

Centre de Referència en Economia Analítica

Barcelona Economics Working Paper Series

Working Paper n° 27

**Causal Assessment in Finite-Length
Extensive-Form Games**

Jose S. Penalva Zuasti and Michael D. Ryall

September 2003

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

Barcelona Economics WP No 27

September 2003

*We thank Adam Brandenburger, Colin Camerer, David M. Chickering, Eddie Dekel, Bryan Ellickson, Ehud Kalai, Uffe Kjærulff, David Levine, Steven Lippman, Glenn MacDonald, Mark Machina, Leslie Marx, Joseph Ostroy, Scott Page, Judea Pearl, Larry Samuelson, Joel Watson, Eduardo Zambrano, Bill Zame, participants in the MEDS, Georgetown, UCLA, and UCSD theory workshops, the Stanford SITE and Stonybrook theory conferences, three anonymous referees and an anonymous co-editor, 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 consider extensive-form games in which the information structure is not known and ask how much of that structure can be inferred from the distribution on action profiles generated by player strategies. One game is said to *observationally imitate* another when the distribution on action profiles generated by every behavior strategy in the latter can also be generated by an appropriately chosen behavior strategy in the former. The first part of the paper develops analytical methods for testing observational imitation. The central idea is to relate a game's information structure to the conditional independencies in the distributions it generates on action profiles. We present a new analytical device, the *influence opportunity diagram of a game*, describe how such a diagram is constructed for a given finite-length extensive-form game, and demonstrate that it provides, for a large class of economically interesting games, a simple test for observational imitation. The second part of the paper shifts the focus to the influence assessments of players within a game. 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: causality, information structure, extensive form game, observational imitation, Bayesian network

JEL classification: C72

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 social scientist, 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 *observational imitation*. One game is said to observationally imitate another when the empirical distribution induced by any behavior strategy profile in the latter can also be induced by an appropriately chosen behavior strategy profile in the former. Thus, even under infinite repetition, it is impossible to distinguish a game from its imitators solely on the basis of the observed behavior of its players.

Our analysis is facilitated by the introduction of a new graphical device, the

influence opportunity diagram of a game (hereafter, IOD). The IOD is defined constructively for a broad class of finite-length extensive-form games, including those with infinite action sets. As we show, the IOD summarizes certain information about the conditional independencies that must be observed in all empirical distributions arising from play of the underlying game. This feature allows us to apply some results from the extensive literature on probabilistic networks in artificial intelligence to address the issues raised above. This literature explores the use of graphs to model uncertainty and decisions in complex domains. Since most economists are not familiar with it, we have included a condensed, self-contained discussion of the relevant results and references in Appendix A.

A basic finding is that a necessary requirement for one game to imitate another is that their IODs imply a consistent set of conditional independencies. This requirement is not, in general, sufficient because differences in the specific information upon which players condition their behavior may imply additional restrictions on empirical distributions that are not picked up by the IOD. However, we do identify a broad class of games, termed games of *exact information*, for which this condition is also sufficient. For games of this type, we use results from the literature on probabilistic networks and develop new ones to show that observational imitation can be identified by simple visual comparison of the IODs of the games in question.

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 observed behavior. 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 concepts.

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, 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 observational imitation. Section 3 lays out the definition of a finite-length extensive-form game (which differs in some ways from the usual setup) and defines observational imitation and observational indistinguishability. In Section 4, we present our main results regarding the analysis of observational imitation (from the perspective of an outside observer). Section 4.1 shows how to construct an IOD from an extensive-form game. Section 4.2 connects information structure to observational imitation through the IOD. Section 4.3 shows how to test whether one game observationally imitates another 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 two 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

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

Of course, m_{θ^B} and m_{θ^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 .

On the other hand, it should be clear that Γ^C cannot imitate all strategies from 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 . For example, any θ^A in which player I conditions her behavior on Nature's play results in a m_{θ^A} that cannot be duplicated 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 whose

members are ordered profiles are indicated by bold capitals (\mathbf{A}, \mathbf{E}) with small bold (\mathbf{a}, \mathbf{e}) denoting typical elements. 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 . Let X_r denote the set of all nodes associated with the r^{th} move. 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* $\mathbf{a}_z = (a_{1z}, \dots, a_{tz})$. 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 (complete) *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* $\tilde{c}_r : \mathbf{A} \rightarrow \mathcal{A}_r$. Assume that feasible action sets are equal for all nodes 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 $\tilde{h}_r(\mathbf{a})$. 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 $\theta_i \equiv (\theta_r)_{r \in o^{-1}(i)}$. Σ_i is the set of all behavior strategies available to i . A *strategy profile* is an element $\theta \in \Sigma \equiv \times_{i \in N} \Sigma_i$. When convenient, we use the familiar shorthand $\theta = (\theta_i, \theta_{-i})$.

3.2 Empirical distribution

Given a game meeting the conditions of the previous section, every behavior strategy profile θ induces a probability space, denoted $(\mathbf{A}, \mathcal{A}, m_\theta)$. The measure m_θ can be

constructed directly from θ as follows: for all $\mathbf{F} \in \mathcal{A}$,

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

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 $\mathbf{a} \in \mathbf{A}$, $F_r \in \mathcal{A}_r$,

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

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

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

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

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 can 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 Observational imitation and indistinguishability

A game Γ' is said to observationally imitate 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., an outcome is 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' observationally imitates 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 observational imitation 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 . Note that $\Gamma' \in \mathbb{O}_\Gamma$ implies $\Gamma \in \mathbb{O}_{\Gamma'}$. 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 Γ' observationally imitates 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 *observationally indistinguishable*, denoted $\Gamma \sim \Gamma'$. The interpretation is that when Γ' and Γ are indistinguishable, any behavior observed under Γ (“observed” in the sense of knowing m_θ) could also be observed under Γ' and vice versa. When $\Gamma \sim \Gamma'$, Γ differs from Γ' in terms of its information and, possibly, payoff structures. Note that observational imitation 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 observational imitation defined only for specific (e.g., equilibrium) profiles. Hereafter, we drop the “observationally” and simply say that one game imitates another or that two games are indistinguishable.

Lemma 1 *Indistinguishability is an equivalence relation on the space of finite-length extensive form games.*

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 $\theta^A \in \Sigma^A$,

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

where, $\mathcal{I}_N^A = \{\emptyset, \mathbf{A}\}$, $\mathcal{I}_I^A = \{\emptyset, \mathbf{A}_U, \mathbf{A}_D, \mathbf{A}\}$ and $\mathcal{I}_{II}^A = \{\emptyset, \mathbf{A}_L, \mathbf{A}_R, \mathbf{A}\}$.

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

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

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 imitation between games

In this section we analyze imitation from the perspective of an outside observer who knows the empirical distribution generated by an arbitrary behavior strategy profile in a game with unknown information structure. To what extent does this knowledge illuminate the game's underlying information structure? To check a candidate game, one might attempt 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, five 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.

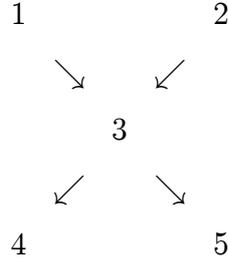
Definition 2 *The influence opportunity diagram of Γ is a graph (T, \rightarrow) such that $r \rightarrow s$ if and only if $r < s$, and there exist $\mathbf{a}, \mathbf{a}' \in \mathbf{A}$ satisfying each of the following conditions: (1) $\tilde{h}_r(\mathbf{a}) = \tilde{h}_r(\mathbf{a}')$; (2) $\exists \mathbf{F} \in \mathcal{I}_s$ 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) $\tilde{a}_I(\mathbf{a}_5) = (R) \neq \tilde{a}_I(\mathbf{a}_4) = (L)$. Item (4) is automatically satisfied since there are no moves between I 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 IOD 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 $1 \rightarrow 3$.

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 IOD (T, \rightarrow) ,*

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

4.2 Imitation and independence in probability

In the following sections we compare outcome compatible games. In order to reduce the notational burden throughout the rest of the paper, when comparing two games Γ and Γ' , $\Gamma' \in \mathbb{O}_\Gamma$, we drop the reference to the permutation function f . Thus, when comparing two games, the “ r^{th} move” refers to the same player and same set of actions in both games, with the order in Γ as the point of reference. This convention will apply to the IOD generated by Γ' too, so that $r \rightarrow' s$ in Γ' means that – in Γ' – the agent who has the r^{th} move in Γ has the potential to influence the agent who has the s^{th} move in Γ .⁸

If $\Gamma \preceq \Gamma'$, then 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, per our notational convention, $\mathcal{I}'_r \subset \mathcal{A}$ corresponds to the information at move $f(r)$ in Γ'). In other words, testing whether Γ' imitates 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 $\Gamma_2 \in \mathbb{O}_{\Gamma_1}$. 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 *exact information* if, for all $r \in T$, $\mathcal{I}_r = \sigma(\tilde{\pi}_r)$; a game of exact information 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. Note that games of perfect information are games of exact information in which *every* player exactly observes the move of *every* predecessor.

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

4.3 Testing for imitation

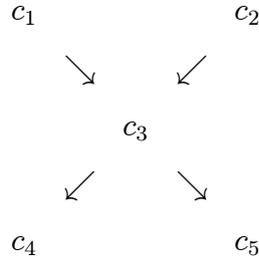
Suppose, given a game Γ and another Γ' of exact information, one wishes to determine whether $\Gamma \preceq \Gamma'$. One might be tempted to use Proposition 2 for this purpose. However, this is not practical for complex games since condition (6) must be checked for all strategy profiles. In this section, we develop results that permit this determination by simple visual inspection of the games' IODs.

For the following proposition, given an IOD (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 3 *Given two games Γ and Γ' with $\Gamma \in \mathbb{O}_{\Gamma'}$ and Γ' a game of exact information, if (T, \rightarrow) and (T, \rightarrow') are such that $\mathcal{E} = \mathcal{E}'$ and $\mathbf{S} = \mathbf{S}'$, then $\Gamma \preceq \Gamma'$.*

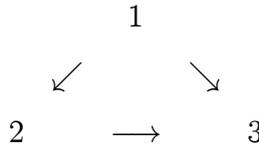
Corollary 2 *Assume Γ and Γ' are both games of exact information with $\Gamma \in \mathbb{O}_{\Gamma'}$, $\mathcal{E} = \mathcal{E}'$, and $\mathbf{S} = \mathbf{S}'$. Then, $\Gamma \sim \Gamma'$.*

To see how these results are used, consider once again games Γ^A and Γ^B in Section 2. These are both games of exact information. The IODs are $N \rightarrow I \rightarrow II$ and $N \leftarrow I \leftarrow II$, respectively. By Corollary 2, it is virtually immediate that these games are indistinguishable: $\mathcal{E}^A = \mathcal{E}^B = \{\{N, I\}, \{I, II\}\}$ and $\mathbf{S}^A = \mathbf{S}^B = \emptyset$. Recall that the IOD for the Gatekeeper game is:



Since this is a game of exact information, Corollary 2 tells us that there are no other IODs from which it is indistinguishable. An indistinguishable game requires an IOD with 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.¹⁰

The requirement that elements of \mathbf{S} not include convergent arrows with adjacent 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 below. Namely, all exact information, outcome compatible games whose IODs are a variation of the fully connected graph



are indistinguishable.

We want to extend Proposition 3 to allow us to establish imitation more generally. One approach is to use the “ d -separation” criterion for graphs defined by Pearl (1986); see the discussion in Appendix A and Proposition 7 in Appendix B. Here, we show that a very simple graph operation, edge deletion, together with the concepts just used in Proposition 3, suffice to determine imitability more generally. Let $rem(T, \rightarrow)$ be the set of IODs obtained from (T, \rightarrow) by the operation of edge deletion, and let $eq(T, \rightarrow)$ be the set of IODs that satisfy the second condition of Proposition 3; i.e., $(T, \rightarrow') \in eq(T, \rightarrow)$ if and only if $\mathcal{E} = \mathcal{E}'$ and $\mathbf{S} = \mathbf{S}'$.

Proposition 4 *Assume $\Gamma \in \mathbb{O}_{\Gamma'}$ and Γ' is a game of exact information. If $(T, \rightarrow) \in rem(T, \rightarrow^*)$ for some $(T, \rightarrow^*) \in eq(T, \rightarrow')$ then $\Gamma \preceq \Gamma'$.*

That is, when comparing two compatible games, Γ and Γ' , the latter of exact information, $\Gamma \preceq \Gamma'$ if the IOD of Γ can be obtained by edge deletion from any IOD (T, \rightarrow^*) in which $\mathcal{E}^* = \mathcal{E}'$ and $\mathbf{S}^* = \mathbf{S}'$. Finally, a much stronger result holds for games of perfect information. Specifically, a game of perfect information imitates *every* game with which it is outcome compatible.

Corollary 3 *Let Γ' be a game of perfect information. For every $\Gamma \in \mathbb{O}_{\Gamma'}$, $\Gamma \preceq \Gamma'$. If Γ is also a game of perfect information, $\Gamma \sim \Gamma'$.*

5 Influence from the player’s perspective

In this section, we consider the implications of observational imitation 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 case is when payoffs are correlated with game structure. In such cases, 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 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 random 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 the 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 2. 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 17, since Stackleberg is a game of exact information, this is indistinguishable from the stage game in which the entrant is the leader. Specifically, suppose the challenger has initial prior $\mu \in [0, 1]$ that Γ_1 is the true stage game. If $\mu \geq \frac{1}{2}$, 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. Apparently, some solution concept other than Nash or BNE is required for the outcome suggested by the preceding example. 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 have beliefs about the game's true influence relationships, they choose strategies that are optimal with respect to these beliefs and, as play unfolds, observe nothing that refutes them. 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 $\mathcal{C}_r^k \in \mathcal{A}$ denote the partition of \mathbf{A} that corresponds to player $o(r)$'s move- r information sets in Γ^k . We assume that players observe their own payoffs at the conclusion of play, which will require a consistency condition in our equilibrium definition. Therefore, let $\mathcal{C}_{i,t}^k$ be the partition of \mathbf{A} that corresponds to what player i learns at the conclusion of play; e.g., this may equal the partition implied by v_i^k , player i 's payoff in Γ^k .

Given $\hat{\mu}_i$, the set of all events that i believes could be observed during play is $\mathcal{C}_i \equiv \cup_{\Gamma^k \in \Lambda_i} \cup_{r \in \{o^{-1}(i), t\}} \mathcal{C}_r^k$. \mathcal{C}_i is termed *i 's set of consequences under $\hat{\mu}_i$* . Note that $\mathcal{C}_i \subset \mathcal{A}$ and, since N may also contain a nature player, this formulation allows for the inclusion of a rich set of environmental observables as well as partial-to-full knowledge of competitor actions. It should also be pointed out that $\Gamma^k \in \Lambda_i$ may be the extensive form of a finitely repeated stage game. For each $\mathbf{C} \in \mathcal{C}_i$, there is a corresponding set of feasible actions for player i , denoted $A_{\mathbf{C}}$ (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 during play is equivalent to being told $(\mathbf{C}, A_{\mathbf{C}})$.

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 true underlying 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 $(\mathbf{C}_L, \{u, d\})$, $(\mathbf{C}_R, \{u, d\})$, or $(\mathbf{C}_{\emptyset}, \{u, d\})$.

A *subjective behavior strategy for player i given μ_i* is a function ϕ_i such that, for all $\mathbf{C} \in \mathcal{C}_i$, $\phi_i(a|\mathbf{C})$ is the probability that $a \in A_{\mathbf{C}}$ is played given i 's arrival at the information set corresponding to \mathbf{C} . It is easy to see that ϕ_i restricted to the information sets of a particular game, e.g. Γ^k , corresponds to a unique behavior strategy for i in that game, written $\phi_i^k \in \Sigma_i^k$. Thus, given a game Γ^k , a profile of subjective strategies $\phi = (\phi_1, \dots, \phi_n)$ implies a probability space $(A, \mathcal{A}, m_{\phi}^k)$ where m_{ϕ}^k is the measure induced by $(\phi_1^k, \dots, \phi_n^k) \in \Sigma^k$. When referring explicitly to the true game, Γ , we simply write m_{ϕ} to refer to the actual measure on (A, \mathcal{A}) induced by ϕ .

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

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

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

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

where Φ_i is the set of all subjective strategies for i (under $\hat{\mu}_i$). 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. Lastly, for the upcoming influence-consistency condition, let Ξ_Γ denote the set of games that imitate 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 $\mathbf{B} \in \sigma(\mathcal{C}_i)$, $\Gamma^k \in \Lambda_i$, $m_\phi(\mathbf{B}) = m_{\phi_i, \hat{\phi}_{-i}}^k(\mathbf{B})$, (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. 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 that imitate Γ . The interpretation of this is that, as players grope their way toward equilibrium during the (unmodelled) pre-equilibrium 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 with positive probability in equilibrium.

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 $i \in N$, $\hat{\Theta}_i = \phi_{-i}$. 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 $\hat{\Theta}_i = \phi_{-i}$. So, $BNE_\Gamma \subseteq SCE_\Gamma$. Summing up:

Proposition 5 *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), hereafter KL, present the notion of a “subjective game” and a corresponding definition of subjective Nash equilibrium (SNE). In this formulation, each player chooses a best response to his “environment response function,” a mapping from his available actions to probability distributions on the consequences he experiences as a result of those actions. We wish to show that CNE is a refinement of SNE. In order to make the comparison formal we must introduce some new concepts and the corresponding notation. As in KL, we now restrict attention to games with countable action sets¹¹.

From player i 's perspective, upon reaching the information set corresponding to $\mathbf{C} \in \mathcal{C}_i$, player i chooses some action $a_{\mathbf{C}} \in A_{\mathbf{C}}$ and then, depending upon the true game and the true strategies of his opponents, observes some new consequence $\mathbf{C}' \in \mathcal{C}_i$, and so on until the conclusion of play. This process can be summarized by an *environment response function* for player i , a device that summarizes his individual decision problem. Formally, for a given game Γ and opponent strategies ϕ_{-i} ,

$e_i|_{\mathbf{C},a_{\mathbf{C}}}(\mathbf{C}')$ denotes the true probability of player i observing consequence \mathbf{C}' given his current information \mathbf{C} and his play of action $a_{\mathbf{C}}$. The computation of i 's environment response function is straightforward: for all $\mathbf{C}, \mathbf{C}' \in \mathcal{C}_i, a_{\mathbf{C}} \in A_{\mathbf{C}}$,

$$e_i|_{\mathbf{C},a_{\mathbf{C}}}(\mathbf{C}') \equiv m_{\phi}(\mathbf{C}'|\mathbf{C})$$

where ϕ_i is any strategy that chooses $a_{\mathbf{C}}$ with probability one conditional on reaching \mathbf{C} and does not make \mathbf{C} impossible. If ϕ_{-i} is such that \mathbf{C} is impossible no matter what strategy i chooses, then $e_i|_{\mathbf{C},a_{\mathbf{C}}}$ 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.

Players may not know their true environment response function. Instead, player i assesses e_i by a *subjective* environment response function \hat{e}_i . That is, $\hat{e}_i|_{\mathbf{C},a}(\mathbf{C}')$ is i 's subjective assessment that \mathbf{C}' occurs after having been told \mathbf{C} and having taken action $a \in A_{\mathbf{C}}$. Then, m_{θ_i, \hat{e}_i} represents i 's beliefs on observable events given his choice of ϕ_i and his assessment \hat{e}_i . Specifically, given beliefs $(\hat{\mu}_i, \hat{\Theta}_i)$

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

where ϕ_i is any strategy that chooses $a_{\mathbf{C}}$ with probability one conditional on reaching \mathbf{C} and that does not make \mathbf{C} impossible.

The last piece of the analysis is to define expected payoffs and best responses given ϕ_{-i} . Since all games in Λ_i are of perfect recall, \mathcal{C}_i can be partitioned into subsets ordered by the period in which their elements are observed by i . Suppose i moves k times in the true game Γ . Then, under the assumption that $\Lambda_i \subset \mathbb{O}_{\Gamma}$, i moves k times in every game contained in Λ_i . Let $\mathcal{C}_{i,0}, \mathcal{C}_{i,1}, \dots, \mathcal{C}_{i,t-1}, \mathcal{C}_{i,t}$ be the sets that partition \mathcal{C}_i in this fashion. So, $\mathcal{C}_{i,r}$ is the set of consequences i thinks could be reported at the start of his r^{th} turn. Then, the probability measure on \mathcal{C}_i implied by ϕ_i and e_i given ϕ_{-i} can be constructed inductively: for all $\mathbf{C}' \in \mathcal{C}_{i,1}$, define

$$m_{\phi_i, e_i}(\mathbf{C}') \equiv \sum_{\mathbf{C} \in \mathcal{C}_{i,0}} \sum_{a \in A_{\mathbf{C}}} e_i|_{\mathbf{C},a}(\mathbf{C}') \phi_i(a|\mathbf{C}) m_{\phi}(\mathbf{C}),$$

for all $\mathbf{C}'' \in \mathcal{C}_{i,2}$,

$$m_{\phi_i, e_i}(\mathbf{C}'') \equiv \sum_{\mathbf{C}' \in \mathcal{C}_{i,1}} \sum_{a \in A_{\mathbf{C}'}} e_i |_{\mathbf{C}, a}(\mathbf{C}'') \phi_i(a | \mathbf{C}') m_{\phi_i, e_i}(\mathbf{C}'),$$

and so on. The expected payoff to a strategy ϕ_i given e_i is

$$E_v(\phi_i | e_i) \equiv \sum_{\mathbf{C} \in \mathcal{C}_{i,t}} v_i(\mathbf{C}) m_{\phi_i, e_i}(\mathbf{C}).$$

Given an e_i , the best-response correspondence is

$$BR(e_i) \equiv \{\phi_i \in \Phi_i \mid \forall \phi'_i \in \Phi_i, E_v(\phi_i | e_i) \geq E_v(\phi'_i | e_i)\}.$$

Using \hat{e}_i , we define m_{ϕ_i, \hat{e}_i} , $E_v(\phi_i | \hat{e}_i)$ and $BR(\hat{e}_i)$ in the obvious way.

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

Proposition 6 *Given a game Γ with \mathbf{A} countable, $CNE_{\Gamma} \subseteq SNE_{\Gamma}$.*

The proof of this is almost immediate. Items (1) and (2) in Definition 4 are implied by items (1) and (2) in Definition 3 (player's payoff information is included in the description of the consequences in a subjective game). So, the only difference is that CNE has the learned influence requirement, item (3), that is not imposed in SNE.

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 take 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 observationally indistinguishable. 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 2. 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 observational imitation is new, the idea of observationally indistinguishable *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 indistinguishable from its reduced normal form. Thus, any two games with the same reduced normal form are indistinguishable.

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 might be useful in refining cost estimation in empirical work in industrial organization (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 their 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; since $CNE \subseteq SNE$ in finitely repeated subjective games, 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 imitative games rather than, say, the set of games capable of imitating a specific equilibrium strategy profile (typically, a much larger set). The stronger condition is appropriate, e.g., 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 literature in artificial intelligence 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*. 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 (or, for example, if correlated strategies are allowed without inclusion of a specific correlating device).

Until recently, work on probabilistic networks focused upon the decision problem of a single individual. Thus, another aspect that separates our work from the existing

literature in this field 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, though along a different line. Both of these papers develop alternative representations for interactive decision problems (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. Our work clearly complements this line of research.

A Dependency models

Here, we provide a condensed discussion of the relevant underlying theory of probabilistic networks. Since most economists are unfamiliar with this literature, we wish to: (i) give readers a sense of its theoretical content, and (ii) provide sufficient technical detail to support the development of our proofs. For those interested in pursuing these ideas further, we suggest starting with the texts by Pearl (1988, 2000) and Cowell et al. (1999).

Definition 5 *A dependency model M over a finite set of elements T is a collection of independence statements of the form $(C \perp D|E)$ in which C, D and E are disjoint subsets of T and which is read “ C is independent of D given E .” The negation of an independency is called a dependency.*

The notion of a general dependency model was originated by Pearl and Paz (1985), who were motivated to develop a set of axiomatic conditions on general dependency models that would include probabilistic and graphical dependencies as special cases. These axioms are known as the *graphoid axioms*.¹³ We are interested in *graphoids*, which are defined as dependency models that are closed under the graphoid axioms.

For example, given a probability space (A, \mathcal{A}, μ) and an associated, finite set of random variables X indexed by $T = \{1, \dots, t\}$ with typical element \tilde{x}_r , M_μ is the list of conditional independencies that hold under μ . For all $W \subseteq T$, let $\tilde{x}_W \equiv (\tilde{x}_r)_{r \in W}$. Then, for all disjoint $C, D, E \subset T$, $(C \perp D|E) \in M_\mu$ if and only if \tilde{x}_C is μ -conditionally independent of \tilde{x}_D given \tilde{x}_E . A proof that the graphoid axioms hold for conditional independence in all probability distributions can be found in Spohn (1980).

Alternatively, if \mathcal{G} is a graph whose vertices are T , then for all disjoint $C, D, E \subset T$, $(C \perp D|E) \in M_{\mathcal{G}}$ if and only if E is a cutset separating C from D . Of course, in this case, the meaning of $(C \perp D|E)$ depends upon how one defines “cutset.” The literature on probabilistic networks contains several such definitions, depending upon

whether the graph is undirected, directed or some mixture of the two (i.e., a chain graph). Since our IODs are directed, acyclic graphs (hereafter, DAGs), we proceed with Pearl’s (1986) notion of *d-separation* (the *d* stands for “directed”).

Given a DAG $\mathcal{G} \equiv (T, \rightarrow)$, a *path* is an ordered set of nodes $P \subseteq T$ such that, for all $\alpha_r, \alpha_{r+1} \in P$, either $\alpha_r \rightarrow \alpha_{r+1}$ or $\alpha_r \leftarrow \alpha_{r+1}$. A node $\alpha_r \in P$ is called *head-to-head with respect to P* if $\alpha_{r-1} \rightarrow \alpha_r$ and $\alpha_r \leftarrow \alpha_{r+1}$ in P . A node that starts or ends a path is not head-to-head. A path $P \subset T$ is *active by* $E \subset T$ if: (i) every head-to-head node is in or has a descendant in E , and (ii) every other node in P is outside E . Otherwise, P is said to be *blocked* by E .

Definition 6 *If $\mathcal{G} = (T, \rightarrow)$ is a DAG and C, D and E are disjoint subsets of T , then E is said to *d-separate* C from D if and only if there exists no active path by E between a node in C and a node in D .*

Examples of *d-separation* can be found in the Pearl references cited above. Thus, given a DAG \mathcal{G} we define $M_{\mathcal{G}}$ such that $(C \perp D | E) \in M_{\mathcal{G}}$ if and only if E *d-separates* C from D in \mathcal{G} .

We wish to characterize the relationship between probabilistic and graphical dependency models. This is done through the general notion of an *independence map* (or, *I-map*).

Definition 7 *An I-map of a dependency model M is any model M' such that $M' \subseteq M$.*

Given a probability space (A, \mathcal{A}, μ) and an associated, finite set of random variables $X \equiv \{\tilde{x}_1, \dots, \tilde{x}_t\}$, the task of constructing a DAG (T, \rightarrow) , $T = \{1, \dots, t\}$, such that $M_{(T, \rightarrow)}$ is an *I-map* of M_{μ} is straightforward (see Geiger et al., 1990, p. 514). First, for all $r \in T$, let $U_r \equiv \{1, \dots, r - 1\}$ index the predecessors of \tilde{x}_r according to T . Next, identify a minimal set of predecessors $\Pi_r \subset T$ such that $(\{r\} \perp U_r \setminus \Pi_r | \Pi_r)_{\mu}$ where the “ μ ” subscript indicates probabilistic independence under μ . This results

in a set of t independence statements known as a *recursive basis drawn from M_μ* and denoted B_μ . Now, construct (T, \rightarrow) such that $s \rightarrow r$ if and only if $s \in \Pi_r$. The resulting graph \mathcal{G} , a DAG, is said to be *generated* by B_μ and $\Pi_r = \{s \in T \mid s \rightarrow r\}$ is the set of parents of r in \mathcal{G} .

The following theorems are from Geiger et al. (1990, Theorems 1 and 2). First, an independence statement $(C \perp D \mid E)$ is a *semantic consequence* (with respect to a class of dependency models \mathcal{M} – e.g., those that satisfy the graphoid axioms) of a set B of such statements if $(C \perp D \mid E)$ holds in every dependency model that satisfies B ; i.e., $(C \perp D \mid E) \in M$ for all M such that $B \subseteq M \in \mathcal{M}$.

Theorem 1 (soundness) *If M is a graphoid and B is any recursive basis drawn from M , then the DAG generated by B is an I -map of M .*

So, given (A, \mathcal{A}, μ) , the DAG \mathcal{G} constructed in the fashion outlined above is an I -map of M_μ . That is, every independence statement implied by (T, \rightarrow) under d -separation corresponds to a valid μ -conditional independency.

Theorem 2 (closure) *Let D be a DAG generated by a recursive basis B . Then M_D , the dependency model generated by D , is exactly the closure of B under the graphoid axioms.*

Following Pearl (2000, pp. 16-20), two DAGs (T, \rightarrow) and (T, \rightarrow') are said to be *observationally equivalent* if every probability distribution that can be factored in accordance with the recursive basis $B_{(T, \rightarrow)} \equiv \{(\{r\} \perp U_r \setminus \Pi_r \mid \Pi_r) \mid r \in T\}$ can also be factored in accordance with $B_{(T, \rightarrow')} \equiv \{(\{r\} \perp U'_r \setminus \Pi'_r \mid \Pi'_r) \mid r \in T\}$. The following theorem is from Pearl (2000, Theorem 1.2.8). It originally appears in Verma and Pearl (1990, Theorem 1) and is generalized by Andersson et al. (1997, Theorems B.1 and 2.1).

Theorem 3 *Two DAGs (T, \rightarrow) and (T, \rightarrow') are observationally equivalent if and only if $\mathcal{E} = \mathcal{E}'$ and $\mathbf{S} = \mathbf{S}'$.*

B Proofs

B.1 Preliminary results

We begin with four lemmas that are used later. Given a game Γ and its IOD $\mathcal{G} \equiv (T, \rightarrow)$. Define $B_\Gamma \equiv \{(\{r\} \perp U_r \setminus \Pi_r | \Pi_r) | r \in T\}$ to be the collection of t independence statements associated with the moves in the game. Let M_Γ be the closure of B_Γ under the graphoid axioms. Notice that B_Γ is a recursive basis drawn from M_Γ and, by construction of B_Γ , \mathcal{G} is the DAG generated by B_Γ . Hence, by the soundness theorem, \mathcal{G} is an I -map of M_Γ and by the closure theorem, $M_\mathcal{G} = M_\Gamma$. By Proposition 1, for all $\theta \in \Sigma$, the conditional independencies in B_Γ hold in m_θ . Thus, the independence statements in M_Γ hold in m_θ for all $\theta \in \Sigma$; i.e., for all $\theta \in \Sigma$, $M_\Gamma \subseteq M_{m_\theta}$.

Lemma 2 *Given an exact information game Γ and a probability distribution μ that can be factored according to B_Γ , there exists a strategy $\theta \in \Sigma$ such that $m_\theta = \mu$, μ -a.s.*

Proof. By the premise, μ can be factored according to B_Γ , so that for all $r \in T$, $(\{r\} \perp U_r \setminus \Pi_r | \Pi_r) \in M_\mu$. Thus, for all $\mathbf{F} \in \mathcal{A}$,

$$\mu(\mathbf{F}) = \int_{\mathbf{F}} \prod_{r \in T} \mu(\tilde{a}_r | \tilde{\pi}_r)(\mathbf{a}) \, d\mathbf{a}.$$

Define, for all $F_r \in \mathcal{A}_r$, $\theta_r(F_r | \tilde{h}_r) \equiv \mu(F_r | \tilde{\pi}_r)$. Because Γ is a game of exact information, $\sigma(\tilde{\pi}_r) = \mathcal{I}_r$. Thus, the resulting profile, $(\theta_r)_{r \in T}$, is a strategy profile of Γ and $m_\theta = \mu$, μ -a.s., by construction. ■

Lemma 3 *If Γ and Γ' , $\Gamma \in \mathcal{O}_{\Gamma'}$, are games of exact information and $\Gamma \in \text{rem}(T', \rightarrow')$ then $\Gamma \preceq \Gamma'$. (Recall the definition of $\text{rem}(T, \rightarrow)$ in section 4.3.)*

Proof. First, imitation is a transitive (see proof of Lemma 1 below): i.e., for any games Γ_1, Γ_2 and Γ_3 , if $\Gamma_1 \preceq \Gamma_2$ and $\Gamma_2 \preceq \Gamma_3$, then $\Gamma_1 \preceq \Gamma_3$. Let $E_{\Gamma_1} \equiv$

$\{(a, b) \mid a \rightarrow b \in (T, \rightarrow_1)\}$ denote the set of edges in (T, \rightarrow_1) . Define

$$E_{\Gamma_2 - \Gamma_1} \equiv \{(a, b) \in T \times T \mid (a, b) \in E_{\Gamma_2}, (a, b) \notin E_{\Gamma_1}\};$$

i.e., the set of edges that are in (T, \rightarrow_2) but not in (T, \rightarrow_1) .

We now continue with the proof in steps:

1. Delete one edge from Γ' . Let $\Gamma_1 \in \mathbb{O}_{\Gamma'}$ be a game of exact information such that $E_{\Gamma_1} \subset E_{\Gamma'}$ and $E_{\Gamma' - \Gamma_1} = \{(a, b)\}$ where $(a, b) \in E_{\Gamma' - \Gamma}$. As both Γ' and Γ_1 are games of exact information, this implies that $\forall r \in T, r \neq b, \Pi'_r = \Pi_r^1$ and that $\Pi_b^1 = \Pi'_b \setminus \{a\}$. Let $\theta^1 \in \Sigma^1$ be an arbitrary strategy profile in Γ_1 . By Proposition 1,

$$\begin{aligned} m_{\theta^1} &= \prod_{r \in T} m_{\theta^1}(\tilde{a}_r \mid \tilde{\pi}_r^1) \\ &= m_{\theta^1}(\tilde{a}_b \mid \tilde{\pi}_b^1) \prod_{r \in T \setminus \{b\}} m_{\theta}(\tilde{a}_r \mid \tilde{\pi}_r^1). \end{aligned}$$

Since $\Pi_b^1 = \Pi'_b \setminus \{a\}$, $\sigma(\tilde{\pi}_b^1) \subseteq \sigma(\tilde{\pi}'_b)$. Thus, $m_{\theta^1}(\tilde{a}_b \mid \tilde{\pi}_b^1) = m_{\theta}(\tilde{a}_b \mid \tilde{\pi}'_b)$. Hence, condition (6) holds. By Proposition 2, $\Gamma_1 \preceq \Gamma'$. If $E_{\Gamma' - \Gamma} = \{(a, b)\}$, go to step 4. Otherwise, let $n = 2$.

2. If $E_{\Gamma' - \Gamma} \neq E_{\Gamma' - \Gamma_{n-1}}$, we proceed by deleting an additional edge. Let $\Gamma_n \in \mathbb{O}_{\Gamma'}$ (and hence $\Gamma_n \in \mathbb{O}_{\Gamma_{n-1}}$) be a game of exact information such that $E_{\Gamma_n} \subset E_{\Gamma_{n-1}}$ and $E_{\Gamma_{n-1} - \Gamma_n} = \{(c, d)\}$ where $(c, d) \in E_{\Gamma' - \Gamma}$. By the same argument as in step 1, $\Gamma_n \preceq \Gamma_{n-1}$. As, from the previous step, $\Gamma_{n-1} \preceq \Gamma'$, by the transitivity of \preceq (see proof of Lemma 1 below), so that $\Gamma_n \preceq \Gamma'$. If $E_{\Gamma' - \Gamma} = E_{\Gamma' - \Gamma_n}$ go to step 4.
3. Otherwise, repeat step 2, increasing the value of n by one. Eventually, as the set \mathcal{E}' is finite and $\mathcal{E} \subset \mathcal{E}'$, we will find an n such that $E_{\Gamma' - \Gamma} = E_{\Gamma' - \Gamma_n}$.
4. Since Γ is a game of exact information, it has the same players, action spaces and information sets as Γ_n so that $\Sigma = \Sigma^n$, hence $\Gamma \preceq \Gamma'$.

■

Lemma 4 *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 $\tilde{h}_{s \setminus r}(\mathbf{a}) \equiv (a_1, \dots, a_{r-1}, a_{r+1}, \dots, a_{s-1})$.*

Proof.

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 $\mathbf{F}' = \mathbf{G}_{r-} \cap \mathbf{G}_r \cap \mathbf{G}_{r+}$ where $\mathbf{G}_{r-} \in \sigma(\tilde{h}_r)$, $\mathbf{G}_r \in \sigma(\tilde{a}_r)$, $\mathbf{G}_{r+} \in \sigma(\tilde{a}_{r+1}, \dots, \tilde{a}_{s-1})$.
2. Let $\mathbf{H}_{r+1} \equiv \{(\mathbf{h}_r, a_r) \in \tilde{h}_{r+1}(\mathbf{A}) \mid \mathbf{h}_r \in \tilde{h}_r(\mathbf{A}), a_r \in A_r\}$. By the definition of an IOD, given $r, s \in T$, $r < s$, if $r \not\rightarrow s$, then for all $(\mathbf{h}_r, a_r), (\mathbf{h}_r, a'_r) \in \mathbf{H}_{r+1}$ we have that for all $\mathbf{F} \in \mathcal{I}_s$ such that $\tilde{h}_{r+1}^{-1}(\mathbf{h}_r, a_r) \subset \mathbf{F}, \tilde{h}_{r+1}^{-1}(\mathbf{h}_r, a'_r) \subset \mathbf{F}$. Thus, for all $\mathbf{h}_r \in \tilde{h}_r(\mathbf{A})$, there exist $\mathbf{G}_{r+} \in \sigma(\tilde{a}_{r+1}, \dots, \tilde{a}_{s-1})$ and $\mathbf{F} \in \mathcal{I}_s$ such that $\tilde{h}_r^{-1}(\mathbf{h}_r) \cap \tilde{a}_r^{-1}(\tilde{c}_r(\mathbf{h}_r)) \cap \mathbf{G}_{r+} \subset \mathbf{F}$. But, the tree structure of the game implies $\tilde{a}_r^{-1}(\tilde{c}_r(\mathbf{h}_r)) = \tilde{h}_r^{-1}(\mathbf{h}_r)$.
3. Taken together, items 1) and 2) imply that, for all $\mathbf{F}' \in \mathcal{I}_s$ there exist $\mathbf{G}_{r-} \in \sigma(\tilde{h}_r)$ and $\mathbf{G}_{r+} \in \sigma(\tilde{a}_{r+1}, \dots, \tilde{a}_{s-1})$ such that

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

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

Lemma 5 *Every influence opportunity diagram is a DAG.*

Proof. This follows by construction: the graph is defined using only directed edges. It is acyclic because nodes are fully ranked by the order of play and directed edges only appear from earlier moves to strictly later moves. ■

Finally, the following proposition implies a straightforward procedure for testing whether $\Gamma \preceq \Gamma'$ where Γ is an arbitrary game and Γ' is a game of exact information: (i) construct the IODs for each game; then, (ii) check (using d -separation) to see

whether the t independence statements in $B_{\Gamma'}$ hold in (T, \rightarrow) . While this result is more general than Proposition 4, it requires an understanding of the material presented in Appendix A and may be quite cumbersome to implement in games with many moves.

Proposition 7 *Given two games Γ and Γ' , the latter a game of exact information, if $\Gamma' \in \mathbb{O}_{\Gamma}$ and $B_{\Gamma'} \subseteq M_{\Gamma}$, then $\Gamma \preceq \Gamma'$.*

Proof. As discussed above, for all $\theta \in \Sigma$, $M_{\Gamma} \subseteq M_{m_{\theta}}$. By the premise, $B_{\Gamma'} \subseteq M_{\Gamma}$. It follows that, for all $\theta \in \Sigma$, m_{θ} can be factored in accordance with $B_{\Gamma'}$; i.e., as in condition (6). By Lemma 2 and Proposition 2, $\Gamma \preceq \Gamma'$. ■

B.2 Lemma 1

1. (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.
2. $((f(\mathbf{A}'), f(\mathcal{A}')) = (\mathbf{A}, \mathcal{A}))$ This is immediate from $\Gamma \in \mathbb{O}_{\Gamma'}$.
3. Let g and \hat{g} be functions meeting the conditions of $\Gamma \preceq \hat{\Gamma}$ and $\hat{\Gamma} \preceq \Gamma$, respectively. Define $\tilde{g} : \Sigma \rightrightarrows \hat{\Sigma}$ as follows:

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

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

B.3 Proposition 1

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

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

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

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

We need to demonstrate that $m_\theta(\tilde{a}_r | \mathcal{I}_r)$ also satisfies these conditions. For all $r \in T$, Lemma 4 implies that $\mathbf{F} \in \mathcal{I}_r \Rightarrow \mathbf{F} \in \sigma(\tilde{a}_k)_{k \in \{s \in T | s \rightarrow r\}^c}$. Of course, $\sigma(\tilde{a}_k)_{k \in \{s \in T | s \rightarrow r\}^c} = \sigma(\tilde{\pi}_r)$, so $\mathcal{I}_r \subset \sigma(\tilde{\pi}_r)$. Therefore, $m_\theta(\tilde{a}_r | \mathcal{I}_r)$ is $\sigma(\tilde{\pi}_r)$ -measurable. But this (and the definition of conditional probability) implies that, for all $\mathbf{F} \in \mathcal{A}$, $\mathbf{G} \in \sigma(\tilde{\pi}_r)$,

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

B.4 Proposition 2

The necessity of (6) follows from Corollary 1. To prove sufficiency, consider an arbitrary $\theta \in \Sigma$. By the premise of the proposition, m_θ can be factored according to $B_{(T, \rightarrow')}$. Since Γ' is a game of exact information, by Lemma 2, there exists a $\theta' \in \Sigma'$ such that $m_\theta = m_{\theta'}$. Therefore, $\Gamma \preceq \Gamma'$.

B.5 Proposition 3

By Theorem 3, (T, \rightarrow) and (T, \rightarrow') are observationally equivalent. Hence, for all $\theta \in \Sigma$, m_θ can be factored in accordance with $B_{(T, \rightarrow')}$. By Proposition 2, $\Gamma \preceq \Gamma'$. By identical reasoning, if Γ is also a game of exact information, $\Gamma' \preceq \Gamma$ and, thus, $\Gamma \sim \Gamma'$.

B.6 Proposition 4

Suppose Γ is not of exact information. Let Γ_1 be the same game as Γ in all respects except that Γ_1 is of exact information. As both Γ and Γ_1 admit (T, \rightarrow) condition (6) holds so that by Proposition 2, $\Gamma \preceq \Gamma_1$. Let $\Gamma^*, \Gamma' \in \mathbb{O}_{\Gamma'}$, be a game of exact information with IOD (T, \rightarrow^*) . Per Corollary 2, $\Gamma^* \sim \Gamma'$. Per Lemma 3 above $\Gamma_1 \preceq \Gamma^*$. By the transitivity of imitation and indistinguishability (see proof of Lemma 1) $\Gamma \preceq \Gamma_1 \preceq \Gamma^* \sim \Gamma$.

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 dependent upon the structure implied by the former. In particular, correlated strategies are not allowed without explicit correlating devices. The perfect recall assumption implies that a player’s own behavior at different moves may be correlated.
8. To illustrate the notational convention take the Games Γ^A and Γ^B in Figures 1 and 2. Their IODs are $N \rightarrow I \rightarrow II$ and $II \rightarrow I \rightarrow N$ respectively. Strictly

according to the definition, both IODs are $1 \rightarrow 2 \rightarrow 3$. Using Γ^A as the reference point, the relevant permutation is $f(1) = 3$, $f(2) = 2$ and $f(3) = 1$ so that $o^B(f(1)) = o^A(1) = N$, $o^B(f(2)) = o^A(2) = I$, and $o^B(f(3)) = o^A(3) = II$. Under our notational convention, the IOD for Γ^B is $3 \rightarrow^B 2 \rightarrow^B 1$, which matches the intuitive description used in the discussion of the example: $II \rightarrow I \rightarrow N$.

9. Note that feasible action consistency is implied by $\Gamma' \in \mathbb{O}_\Gamma$.
10. This last result raised a question that was put to us by E. Dekel in correspondence. Given the well-known work 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 2)? Due to space limitations, we do not provide a formal reply. Clearly, however, Corollary 2 does suggest a step-wise transformation that will yield observationally indistinguishable 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 2) in the original IOD.
11. Although KL allow for a more general set of possible consequences, we restrict attention to those defined above, \mathcal{C}_i
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 (1997). Howitt and McAfee (1992) employ a similar idea in a macroeconomic application.

13. The graphoid axioms are (Geiger et al. (1990) p. 515):

1. *Symmetry* $(C \perp D|E) \Rightarrow (D \perp C|E)$
2. *Decomposition* $(C \perp D \cup F|E) \Rightarrow (C \perp D|E)$
3. *Weak union* $(C \perp D \cup F|E) \Rightarrow (C \perp D|E \cup F)$
4. *Contraction* $(C \perp D|E) \wedge (C \perp F|E \cup D) \Rightarrow (C \perp D \cup F|E)$

References

- [1] Abreu, D., Pearce, D., Stacchetti, E., 1990. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5), 1041-63.
- [2] Andersson, S.A., Madigan D., Perlmann, M.D., 1997. A Characterization of Markov Equivalence Classes for Acyclic Digraphs. *The Annals of Statistics*, 25(2), 505-541.
- [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. *Istituto Economia Politica*, Milan.
- [5] Blume, L. E., Easley, D., 1992. *Rational expectations and rational learning*. Cornell University, Cornell.
- [6] Brandenburger, A., Dekel, E., 1993. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59, 189-98.
- [7] Cowell, R. G., Dawid, A. P., Lauritzen, S. L., Spiegelhalter, D. J., 1999. *Probabilistic Networks and Expert Systems*. Springer, New York.
- [8] Elmes, S., Reny, P., 1994. On the strategic equivalence of extensive form games. *Journal of Economic Theory*, 62(1), 1-23.
- [9] Fristedt, B., Gray, L., 1997. *A Modern Approach to Probability Theory*. Birkhäuser, Boston.
- [10] Fudenberg, D., Levine, D. K., 1993. Self-confirming equilibrium. *Econometrica*, 61, 523-45.

- [11] Geanakoplos, J., 1994. Common Knowledge. Handbook of Game Theory. Elsevier.
- [12] Geiger, D., Verma, T. S., J. Pearl, J., 1990. Identifying independence in Bayesian networks. *Networks*, 20, 507-34.
- [13] Harsanyi, J., 1967-68. Games with incomplete information played by ‘Bayesian’ players,” *Management Science*, 8, Parts I-III, 159-82, 320-34, 486-502.
- [14] Howitt, P., McAfee, R. P., 1992. Animal Spirits. *American Economic Review*, 82, 493-507.
- [15] Kalai, E., Lehrer, E., 1993. Subjective equilibrium in repeated games. *Econometrica*, 61, 1231-1240.
- [16] ——— 1995. Subjective games and equilibria. *Games and Economic Behavior*, 8, 123-63.
- [17] Koller, D. Milch, B., 2002. Multi-agent influence diagrams for representing and solving games. Seventeenth International Joint Conference on Artificial Intelligence, Seattle, WA.
- [18] La Mura, P., 2002. Game Networks. Unpublished manuscript.
- [19] Mackey, G. W., 1957. Borel structures in groups and their duals. *Trans. Amer. Math. Soc.*, 85, 283-95.
- [20] Milgrom, P. Roberts, J., 1990. Rationalizability, learning and equilibrium in games with strategic complementarities. *Econometrica*, 58, 1255-1277.
- [21] Nachbar, J. H., 1997. Prediction, optimization and learning in repeated games. *Econometrica*, 65, 275-309.
- [22] Pearl, J., 1986. Fusion, propagation and structuring in belief networks. *Artificial Intelligence* 29, 241-88.

- [23] —, 1988. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. North Holland, Amsterdam.
- [24] —, 2000. Causality: Models, Reasoning, and Inference. Cambridge University Press, Cambridge.
- [25] —, and Paz 1985. Graphoids: A graph-based logic for reasoning about relevance relations. Technical report 850038 (R-53-L), Cognitive Systems Laboratory, UCLA. Short version in *Advances in artificial intelligence 2*, ed. Du Boulay, B., Hogg, D., Steels, L. North Holland, Amsterdam.
- [26] Spohn, W., 1980. Stochastic independence, causal independence and shieldability. *J. Phil. Logic*, 9, 73-99.
- [27] Thompson, F. B., 1952. Equivalence of games in extensive form. The Rand Corporation, Research Memorandum 759.
- [28] 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

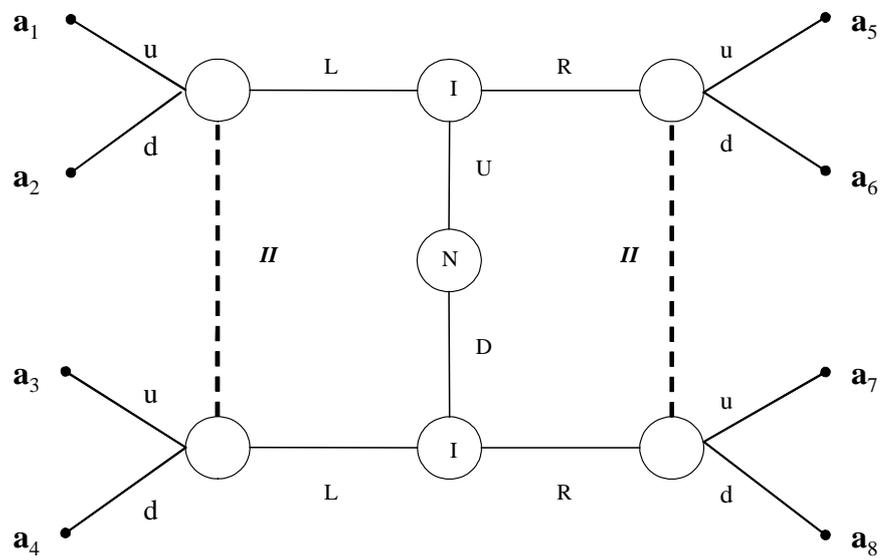


Figure 1: game Γ^A

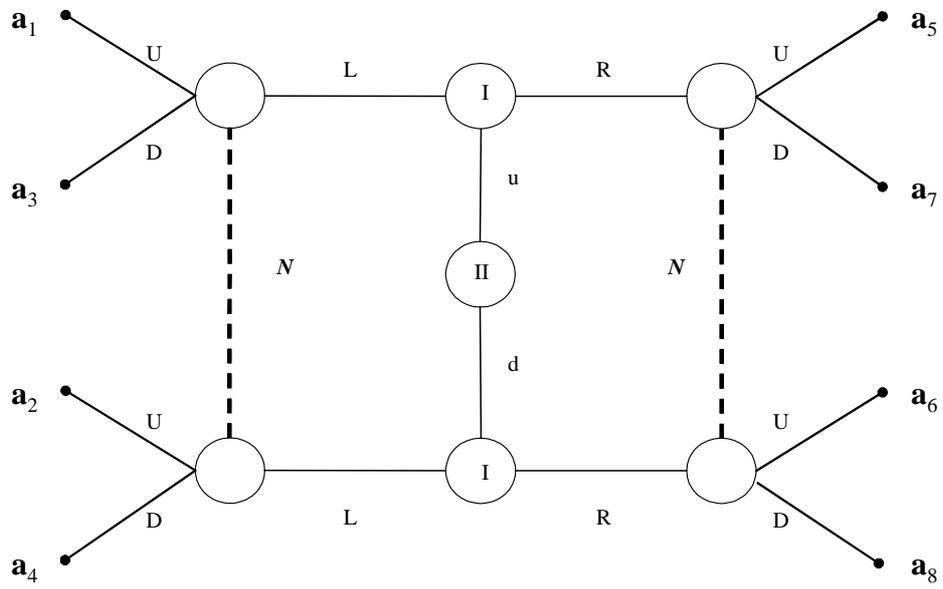


Figure 2: game Γ^B .

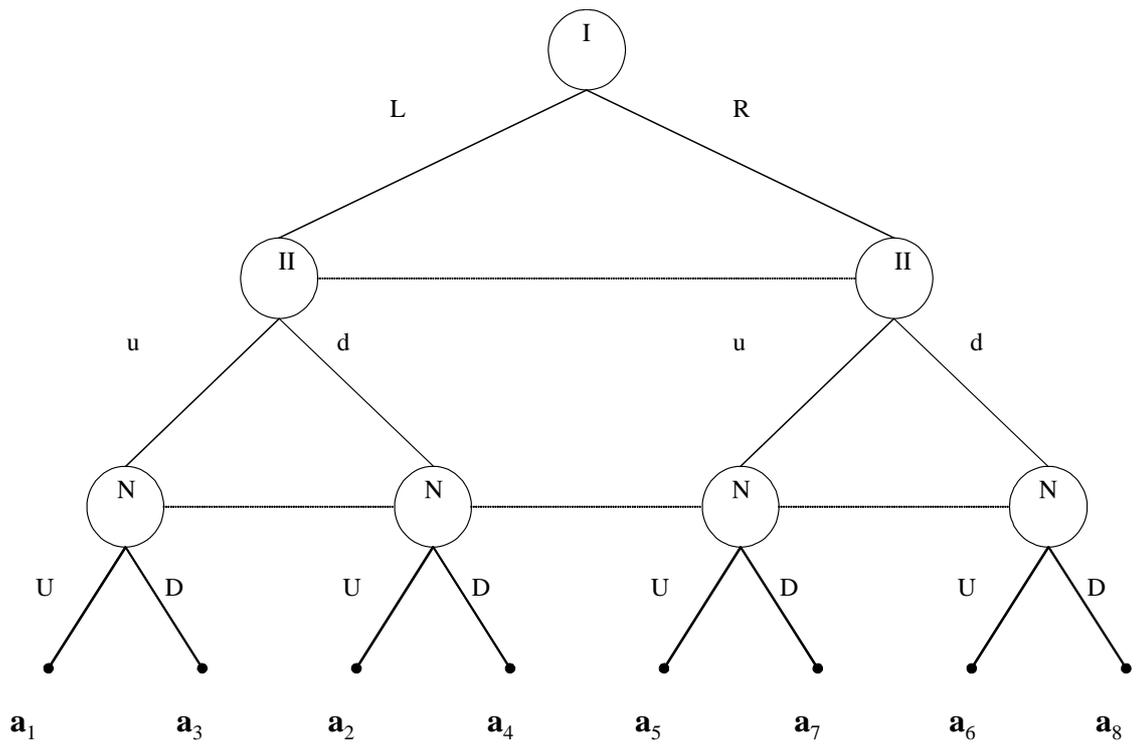


Figure 3: game Γ^C .

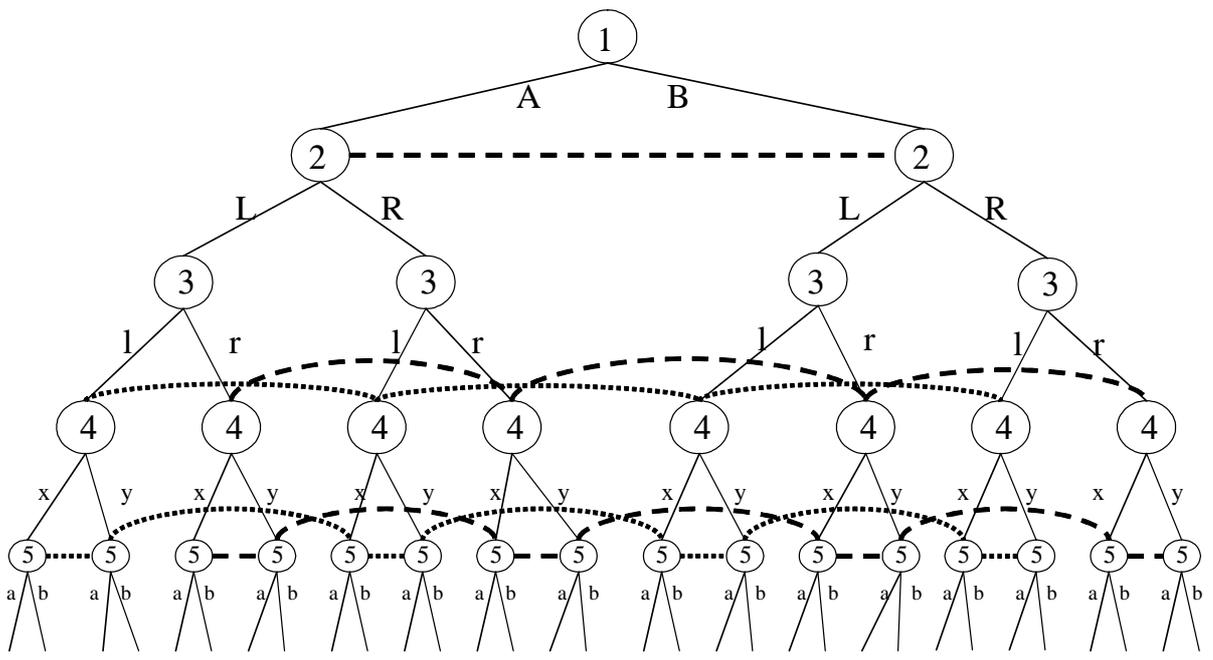


Figure 4: the Gatekeeper game.

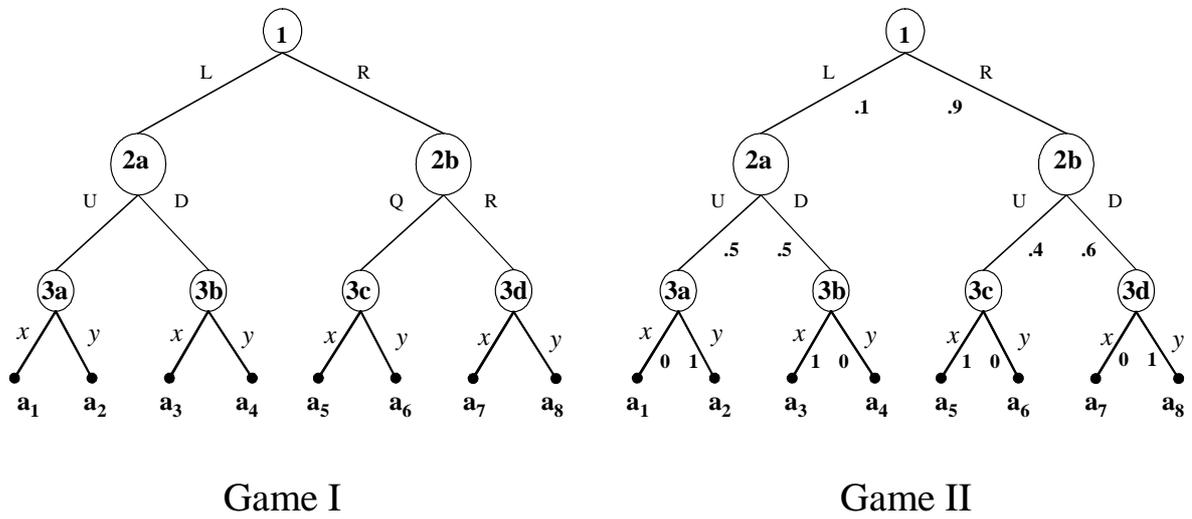


Figure 5: illustration of IOD condition (4).

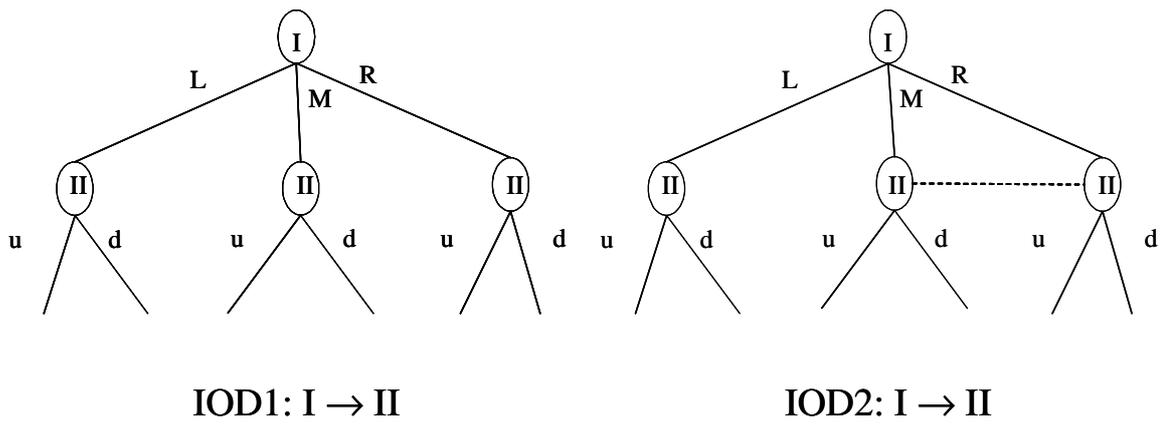


Figure 6: necessity of condition (6).

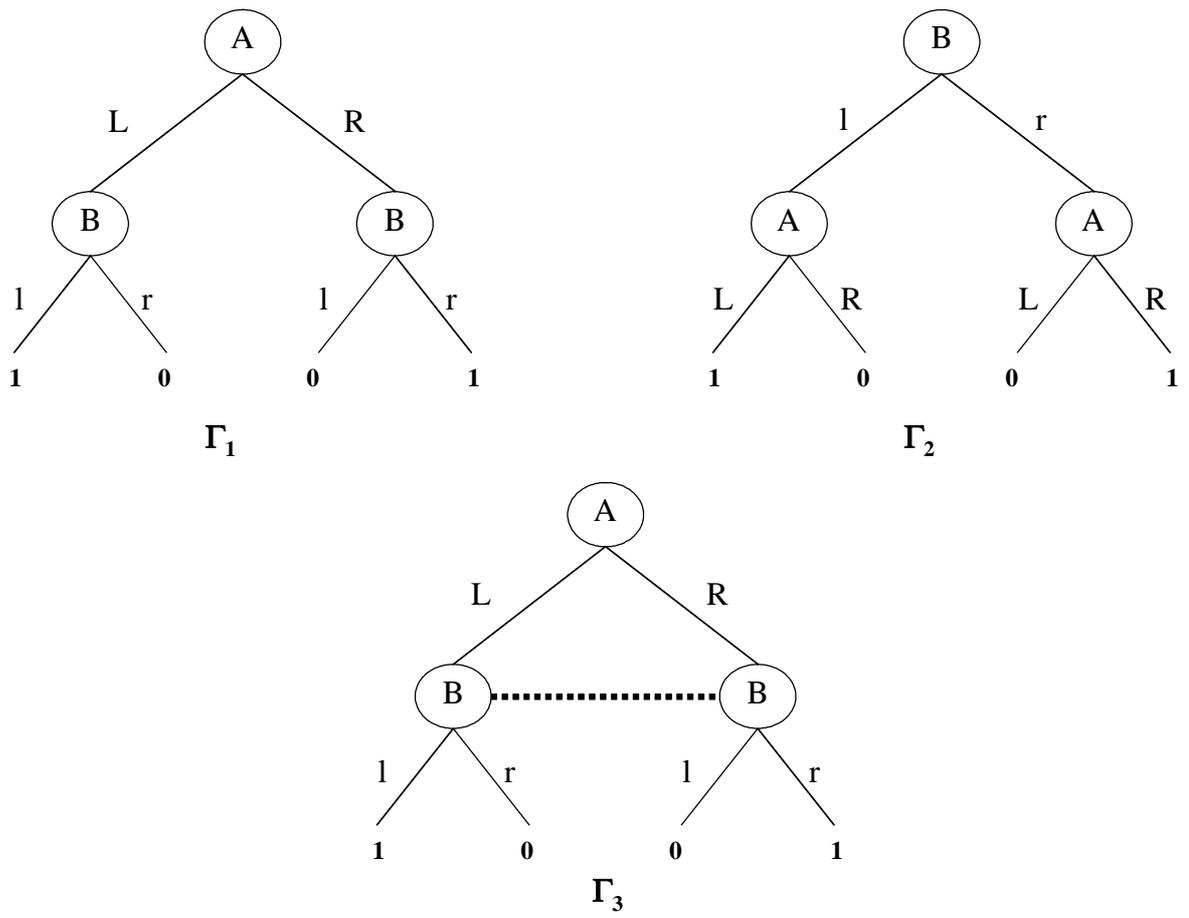


Figure 7: possible departmental subgames.