, $\Gamma' \preceq \Gamma$. Therefore, $\Gamma \sim \Gamma'$. Symmetry: This is immediate from the definition. **Part II** $((f(\mathbf{A}'), f(\mathcal{A}')) = (\mathbf{A}, \mathcal{A}))$ This is immediate from $\Gamma \in \mathbb{O}_{\Gamma'}$ and $\Gamma' \in \mathbb{O}_{\Gamma}$. **Part III** Let g and \hat{g} be functions meeting the conditions of $\Gamma \preceq \hat{\Gamma}$ and $\hat{\Gamma} \preceq \Gamma$, respectively.

Define $\tilde{g} : \Sigma \rightrightarrows \hat{\Gamma}$ as follows:

$$\forall \boldsymbol{\theta} \in \Sigma, \tilde{g}(\boldsymbol{\theta}) \equiv \begin{cases} g(\boldsymbol{\theta}) & \text{if } \boldsymbol{\theta} \notin \hat{g}(\hat{\Sigma}) \\ \hat{g}^{-1}(\boldsymbol{\theta}) & \text{if } \boldsymbol{\theta} \in \hat{g}(\hat{\Sigma}) \end{cases}$$

Clearly, \tilde{g} is onto and has the desired property for $\boldsymbol{\theta} \notin \hat{g}(\hat{\Sigma})$. Suppose $\boldsymbol{\theta} \in \hat{g}(\hat{\Sigma})$. Then, for all $\hat{\boldsymbol{\theta}} \in \hat{g}^{-1}(\boldsymbol{\theta})$, $(\hat{\mathbf{A}}, \hat{\mathcal{A}}, \hat{m}_{\hat{\boldsymbol{\theta}}}) = (\hat{\mathbf{A}}, \hat{\mathcal{A}}, \hat{m}_{\boldsymbol{\theta}})$ by the definition of \hat{g} . By the equality of measurable spaces (Part II), it is also the case that $(\mathbf{A}, \mathcal{A}, m_{\hat{\boldsymbol{\theta}}}) = (\mathbf{A}, \mathcal{A}, m_{\boldsymbol{\theta}})$.

B.2 Lemma 2

1. Let $\mathbf{F}' \in \mathcal{I}_s$ be an arbitrary element of the partition of \mathbf{A} that generates \mathcal{I}_s . Recall, $\mathcal{I}_s \subseteq \sigma(\tilde{h}_s)$, so we can trivially construct an index set J , such that $\mathbf{F}' = \bigcup_{j \in J} \mathbf{h}_r^j \cap \mathbf{G}_r^j \cap \mathbf{G}_{r+}^j$ where $\mathbf{h}_r^j \in \sigma(\tilde{h}_r)$, $\mathbf{G}_r^j \in \sigma(\tilde{a}_r)$, $\mathbf{G}_{r+}^j \in \sigma(\tilde{a}_{r+1}, \dots, \tilde{a}_{s-1})$.
2. We claim that for all $\mathbf{F}' \in \mathcal{I}_s$ there exist an index set J , $\mathbf{h}_r^j \in \sigma(\tilde{h}_r)$ and $\mathbf{G}_{r+}^j \in \sigma(\tilde{a}_{r+1}, \dots, \tilde{a}_{s-1})$ such that

$$\mathbf{F}' = \bigcup_{j \in J} (\mathbf{h}_r^j \cap \tilde{a}_r^{-1}(\tilde{c}_r(\mathbf{h}_r^j)) \cap \mathbf{G}_{r+}^j)$$

Suppose not. Then, there exist \mathbf{h}_r , $B \subset \tilde{c}_r(\mathbf{h}_r)$, and \mathbf{G}_{r+} such that $B \neq \emptyset$, $B^c \neq \emptyset$, and

$$\begin{aligned} \mathbf{F}' &\supset \mathbf{h}_r \cap B \cap \mathbf{G}_{r+}, \\ \mathbf{F}' &\not\subset \mathbf{h}_r \cap B^c \cap \mathbf{G}_{r+}. \end{aligned}$$

But this clearly contradicts $r \leftrightarrow s$.

3. Using, items 1) and 2)

$$\begin{aligned} \mathbf{F}' &= \bigcup_{j \in J} (\mathbf{h}_r^j \cap \tilde{a}_r^{-1}(\tilde{c}_r(\mathbf{h}_r^j)) \cap \mathbf{G}_{r+}^j) \\ &= \bigcup_{j \in J} (\mathbf{h}_r^j \cap \mathbf{G}_{r+}^j) \end{aligned}$$

But, this implies that $\mathbf{F}' \in \sigma(\tilde{h}_{s \setminus r})$. Thus, $\mathcal{I}_s \subseteq \sigma(\tilde{h}_{s \setminus r})$.

B.3 Proposition 1

Given equation (3), equation (4) holds if, for all $\boldsymbol{\theta} \in \boldsymbol{\Sigma}$, $r \in T$,

$$m_{\boldsymbol{\theta}}(\tilde{a}_r | \tilde{\pi}_r) = m_{\boldsymbol{\theta}}(\tilde{a}_r | \mathcal{I}_r). \quad (7)$$

Recall, for all $\mathbf{F} \in \mathcal{A}$, $m_{\boldsymbol{\theta}}(\tilde{a}_r(\mathbf{F}) | \tilde{\pi}_r)$ is the conditional probability of $\tilde{a}_r^{-1}(F_r)$ given $\sigma(\tilde{\pi}_r)$ where $F_r \equiv \{a \in A_r | \exists \mathbf{a} \in \mathbf{F}, \tilde{a}_r(\mathbf{a}) = a\}$. Thus, the two conditions characterizing $m_{\boldsymbol{\theta}}(\tilde{a}_r | \tilde{\pi}_r)$ are: (i) $m_{\boldsymbol{\theta}}(\tilde{a}_r | \tilde{\pi}_r)$ is $\sigma(\tilde{\pi}_r)$ -measurable and (ii) for all $\mathbf{F} \in \mathcal{A}$, $\mathbf{G} \in \sigma(\tilde{\pi}_r)$,

$$\int_{\mathbf{G}} m_{\boldsymbol{\theta}}(\tilde{a}_r(\mathbf{F}) | \tilde{\pi}_r)(\mathbf{a}) m_{\boldsymbol{\theta}}(d\mathbf{a}) = m_{\boldsymbol{\theta}}(\tilde{a}_r^{-1}(F_r) \cap \mathbf{G}).$$

We need to demonstrate that $m_{\boldsymbol{\theta}}(\tilde{a}_r | \mathcal{I}_r)$ also satisfies these conditions. For all $r \in T$, Lemma 2 implies that $\mathbf{F} \in \mathcal{I}_r \Rightarrow \mathbf{F} \in \sigma(\tilde{a}_k)_{k \in \{s \in T | s \leftrightarrow r\}^c}$. Of course, $\sigma(\tilde{a}_k)_{k \in \{s \in T | s \leftrightarrow r\}^c} = \sigma(\tilde{\pi}_r)$, so $\mathcal{I}_r \subset \sigma(\tilde{\pi}_r)$. Therefore, $m_{\boldsymbol{\theta}}(\tilde{a}_r | \mathcal{I}_r)$ is $\sigma(\tilde{\pi}_r)$ -measurable. But this (and the definition of conditional probability) implies that, for all $\mathbf{F} \in \mathcal{A}$, $\mathbf{G} \in \sigma(\tilde{\pi}_r)$,

$$\int_{\mathbf{G}} m_{\boldsymbol{\theta}}(\tilde{a}_r(\mathbf{F}) | \mathcal{I}_r)(\mathbf{a}) m_{\boldsymbol{\theta}}(d\mathbf{a}) = m_{\boldsymbol{\theta}}(\tilde{a}_r^{-1}(F_r) \cap \mathbf{G}).$$

B.4 Proposition 2

The necessity of (6) follows from Corollary 1. Assume $\Gamma' \in \mathbb{O}_{\Gamma}$ and let f be the permutation guaranteed by this relationship. Consider an arbitrary $\boldsymbol{\theta} \in \boldsymbol{\Sigma}$. By the premise, $m_{\boldsymbol{\theta}} = \prod_{r \in T} m_{\boldsymbol{\theta}}(\tilde{a}_r | \tilde{\pi}'_r)$. By the definition of perfect observation, for all $r \in T$, $\mathcal{I}'_r = \sigma(\tilde{\pi}'_r)$. Therefore, for all $\boldsymbol{\theta} \in \boldsymbol{\Sigma}$, $r \in T$, choose $\boldsymbol{\theta}' \in \boldsymbol{\Sigma}'$ such that $\forall F \in \mathcal{A}_r$, $\theta'_{f(r)}(F | \tilde{h}_{f(r)}(f(\mathbf{a}))) = m_{\boldsymbol{\theta}}(F | \tilde{\pi}'_r)(\mathbf{a})$, $r \in T$.

B.5 Proposition 3

Let $\boldsymbol{\theta} \in \boldsymbol{\Sigma}$, $r \in T$ satisfy the premise. Since $m_{\boldsymbol{\theta}}(\tilde{a}_r | \tilde{\pi}_r)$ is not $\sigma(\tilde{\pi}'_r)$ -measurable, condition (6) fails and hence, by the contrapositive of Corollary 1, $\Gamma \not\leq \Gamma'$.

B.6 Proposition 4

Let \mathcal{C}_r denote the partition generating \mathcal{I}_r , $|G|$ denote the cardinality of the set G and, for all $s \in T$, let $\Pi_s \equiv \{r \in T | r \rightarrow s\}$. We prove Proposition 4 by constructing a strategy for the game and showing it is maximally revealing. We construct the strategy as follows:

1. For all s such that $|\Pi_s| = 0$, pick arbitrary $\mathbf{a} \in \mathbf{A}$ and let $\mathbf{H}_s = \{\mathbf{a}\}$ and $\mathbf{G}_s = \emptyset$.
2. Otherwise, by the definition of an IOD (Definition 2), for each $r \in \Pi_s$, there exist $\mathbf{F}_r \in \mathcal{C}_s$ and $(\mathbf{a}_r, \mathbf{a}'_r) \in \mathbf{F}_r \times \mathbf{F}_r^c$ such that $\tilde{h}_r(\mathbf{a}_r) = \tilde{h}_r(\mathbf{a}'_r)$, $\tilde{a}_r(\mathbf{a}_r) \neq \tilde{a}_r(\mathbf{a}'_r)$, $\tilde{a}_s(\mathbf{a}_r) \neq \tilde{a}_s(\mathbf{a}'_r)$, and $\tilde{a}_k(\mathbf{a}_r) = \tilde{a}_k(\mathbf{a}'_r)$ for $k \in \Pi_s \setminus \{r\}$. For each $r \in \Pi_s$, choose such a pair $(\mathbf{a}_r, \mathbf{a}'_r)$ and define $\mathbf{H}_s \equiv \bigcup_{r \in \Pi_s} \{\mathbf{a}_r, \mathbf{a}'_r\}$, and $\mathbf{G}_s \equiv \{\mathbf{G} \in \mathcal{C}_s | \mathbf{H}_s \cap \mathbf{G} \neq \emptyset\}$.
3. Let $\mathbf{H} \equiv \bigcup_{s=1}^t \mathbf{H}_s$. Note that \mathbf{H} is finite.
4. Define behavior strategies of the agent at move $s \in T$, θ_s , as follows:
 - (a) For each $\mathbf{F} \in \mathcal{C}_s$, $\mathbf{H} \cap \mathbf{F} \neq \emptyset$, let $\theta_s(a|\mathbf{F})$ be such that: (i) $\theta_s(a|\mathbf{F}) > 0$ if $a \in \tilde{a}_s(\mathbf{H} \cap \mathbf{F})$ and zero otherwise, and (ii) for any two distinct $\mathbf{F}, \mathbf{F}' \in \mathcal{C}_s$, $\mathbf{H} \cap \mathbf{F} \neq \emptyset$, $\mathbf{H} \cap \mathbf{F}' \neq \emptyset$, $\theta_s(a|\mathbf{F}) \neq \theta_s(b|\mathbf{F}')$ for any distinct $a, b \in \tilde{a}_s(\mathbf{H})$. If $\tilde{a}_s(\mathbf{H} \cap \mathbf{F})$ is a singleton, condition (ii) may require the need to find a second $a' \in \tilde{c}_s(\mathbf{F}) \setminus \tilde{a}_s(\mathbf{H})$ with which to construct a distinct $\theta_s(a|\mathbf{F})$ (as long as $|\tilde{c}_s(\mathbf{H})| > 1$). In such a case a' can be an arbitrary element of $c_s(\mathbf{F}) \setminus \tilde{a}_s(\mathbf{H})$.
 - (b) For each $\mathbf{F} \in \mathcal{C}_s$, $\mathbf{H} \cap \mathbf{F} = \emptyset$, define $\theta_s(a|\mathbf{F})$ such that it assigns full probability to some $a \in \tilde{c}_s(\mathbf{F})$.

Clearly, this procedure generates a strategy $\theta \in \Sigma$. Furthermore, $m_\theta(\mathbf{a}) > 0$ for all $\mathbf{a} \in \mathbf{H}$. Let $\mathbf{H}' = \{\mathbf{a} \in \mathbf{A} | m_\theta(\mathbf{a}) > 0\}$ so that $\mathbf{H} \subseteq \mathbf{H}'$ and by construction $|\mathbf{H}'| < \infty$. We need to confirm that this strategy satisfies the condition of being

maximally revealing; i.e., there does not exist $s \in T$ and $S \subset \{r \in T | r \rightarrow s\}$ such that $m_\theta(\tilde{a}_s | \tilde{\pi}_S) = m_\theta(\tilde{a}_s | \tilde{\pi}_s)$, where $\tilde{\pi}_S()$ is the projection of \mathbf{a} to the dimensions indexed by S . For all (r, s) such that $r \rightarrow s$, and all $\mathbf{a} \in \mathbf{A}$, define the events $\mathbf{G}(\mathbf{a})$, $\mathbf{F}_r^{\mathbf{a}|s}$, $\mathbf{F}_{-r}^{\mathbf{a}|s}$ as follows: $\mathbf{G}(\mathbf{a})$ is the intersection of \mathbf{H}' with the smallest set $B \in \mathcal{I}_s$ such that $\tilde{h}_{r+1}^{-1}(\tilde{h}_{r+1}(\mathbf{a})) \subseteq B$; $\mathbf{F}_r^{\mathbf{a}|s} \equiv \mathbf{H}' \cap_{\{j|j \rightarrow s\}} \tilde{a}_j^{-1}(\mathbf{a})$ and $\mathbf{F}_{-r}^{\mathbf{a}|s} \equiv \mathbf{H}' \cap_{\{j|j \neq r, j \rightarrow s\}} \tilde{a}_j^{-1}(\mathbf{a})$.

The strategy θ , as constructed above, is maximally revealing if there does not exist (r, s) such that $r \rightarrow s$ and for all $\mathbf{a} \in \mathbf{H}'$

$$\frac{m_\theta(\mathbf{F}_r^{\mathbf{a}|s} \cap \tilde{a}_s^{-1}(\mathbf{a}))}{m_\theta(\mathbf{F}_r^{\mathbf{a}|s})} = \frac{m_\theta(\mathbf{F}_{-r}^{\mathbf{a}|s} \cap \tilde{a}_s^{-1}(\mathbf{a}))}{m_\theta(\mathbf{F}_{-r}^{\mathbf{a}|s})} \quad (8)$$

By construction above, for all (r, s) there exist $\mathbf{a}_{rs}, \mathbf{a}'_{rs} \in \mathbf{H}$, such that $\mathbf{a}_{rs}, \mathbf{a}'_{rs}$ satisfy the conditions of definition 2, and $\mathbf{a}'_{rs} \notin \mathbf{G}(\mathbf{a}_{rs})$. Recall that we chose \mathbf{a}_{rs} and \mathbf{a}'_{rs} such that $\tilde{a}_k(\mathbf{a}_{rs}) = \tilde{a}_k(\mathbf{a}'_{rs})$ for all $k \in \Pi_s \setminus \{r\}$. Let $\tilde{a}_r(\mathbf{a}'_{rs}) = b$ and define $S = \Pi_s \setminus \{r\}$. By construction $\theta_s(b | \tilde{h}_s(\mathbf{a}_{rs})) \neq \theta_s(b | \tilde{h}_s(\mathbf{a}'_{rs}))$. This implies $m_\theta(b | \tilde{h}_s)(\mathbf{a}_{rs}) \neq m_\theta(b | \tilde{h}_s)(\mathbf{a}'_{rs})$ which is true if and only if

$$\frac{m_\theta(\mathbf{F}^{\mathbf{a}_{rs}|s} \cap \tilde{a}_s^{-1}(\mathbf{a}'_{rs}))}{m_\theta(\mathbf{F}^{\mathbf{a}_{rs}|s})} \neq \frac{m_\theta(\mathbf{F}^{\mathbf{a}'_{rs}|s} \cap \tilde{a}_s^{-1}(\mathbf{a}'_{rs}))}{m_\theta(\mathbf{F}^{\mathbf{a}'_{rs}|s})}.$$

Since $\mathbf{F}_{-r}^{\mathbf{a}_{rs}|s} = \mathbf{F}_{-r}^{\mathbf{a}'_{rs}|s}$, Equation (8) holds for, at most, one of $\{\mathbf{a}_{rs}, \mathbf{a}'_{rs}\}$.

B.7 Proposition 5

Once Proposition 2 is established, we need only show that condition (6) holds. Take an arbitrary $\hat{\theta} \in \Sigma$ and $r \in T$ and consider $\hat{\theta}_r(a_r | \mathcal{I}_r)$. By Proposition 1, $\hat{\theta}_r(a_r | \mathcal{I}_r) = \hat{\theta}_r(a_r | \tilde{\pi}_r)$. Using equation (2), we can write

$$\forall F \in \mathcal{A}_r, \hat{\theta}_r(F | \tilde{\pi}_r) = m_{\hat{\theta}}(F | \tilde{\pi}_r)$$

Since θ is maximally revealing for Γ , for all $r \in T$, there does not exist $S \subset \{k | k \rightarrow r\}$ such that $m_\theta(a_r | \tilde{\pi}_S) = m_\theta(a_r | \tilde{\pi}_r)$. This combined with $m_\theta(a_r | \tilde{\pi}_r) = m_\theta(a_r | \tilde{\pi}'_r)$ implies that $\sigma(\tilde{\pi}_r) \subseteq \sigma(\tilde{\pi}'_r)$. Hence, for all $F \in \mathcal{A}_r$, $m_{\hat{\theta}}(F | \tilde{\pi}_r) = m_{\hat{\theta}}(F | \tilde{\pi}'_r)$. This implies $m_{\hat{\theta}} = \prod_{r \in T} m_{\hat{\theta}}(\tilde{a}_r | \tilde{\pi}'_r)$.

B.8 Proposition 6

Let $v_j = \tilde{a}_j = \tilde{a}'_{f(j)}$ in Corollary 4.

(\Rightarrow) Given $\Gamma \in \mathbb{O}_{\Gamma'}$, pick a fully revealing $\theta \in \Sigma$. As (T, \rightarrow) is minimal apply Corollary 4 to get $\mathcal{E} \subseteq \mathcal{E}'$ and $\mathbf{S} \subseteq \mathbf{S}'$.

(\Leftarrow) Given $\Gamma \in \mathbb{O}_{\Gamma'}$, pick a fully revealing $\theta \in \Sigma$. As $\mathcal{E} \subseteq \mathcal{E}'$ and $\mathbf{S} \subseteq \mathbf{S}'$, and (T, \rightarrow) is minimal, apply Corollary 4 to obtain $m_\theta(v_j|\tilde{\pi}_j) = m_\theta(v_j|\tilde{\pi}'_j)$. Use this equation to define m_θ on (T, \rightarrow') . As Γ' is a game of perfect observation, the strategy associated with $m_\theta(\tilde{a}_r|\tilde{\pi}_r)$ is a valid strategy on Γ' . Define $g(\theta) \equiv \theta'(\tilde{a}'_r|\tilde{\pi}'_r) = \theta(\tilde{a}_r|\tilde{\pi}_r)$. As this procedure generates a valid strategy for a fully revealing strategy it also works for any arbitrary $\theta \in \Sigma$ (see proof of Prop 5).

B.9 Proposition 8

In a subjective game, each player knows his own consequences and payoffs. This implies, $\Omega_i = \{(\mathbf{C}, A_i) \mid \mathbf{C} \in \mathcal{C}_i\}$. Therefore, for all $\Gamma^k \in \Lambda_i$, $\phi_i|_{\Gamma^k} = \phi_i|_{\Gamma}$; so, we write ϕ_i without ambiguity. Let ϕ be a causal Nash Equilibrium in a subjective game supported by beliefs $(\hat{\mu}, \hat{\Theta})$. For all $i \in N$, for all $r \in T$, $a_{i,r} \in A_i$ and $\mathbf{C}_{r-1}, \mathbf{C}_r \in \mathcal{C}_i$, define,

$$\hat{e}_i|_{\mathbf{C}_{r-1}a_{i,r}}(\mathbf{C}_r) \equiv \sum_{\Gamma^k \in \Lambda_i} \hat{\mu}_i(\Gamma^k) \sum_{\mathbf{C}_r \in \mathcal{C}_i} m_{(\phi'_i, \theta_{-i}^k)}(\mathbf{C}_r|\mathbf{C}_{r-1})$$

where, as before, ϕ'_i is chosen such that $\mathbf{C}_{r-1} \cap \{\mathbf{a}' \in \mathbf{A} \mid a'_{i,r} = a_{i,r}\}$ occurs with positive probability. If ϕ_{-i} is such that \mathbf{C}_{r-1} is impossible, define $\hat{e}_i|_{\mathbf{C}_{r-1}a_{i,r}}$ arbitrarily. For all $\mathbf{C}_t \in \mathcal{C}_t$, define

$$m_{\phi_i, \hat{e}_i}(\mathbf{C}_t) = \sum_{a_{i,r} \in A_i} \hat{e}_i|_{\mathbf{C}_{t-1}a_{i,r}}(\mathbf{C}_t) \phi_i((\mathbf{C}_{t-1}, A_i))(a_{i,r}),$$

where, given the perfect recall assumption, $\mathbf{C}_{t-1} \in \mathcal{C}_i$ is the unique consequence such that $\mathbf{C}_{t-1} \supset \mathbf{C}_t$. Therefore,

$$E_v(\phi_i|\hat{\mu}_i, \hat{\Theta}_i) = \sum_{\mathbf{C} \in \mathcal{V}_i} v_i(\mathbf{C}) m_{\phi_i, \hat{e}_i}(\mathbf{C}) = E_v(\phi_i|\hat{e}_i).$$

Thus, $\phi_i \in BR(\hat{\mu}_i, \hat{\Theta}_i) \Rightarrow \phi_i \in BR(\hat{e}_i)$; i.e., item (1) of Definition 3 implies item (1) of Definition 4. Since, for all $\Gamma^k \in \Lambda_i$, $r = T$, $(\mathbf{C} \in \mathcal{C}_r^k) \Rightarrow (\mathbf{C} \in \mathcal{C}_i)$, item (2.i) of Definition 3 implies $m_{\phi_i, \hat{e}_i} = m_{\phi_i, e_i}$. Since, in a subjective game, player i knows v_i , condition 2.ii is automatically satisfied (i.e., for all subjective games). Thus, item (2) of Definition 3 implies item (2) of Definition 4. This proves $(\phi \in CNE_\Gamma) \Rightarrow (\phi \in SNE_\Gamma)$.

C Footnotes

1. Many real-world managerial situations, for example, appear to be characterized by this structure.
2. If our results are to be interpreted as relevant to situations in which players learn about their ability to influence others, it seems reasonable to assume that they do so in an environment in which such influence is a stationary aspect of the game.
3. These conditions are less restrictive than they may at first appear since players may make multiple moves and/or may be limited to a single, ‘null’ action at certain information sets (see, e.g., Elmes and Reny, 1994).
4. That is, they are finite, denumerable or isomorphic with the unit interval. In particular, this assumption implies the points in each set are measurable. The use of this word is due to Mackey (1957).
5. Both (1) and (2) follow from a standard result in probability theory. See, e.g., Fristedt and Gray (1997, p. 430-31).
6. Note that N has some hope of influencing II indirectly through I . Even so, II may choose to ignore the move of I (e.g., pick u at both information sets).
7. In what follows, keep in mind the distinction between probability measures on $(\mathbf{A}, \mathcal{A})$ induced by a strategy profile in the underlying game versus generic elements of the much larger space $\Delta(\mathbf{A}, \mathcal{A})$. Our results are critically dependant upon the structure implied by the former.
8. That is, if $\mathcal{I}'_{f(r)} \subset \mathcal{A}'$ is the information at move $f(r)$ in Γ' (where f is the permutation guaranteed by $\Gamma' \in \mathbb{O}_\Gamma$), then $\mathbf{F} \in \mathcal{I}'_r$ if and only if $f(\mathbf{F}) \in \mathcal{I}'_{f(r)}$.
9. Note that feasible action consistency is implied by $\Gamma' \in \mathbb{O}_\Gamma$.

10. We thank an anonymous referee for pointing out the relationship between Proposition 1 and an earlier result which enabled us to generalize it (the following proposition) and simplify the proof.
11. This last result raised a question that was put to us by E. Dekel in correspondence. Given the well-known works by Thompson (1952) and Elmes and Reny (1994) that identify transformations on extensive form games that yield the equivalence class of games with the same strategic form, is there a set of operations, similar to these in spirit, that yield games with equivalent IODs (in the sense of Corollary 3)? Due to space limitations, we do not provide a formal reply. Clearly, however, Corollary 3 does suggest a step-wise transformation that will yield empirically equivalent extensive forms with different IODs. The transformation, while difficult to formalize in the context of an extensive form game, is easy to describe: it is the transformation that flips an “allowed” arrow (per Corollary 3) in the original IOD.
12. A few of the more important contributions include Battigalli and Guatoli’s (1988) conjectural equilibrium, Fudenberg and Levine’s (1993) self-confirming equilibrium and, of course, Kalai and Lehrer’s (1993, 1995) subjective equilibrium. For related work, see Abreu, Pearce, and Stacchetti (1990), Blume and Easley (1992), Brandenburger and Dekel (1993), Geanakoplos (1994), Milgrom and Roberts (1990), and Nachbar (1996). Howitt and McAfee (1992) employ a similar idea in a macroeconomic application.
13. Cowell et al. (1999), Jensen (2001) and Pearl (1988, 2000) provide nice introductions for those interested in pursuing this material further.

References

- [1] Armbruster, W., Boge, W., 1979. Bayesian game theory, in: Moeschlin, O., Pallaschke, D. (Eds.), *Game Theory and Related Topics*. North-Holland, Amsterdam, pp. 17-28.
- [2] Abreu, D., et al., 1987. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58, 1041-63.
- [3] Aumann, R. J., 1964. Mixed and behavior strategies in infinite extensive games, in: Dresher, O., Shapley, L. S., Tucker, A. W. (Eds.), *Advances in Game Theory*. Princeton University Press, Princeton, pp. 627-650.
- [4] Battigalli, P., Guatoli, D., 1988. Conjectural equilibria and rationalizability in a macroeconomic game with incomplete information. *Instituto Economia Politica*, Milan.
- [5] Blume, L. E., Easley, D., 1992. *Rational expectations and rational learning*. Cornell University, Cornell.
- [6] Bollobas, B., 1998. *Modern Graph Theory*. Springer, New York.
- [7] Brandenburger, A., Dekel, E., 1993. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59, 189-98.
- [8] Cowell, R. G., Dawid, A. P., Lauritzen, S. L., Spiegelhalter, D. J., 1999. *Probabilistic Networks and Expert Systems*. Springer, New York.
- [9] Elmes, S., Reny, P., 1994. On the strategic equivalence of extensive form games. *Journal of Economic Theory*, 62(1), 1-23.
- [10] Fristedt, B., Gray, L., 1997. *A Modern Approach to Probability Theory*. Birkhäuser, Boston.

- [11] Fudenberg, D., Levine, D. K., 1993. Self-confirming equilibrium. *Econometrica*, 61, 523-45.
- [12] ——— 1998. *Theory of Learning in Games*. MIT Press, Cambridge.
- [13] Geanakoplos, J., 1994. Common Knowledge. *Handbook of Game Theory*. Elsevier.
- [14] Geiger, D., Verma, T. S., J. Pearl, J., 1990. Identifying independence in Bayesian networks. *Networks*, 20, 507-34.
- [15] Harsanyi, J., 1967-68. Games with incomplete information played by ‘Bayesian’ players,” *Management Science*, 8, Parts I-III, 159-82, 320-34, 486-502.
- [16] Howitt, P., McAfee, R. P., 1992. Animal Spirits. *American Economic Review*, 82, 493-507.
- [17] Kalai, E., Lehrer, E., 1993. Subjective equilibrium in repeated games. *Econometrica*, 61, 1231-1240.
- [18] ——— 1995. Subjective games and equilibria. *Games and Economic Behavior*, 8, 123-63.
- [19] Koller, D. Milch, B., 2002. Multi-agent influence diagrams for representing and solving games. *Seventeenth International Joint Conference on Artificial Intelligence*, Seattle, WA.
- [20] La Mura, P., 2002. *Game Networks*. Unpublished manuscript.
- [21] Mackey, G. W., 1957. Borel structures in groups and their duals. *Trans. Amer. Math. Soc.*, 85, 283-95.
- [22] Milgrom, P. Roberts, J., 1990. Rationalizability, learning and equilibrium in games with strategic complementarities. *Econometrica*, 58, 1255-1277.

- [23] Nachbar, J. H., 1997. Prediction, optimization and learning in repeated games. *Econometrica*, 65, 275-309.
- [24] Shafer, G., 1996. *The Art of Causal Conjecture*. Cambridge University Press, Cambridge.
- [25] Thompson, F. B., 1952. Equivalence of games in extensive form. The Rand Corporation, Research Memorandum 759.
- [26] Verma, T. S., Pearl, J., 1990. Equivalence and synthesis of causal models, in: *Proceedings of the 6th Conference on Uncertainty in Artificial Intelligence*. Cambridge, pp. 220-7. Reprinted in: Bonissone, P., Henrion, M., Kanal, L. N., Lemmer, J. F. (Eds.), *Uncertainty in Artificial Intelligence*, vol. 6, 255-68.

D Diagrams

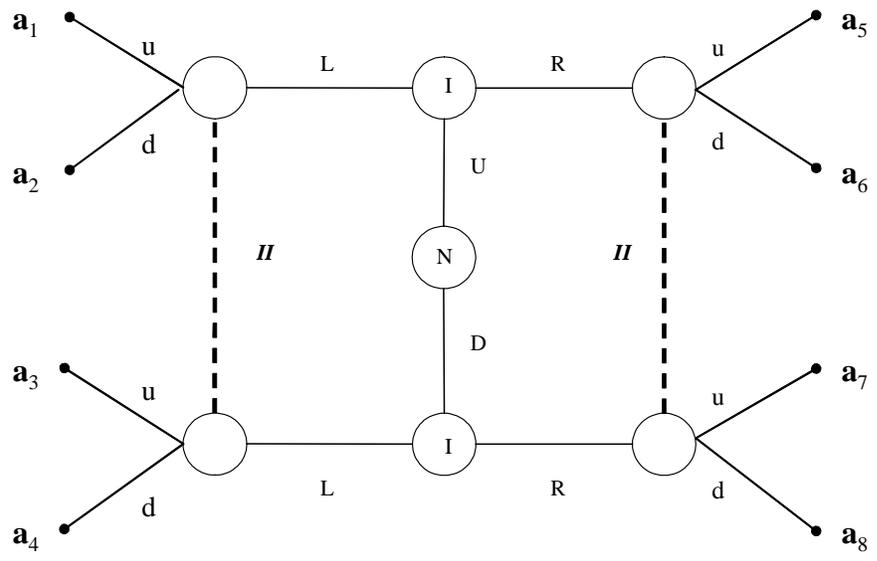


Figure 1: game Γ^A

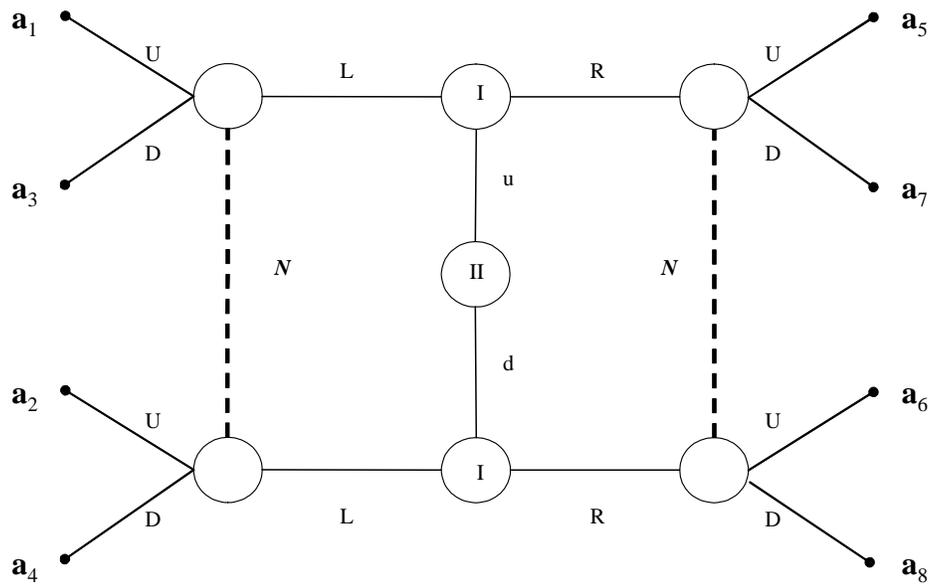


Figure 2: game Γ^B .

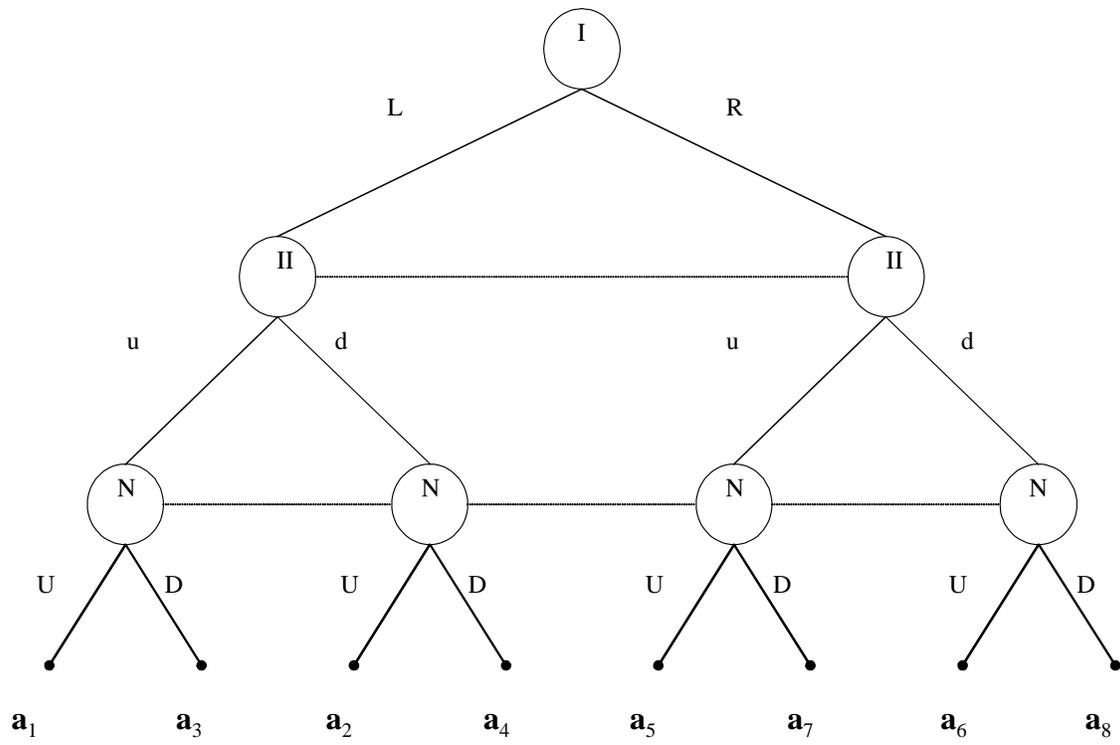


Figure 3: game Γ^C .

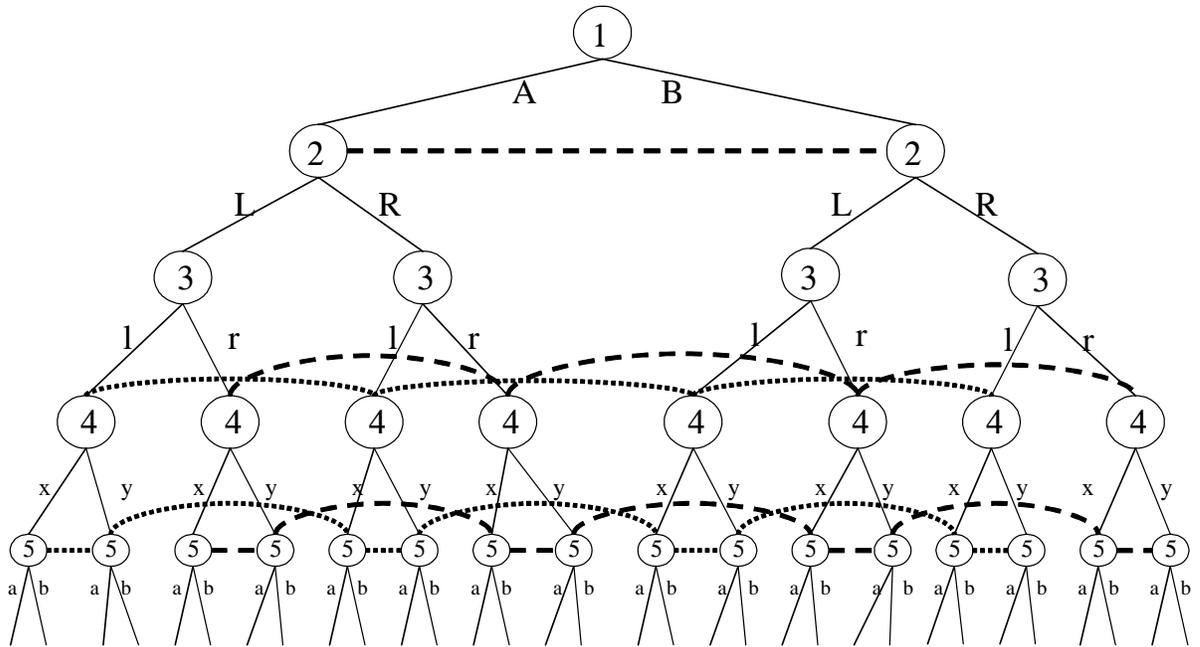


Figure 4: the Gatekeeper game.

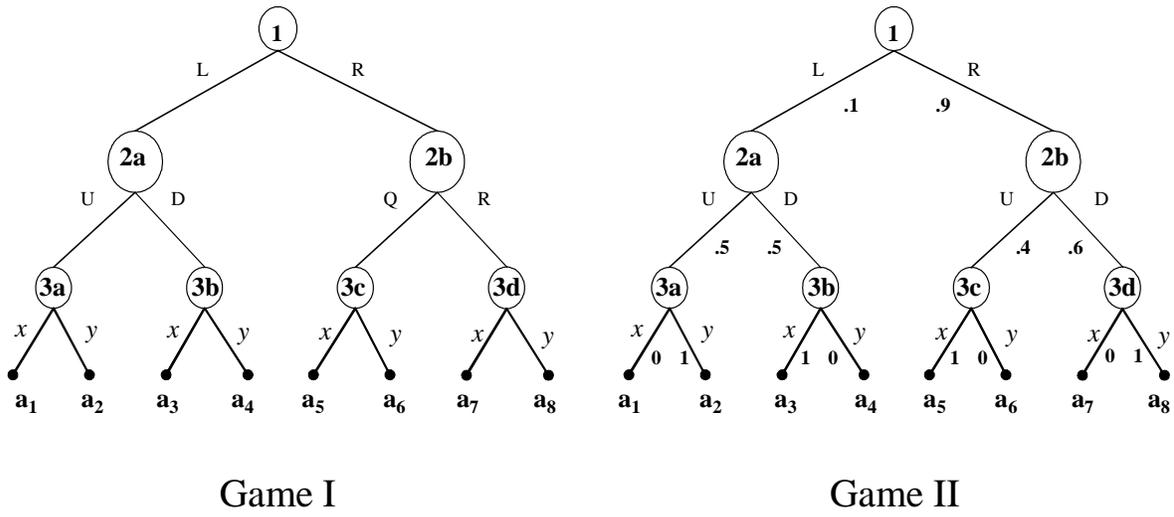


Figure 5: illustration of IOD condition (4).

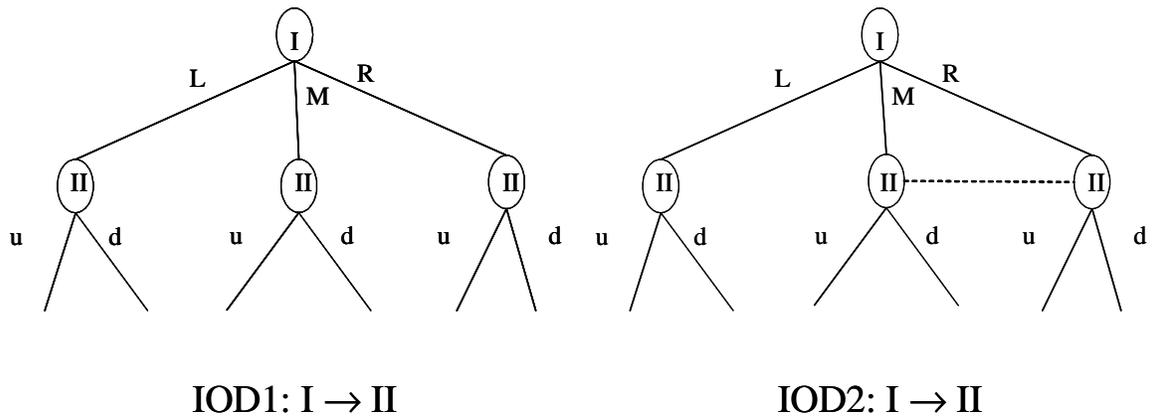


Figure 6: necessity of condition (6).

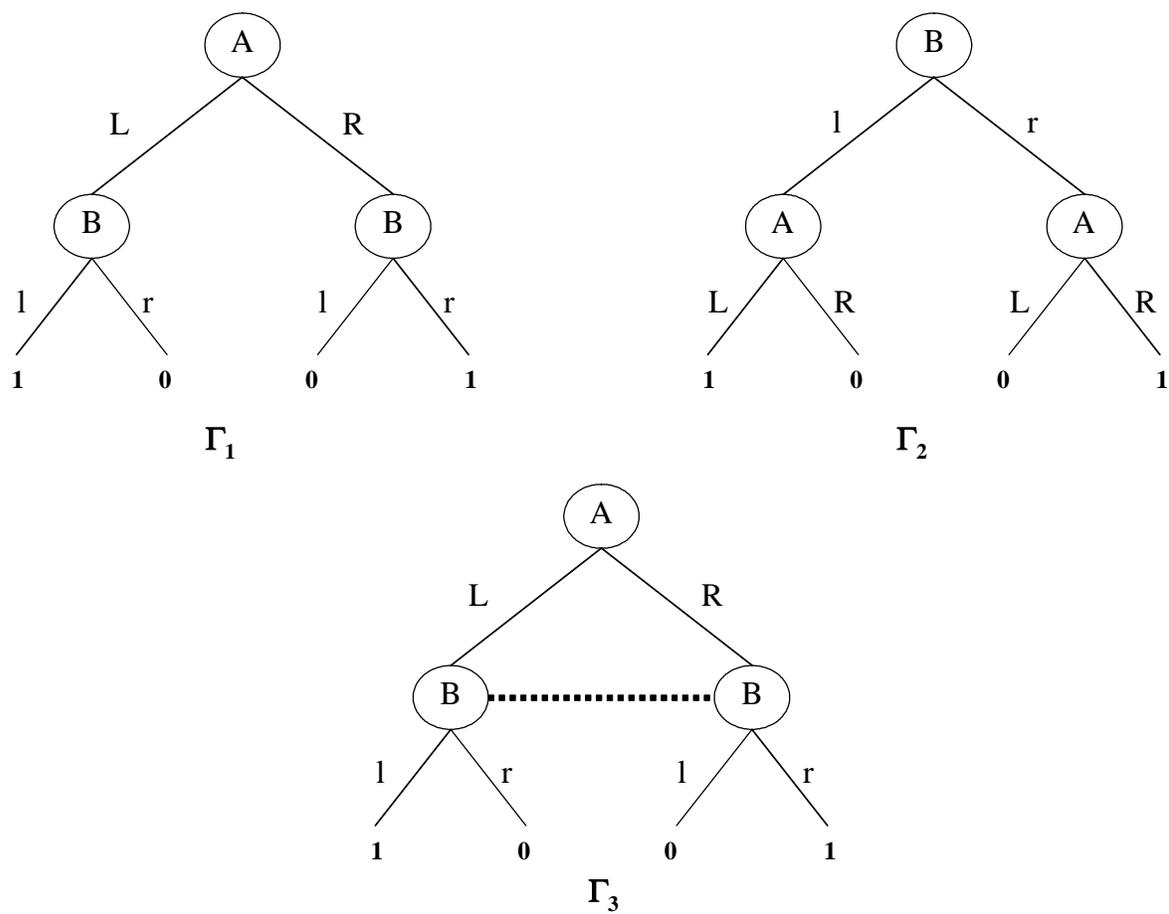


Figure 7: possible departmental subgames.