The B.E. Journal of Theoretical Economics

Contributions

Volume 8, *Issue* 1 | 2008 | *Article* 2

Empirical Implications of Information Structure in Finite Extensive Form Games

Jose Penalva* Michael D. Ryall†

*Universidad Carlos III, Madrid, jpenalva@emp.uc3m.es

†University of Melbourne, m.ryall@mikeryall.com

Recommended Citation

Jose Penalva and Michael D. Ryall (2008) "Empirical Implications of Information Structure in Finite Extensive Form Games," *The B.E. Journal of Theoretical Economics*: Vol. 8: Iss. 1 (Contributions), Article 2.

Available at: http://www.bepress.com/bejte/vol8/iss1/art2

Empirical Implications of Information Structure in Finite Extensive Form Games*

Jose Penalva and Michael D. Ryall

Abstract

We analyze what can be inferred about a game's information structure solely from the probability distributions on action profiles generated during play; i.e., without reference to special behavioral assumptions or equilibrium concepts. Our analysis focuses on deriving payoff-independent conditions that must be met for one game form to be empirically distinguished from another. We define empirical equivalence and independence equivalence. The first describes when two game forms can never be distinguished based solely on the empirical distribution of player actions. As this turns out to be difficult to characterize, we introduce the latter, which describes two game forms that imply the same minimal sets of conditional independencies in every one of their empirical distributions. Our main contribution is to identify, for an arbitrary game form, the minimal set of conditional independencies that must arise in every one of its empirical distributions. We also introduce a new graphical device, the influence opportunity diagram of a game form which facilitates verifying independence equivalence, and hence provides a simple necessary condition for empirical equivalence.

KEYWORDS: empirical inference, information structure, extensive form, compatibility, Bayesian network, causality

*Our thinking on the issues analyzed in this paper have been substantially influenced by many fine colleagues, including Adam Brandenburger, Colin Camerer, David Chickering, Catherine de Fontenay, Eddie Dekel, Bryan Ellickson, Joshua Gans, Fabrizio Germano, Ehud Kalai, Uffe Kj, David Levine, Steven Lippman, Glenn MacDonald, Mark Machina, Leslie Marx, Joseph Ostroy, Scott Page, Judea Pearl, Ben Polak, Larry Samuelson, Joel Watson, Eduardo Zambrano, Bill Zame, participants in the MEDS, Georgetown, UCLA, and UCSD theory workshops as well as the SITE, Stonybrook, Midwest, and Melbourne theory conferences. The authors are, of course, entirely responsible for all shortcomings. Penalva's research was partially supported by the Spanish Ministry of Science grant SEJ2005-03924.

1 Introduction

Applied noncooperative game theory has spawned numerous streams of empirical research in economics and other social sciences. A common feature of virtually all contributions along these lines is that they begin by specifying a game that is presumed to generate the data.[1] An equilibrium concept is then applied and either tested or used to infer unknown parameters, depending upon the type of study at hand. It is well understood that different assumptions about game structure can lead to very different empirical implications. Indeed, a common complaint invariably encountered at some point by every game theorist is that *any* observed economic phenomenon can be described by *some* appropriately constructed game and, hence, the theory has no empirical content. This critique would be neutralized by the invention of a joint test of game structure and equilibrium behavior.

Although this paper does not deliver such a comprehensive test, it does contribute to our understanding of the extent to which game structure can be estimated from data. The key components of game structure are game form and payoffs. A very recent line of work explores situations in which the empiricist faces unobserved payoffs (discussed in greater detail below). To the best of our knowledge, no work has been done along these lines on settings with unobserved game form. In this paper, we consider such settings and develop a novel method by which to identify certain empirical regularities implied by game form. We use this method to derive several results on the extent to which one can empirically discriminate between games with different forms. Our findings do not rely upon behavioral assumptions, such as whether play is consistent with some notion of individual rationality. Rather, we sidestep these issues and operate directly on game form to derive payoff-independent conditions that must be met for one game form to be empirically distinguished from another.

Our analysis concentrates on the probability distributions over action profiles generated by behavior strategies in an extensive form game. We refer to these as the *empirical distributions* induced by their associated behavior strategy profiles. The motivation for this term is that, over a large number of independent plays of a game according to a specific (e.g., equilibrium) behavior strategy profile, the observed action frequencies converge to the empirical distribution induced by that profile. What we have in mind are situations in which the only data available to a social scientist who does not

[1] In the case of the experimental literature and certain areas of mechanism design, the structure of the game actually *is* known to the researcher.

know the underlying game form is a cross-sectional listing of actions taken by each of the players – i.e., data that does not directly indicate what players knew about each others' actions at the times of their own moves. We are interested in understanding the potential of such data to identify certain aspects of the underlying information structure.

To begin, we introduce the notion of *empirical equivalence*. One form is said to be empirically equivalent to 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 and viceversa. This is a strong condition in the sense that, even under infinitely repeated play, it is impossible to distinguish one form from any within its empirical equivalence class based solely on the observed actions of the players.

For reasons we discuss below, it turns out that empirical equivalence is difficult to characterize. Hence, we introduce a weaker equivalence concept that we term *independence equivalence*. Two game forms are independence equivalent when they imply the same minimal sets of conditional independencies in every one of their empirical distributions. Clearly, empirical equivalence implies independence equivalence, but not conversely. This leads to the idea that one way to test a game form hypothesis independent of any behavioral assumptions is to conduct an empirical consistency check against that form's conditional independence implications.

Our main contribution is to identify, for an arbitrary game form, the minimal set of conditional independencies that must arise in every one of its empirical distributions. Our analysis is facilitated by the introduction of a new graphical device, the *influence opportunity diagram of a game form* (hereafter, IOD). The IOD is a directed, acyclic graph whose nodes correspond to the moves made by players in the underlying game. Given a game form, we illustrate how to construct its IOD. We then demonstrate that the IOD is an "independence map" of every empirical distribution that can be generated in the underlying game. That is, according to a particular notion of conditional independence in graphs (introduced below), every conditional independence between players in a game's IOD will also be present in *all* of the game's empirical distributions.[2]

Thus, for one form to be independence equivalent to another, their IODs must imply the same set of conditional independencies. If not, then there

[2] Our work is closely related to – indeed, contributes to – the rapidly-expanding literature on probabilistic networks. Since most economists are not familiar with this literature, we have included a condensed, self-contained discussion of some relevant results and references in Appendix A.

are probabilistic dependencies that can arise in one game that cannot arise in the other. As we show, the required consistency is easy to determine via direct pair-wise comparison of IODs. Although IOD consistency is both necessary and sufficient to determine whether two game forms imply the same minimum sets of conditional independencies, such consistency is not, in general, sufficient to establish the stronger condition of empirical equivalence. The problem is that 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. Thus, IOD consistency is only a necessary condition for empirical equivalence. We show that for a class we call forms of semiperfect information, independence equivalence implies empirical equivalence. Though special, this class does include many games of economic interest.

Our hope is that these findings will convince readers of the value of assessing game form independent of equilibrium assumptions via the probabilistic independencies observed during play. While we present some new methodological tools and demonstrate some of their uses along this line, we stop well short of implementing an explicit empirical procedure. Remaining issues that must be resolved prior to achieving this goal include: i) extending the results to dynamic strategies under repeated play, ii) admitting hidden variables, iii) identifying reasonable assumptions about the noise process associated with play, and iv) refining the IOD to account for the additional restrictions implied by specific information. Although we do not presently have solutions to these issues, our sense is that any practical empirical test will require the relaxation of certain of our assumptions and the strengthening of others. We discuss these in some detail below.

The layout of the paper is as follows. In the next section, we discuss related work. Section 3 contains two extended examples designed to motivate our results and illuminate the ideas that are generalized later in the paper. In Section 4, we set up the formalities. Section 5 presents and illustrates the definition of a game's IOD. Some useful preliminary results are contained in Section 6, followed by our main results in Section 7. Section 8 discusses these results, their limitations and some thoughts on future directions. We close with brief concluding remarks in Section 9.

2 Related literature

To the best of our knowledge, this is the first paper to consider the empirical implications of game form independent of equilibrium assumptions. Nevertheless, there are some papers that address questions that are similar in spirit and others that use several of the tools also used in this paper.

The line of work that adresses questions similar to ours assumes the game form is known and explores what observed actions can tell us about players' unknown payoffs. Sprumont (2000) considers whether a given joint choice function (combination of actions chosen by the players and normal game forms defined via restrictions on a given set of actions) can be rationalized as a Nash equilibrium or Pareto efficient outcome. Ray & Zhou (2001) and Ray & Snyder (2004) extend Sprumont to extensive game forms of perfect information, while Demuynck & Lauwers (2005) extend the analysis to mixed strategies. Our analysis of unobserved game form is complementary to this line – a synthesis of the two may lead to the implementation of *joint* tests of structure and equilibrium (a topic we consider at greater length in Section 8).

Kalai (2004) is related to the current paper in that he is also interested in situations with uncertainty about game structure. He studies the robustness of Nash equilibria to changes in the game form as the number of players increases. He analyzes special classes of Bayesian games with given payoffs. Notably, he finds two conditions on extensive game modifications such that all modifications of a game that satisfy them have Nash equilibria that correspond to constant versions of the Nash equilibria of the original game (if the number of players is large enough).

Several of our results build upon work in the literature on probabilistic networks. This literature focuses overwhelmingly upon the decision problem of a single individual. There are some notable exceptions, including Ryall (1997), Kearns et al. (2001), Kakade et al. (2003), Koller and Milch (2002), and La Mura (2002). These works have the common feature of using ideas in the probabilistic network literature to develop alternative representations for interactive decision problems (i.e., to the conventional normal or extensive forms). In this paper, we also apply these ideas to develop a graphical device, the influence opportunity diagram (IOD) of an extensive form. As a result, this paper is loosely related to this literature. The primary difference is that, here, we take the extensive form as the essential primitive and, from it, construct its IOD. We do not view the IOD as a substitute for the extensive form but, rather, as a device that summarizes useful information about it.[3] Another difference is our motivation: whereas the latter four

[3] In fact, as we demonstrate in Proposition 4, there is some reason to doubt that IOD-style representations can ever fully substitute for the standard extensive form outside of a limited class of games.

papers focus primarily upon efficient algorithms for approximating equilibria, we explore the behavioral regularities implied by game form as a basis for empirical refutation of game form hypotheses.

In addition to the IOD, this paper formally introduces several new ideas including independence equivalent game forms, empirically equivalent forms, and forms of semiperfect information. We are not aware of any precedent to the notion of independence equivalent game forms. Although our definition of empirically equivalent game forms is also new, the idea of empirically equivalent *strategies* is introduced at least as early as Kuhn (1953). There, two strategies, behavior or mixed, are said to be equivalent if they lead to the same probability distribution over outcomes for all strategies of one's opponents. Our notion of semiperfect information games is a mild generalization of perfect information games. In the latter, each player perfectly observes the actions of all predecessors while, in our formulation, an arbitrary subset of predecessors is allowed to remain unobserved.

Game theory itself has long been interested in the descriptive equivalence of alternative game forms. Kuhn (1953) demonstrates that, in games of perfect recall, every mixed strategy (in an appropriately defined normal form) is equivalent to the unique behavior strategy it generates (in the related extensive form) and each behavior strategy is equivalent to every mixed strategy that generates it (see Aumann (1964) for an extension to infinite games). Thompson (1952) and Elmes & Reny (1994) identify transformations on extensive form games that yield the equivalence class of games with the same strategic form. Glazer and Rubinstein (1996) construct an extensive form game from a normal form game such that the strategies that are obtained from the former using backward induction is the same as the strategies that survive iterative elimination of dominated strategies in the latter. These are very strong notions of equivalence, as is our definition of empirical equivalence. Classifying equivalence on the basis of the empirical regularities implied by game form is new to this paper.

3 Examples and intuition

In this section, we present two examples, one in which the issue of interest is a test of equilibrium refinement and another in which the issue is estimation of cost parameters in a market game. We wish to highlight the issues that arise when game structure is unknown and also to foreshadow the analytical approach that follows. Throughout this section, we assume that the data observed by the social scientist is generated by an extensive form game in which the participating players know all the relevant structural details. Since the interest is in empirical work, we also assume that players take noisy actions and that the noise is independent across player moves. To keep things simple in this section, we do not add explicit "nature" nodes to the extensive forms to represent the noise terms but, instead, embed them directly into player strategies in the form of independent "trembles."

3.1 A test of equilibrium

Suppose a social scientist observes a two player interaction in which player I plays either L or R and player II either u or d. The scientist hypothesizes that the true underlying game is the one shown in Figure 1 below. Although she is certain that this model captures all the relevant players (e.g., there are no unobserved factors influencing outcomes outside the choices of the two players shown) and their payoffs, she is unsure of the game's information structure (who moves when and what they know at the time).

The social scientist's data set represents a tally of the outcomes observed over a *large* number of repetitions. She would like to know whether the data are consistent with one of the game's two pure strategy Nash equilibria, (L, u) and (R, d). Assuming player I chooses an action with noise ϵ_I and player II with noise ϵ_{II}, the data set should be consistent with one of the following frequency tables (the one on the left corresponding to (L, u) and the one on the right to (R, d)).

Table 1

Outcome	Frequency
(L, u)	$(1 - \epsilon_I)(1 - \epsilon_{II})$
(L, d)	$(1 - \epsilon_I)(\epsilon_{II})$
(R, u)	$(\epsilon_I)(1 - \epsilon_{II})$
(R, d)	$(\epsilon_I)(\epsilon_{II})$

Table 2

Outcome	Frequency
(L, u)	$(\epsilon_I)(\epsilon_{II})$
(L, d)	$(\epsilon_I)(1 - \epsilon_{II})$
(R, u)	$(1 - \epsilon_I)(\epsilon_{II})$
(R, d)	$(1 - \epsilon_I)(1 - \epsilon_{II})$

Suppose the scientist's best estimate of the true distribution generating the observed outcomes is as shown in the following table.

Table 3

Outcome	Frequency
(L, u)	.81
(L, d)	.09
(R, u)	.01
(R, d)	.09

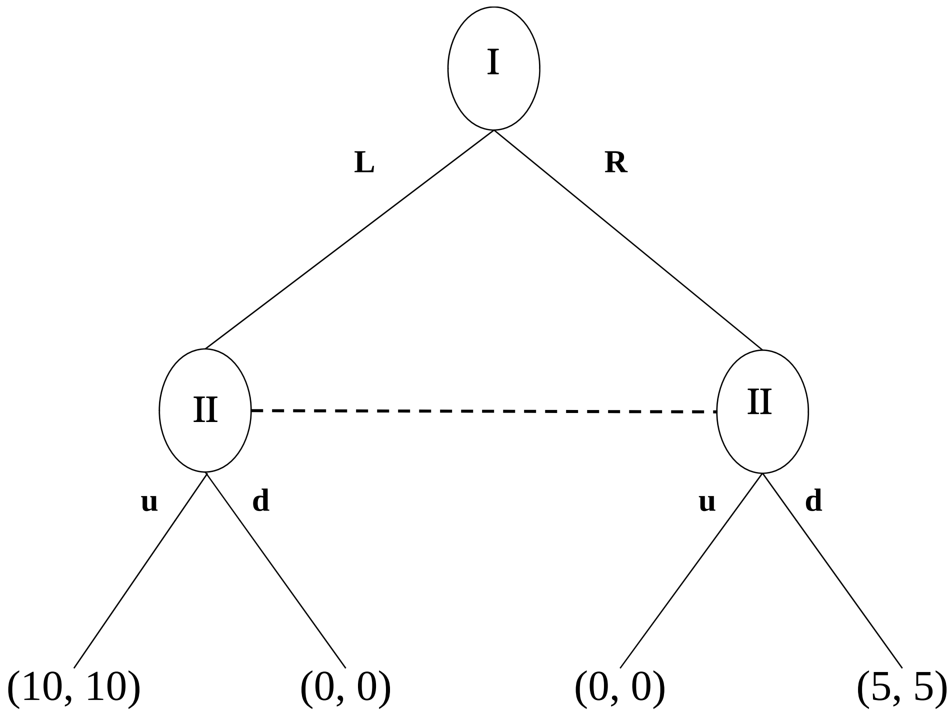

Figure 1: two-player, simultaneous move game

We call this the empirical distribution generated by play of a behavior strategy profile. It is not difficult to see that these frequencies do not appear to conform to either of the equilibrium distributions hypothesized above. There are no parameters ϵ_I and ϵ_{II} that deliver the observed values. Apparently, then, the Nash hypothesis should be rejected.

Such a conclusion is premature, however. What the data reject is not the Nash hypothesis but, more fundamentally, the assumption that the observed actions arise from strategies – any strategies – arising from a game form in which the players must choose their actions simultaneously. Thus, *independent of payoff and equilibrium hypotheses*, the social scientist discovers that the game form in Figure 1 is mispecified. In fact, the data *is* consistent with a Nash equilibrium (a subgame perfect one) of a different game.

To see this, first note that the structure of the hypothesized game (simultaneous move) implies $\Pr(\tilde{a}_I, \tilde{a}_{II}) = \Pr(\tilde{a}_I)\Pr(\tilde{a}_{II})$ for *all* true empirical distributions generated by play (where $\tilde{a}_i$ is the action taken by player i).[4] This restriction is violated by the distribution in Table 3. By the definition of conditional probability, every empirical distribution from *any* two player game can be factored as

$$\Pr(\tilde{a}_I, \tilde{a}_{II}) = \Pr(\tilde{a}_{II}|\tilde{a}_I)\Pr(\tilde{a}_I). \quad (1)$$

The parameters that deliver the distribution in Table 3 according to the factorization in (1) are as follows.

	L	R
$\Pr(\tilde{a}_I)$	.9	.1

$\tilde{a}_I$	$\Pr(u\|\tilde{a}_I)$	$\Pr(d\|\tilde{a}_I)$
L	.9	.1
R	.1	.9

This is significant because these parameters correspond exactly to (noisy) behavior strategies in a different game; namely, the one in which player I moves R or L and *is observed by player II,* who then moves u or d. The information structure is intuitively captured by the directed graph

$$1 \rightarrow 2.$$

Indeed, the behavior strategy profile implied by these parameters is a subgame perfect equilibrium in which $\epsilon_I = \epsilon_{II} = .1$. Hence, the data summarized in Table 3 is consistent with the joint hypothesis that the underlying game is the perfect information game with player I moving first and in which play proceeds according to the subgame perfect equilibrium of that game.

Are there other joint hypotheses with which this data is consistent? The answer is yes. To see this, note that by applying the definition of conditional probability once again, we can equally well consider the factorization

$$\Pr(\tilde{a}_I, \tilde{a}_{II}) = \Pr(\tilde{a}_I|\tilde{a}_{II})\Pr(\tilde{a}_{II}). \quad (2)$$

The corresponding parameters are

	u	d
$\Pr(\tilde{a}_{II})$	.82	.18

$\tilde{a}_{II}$	$\Pr(L\|\tilde{a}_{II})$	$\Pr(R\|\tilde{a}_{II})$
u	.99	.01
d	.50	.50

[4] Here, and throughout, we assume that the player set is complete; i.e., there are no hidden correlating devices in the form of unspecified nature moves or historical player moves.

These parameters correspond exactly to the game of complete information in which player II moves first and which is described by the directed graph

$$1 \leftarrow 2.$$

In this case, play is consistent with a subgame perfect equilibrium with $\epsilon_{II} = .18$ and error parameters for player I at his respective information sets of $\epsilon_{I,u} = .01$ and $\epsilon_{I,d} = .5$.

Notice that an implication of equations (1) and (2) is that any empirical distribution arising as a result of play in the perfect information game in which I moves first can be empirically imitated by appropriately chosen behavior strategies (with error terms) in the perfect information game in which II moves first, and viceversa. When two games have this property, we say they are *empirically equivalent*.

Which, if either, of the two data-consistent joint hypotheses seems reasonable depends upon what the social scientist wishes to assume about the relative magnitude of the error terms, their consistency across information sets, and so on. The important point is that the data in Table 3 appears to rule out the simultaneous move game in Figure 1 and, hence, any equilibrium implications that might be drawn from it.

3.2 Inferring firm costs in market games

Now, suppose the social scientist is consulting with the FTC and wishes to estimate costs for a 3-firm industry, say automobiles. Average market shares are estimated at 40%, 40%, and 20% of aggregate industry unit volume. As before, we suppose these estimates are taken from a large number of observations. It is assumed that the industry operates as a Cournot oligopoly in which firm costs are noisy and inverse demand is given by $p \equiv 200 - \sum_{i=1}^{3} q_i$. This implies average costs are $c_1 = c_2 = 67$ and $c_3 = 93$. Hence, the consultant concludes that firm 3's low market share is due to a substantial cost disadvantage. Is such a conclusion appropriate?

Suppose that, upon deeper inspection and controlling for external correlating factors (e.g., industry-wide shocks) and past production (to rule out history-dependent coordination), our consultant discovers, with a high degree of confidence, that the cross-sectional empirical distribution on firm production choices can be factored as

$$\Pr(q_1, q_2, q_3) = \Pr(q_3 | q_1, q_2) \Pr(q_1) \Pr(q_2) \tag{3}$$

but *not*

$$\Pr(q_1, q_2, q_3) = \Pr(q_3) \Pr(q_1) \Pr(q_2). \tag{4}$$

The finding of (3) and not (4) is highly problematic with respect to the Cournot hypothesis for, under the simultaneous-move assumption of Cournot, firm quantity choices should be independent. The dependencies in (3) suggest a different, Stackleberg-type mechanism, one in which firms 1 and 2 move first and are observed by 3 prior to its production decision. This latter information structure is intuitively captured by the directed graph

$$1 \rightarrow 3 \leftarrow 2. \tag{5}$$

As we will show momentarily, there is a close connection between graphs of this type, which we call influence opportunity diagrams (IODs), and the conditional independencies implied by a game's information structure. Estimating costs according to the Stackleberg game consistent with (5) leads to the conclusion

$$c_1 = c_2 = c_3 = 80,$$

which is both quantitatively and qualitatively different than the conclusion reached under the Cournot assumption.

4 The framework

Wherever possible, capital letters (X, Z) denote sets, small letters (a, w) elements of sets, small letters with a tilde $(\tilde{a}, \tilde{\pi})$ denote functions and random variables, and script letters $(\mathcal{A}, \mathcal{F})$ collections of sets and graphs. Sets with product structure are indicated by bold capitals $(\mathbf{A}, \mathbf{E})$ and small bold $(\mathbf{a}, \mathbf{e})$ denotes typical elements (ordered profiles) in such sets. Graphs and probability spaces play a large role in the following analysis – standard notation and definitions are adopted wherever possible.

4.1 Game form

A game form, denoted Γ, contains a set of players, the order of moves, each player's set of choices when he or she moves, what each player knows when he makes a move (the information structure), and the probability distribution over exogenous events. Specifically, let $N \equiv \{1, ..., n\}$, $n < \infty$, index the set of players, which includes a Nature player that describes exogenous events if necessary. The game tree is denoted $(X, \mathbf{E})$ with nodes X, edges $\mathbf{E} \subset X \times X$, and terminal nodes $Z \subset X$. The order of play is a partial order over the set of nodes and there is a function that maps nodes into the set of players that identifies which player moves at that node.

We assume that the game form is one in which each player makes one move.[5] More substantially, we also assume that players always move in the same order.[6] Thus, we can identify the player's index with the move order: player k is the one who has the k^{th} move. Let X_k denote the set of move-k nodes, $k = 1, \ldots, n$. We also assume that the set of actions available for the k-th move is a finite set, A_k, and does not depend on the actions of other players so that the set of *action profiles*, $\mathbf{A} = \times_{k \in N} A_k$, has a product structure. An action is represented in the game tree as an edge so that each edge connecting any two nodes, x and x', where $x \in X_k$ and $x' \in X_{k+1}$, corresponds to some $a_k \in A_k$.[7] Given this structure, the relationship between Z and $\mathbf{A}$ is bijective.

The final object needed to characterize the game form is the information structure, what each player knows at the time of his move. The *information structure* is represented by $\mathcal{X} \equiv \{\mathcal{X}_1, ..., \mathcal{X}_n\}$ where $\mathcal{X}_k \equiv \left\{X_{k,1}, ..., X_{k,m_k}\right\}$ is the partition of the move-k game tree nodes into their corresponding information sets $X_{k,j}$. Thus, a game form, Γ, is fully characterized by $(N, \mathcal{X}, (X, \mathbf{E}), \mathbf{A})$.[8]

4.2 Actions and strategies

We adopt the convention that functions on Z are denoted with tildes to both emphasize this fact and to distinguish them from the specific values they take. For example, the function $\tilde{x}_k : Z \to X$ indicates the k^{th} node on the game tree path terminating at z; that is, the function $\tilde{x}_k$ evaluated at terminal node z, $\tilde{x}_k(z)$, indicates a node, x, where x is a node in the set X_k such that there is a path in $(X, \mathbf{E})$ from x to z. Similarly, $\tilde{a}_k : Z \to A_k$ represents the action corresponding to the k^{th} edge on the game tree which connects the

[5]Our results extend to games in which players have multiple moves, though doing so requires an additional layer of notational bookkeeping.

[6]There is a narrow class of game forms for which this assumption fails, namely games in which one player conditions the order of moves by subsequent players in non-trivial ways; i.e., player 1 decides whether player 3 observes player 2 or visa versa. We ignore such games: in order to identify influence relationships from empirical data, it seems reasonable to limit our attention to situations in which such relationships can be assumed stationary.

[7]Earlier versions of this paper allowed for infinite action sets. Because this generalization added substantial technical detail without much additional insight, it was dropped. However, readers may wish to keep in mind that all of our results apply to forms with "well-behaved" infinite action sets (e.g., choosing points in closed intervals of real numbers).

[8]To simplify notation we obviate the probability distribution over exogenous events in the description of the game and include it as the strategy of the Nature player.

nodes $\tilde{x}_k(z)$ and $\tilde{x}_{k+1}(z)$ on the path terminating at z. Let $\tilde{a} \equiv (\tilde{a}_1, ..., \tilde{a}_n)$ so that $\tilde{a}(z) \in \mathbf{A}$ identifies the action profile associated with $z \in Z$. The *move-k history* is given by $\tilde{h}_k \equiv (\tilde{a}_1, ..., \tilde{a}_{k-1})$, a list of *all* of the actions taken prior to move k (both the ones observed by player k and those not observed by the player at the time of his move). The null history, $\tilde{h}_\emptyset$, is an arbitrary constant and we set $\tilde{h}_1 = \tilde{h}_\emptyset$. Finally, let $\tilde{X}_k(z) = X_{k,i}$ indicate the move-k information set that intersects the path terminating in z; i.e., $\mathcal{X}_k \ni X_{k,i} \ni \tilde{x}_k(z)$.

A *behavior strategy* for player k is a plan of what player k will do when he gets to move, conditional on the information he has available at the time of his move. Let $l_k \equiv |A_k|$ denote the number of actions in A_k. Then, every node in X_k has l_k edges emanating from it. In our simplified setup, a *behavior strategy* for player $k \in N$ is a matrix of conditional probabilities

$$\Theta_k \equiv \begin{bmatrix} \theta^k_{1,1} & \cdots & \theta^k_{1,l_k} \\ \vdots & \ddots & \vdots \\ \theta^k_{m_k,1} & \cdots & \theta^k_{m_k,l_k} \end{bmatrix}$$

in which $\theta^k_{i,j}$ is the probability player k places on action $a_{k,j} \in A_k$ upon reaching information set $X_{k,i} \in \mathcal{X}_k$. Hence, each row of Θ_k corresponds to an information set and constitutes a probability distribution on the elements of A_k. Let $\mathbf{\Theta} \equiv (\Theta_1, ..., \Theta_n)$ denote a *strategy profile* and $\mathbf{\Sigma}$ the set of all such profiles.

To keep track of strategy parameters, let $\tilde{\theta}_k(z) \equiv \theta^k_{i,j}$ where $\tilde{X}_k(z) = X_{k,i}$ and $\tilde{a}_k(z) = a_{k,j}$. So, $\tilde{\theta}_k(z)$ indicates the probability associated with the k^{th} edge in the path terminating at z given player k's behavior strategy. As usual, the set of behavior strategy profiles, $\mathbf{\Sigma}$, is restricted such that players assign equal probabilities to identical actions in the same information set; that is,

$$\forall z, z' \in Z \; s.t. \; \tilde{X}_k(z) = \tilde{X}_k(z') \; and \; \tilde{a}_k(z) = \tilde{a}_k(z'), \; \tilde{\theta}_k(z) = \tilde{\theta}_k(z'). \quad (6)$$

The function $\tilde{\theta} \equiv (\tilde{\theta}_1, ..., \tilde{\theta}_n)$ identifies the edge probabilities associated with each branch of the game tree given $\mathbf{\Theta}$.

Finally, we only allow correlated strategies when the correlating device is explicitly included in the description of the game form (i.e., as an additional player whose actions correspond to the possible states of the device). This is, from the empirical perspective, a substantial assumption. In practice, it means the social scientist must either have strong priors that correlated strategies are unlikely in the situation under analysis or, alternatively, he must be able to control for any variables likely to induce them.

4.3 Measures and empirical distributions

At this point, consider things from the perspective of the intrepid social scientist faced with a game of unknown information structure. For example, imagine she has a data set in which each point is simply a list of actions taken by each player with no indication of timing or conditioning information. To her, the strategies adopted by the players, $\mathbf{\Theta}$, are unobserved parameters that determine an empirical probability distribution over action profiles, $\mathbf{A}$. Suppose her data set is sufficiently rich to permit an accurate estimate of the true empirical distribution. We wish to formalize this case.

Notationally, it will be easier if we use Z to describe the underlying probability space.[9] Then, given a behavior strategy profile $\mathbf{\Theta} \in \mathbf{\Sigma}$, let $\mu_{\mathbf{\Theta}}$ denote the probability measure on $\left(Z, 2^Z\right)$ defined in the following way:

$$\forall W \in 2^Z,\ \mu_{\mathbf{\Theta}}\left(W\right) \equiv \sum_{z \in W} \prod_{k \in N} \tilde{\theta}_k\left(z\right). \tag{7}$$

We refer to $\mu_{\mathbf{\Theta}}$ as the *empirical distribution induced by* $\mathbf{\Theta}$. Note that the functions earlier defined on Z (e.g., $\tilde{a}_k$) are measurable random variables on $\left(Z, 2^Z\right)$.

It is useful to observe that, for all $\mathbf{\Theta} \in \mathbf{\Sigma}$, the definition of conditional probability permits us to factor $\mu_{\mathbf{\Theta}}$ as follows

$$\forall z \in Z,\ \mu_{\mathbf{\Theta}}\left(z\right) \equiv \prod_{k=1}^{n} \mu_{\mathbf{\Theta}}\left(\tilde{a}_k | \tilde{h}_k\right)\left(z\right)\ \ \mu_{\mathbf{\Theta}}\text{-}a.e., \tag{8}$$

where $\mu_{\mathbf{\Theta}}\left(\tilde{a}_k | \tilde{h}_k\right)$ is the $\mu_{\mathbf{\Theta}}$-conditional distribution of $\tilde{a}_k$ given $\tilde{h}_k$. To reduce clutter from this point on, we omit the "$\mu_{\mathbf{\Theta}}$-*a.e.*" qualifier whenever it is clear from the context.

Finally, we need some additional notation for what is to follow: for all $i, k \in N$ such that $i < k$, let $\tilde{h}_{k\setminus i}(z) \equiv \{(a_1, ..., a_{k-1})\,|a_j = \tilde{a}_j(z), \forall j \in \{1, \ldots, k-1\} \setminus \{i\}\}$; i.e., the set of histories that coincides with $\tilde{h}_k(z)$ except (possibly) for the i-th action. Given $\mathbf{\Theta} \in \mathbf{\Sigma}$, from the definition of conditional probability and condition (8), for all $i, k \in N$ such that $i < k$, $\tilde{a}_k$ is $\mu_{\mathbf{\Theta}}$-independent of $\tilde{a}_i$ given $\tilde{h}_{k\setminus i}$ if

$$\mu_{\mathbf{\Theta}}\left(\tilde{a}_k | \tilde{h}_k\right) = \mu_{\mathbf{\Theta}}\left(\tilde{a}_k | \tilde{h}_{k\setminus i}\right). \tag{9}$$

[9] Note that we are assuming the social scientist has no information on the order of moves. Thus, we use the notation associated to terminal nodes to refer to states of the world (sequences of actions $(a_1, \ldots, a_n)$) without associating any meaning to the way the actions are ordered.

5 Influence opportunity diagrams

We are now equipped with two objects that are essential to the social scientist's problem: the (unknown) game form and the (known) empirical distribution. These are linked by players' strategies (Equation 7). Moreover, there is information about game form embedded in the empirical distribution stemming from the fact that it can be factored in a particular way (Equation 8). In order to exploit this information we introduce a new tool, the influence opportunity diagram. As we see from the following definition, the construction is quite intuitive.

Definition 1 *The* influence opportunity diagram *(hereafter, IOD) implied by a game form* Γ *is a directed graph* $(N, \rightarrow)$ *such that* $i \rightarrow k$ *if and only if* $i < k$, $l_k > 1$ *and there exist* $z \in Z$ *and* $z' \in \tilde{h}_{k\backslash i}^{-1}\left(\tilde{h}_{k\backslash i}(z)\right)$*, such that* $\tilde{X}_k(z) \neq \tilde{X}_k(z')$*.*

Hence, $i \rightarrow k$ if and only if there exists at least one move-k history $(a_1, ..., a_i, ..., a_{k-1})$ such that the player moving at i has the ability to – unilaterally – place the player moving at k into a different information set simply by deviating to some other action (i.e., holding other players' actions constant). When $i \rightarrow k$, influence (in the sense that k is able to condition his move on i's behavior) is only an "opportunity" because player k may always choose to ignore the actions of player i. Neither is $i \rightarrow k$ required for the appearance of correlation in an empirical distribution since player i may affect player k's behavior indirectly through some other player or players. In Section 3 we introduced intuitive graphical depictions of the interactions between players. These intuitive depictions are the IODs of the corresponding game forms.

Example: IOD corresponding to the standard signalling game form

To further illustrate, consider the extensive form of a standard signalling-style game (Figure 2). The game proceeds in the usual way: player 2 observes player 1's move (usually, a "Nature" player) before he chooses his own; subsequently, player 3, after observing player 2's move, chooses her action. The set of possible action profiles is $\mathbf{A} = \{U, D\} \times \{L, R\} \times \{u, d\}$ with elements $\tilde{a}(z_1) = (U, L, u)$, $\tilde{a}(z_6) = (U, R, d)$, etc. Player 1 moves first and no one has the opportunity to influence him. Player 2's information is defined by the partition $\mathcal{X}_2 = \{X_{2,1}, X_{2,2}\}$ where $X_{2,1} = \{x_2\}$ and $X_{2,2} = \{x_3\}$. To construct the IOD, pick a terminal node, say z_1, and note the following: $\tilde{X}_2(z_1) = X_{2,1}$, $\tilde{h}_2(z_1) = (U)$ and $\tilde{h}_{2\backslash 1} = \tilde{h}_{\emptyset}$. Since $\tilde{h}_{\emptyset}$ is a constant, $\tilde{h}_{\emptyset}^{-1}\left(\tilde{h}_{\emptyset}(z_1)\right) = Z$. In this case we merely need to find a terminal node z' such that $\tilde{X}_2(z') \neq X_{2,1}$. Hence, any $z' \in \{z_3, z_4, z_7, z_8\}$ suffices to show that $1 \to 2$.

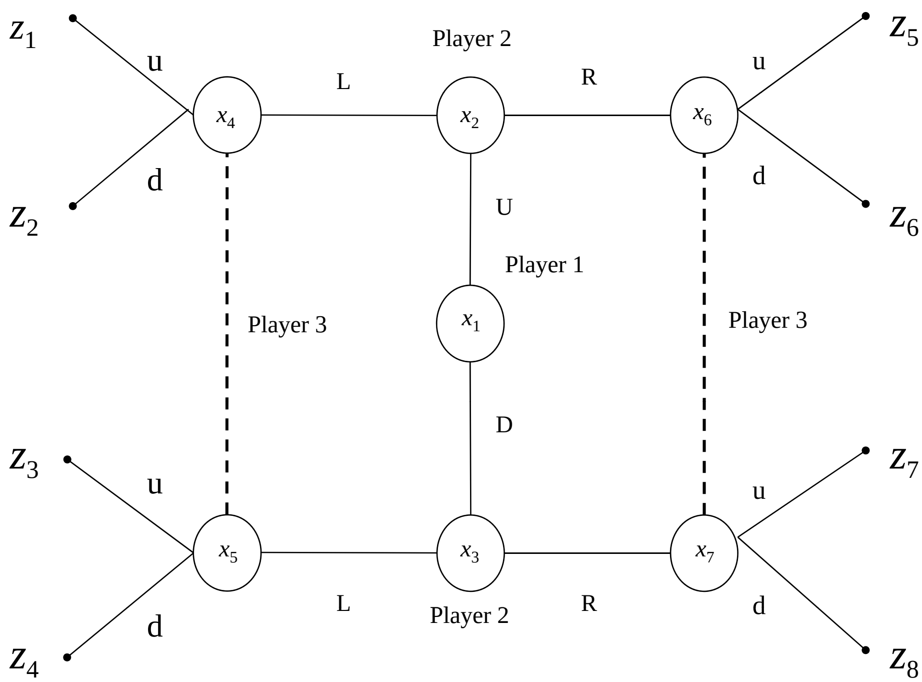

Figure 2: a signalling-style game form

Now, consider player 3. This player has two information sets: $X_{3,1} = \{x_4, x_5\}$ and $X_{3,2} = \{x_6, x_7\}$. First, check $2 \to 3$. Pick the same terminal node as before, z_1. We have $\tilde{X}_3(z_1) = \{x_4, x_5\}$, $\tilde{h}_3(z_1) = (U, L)$, $\tilde{h}_{3\setminus 2}(z_1) = (U)$ and $\tilde{h}^{-1}_{3\setminus 2}\left(\tilde{h}_{3\setminus 2}(z_1)\right) = \{z_1, z_2, z_5, z_6\}$. Choosing z_5 from the previous set and noting that $\tilde{X}_3(z_5) = \{x_6, x_7\} \neq \tilde{X}_3(z_1)$, we establish $2 \to 3$. Finally, check $1 \to 3$. Pick an arbitrary $z \in Z$. Either $\tilde{h}^{-1}_{3\setminus 1}\left(\tilde{h}_{3\setminus 1}(z)\right) = \{z_1, z_2, z_3, z_4\}$ or $\tilde{h}^{-1}_{3\setminus 1}\left(\tilde{h}_{3\setminus 1}(z)\right) = \{z_5, z_6, z_7, z_8\}$. In both cases, for all $z' \in \tilde{h}^{-1}_{3\setminus 1}\left(\tilde{h}_{3\setminus 1}(z)\right)$, $\tilde{X}_3(z) = \tilde{X}_3(z')$. Hence, $1 \not\to 3$. Thus, the IOD for this form is:

$$1 \to 2 \to 3.$$

6 Empirical distributions and the IOD

The following lemma establishes the rather obvious result that the conditional probability of $\tilde{a}_k$ given a history $\tilde{h}_k$ is simply the value of the probability parameter on the edge associated with $\tilde{a}_k$ following the history $\tilde{h}_k$.

Lemma 1 *Given a game form Γ, for all $k \in N$, $\mathbf{\Theta} \in \mathbf{\Sigma}$, $z \in Z$,*

$$\mu_{\mathbf{\Theta}}\left(\tilde{a}_k|\tilde{h}_k\right)(z) = \tilde{\theta}_k(z).$$

The next lemma demonstrates a key fact linking pair-wise relations between players in the IOD and conditional independence of their actions in all empirical distributions. The proof relies on an extension of the ideas used to prove Lemma 1. Given random variables $\tilde{x}, \tilde{y}, \tilde{z}$ on $\left(Z, 2^Z, \mu_{\mathbf{\Theta}}\right)$, the notation $(\tilde{x} \perp \tilde{y}|\tilde{z})_{\mathbf{\Theta}}$ means: $\tilde{x}$ is $\mu_{\mathbf{\Theta}}$-independent of $\tilde{y}$ given $\tilde{z}$.

Lemma 2 *Given a game form Γ and its IOD $(N, \rightarrow)$, for all $k \in N$ and all $i \in \{1, ..., k-1\}$,*

$$(i \nrightarrow k) \Leftrightarrow (\forall \mathbf{\Theta} \in \mathbf{\Sigma})\left(\tilde{a}_k \perp \tilde{a}_i|\tilde{h}_{k\setminus i}\right)_{\mathbf{\Theta}}.$$

The third lemma is critical to all that follows. It establishes a direct connection between a game's IOD and certain conditional independencies that must hold in every one of its empirical distributions. Given an IOD $\mathcal{G} = (N, \rightarrow)$, let $Pa_k^{\mathcal{G}} \subset N$ denote the *parents of* k (i.e., the set of all $i \in N$ such that $i \rightarrow k$) in $\mathcal{G}$ and, let $\widetilde{pa}_k^{\mathcal{G}}$ represent the projection of $\tilde{h}_k$ into the dimensions indexed by $Pa_k^{\mathcal{G}}$—the superscript denoting the IOD will be dropped whenever there is no ambiguity about the corresponding IOD. The notation $\tilde{h}_{k\setminus Pa_k}$ has the obvious meaning.

Lemma 3 *Given a game form Γ and its IOD $(N, \rightarrow)$, for all $k \in N$ and all $\mathbf{\Theta} \in \mathbf{\Sigma}$,* $\left(\tilde{a}_k \perp \tilde{h}_{k\setminus Pa_k}|\widetilde{pa}_k\right)_{\mathbf{\Theta}}$.

This brings us to one of our main results. Specifically, every empirical distribution generated in the play of a game can be factored in accordance with the parental relations embodied in its IOD. This result provides the key that allows us to compare the conditional independencies implied by one game form with those of another simply by referring to their IODs (i.e., rather than trying to analyze the empirical distributions directly).

Proposition 1 *Given a game form* Γ, *for all* $\mathbf{\Theta} \in \mathbf{\Sigma}$, $\mu_{\mathbf{\Theta}}$ admits recursive factorization according to $(N, \rightarrow)$. *That is,*

$$\forall \mathbf{\Theta} \in \mathbf{\Sigma}, \mu_{\mathbf{\Theta}} = \prod_{k=1}^{n} \mu_{\mathbf{\Theta}} \left(\tilde{a}_k | \widetilde{pa}_k \right). \tag{10}$$

Proof. Apply Lemma 3 to condition (8).[10] ■

A key property of the IOD is that it is "minimal" in the following sense: there exist strategies that generate empirical distributions for condition (10) which fail if one removes any element of Pa_k for any $k \in N$. More plainly, there are strategy profiles that generate empirical distributions with the following property: remove any edge from the IOD and the new graph is not consistent with the empirical distribution in the sense of condition (10). Such distributions are said to be *maximally revealing*.[11]

Most empirical distributions have more independencies than the ones implied by (10). This is because players may choose to ignore some of the information available to them, thereby generating independencies above and beyond those enforced by the game form. The canonical example of a distribution that has more independencies than those captured by the IOD is the one in which all players ignore whatever information they have when making their move and behave as if they are in a simultaneous game, in which case $\mu_{\mathbf{\Theta}} = \prod_{k=1}^{n} \mu_{\mathbf{\Theta}} \left(\tilde{a}_k \right)$.

7 Refutation of game form hypotheses

Given data on a strategic interaction it would be ideal to find some feature that would allow one to uniquely identify the game form from which it arose without resorting to specific behavioral assumptions (e.g., Nash equilibrium). Unfortunately, this goal is too ambitious. Rather, the best we can hope for in this context is to identify a *class* of game forms that are capable of generating the observed behavior.

We now analyze this issue using the concepts and results above. We start from a given game form and study the properties of the set of empirical dis-

[10] See Appendix A, for additional details on recursive factorization.

[11] A maximally revealing empirical distribution is one for which $(N, \rightarrow)$ is said to be a *perfect map* (Geiger et al. 1990). It can be shown (using results in Meek 1997) that for any $(N, \rightarrow)$ corresponding to the game form of an extensive form game, there is a distribution for which its IOD is a perfect map. The corresponding strategy profile is identified by the conditional probability parameters in (10).

tributions it could generate. One form is said to be "empirically compatible" to another when any empirical distribution generated by a strategy profile in the latter can be replicated by an appropriately chosen strategy profile in the former. Clearly – even with a perfectly accurate estimate obtained from a perfectly revealing empirical distribution – one can do no better than to narrow things down to the generating form's empirical compatibility class. As we illustrate below, certain measurability issues inherent in game structure conspire to make this classification problem a messy one to solve for the general case.

On the brighter side, Proposition 1 suggests the simpler possibility of refuting game form hypotheses by testing estimated empirical distributions for inconsistencies in the required set of conditional independencies. Moreover, because a game form's IOD provides a very compact representation of these independencies, the IOD appears to hold some promise as a useful empirical device. Of course, the very simplicity of the IOD also defines the limits of its ability to completely identify the empirical compatibility class. Exploring these issues is the task to which we now turn.

7.1 Empirical relationships between game forms

A game form Γ' is said to be *empirically compatible to another form* Γ when the empirical distribution induced by any strategy profile in Γ can also be induced by an appropriately chosen strategy profile in Γ'. The social scientist observing outcomes generated by independent repeated play of $\mathbf{\Theta}$ in Γ will, eventually, develop a fairly precise estimate of $\mu_{\mathbf{\Theta}}$. However, if Γ' is empirically compatible to Γ, then there is no collection of Γ-generated data capable of ruling out Γ' as the true underlying game – that is, no data from Γ will allow us to rule out the possibility that the true game form is Γ'.

An obvious necessary condition for empirical compatibility is that the games have consistent player sets and outcome profiles. Given a permutation $f : N \to N$, let $f(\mathbf{a})$ denote the permuted profile $\left(a_{f(k)}\right)_{k \in N}$. Then, Γ and Γ' are said to be *outcome equivalent* if there exists an f such that $f(\mathbf{A}) = \mathbf{A}'$. Note that when Γ and Γ' are outcome equivalent there exists a bijection between $\mathbf{A}$ and $\mathbf{A}'$ and, hence, between Z and Z'. Thus, every $\mathbf{\Theta}' \in \mathbf{\Sigma}'$ implies a distribution on $\left(Z, 2^Z\right)$ which we denote by $\mu_{\mathbf{\Theta}'}$; this distribution is constructed as in (7) but using the appropriate permutation. Outcome equivalence is an equivalence relation.

Definition 2 *Game form Γ' is* empirically compatible to form Γ, *denoted* $\Gamma \preceq \Gamma'$, *if* Γ *and* Γ' *are outcome equivalent and there exists a function* $g : \mathbf{\Sigma} \to \mathbf{\Sigma}'$ *such that* $\forall \mathbf{\Theta} \in \mathbf{\Sigma}$, $\mu_{\mathbf{\Theta}} = \mu_{g(\mathbf{\Theta})}$.

If $\Gamma \preceq \Gamma'$ and $\Gamma' \preceq \Gamma$, then Γ and Γ' are said to be *empirically equivalent*, denoted $\Gamma \sim \Gamma'$. When Γ' and Γ are empirically equivalent, any behavior observed under Γ can also be observed under Γ' and viceversa. Empirical equivalence does *not* imply identical information structures; e.g., in Section 3.1, the complete information game in which I is observed by II is empirically equivalent to the one in which II is observed by I.

Proposition 2 *Empirical equivalence is an equivalence relation on the space of finite-length extensive forms. Moreover, if* $\Gamma \sim \Gamma'$, *there exists an onto correspondence* $g : \mathbf{\Sigma} \rightrightarrows \mathbf{\Sigma}'$ *such that* $\forall \mathbf{\Theta} \in \mathbf{\Sigma}, \mathbf{\Theta}' \in g(\mathbf{\Theta}), \mu_{\mathbf{\Theta}} = \mu_{\mathbf{\Theta}'}$.

Note that all extensive form games with the same game form are empirically equivalent so that empirical equivalence is also an equivalence relation on the space of finite-length extensive form games.

Empirical compatibility is a "strong" notion in two senses. First, the condition on empirical distributions must hold for *all* behavior strategy profiles in the target game form. Second, the definition requires measure *equivalence*; i.e., every empirical distribution in the target form must be exactly replicable in every detail. The first requirement is useful in the context of testing game form hypotheses in lieu of specific behavioral assumptions. However, we can imagine replacing the measure equivalence requirement with a weaker one pertaining to conditional independencies. This is the motivation behind the following definition.[12]

Definition 3 *Given two game forms,* Γ *and* Γ', *let* $\mathcal{G}$ *be the IOD of* Γ *and* $\mathcal{F}$ *that of* Γ'. *Then,* Γ *and* Γ' *are* independence equivalent, *denoted* $\Gamma \doteqdot \Gamma'$, *if* Γ *and* Γ'

[12] Note that $\widetilde{pa}_k^{\mathcal{F}}$ in the first equation of the definition of independence equivalence is shorthand for the actions of the players who are the parents of move $f(k)$ in $\mathcal{F} \equiv (N, \to')$, appropriately indexed by the move made in Γ (where f is the player permutation required by outcome equivalence: $f(\mathbf{A}) = \mathbf{A}'$). Formally, $\widetilde{pa}_k^{\mathcal{F}}$ is the profile of actions in Γ indexed by $f^{-1}\left(Pa_{f(k)}^{\mathcal{F}}\right)$. The same shorthand is used for $\widetilde{pa}_k^{\mathcal{G}}$ in the second equation of the definition and in every other equation where the parents of a move in one game form are substituted for the parents of the equivalent move in another, outcome equivalent, game form.

are outcome equivalent and

$$\begin{aligned} \forall \boldsymbol{\Theta} &\in \boldsymbol{\Sigma}, \mu_{\boldsymbol{\Theta}} = \prod_{k=1}^{n} \mu_{\boldsymbol{\Theta}} \left(\tilde{a}_k | \widetilde{pa}_k^{\mathcal{F}} \right), \textit{ and} \\ \forall \boldsymbol{\Theta}' &\in \boldsymbol{\Sigma}', \mu_{\boldsymbol{\Theta}'} = \prod_{k=1}^{n} \mu_{\boldsymbol{\Theta}'} \left(\tilde{a}_k | \widetilde{pa}_k^{\mathcal{G}} \right). \end{aligned}$$

Independence equivalence is clearly an equivalence relation. Two game forms are in the same independence equivalence class when their IODs imply the same minimal sets of conditional independencies in all empirical distributions (see Appendix A). Obviously, when this condition holds, one form cannot be distinguished from another simply by testing an observed empirical distribution for inconsistent dependencies. When it does not, there then arises at least the possibility that such a distinction can be made.

7.2 Empirical vs. independence equivalence

Before returning to our analysis, let us illustrate these concepts with some examples. Recall the signalling game form (Figure 1) with IOD $1 \rightarrow 2 \rightarrow 3$. Forms that are outcome equivalent to this one are those that have three players, one with action set $\{U, D\}$, another with $\{u, d\}$ and a third with $\{L, R\}$. Among all the possible outcome equivalent game forms there is at least one that is empirically compatible: it is the signalling game form with the addition that player 3 also observes player 1. This game has IOD

$$\begin{array}{ccc} 1 & \rightarrow & 2 \\ & \searrow & \downarrow \\ & & 3 \end{array}$$

Clearly, any strategy in this form can be replicated in this game as player 3 can always ignore player 1. The converse is not true.

A more subtle situation occurs with the game form shown in Figure 3 in which both player 1 and player 3 observe player 2 ($1 \leftarrow 2 \rightarrow 3$). Consider an arbitrary strategy in this game, $\boldsymbol{\Theta}'$. Equating moves with player labels, and replacing histories with the appropriate actions, Proposition 1 says $\mu_{\boldsymbol{\Theta}'}$ can be factored as follows:

$$\mu_{\boldsymbol{\Theta}'} = \mu_{\boldsymbol{\Theta}'}(\tilde{a}_2) \mu_{\boldsymbol{\Theta}'} \left(\tilde{a}_1 | \tilde{a}_2 \right) \mu_{\boldsymbol{\Theta}'} \left(\tilde{a}_3 | \tilde{a}_2 \right).$$

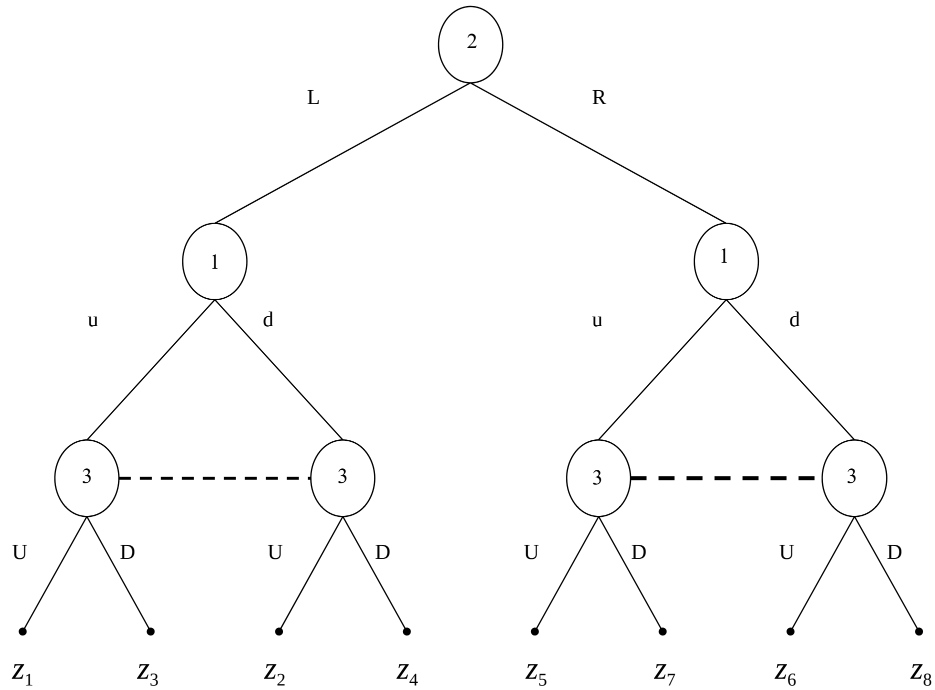

Figure 3: this form is outcome compatible with the signalling form

Applying Bayes rule and multiplying by $\frac{\mu_{\Theta'}(\tilde{a}_1)}{\mu_{\Theta'}(\tilde{a}_1)}$,

$$\begin{aligned} \mu_{\Theta'} &= \mu_{\Theta'}(\tilde{a}_2)\mu_{\Theta'}(\tilde{a}_1|\tilde{a}_2)\,\mu_{\Theta'}(\tilde{a}_3|\tilde{a}_2)\,, \\ &= \mu_{\Theta'}(\tilde{a}_2)\frac{\mu_{\Theta'}(\tilde{a}_1,\tilde{a}_2)}{\mu_{\Theta'}(\tilde{a}_2)}\mu_{\Theta'}(\tilde{a}_3|\tilde{a}_2)\,, \\ &= \mu_{\Theta'}(\tilde{a}_1)\frac{\mu_{\Theta'}(\tilde{a}_1,\tilde{a}_2)}{\mu_{\Theta'}(\tilde{a}_1)}\mu_{\Theta'}(\tilde{a}_3|\tilde{a}_2)\,, \\ &= \mu_{\Theta'}(\tilde{a}_1)\mu_{\Theta'}(\tilde{a}_2|\tilde{a}_1)\,\mu_{\Theta'}(\tilde{a}_3|\tilde{a}_2)\,. \end{aligned} \tag{11}$$

In the signalling game form $(1 \to 2 \to 3)$, the empirical distribution generated by an arbitrary strategy $\mathbf{\Theta}$ can always be factored as

$$\mu_{\Theta} = \mu_{\Theta}(\tilde{a}_1)\mu_{\Theta}(\tilde{a}_2|\tilde{a}_1)\,\mu_{\Theta}(\tilde{a}_3|\tilde{a}_2)\,. \tag{12}$$

The joint implication of (11) and (12) is that any empirical distribution generated under the form in Figure 3 can also be generated in the game form of the signalling game. Specifically, given $\mathbf{\Theta}'$, construct $\mathbf{\Theta}$ by setting, for all $z \in Z$,

$$\tilde{\theta}(z) = (\mu_{\mathbf{\Theta}'}(\tilde{a}_1)(z), \mu_{\mathbf{\Theta}'}(\tilde{a}_2|\tilde{a}_1)(z), \mu_{\mathbf{\Theta}'}(\tilde{a}_3|\tilde{a}_2)(z)).$$

Then, $\mu_{\mathbf{\Theta}'} = \mu_{\mathbf{\Theta}}$. Similarly, for any strategy, $\mathbf{\Theta}$ in the signalling game form, the converse is also true. This implies the empirical *equivalence* of the two forms (and, hence, independence equivalence as well).

Using this reasoning, it can be shown that the empirical equivalence class of the signalling game form includes three elements: signalling; the form in Fig. 3; and, the form in which 3 moves first, 2 observes 3 and 1 observes 2 (IOD $1 \leftarrow 2 \leftarrow 3$).[13] Notice that, here, the class of empirically equivalent game forms coincides with the class of game forms that are independence equivalent.

To see a case in which independence equivalence holds but empirical equivalence fails, consider the outcome compatible games in Figure 4. Clearly, every empirical distribution arising in Γ' can be replicated by one in Γ (by having player 2 play identically to M and R). The converse is not true. Therefore, $\Gamma' \preceq \Gamma$ but $\Gamma \npreceq \Gamma'$. Both have the same IOD: $1 \rightarrow 2$. This means that, even under maximally revealing strategy profiles, both forms imply the same set of conditional independencies (none). That is, for all $\mathbf{\Theta}$ and $\mathbf{\Theta}'$

$$\begin{aligned} \mu_{\mathbf{\Theta}} &= \mu_{\mathbf{\Theta}}(\tilde{a}_2|\tilde{a}_1)\,\mu_{\mathbf{\Theta}}(\tilde{a}_1), \textit{ and} \\ \mu_{\mathbf{\Theta}'} &= \mu_{\mathbf{\Theta}'}(\tilde{a}_2|\tilde{a}_1)\,\mu_{\mathbf{\Theta}'}(\tilde{a}_1). \end{aligned}$$

Thus, the two games *are* independence equivalent.

The main issue that makes empirical equivalence difficult to characterize is the relative measurability of the strategy variables $\tilde{\theta}_2$ and $\tilde{\theta}'_2$. To see this, let $\sigma(\cdot)$ indicate the σ-algebra generated by either a collection of sets or a random variable. Define $\mathcal{I}_2 \equiv \{\{z_1, z_2\}, \{z_3, z_4\}, \{z_5, z_6\}\}$ and $\mathcal{I}'_2 \equiv \{\{z_1, z_2\}, \{z_3, z_4, z_5, z_6\}\}$. Note that $\sigma(\mathcal{I}'_2) \subset \sigma(\mathcal{I}_2)$. In both games $1 \rightarrow 2$ so that $\widetilde{pa}_2^{\mathcal{G}} = \widetilde{pa}_2^{\mathcal{F}} = \tilde{a}_1$ and, hence, $\sigma(\widetilde{pa}_2^{\mathcal{G}}) = \sigma(\widetilde{pa}_2^{\mathcal{F}})$. The following is true for Γ':

$$\forall \mathbf{\Theta}' \in \mathbf{\Sigma}', \mu_{\mathbf{\Theta}'} = \mu_{\mathbf{\Theta}'}\left(\tilde{a}_2|\sigma\left(\mathcal{I}'_2\right)\right)\mu_{\mathbf{\Theta}'}(\tilde{a}_1). \tag{13}$$

[13] Keep in mind that behavioral assumptions typically refine the observational equivalence class. For example, if it is known in advance that player 1 cannot condition his behavior on any other (e.g., he represents "nature" whose move is required to be independent of all player actions), then the form in Figure 3 is immediately eliminated from consideration.

Since $\sigma(\mathcal{I}_2) = \sigma(\tilde{\pi}_2) \supset \sigma(\mathcal{I}_2')$, for any $\tilde{\theta}'$ we can define a $\tilde{\theta}$ by setting $\tilde{\theta}_2(z) = \tilde{\theta}_2'(z)$ and $\tilde{\theta}_1(z) = \tilde{\theta}_1'(z)$. The resulting $\tilde{\theta}$ is a valid strategy in Γ and generates the same empirical distribution as $\tilde{\theta}'$, so $\Gamma' \preceq \Gamma$. On the other hand, the argument does not work in reverse: $\sigma(\mathcal{I}_2') \subset \sigma(\mathcal{I}_2)$ implies that there are strategies available to player 2 in Γ that are not valid strategies for player 2 in Γ', so $\Gamma \not\preceq \Gamma'$.

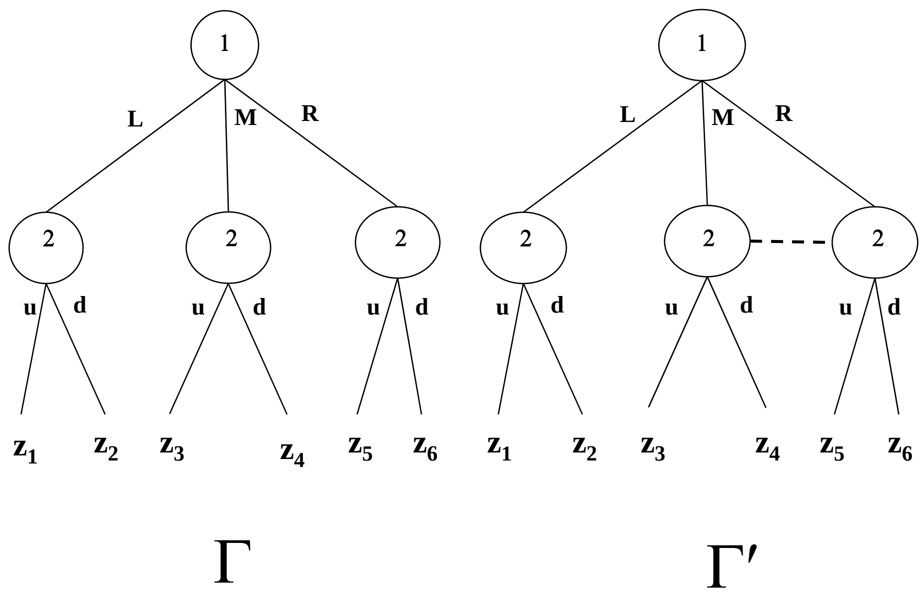

Figure 4: strategy measurability issue

Thus, a social scientist who observes outcomes generated under independent repetitions of a behavior strategy in Γ' cannot – without additional behavioral assumptions – refute Γ no matter how large the sample size. Conversely, rejecting Γ' in favor of Γ by observing differences in player 2's responses to M and R is a distinct possibility.

Identifying measurability differences of the kind arising in Figure 4 is too much to ask of our simple IODs. That being said, IODs do provide a necessary (and, hence, testable) condition for empirical compatibility. The following proposition is an almost immediate corollary to Proposition 1.

Proposition 3 *Consider two outcome equivalent game forms,* Γ *and* Γ'*, the latter with IOD* $\mathcal{F}$*. If* $\Gamma \preceq \Gamma'$ *then* Γ *admits recursive factorization according to* $\mathcal{F}$*. That is,*

$$\forall \mathbf{\Theta} \in \mathbf{\Sigma}, \mu_{\mathbf{\Theta}} = \prod_{k \in N} \mu_{\mathbf{\Theta}} \left(\tilde{a}_k | \widetilde{pa}_k^{\mathcal{F}} \right). \tag{14}$$

There remains the interesting question of whether there is some class of game forms for which the information contained in the IOD is both necessary and sufficient to determine empirical equivalence. The following proposition answers this question. We begin with a definition in which the phrase "k observes i perfectly" means: every pair of k-length histories that differ only on the i-th action belong to different information sets: $\forall z, z' \in Z$ such that $z' \in \tilde{h}_{k\setminus i}^{-1}\left(\tilde{h}_{k\setminus i}(z)\right)$ and $\tilde{a}_i(z) \neq \tilde{a}_i(z')$, $\tilde{X}_k(z) \neq \tilde{X}_k(z')$.

Definition 4 $\Gamma = (N, \mathcal{X}(X, \mathbf{E}), \mathbf{A})$ *is a game form of* semi-perfect information *if* $i \rightarrow k$ *implies* k *observes* i *perfectly.*

This is a generalization of the notion of perfect information. In perfect information forms, each player perfectly observes the move of *every* player who precedes her. Semi-perfect game forms allow for cases in which players do not observe the actions of arbitrary predecessors. From a purely theoretical perspective, the semi-perfect information condition is fairly restrictive. In practical terms, however, this class contains forms from many games of economic interest, including: market games such as Cournot, Bertrand, and Stackleberg; all simultaneous-move games; perfect information games; finite-horizon bargaining, war of attrition, and preemption games; and all finite-horizon Bayesian games with observed actions such as signalling, public good games, simultaneous-bid double auction games, first-price and other auctions, and principle-agent games.[14] All the forms in Section 3 are of semi-perfect information.

Proposition 4 *Given an n-dimensional space of actions,* $\mathbf{A}$*, and an IOD,* $\mathcal{G} = (N, \rightarrow)$*, there exists a game form* $\Gamma_{\mathcal{G}} = (N, \mathcal{X}, (X, \mathbf{E}), \mathbf{A})$ *such that a probability measure* μ *on* $\left(Z, 2^Z\right)$ *admits recursive factorization according to its IOD if and only if there exists a strategy,* $\mathbf{\Theta} \in \mathbf{\Sigma}_{\mathcal{G}}$*, such that* $\mu = \mu_{\mathbf{\Theta}}$*. Moreover,* $\Gamma_{\mathcal{G}}$ *is a game form of semi-perfect information.*

[14] As noted earlier, the following proposition also applies to forms in which players choose points in real intervals.

This result says that, given a space of action profiles and a corresponding (i.e., n-node) IOD, there exists a game form with the property that every empirical distribution it generates can be factored according to the (given) IOD and every empirical distribution that is factored according to the IOD is an empirical distribution generated by some strategy from that game form. Moreover, this form is one of semi-perfect information.[15] It follows immediately that the set of forms for which independence equivalence is necessary and sufficient for empirical equivalence is the one characterized by semi-perfect information.

Corollary 1 *Suppose* Γ *and* Γ' *are outcome equivalent forms of semi-perfect information. Then,* $\Gamma \sim \Gamma'$ *if only if* $\Gamma \doteqdot \Gamma'$.

7.3 Identifying independence equivalence via the IOD

Let us now return to the general case and focus entirely upon the idea of refuting game form hypotheses on the basis of the minimal set of conditional independencies they imply. Presumably, because the IOD provides a complete summary of this information, one should be able to determine whether two game forms belong to the same independence equivalence class by comparing their IODs. This is true and, as it turns out, the procedure is remarkably simple.

First, given an IOD $(N, \rightarrow)$, let $\mathcal{E} \subset N^2$ be the set of undirected edges; i.e., $(i, j) \in \mathcal{E}$ if and only if $(i \rightarrow j)$ or $(j \rightarrow i)$. Let $\mathbf{S} \subset N^3$ 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 $(i, k) \notin \mathcal{E}$ (i.e., i and k can influence j but not each other). The following proposition confirms that independence equivalence is fully characterized by the IOD. Moreover, such equivalence can be ascertained simply via pairwise inspection of the IODs.[16]

Proposition 5 *Assume* Γ *and* Γ' *are outcome equivalent forms with IODs* $(N, \rightarrow)$ *and* $(N, \rightarrow')$, *respectively. Then,* $\Gamma \doteqdot \Gamma'$ *if only if* $(N, \rightarrow)$ *and* $(N, \rightarrow')$ *satisfy* $\mathcal{E} = \mathcal{E}'$ *and* $\mathbf{S} = \mathbf{S}'$.

[15] It is worth adding that the resulting form is unique once the order of play is fixed.

[16] This last result raised a question that was put to us by E. Dekel in correspondence. Is there a set of operations, similar in spirit to Thompson (1952) and Elmes & Reny (1994), that yields games that are independence equivalent? Due to space limitations, we do not provide a formal reply. Clearly, however, Proposition 5 does suggest a step-wise transformation to the extensive form that causes an "allowed" arrow flip.

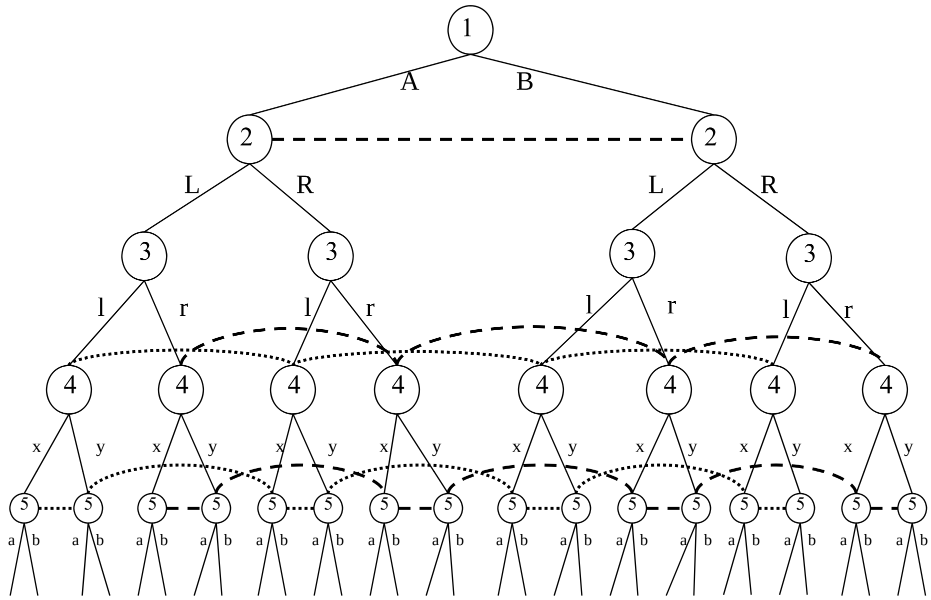

Figure 5: the Gatekeeper game form

In other words, two game forms are independence equivalent if and only if their IODs have the same edges and pairs of converging arrows from unlinked nodes. In simple cases like those in our preceding examples, independence equivalence can be checked via "brute-force" comparisons using Bayes' rule (i.e., without reference to the IOD). On the other hand, consider the Gatekeeper form shown in Figure 5. Here, 5 players interact under a relatively complex information structure. The implications of this form with respect to independence equivalence are not immediately obvious. Presumably, most analysts would not relish the thought of grinding pair-wise through the set of outcome-equivalent games in order to identify Gatekeeper's independence equivalence class. In a case such as this, we see the value of the IOD. The IOD for Gatekeeper is shown below.

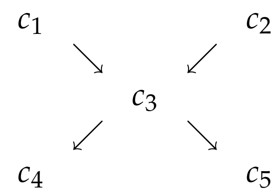

We see where the game gets its name: player 3 is the "gatekeeper" of information flowing from players 1 and 2 to players 4 and 5. By Proposition 5, any form with which this one is independence equivalent must have exactly the same IOD. To see this, note that any independence equivalent game form would have an IOD with the same set of edges, some with different directions. However, reversing any arrow in the Gatekeeper IOD either breaks a converging pair of arrows or creates a new one.[17] Therefore, Gatekeeper's independence equivalence class is uniquely characterized by its own IOD. All independence equivalent forms have the exactly the influence relations: 1 and 2 influence 3, 3 influences 4 and 5. Because each player only has two moves, these relations are uniquely arranged via the (semi-perfect information) form shown in Figure 5 – Gatekeeper has no independence equivalent counterparts.

8 Discussion

The preceding analysis is motivated by the idea that an outside observer with sufficiently informative empirical data should be able to test a game form hypotheses on the basis of the stochastic regularities it implies. In particular, we focus upon the minimal set of conditional independencies determined by a game form in every one of its empirical distributions. Can these results be extended to implement useful empirical tools? While we are optimistic that this goal is attainable, we must also acknowledge that several substantive issues remain ahead. Among these, first, is the problem of hidden variables and, second, is the need to be specific about the mechanism assumed to be generating actual data. Let us discuss these in turn.

We take the standard theoretical view that the game form fully describes the physical rules of a given strategic interaction (e.g., Rubinstein, 1991) and, crucially, that any random event observed by the players is explicitly

[17] Although not required by Proposition 5, it should be noted that Gatekeeper is also a form of semi-perfect information.

represented in the game tree (e.g., Kohlberg & Mertens, 1986). This means that, technically, our findings only hold in settings where the social scientist both identifies and observes all significant variables upon which players condition their behavior. In the real world, such situations are, at best, exceedingly rare. The problem with this is that missing or hidden variables have the potential to introduce spurious dependencies among the observed variables and cause the unwarranted rejection of the true game form.

To see this, reconsider the quantity market game (p. 10) in which firm 3 observes the choices of firms 1 and 2 ($1 \rightarrow 3 \leftarrow 2$). Suppose this game is repeated with the added twist that, at the conclusion of each period, firms observe each other's payoffs (say, from public financial reports). If previous period payoffs enter current period strategies, failing to account for them can give rise to an empirical distribution consistent with the alternative IOD:

$$\begin{array}{ccc} 1 & \rightarrow & 2 \\ & \searrow & \downarrow \\ & & 3 \end{array}$$

Applying Proposition 5 one can see that the independence equivalence class associated with this IOD contains all 8 fully connected graphs – each of which exhibits dependencies that are at odds with the true game form.

Hidden variable complications arise in the empirical investigation of most game theory models – our work is not unique in this regard. Clearly, empirical specifications should include all potentially significant, observable influences. In the preceding example, the sensible IOD is

$$\begin{array}{ccccc} & & \pi_{t-1} & & \\ & \swarrow & \downarrow & \searrow & \\ q_{1,t} & \rightarrow & q_{3,t} & \longleftarrow & q_{2,t} \end{array}$$

where π_{t-1} is lagged public information and $q_{i,t}$ the quantity decision of firm i at date t. Since the direction of influence from π_{t-1} can be only one way, Proposition 5 says that this IOD uniquely defines the independence equivalence class.[18] Unobserved influences are, of course, more problematic. Nevertheless, in the econometrics literature there are sophisticated empirical techniques for dealing with hidden variable problems. Furthermore, the topic of hidden variables receives substantial attention in the probabilistic network literature—including, for example, the construction of generalized graphs designed to admit them and the development of associated

[18] The direction of arrows from π_{t-1} is fixed and $q_{1,t}$ and $q_{2,t}$ are not directly linked (hence, the converging pair of arrows from these to $q_{3,t}$ cannot be flipped).

results. Both lines of research make us hopeful that these issues are surmountable.

A second substantive issue that arises is our interpretation of the statistical process generating the empirical distribution. In this paper we have rather narrowly interpreted the empirical distribution as the outcome of play of a single (stage game) behavior strategy profile. Thus, the empiricist's data is interpreted as a large number of outcomes generated under independent plays of a stationary strategy profile. Of course, we do not imagine that many real-world situations conform to this ideal. Whether the setting is one-shot play of a game over a cross-section of players or repeated play by the same players over time, stationarity of the underlying strategy profile is very likely to be a poor assumption. In a given role, different players may very well adopt different strategies, as may a given player in the same role over time (e.g., due to learning).

Nonstationarity of stage game strategy profiles is a significant concern for applications aiming to test empirical equivalence. Interestingly, however, behavioral variation may actually help the assessment of independence equivalence. In the context of our analysis, a behavior strategy should be broadly interpreted: it is that which generates a distribution over actions given what other players are doing. A "behavior strategy" could be an introspective contingency plan, some combination of purposeful choices and accidents, a description of behaviors taken from a large population, a conjecture about expected actions, a sequence of gropes toward optimality, and so on. All we require is that the empirical distribution be generated in such a way as to preserve the conditional independencies that we posit for a strategy. Here, the key characteristic of a strategy is its empirical content, not *why* it is being played. Not only is repetition of a static strategy profile unnecessary, but behavioral variation may actually favor the generation of fully revealing 'strategies.' In learning situations, for example, groping toward optimality may cause players to repeatedly visit all paths in the game and, thereby, reveal all the conditional independencies of the underlying game form.

9 Conclusions

Having tools to test game form hypotheses would greatly enrich the empirical content of game theory. Even so, this goal has received surprisingly little attention within the profession.

In the preceding analysis, we have focused on assessing game form solely on the basis of observed behavior without invoking equilibrium assumptions. Ultimately, the most powerful techniques are likely to involve joint tests of game form, preferences over outcomes and equilibrium behavior. Taking these elements together, the number of postulated empirical distributions (i.e., those requiring examination) would drop dramatically and, hence, allow more refined analysis.

For example, a two-player form of semiperfect information with $1 \to 2$ is empirically compatible to its simultaneous-move counterpart. Thus, based on game form alone, the empirical finding that the moves of 1 and 2 are independent does not refute $1 \to 2$. However, independence does refute $1 \to 2$ under the addition of an equilibrium hypothesis that the dependency *is* realized when the influence relation does exist.[19]

Still, we believe the independent study of game form is an interesting endeavor in its own right. First, examination of game form hypotheses in the preliminary stages of a study may be useful in guiding later-stage analysis. For example, in the preceding paragraph, uncovering a dependency rejects simultaneous moves even without the addition of payoff and equilibrium assumptions. Second, game forms in the real world may be subtle. For example, investigating influence relationships in a large market setting may hold some interesting surprises. Our hope is that the ideas and findings in this paper will spur additional activity along this line and, eventually, find their way into real applications.

A Dependency models

Here, we provide a condensed discussion of the relevant underlying theory of probabilistic networks. Some of it will be used in the proof of Proposition 5 below. 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 related propositions. For those interested in pursuing these ideas further, we suggest starting with the texts by Cowell et al. (1999), Pearl (1988, 2000) and Spirtes at al. (2000).

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*

[19] A good example of this, as pointed out by an anonymous referee, is Cournot vs. Stackleberg. Although Stackleberg is empirically compatible to Cournot, a finding of independence rejects Stackleberg on the basis of the equilibrium implications.

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 originated with 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. The graphoid axioms are:

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* $(C \perp D|E) \wedge (C \perp F|E \cup D) \Rightarrow (C \perp D \cup F|E)$
5. *Intersection* $(C \perp E|D \cup F) \wedge (C \perp F|D \cup E) \Rightarrow (C \perp E \cup F|D)$

A *graphoid* is defined as any dependency model that is closed under the five graphoid axioms.

For example, given a probability space $(\Omega, \mathcal{F}, \mu)$ and an associated, finite set of random variables X indexed by $N = \{1, ..., n\}$ with typical element $\tilde{x}_r$, M_μ denotes the list of conditional independencies that hold under μ. For all $W \subseteq N$, let $\tilde{x}_W \equiv (\tilde{x}_r)_{r \in W}$. Then, for all disjoint $C, D, E \subset N$, $(C \perp D|E) \in M_\mu$ if and only if $\tilde{x}_C$ is μ -conditionally independent of $\tilde{x}_D$ given $\tilde{x}_E$ (in the main text we do not work with general dependency models and, hence, write $(\tilde{x}_C \perp \tilde{x}_D|\tilde{x}_E)_\mu$ for clarity). M_μ is a semi-graphoid (Spohn 1980); if μ assigns strictly positive probability to all $\omega \in \Omega$ (μ has full support), then M_μ is a graphoid (Geiger et al. 1990).

Alternatively, if $\mathcal{G}$ is a graph whose vertices are N, then for all disjoint $C, D, E \subset N$, $(C \perp D|E) \in M_{\mathcal{G}}$ 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 (e.g., 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 $\mathcal{G} \equiv (N, \rightarrow)$, a *path* is an ordered set of nodes $P \subseteq N$ 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 $\alpha_{r-1} \rightarrow \alpha_r \leftarrow \alpha_{r+1}$ in P. A node that starts or ends a path is not head-to-head. A path $P \subset N$ is *active by* $E \subset N$ 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* $\mathcal{G} = (N, \rightarrow)$ *is a DAG and* C, D *and* E *are disjoint subsets of* N, *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 at the beginning of the section. 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 $\mathcal{G}$. $M_{\mathcal{G}}$ is a graphoid (Geiger et al. 1990).

It is useful to characterize the relationship between probabilistic and graphical dependency models. This is done through the general notions of an *independence map* (or, I-map) and a *perfect map*.

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

Given a probability space $(\Omega, \mathcal{F}, \mu)$ and an associated, finite set of random variables $X \equiv \{\tilde{x}_1, ..., \tilde{x}_n\}$, the task of constructing a DAG $\mathcal{G} = (N, \rightarrow)$, $N = \{1, ..., n\}$, such that $M_{\mathcal{G}}$ is an I-map of M_μ is straightforward (see Geiger et al., 1990, p. 514). First, for all $r \in N$, let $U_r \equiv \{1, ..., r-1\}$ index the predecessors of $\tilde{x}_r$ according to N. Next, identify a minimal set of predecessors $Pa_r \subset N$ such that $(\{r\} \perp U_r \backslash Pa_r | Pa_r)_\mu$. This results in a set of n independence statements known as a *recursive basis drawn from* M_μ and denoted B_μ. Now, construct $(N, \rightarrow)$ such that $s \rightarrow r$ if and only if $s \in Pa_r$. The resulting graph $\mathcal{G}$, a DAG, is said to be *generated* by B_μ and $Pa_r = \{s \in N | s \rightarrow r\}$ is the set of parents of r in $\mathcal{G}$. The following theorems are from Geiger et al. (1990, Theorems 1 and 2).

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

So, given $(\Omega, \mathcal{F}, \mu)$, the DAG $\mathcal{G}$ constructed in the fashion outlined above implies $M_{\mathcal{G}}$ is an I -map of M_μ. That is, every independence statement implied by $(N, \rightarrow)$ under d -separation corresponds to a valid μ -conditional independency.[20] Moreover, no further dependencies can be deduced from the recursive basis using the semi-graphoid axioms:

Theorem 2 (closure) *Let $\mathcal{G}$ be a DAG generated by a recursive basis B. Then $M_{\mathcal{G}}$, the dependency model generated by $\mathcal{G}$, is exactly the closure of B under the semi-graphoid axioms.*

[20]Note the implication, from our Proposition 1, that a game's IOD is an I-map of every empirical distribution generated during play.

If μ above has full support the DAG $\mathcal{G}$ generated above uniquely determines the edges of every graph that is an I-map of M_μ in the following sense: An *I*-map of a dependency model M is *minimal* if the removal of any of its independence statements implies that it ceases to be an *I*-map of M. A DAG $\mathcal{G}$ is said to be *I-minimal* with respect to M if the graph resulting from the removal of any edge no longer implies an *I*-map of M.

Theorem 3 (minimal I-map) *If the dependency model M satisfies symmetry, decomposition, and intersection, then there exists a unique $\mathcal{G}_0 = (N, E_0)$ that is I-minimal with respect to M produced by connecting only those pairs (s,r) for which $(s \perp r | N \setminus \{s,r\})_M$ is FALSE, i.e.,*

$$(s,r) \notin E_0 \;\Leftrightarrow (s \perp r | N \setminus \{s,r\})_M.$$

Theorem 4 (perfect map) *Let M be a dependency model on N. M is closed under the graphoid axioms iff a graph $\mathcal{G}$ can be constructed such that $M_\mathcal{G}$ is a perfect map of M.*

Theorem 4 is proven in Paz et al. (1993) and uses the same construction as Theorem 3. Thus, given any measure μ, the graph $\mathcal{G}$ constructed from the recursive basis B_μ as defined above is a minimal *I*-map of M_μ. If μ has full support, then $M_\mathcal{G}$ is a perfect map of M_μ.

Given a graph $\mathcal{G}$ and a measure P we want to establish whether P reflects the independence statements in $M_\mathcal{G}$. Let A_k be a compact subset of $\mathbb{R}^{l_k}$, where l_k is a finite integer. Let $N = \{1, \ldots, n\}$, $n < \infty$, $\mathbf{A} \equiv \times_{k \in N} A_k$ and $\mathcal{F}$ the Borel subsets of $\mathbf{A}$ and assume P is a probability distribution on $(\mathbf{A}, \mathcal{F})$. Following Cowell et al. (1999, p. 70), P *admits recursive factorization according to* a DAG $\mathcal{G} = (N, \rightarrow)$ if there exist (σ-finite) measures μ_k over $\mathbf{A}$ and non-negative functions $v_k(\cdot,\cdot)$ defined on $A_k \times \left(\times_{j \in \Pi_k} A_j\right)$ such that $\int v_k(a_k, pa_k)\mu_k(\mathrm{d}a_k) = 1$ and P has density p with respect to the product measure $\mu = \otimes_{k \in n} \mu_k$ given by $p(\mathbf{a}) = \prod_{k \in N} v_k(a_k, pa_k)$. The main objective in Section 6 is to establish that every empirical distribution generated as a result of play in a finite extensive-form game admits recursive factorization according to its DAG (Proposition 1).

We now want to establish whether given $\mathcal{G}$ there exists μ such that M_μ is a perfect map of $M_\mathcal{G}$. The following definition is taken from Meek (1997): a distribution P is *faithful* to the graphical structure $\mathcal{G}$ if and only if exactly the independence facts true in P are entailed by the graphical structure $\mathcal{G}$, i.e. $M_P = M_\mathcal{G}$. His notion of faithfulness is equivalent to Geiger et al.'s notion of a perfect map as defined above. Meek(1997) proves the following two theorems (the second being a new proof of an existing result):

Theorem 5 *For all directed acyclic graphs $\mathcal{G}$ there exists a probability distribution P in the class of multinomial distributions which is faithful to $\mathcal{G}$.*

Theorem 6 *For all directed acyclic graphs $\mathcal{G}$ there exists a probability distribution P in the class of multivariate normal distributions which is faithful to $\mathcal{G}$.*

In Meek(1997) we further find a proof that non-faithful distributions are very rare (the parameters that correspond to non-faithful distributions are a set of 'measure zero' – we refer the interested reader to the original paper for further details).

Following Pearl (2000, p. 16-20), two DAGs $(N, \rightarrow)$ and $(N, \rightarrow')$ are said to be *observationally equivalent* if every probability distribution that can be factored in accordance with the recursive basis

$$B_{(N,\rightarrow)} \equiv \{(\{r\} \perp U_r \backslash Pa_r | Pa_r) \, | r \in N\}$$

can also be factored in accordance with

$$B_{(N,\rightarrow')} \equiv \left\{ \left(\{r\} \perp U'_r \backslash Pa'_r | Pa'_r\right) | r \in N \right\}.$$

Let $\mathcal{E} \subset N^2$ be the set of edges of $(N, \rightarrow)$ without reference to direction and let $\mathbf{S} \subset N^3$ 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 $(i, k) \notin \mathcal{E}$. The following theorem is from Pearl (2000, Theorem 1.2.8). It originally appears in Verma and Pearl (1990, Theorem 1).

Theorem 7 *Two DAGs $(N, \rightarrow)$ and $(N, \rightarrow')$ are observationally equivalent if and only if $\mathcal{E} = \mathcal{E}'$ and $\mathbf{S} = \mathbf{S}'$.*

B Proofs

B.1 Additional notation used in the proofs

The following proofs involve the manipulation of subsets of terminal nodes whose paths in the game tree coincide on certain edge labels (actions). In order to help keep things straight and minimize clutter, we adopt the notational convention that $\{Z| (\cdot)\}$ indicates the subset of Z whose elements are associated with game tree paths that are coincident with the object in $(\cdot)$. For example,

$$\{Z|\tilde{a}_k (z)\} = \left\{z' \in Z|\tilde{a}_k \left(z'\right) = \tilde{a}_k (z)\right\} = \tilde{a}_k^{-1} \left(\tilde{a}_k (z)\right), \; k \in N,$$
$$\left\{Z|a_{k,j}\right\} = \left\{z \in Z|\tilde{a}_k(z) = a_{k,j}\right\}, \; a_{k,j} \in A_k.$$

Other common, self-explanatory cases include
$\left\{Z|\tilde{h}_k(z)\right\}$, $\left\{Z|\tilde{h}_{k\backslash i}(z)\right\}$, $\left\{Z|\tilde{\pi}_k(z)\right\}$, and $\left\{Z|\tilde{h}_{k\backslash Pa_k}(z)\right\}$.

B.2 Lemma 1

We begin with the following lemma.

Lemma 4 *Given $\mathbf{\Theta} \in \mathbf{\Sigma}$, $\forall z$, and all $k \in N$*

$$\sum_{z' \in \left\{Z|\tilde{h}_k(z)\right\}} \left(\prod_{j=k}^{n} \tilde{\theta}_j(z')\right) = 1$$

and for $k = 2, \ldots, n$

$$\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_k(z)\right\}\right) = \prod_{j=1}^{k-1} \tilde{\theta}_j(z).$$

Proof. Given $\mathbf{\Theta} \in \mathbf{\Sigma}$, set $k = n$: $\forall z \in Z$, by Equation (7): $\mu_{\mathbf{\Theta}}(z) = \prod_{j=1}^{n} \tilde{\theta}_j(z)$. Suppose $X_{n,s} = \tilde{X}_n(z)$. For every $z' \in \left\{Z|\tilde{h}_n(z)\right\}$, $\tilde{h}_n(z') = \tilde{h}_n(z)$ and $\tilde{\theta}_j(z') = \tilde{\theta}_j(z)$, $j = 1, \ldots, n-1$. Recalling that l_k denotes the number of elements in A_k and that the behavior strategy at information set $X_{n,s}$ is described by the row $(\theta^n_{s,1}, \ldots, \theta^n_{s,l_n})$ in the matrix Θ_n,

$$\begin{aligned}
\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_n(z)\right\}\right) &= \sum_{z' \in \left\{Z|\tilde{h}_n(z)\right\}} \prod_{j=1}^{n} \tilde{\theta}_j(z') \\
&= \sum_{z' \in \left\{Z|\tilde{h}_n(z)\right\}} \prod_{j=1}^{n-1} \tilde{\theta}_j(z)\tilde{\theta}_n(z') \\
&= \prod_{j=1}^{n-1} \tilde{\theta}_j(z) \sum_{z' \in \left\{Z|\tilde{h}_n(z)\right\}} \tilde{\theta}_n(z') = \prod_{j=1}^{n-1} \tilde{\theta}_j(z) \sum_{i=1}^{l_n} \theta^n_{s,i}.
\end{aligned}$$

The last equality follows from the fact that $\left\{Z|\tilde{h}_n(z)\right\}$ contains all and only those $z' \in Z$ such that there exist $a_{n,i} \in A_n$ satisfying $\tilde{a}(z') = \left(\tilde{h}_n(z), a_{n,i}\right)$. By the definition of behavior strategies $\sum_{i=1}^{l_n} \theta^n_{s,i} = 1$, so that

$$\mu_{\mathbf{\Theta}}(\left\{Z|\tilde{h}_n(z)\right\}) = \prod_{j=1}^{n-1} \tilde{\theta}_j(z) \sum_{i=1}^{l_n} \theta^n_{s,i} = \prod_{j=1}^{n-1} \tilde{\theta}_j(z),$$

and

$$\sum_{z' \in \{Z|\tilde{h}_n(z)\}} \tilde{\theta}_n(z') = 1.$$

Now, set $k = n-1$. Then,

$$\begin{aligned} \{Z|\tilde{h}_{n-1}(z)\} &= \{z'|\tilde{a}(z') = (\tilde{h}_{n-1}(z), \tilde{a}_{n-1}(z'), \tilde{a}_n(z'))\} \\ \mu_{\Theta}(\{Z|\tilde{h}_{n-1}(z)\}) &= \sum_{z' \in \{Z|\tilde{h}_{n-1}(z)\}} \prod_{j=1}^{n-2} \tilde{\theta}_j(z)\tilde{\theta}_{n-1}(z')\tilde{\theta}_n(z'). \end{aligned}$$

For each $a_{n-1,i}$, $i = 1, \ldots, l_{n-1}$, let z_i be any $z' \in \{Z|\tilde{h}_{n-1}(z)\}$ such that $\tilde{h}_n(z_i) = (\tilde{h}_{n-1}(z'), a_{n-1,i})$. Because $z' \in \{Z|\tilde{h}_{n-1}(z)\}$ implies $\tilde{a}(z') = (\tilde{h}_{n-1}(z), a_{n-1,i}, \tilde{a}(z'))$ for some $a_{n-1,i} \in A_{n-1}$, the set $\{Z|\tilde{h}_{n-1}(z)\}$ can be partitioned as

$$\{\{Z|\tilde{h}_{n-1}(z)\} \cap \{Z|a_{n-1,1}\}, \ldots, \{Z|\tilde{h}_{n-1}(z)\} \cap \{Z|a_{n-1,l_{n-1}}\}\}.$$

Note that, for all $i = 1, \ldots, l_{n-1}$ and $z' \in \{Z|\tilde{h}_n(z_i)\}$, $\tilde{\theta}_{n-1}(z') = \tilde{\theta}_{n-1}(z_i)$. Therefore,

$$\begin{aligned} \mu_{\Theta}(\{Z|\tilde{h}_{n-1}(z)\}) &= \sum_{i=1}^{l_{n-1}} \sum_{z' \in \{Z|\tilde{h}_n(z_i)\}} \prod_{j=1}^{n-2} \tilde{\theta}_j(z)\tilde{\theta}_{n-1}(z_i)\tilde{\theta}_n(z') \\ &= \sum_{i=1}^{l_{n-1}} \prod_{j=1}^{n-2} \tilde{\theta}_j(z)\tilde{\theta}_{n-1}(z_i) \sum_{z' \in \{Z|\tilde{h}_n(z_i)\}} \tilde{\theta}_n(z'). \end{aligned}$$

We have just demonstrated that for every z, $\sum_{z' \in \{Z|\tilde{h}_n(z)\}} \tilde{\theta}_n(z') = 1$. This implies that, for $i = 1, ..., l_{n-1}$, $\sum_{z' \in \{Z|\tilde{h}_n(z_i)\}} \tilde{\theta}_n(z') = 1$. Hence,

$$\mu_{\Theta}(\{Z|\tilde{h}_{n-1}(z)\}) = \sum_{i=1}^{l_{n-1}} \prod_{j=1}^{n-2} \tilde{\theta}_j(z)\tilde{\theta}_{n-1}(z_i) = \prod_{j=1}^{n-2} \tilde{\theta}_j(z) \sum_{i=1}^{l_{n-1}} \tilde{\theta}_{n-1}(z_i).$$

Because all the z_is are in the same information node, $\tilde{\theta}_{n-1}$ identifies the parameters of a behavior strategy, $\sum_{i=1}^{l_{n-1}} \tilde{\theta}_{n-1}(z_i) = 1$. This implies

$$\mu_{\Theta}(\{Z|\tilde{h}_{n-1}(z)\}) = \prod_{j=1}^{n-2} \tilde{\theta}_j(z).$$

And it follows that

$$\sum_{z'\in H_{N-1}(z)} \tilde{\theta}_{n-1}(z')\tilde{\theta}_n(z') = 1.$$

By induction, then, for $k = 2, ..., n,$

$$\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_k(z)\right\}\right) = \prod_{j=1}^{k-1}\tilde{\theta}_j(z) \qquad \text{and} \qquad \sum_{z'\in\left\{Z|\tilde{h}_k(z)\right\}}\left(\prod_{j=k}^{n}\tilde{\theta}_j(z')\right) = 1.$$

For $k = 1$, $\left\{Z|\tilde{h}_k(z)\right\} = Z$ so that it is trivially true that $\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_k(z)\right\}\right) = 1$ and

$$\sum_{z'\in\left\{Z|\tilde{h}_k(z)\right\}}\left(\prod_{j=k}^{n}\tilde{\theta}_j(z')\right) = 1.$$

■

Proof of Lemma 1

Given $\mathbf{\Theta} \in \mathbf{\Sigma}$, for all $z \in Z$ and $k \in \{2, \ldots, n\}$ such that $\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_k(z)\right\}\right) > 0$ by the definition of conditional probability,

$$\begin{aligned}
\mu_{\mathbf{\Theta}}(\tilde{a}_k|\tilde{h}_k)(z) &= \frac{\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_k(z)\right\}\cap\left\{Z|\tilde{a}_k(z)\right\}\right)}{\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_k(z)\right\}\right)} \\
&= \frac{\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_{k+1}(z)\right\}\right)}{\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_k(z)\right\}\right)} \\
&= \frac{\prod_{j=1}^{k}\tilde{\theta}_j(z)}{\prod_{j=1}^{k-1}\tilde{\theta}_j(z)} \\
&= \tilde{\theta}_k(z).
\end{aligned}$$

For $k = 1$,

$$\begin{aligned}
\mu_{\mathbf{\Theta}}(\tilde{a}_k|\tilde{h}_k)(z) &= \frac{\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_k(z)\right\}\cap\left\{Z|\tilde{a}_k(z)\right\}\right)}{\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_k(z)\right\}\right)} \\
&= \frac{\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_{k+1}(z)\right\}\right)}{\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_k(z)\right\}\right)} \\
&= \frac{\tilde{\theta}_1(z)}{1} = \tilde{\theta}_1(z).
\end{aligned}$$

B.3 Lemma 2

($\Rightarrow$) By the definition of the IOD, $i \in \{1, \ldots, k-1\}$ and $i \not\rightarrow k$ implies one of the following:
(a) $l_k > 1$ and $\forall z \in Z$ and $\forall z' \in \left\{Z|\tilde{h}_{k\backslash i}(z)\right\}$, $\tilde{X}_k(z) = \tilde{X}_k(z')$, or
(b) $l_k = 1$.
Case (a): The premise implies that for all $z' \in \left\{Z|\tilde{h}_{k\backslash i}(z)\right\} \cap \{Z|\tilde{a}_k(z)\}$, $\tilde{X}_k(z) = \tilde{X}_k(z')$ and, therefore, $\tilde{\theta}_k(z) = \tilde{\theta}_k(z')$. By Lemma 1, $\forall z' \in \left\{Z|\tilde{h}_{k\backslash i}(z)\right\}$, $\mu_{\mathbf{\Theta}}(\tilde{a}_k|\tilde{h}_k)(z') = \tilde{\theta}_k(z')$. Since $\tilde{\theta}_k(z) = \tilde{\theta}_k(z')$,

$$\forall z' \in \left\{Z|\tilde{h}_{k\backslash i}(z)\right\},\ \mu_{\mathbf{\Theta}}\left(\{Z|\tilde{h}_{k+1}(z')\}\right) = \mu_{\mathbf{\Theta}}\left(\{Z|\tilde{h}_k(z')\}\right)\tilde{\theta}_k(z). \quad (15)$$

Given $\mathbf{\Theta} \in \mathbf{\Sigma}$ and $k \in N$, consider $z \in Z$ such that $\mu_{\mathbf{\Theta}}\left(\{Z|\tilde{h}_k(z)\}\right) > 0$. Given this context, let $\mathbf{H}$ denote the set of all $k-1$ length histories associated with $\left\{Z|\tilde{h}_{k\backslash i}(z)\right\}$; that is, $\mathbf{H} \equiv \left\{\tilde{h}_k(z')|\forall z' \in \left\{Z|\tilde{h}_{k\backslash i}(z)\right\}\right\}$. Noting that $\{\{Z|h\}\}_{h\in\mathbf{H}}$ is a partition of $\left\{Z|\tilde{h}_{k\backslash i}(z)\right\}$, we have $\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_{k\backslash i}(z)\right\}\right) = \sum_{h\in\mathbf{H}} \mu_{\mathbf{\Theta}}(\{Z|h\})$. Moreover, for all $z' \in \left\{Z|\tilde{h}_{k\backslash i}(z)\right\} \cap \{Z|\tilde{a}_k(z)\}$,

$$\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_{k\backslash i}(z)\right\} \cap \{Z|\tilde{a}_k(z)\}\right) = \sum_{h\in\mathbf{H}} \mu_{\mathbf{\Theta}}(\{Z|h\})\tilde{\theta}_k(z').$$

Therefore, taking the definition of conditional probability and combining: for all $z' \in \left\{Z|\tilde{h}_{k\backslash i}(z)\right\} \cap \{Z|\tilde{a}_k(z)\}$,

$$\begin{aligned}\mu_{\mathbf{\Theta}}\left(\tilde{a}_k|\tilde{h}_{k\backslash i}\right)(z) &= \frac{\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_{k\backslash i}(z)\right\} \cap \{Z|\tilde{a}_k(z)\}\right)}{\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_{k\backslash i}(z)\right\}\right)} \\ &= \frac{\sum_{h\in\mathbf{H}} \mu_{\mathbf{\Theta}}(\{Z|h\})\tilde{\theta}_k(z')}{\sum_{h\in\mathbf{H}} \mu_{\mathbf{\Theta}}(\{Z|h\})},\end{aligned}$$

so that, $\mu_{\mathbf{\Theta}}\left(\tilde{a}_k|\tilde{h}_{k\backslash i}\right)(z) = \tilde{\theta}_k(z')$. By Equation (15), $\forall h \in \mathbf{H}$, $z' \in \{Z|h\}$, $\tilde{\theta}_k(z') = \tilde{\theta}_k(z)$. This and the fact that $\mu_{\mathbf{\Theta}}(\tilde{a}_k|\tilde{h}_k)(z) = \tilde{\theta}_k(z)$ implies the desired conclusion.

Case (b): If $l_k = 1$ then $\forall z \in Z, \tilde{\theta}_k(z) = 1$, so that for all i, $\mu_{\mathbf{\Theta}}(\tilde{a}_k|\tilde{h}_{k\backslash i})(z) = \mu_{\mathbf{\Theta}}(\tilde{a}_k|\tilde{h}_k)(z) = 1$.

($\Leftarrow$) Assume $\forall \mathbf{\Theta} \in \Sigma$, $(\tilde{a}_k \perp \tilde{a}_i|\tilde{h}_{k\setminus i})_{\mathbf{\Theta}}$. If $l_k = 1$ then there does not exist $i \in N$ such that $i \to k$ by the IOD definition. Instead, assume $l_k > 1$, and $i \to k$. Choose $z \in Z$, $z' \in \left\{Z|\tilde{h}_{k\setminus i}(z)\right\}$ such that $\tilde{X}_k(z') \neq \tilde{X}_k(z)$ and $\tilde{a}_k(z) = \tilde{a}_k(z')$. The existence of such nodes is guaranteed by the definition of $i \to k$ and the product structure of the game. Note that $\tilde{a}_i(z') \neq \tilde{a}_i(z)$ while, for $j \in \{1, \ldots, k-1\} \setminus \{i\}$, $\tilde{a}_j(z') = \tilde{a}_j(z)$.

Consider a strategy profile $\mathbf{\Theta} \in \Sigma$ consistent with the following parameters. For $j \in \{1, \ldots, k-1\} \setminus \{i\}$, $\tilde{\theta}_j(z) = \tilde{\theta}_j(z') = 1$ ($\tilde{a}_j(z)$ chosen with probability 1 at all information sets). This implies $\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_{k\setminus i}(z)\right\}\right) = \mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_{k\setminus i}(z')\right\}\right) = 1$. For i, $\tilde{\theta}_i(z) = 1/4$ and $\tilde{\theta}_i(z') = 3/4$ so that $\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_k(z)\right\}\right) = 1/4$ and $\mu_{\mathbf{\Theta}}\left(\left\{Z|\tilde{h}_k(z')\right\}\right) = 3/4$. Because $\tilde{X}_k(z') \neq \tilde{X}_k(z)$, the behavior strategies at these information sets may be chosen freely. Suppose $\mathbf{\Theta}$ is such that $\tilde{\theta}_k(z) = 1/2$. Without loss of generality, assume $\tilde{a}_k(z) = a_{k,1}$ and $\tilde{X}_k(z) = X_{k,1}$, and let $\theta^k_{1,1}$ be some constant, $\alpha \in [0,1]$. Then, by Lemma 1, $\mu_{\mathbf{\Theta}}(a_{k,1}|\tilde{h}_k)(z) = 1/2$ and $\mu_{\mathbf{\Theta}}(a_{k,1}|\tilde{h}_k)(z') = \alpha$. By the definition of conditional probability,

$$\begin{aligned}\mu_{\mathbf{\Theta}}(a_{k,1}|\tilde{h}_{k\setminus i})(z) &= \frac{(1/4)\cdot(1/2) + (3/4)\cdot\alpha}{1} \\ &= \frac{1}{8}(1+6\alpha).\end{aligned}$$

By the premise $\mu_{\mathbf{\Theta}}(a_{k,1}|\tilde{h}_{k\setminus i})(z) = \mu_{\mathbf{\Theta}}(a_{k,1}|\tilde{h}_k)(z) = \frac{1}{2}$, implying $\alpha = 1/2$. However, α can take any value in [0,1], e.g. 3/5, a contradiction.

B.4 Lemma 3

Consider an arbitrary $k \in N$. Let $W_k \equiv \{1, \ldots, k-1\} \setminus Pa_k$. If $W_k = \emptyset$, or W_k is a singleton the Lemma holds trivially or follows immediately from Lemma 2, respectively. Therefore, assume W_k contains more than one element. The proof then consists of the sequential demonstration that: 1) $\forall z \in Z, z' \in \left\{Z|\tilde{h}_{k\setminus Pa_k}(z)\right\}$, $\tilde{X}_k(z') = \tilde{X}_k(z)$; and, 2) $\forall\, \mathbf{\Theta} \in \mathbf{\Sigma}$, $\mu_{\mathbf{\Theta}}\left(\tilde{a}_k|\tilde{h}_k\right) = \mu_{\mathbf{\Theta}}(\tilde{a}_k|\widetilde{pa}_k)$. Once the first condition is established, the second quickly follows using identical logic to that in the proof of Lemma 2. Therefore, we only elaborate part 1 below.

Start with any pair i, j such that $i, j \in W_k$, $i \neq j$. Assume there exist

$z \in Z, z' \in \left\{Z|\tilde{h}_{k\setminus\{i,j\}}(z)\right\}$ such that $\tilde{X}_k(z') \neq \tilde{X}_k(z)$. Note that

$$\left\{Z|\tilde{h}_{k\setminus i}(z)\right\} = \left\{Z|\tilde{h}_{k\setminus\{i,j\}}(z)\right\} \cap \left\{Z|\tilde{a}_j(z)\right\}, \tag{16}$$

and similarly for $\left\{Z|\tilde{h}_{k\setminus j}(z)\right\}$. Therefore,

$$\left\{Z|\tilde{h}_{k\setminus i}(z)\right\} \cup \left\{Z|\tilde{h}_{k\setminus j}(z)\right\} = \left\{Z|\tilde{h}_{k\setminus\{i,j\}}(z)\right\} \cap \left(\left\{Z|\tilde{a}_i(z)\right\} \cup \left\{Z|\tilde{a}_j(z)\right\}\right).$$

By the definition of the IOD, for all $z \in Z$,

$$z^* \in \left\{Z|\tilde{h}_{k\setminus i}(z)\right\} \cup \left\{Z|\tilde{h}_{k\setminus j}(z)\right\} \Rightarrow \tilde{X}_k(z^*) = \tilde{X}_k(z).$$

Hence, it must be that

$$z' \in \left\{Z|\tilde{h}_{k\setminus\{i,j\}}(z)\right\} \cap \left\{\left\{Z|\tilde{a}_i(z)\right\} \cup \left\{Z|\tilde{a}_j(z)\right\}\right\}^c,$$

where X^c indicates the complement of set X.
That is, for $l \in \{1, ..., k-1\} \setminus \{i, j\}$, $\tilde{a}_l(z') = \tilde{a}_l(z)$ but $\tilde{a}_i(z') \neq \tilde{a}_i(z)$ and $\tilde{a}_j(z') \neq \tilde{a}_j(z)$. Consider

$$z'' \in \left\{Z|\tilde{h}_{k\setminus\{i,j\}}(z')\right\} \cap \left\{Z|\tilde{a}_j(z')\right\} \cap \left\{Z|\tilde{a}_i(z)\right\}, \tag{17}$$

i.e., $h_k(z'')$ is $h_k(z')$ with the i^{th} move switched to $\tilde{a}_i(z)$. By (16), this implies $z'' \in \left\{Z|\tilde{h}_{k\setminus i}(z')\right\}$. Since $i \notin Pa_k$, it follows that $\tilde{X}_k(z'') = \tilde{X}_k(z')$. Now, consider

$$z''' \in \left\{Z|\tilde{h}_{k\setminus\{i,j\}}(z'')\right\} \cap \left\{Z|\tilde{a}_i(z'')\right\} \cap \left\{Z|\tilde{a}_j(z)\right\}.$$

By identical reasoning, $\tilde{X}_k(z''') = \tilde{X}_k(z'')$. Hence, $\tilde{X}_k(z''') = \tilde{X}_k(z')$. From (17), $\tilde{a}_i(z'') = \tilde{a}_i(z)$. By construction, $\left\{Z|\tilde{h}_{k\setminus\{i,j\}}(z'')\right\} = \left\{Z|\tilde{h}_{k\setminus\{i,j\}}(z)\right\}$. Taken together, this implies $h_k(z''') = h_k(z)$ and, as a consequence, $\tilde{X}_k(z) = \tilde{X}_k(z')$, a contradiction. It is not difficult to extend this argument to all of W_k should it have more than two elements.

B.5 Proposition 2

1. (Equivalence relation) Reflexivity: Given Γ and the identity mappings $f(r) = r$ and $g(\boldsymbol{\Theta}) = \boldsymbol{\Theta}$ then $\Gamma \preceq \Gamma$. 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 $\tilde{f} = f \circ \hat{f}$. and $\tilde{g} = g \circ \hat{g}$. By similar reasoning, $\Gamma' \preceq \Gamma$. Therefore, $\Gamma \sim \Gamma'$. Symmetry: This is immediate from the definition.

2. (Correspondence) Γ and Γ' are outcome compatible by definition. Let g' and $\hat{g}$ be the strategy mappings defined by $\Gamma \preceq \hat{\Gamma}$ and $\hat{\Gamma} \preceq \Gamma$, respectively. Let $\hat{g}\left(\hat{\mathbf{\Sigma}}\right) = \{\mathbf{\Theta} \in \mathbf{\Sigma} | \exists \hat{\mathbf{\Theta}} \in \hat{\mathbf{\Sigma}}, \hat{g}(\hat{\mathbf{\Theta}}) = \mathbf{\Theta}\}$. Define $g : \mathbf{\Sigma} \rightrightarrows \hat{\mathbf{\Sigma}}$ as follows:

$$\forall \mathbf{\Theta} \in \mathbf{\Sigma}, g\left(\mathbf{\Theta}\right) = \begin{cases} \hat{g}^{-1}\left(\mathbf{\Theta}\right) & if \quad \mathbf{\Theta} \in \hat{g}\left(\hat{\mathbf{\Sigma}}\right) \\ g'\left(\mathbf{\Theta}\right) & if \quad \mathbf{\Theta} \notin \hat{g}\left(\hat{\mathbf{\Sigma}}\right) \end{cases}.$$

Clearly, $\tilde{g}$ is onto and $\forall \mathbf{\Theta}, \mathbf{\Theta}' \in g(\mathbf{\Theta})$, $\mu_{\mathbf{\Theta}} = \mu_{\mathbf{\Theta}'}$.

B.6 Proposition 3

$\Gamma \preceq \Gamma'$ implies that Γ and Γ' are outcome compatible; let f denote the corresponding permutation and $\mathcal{G}$ denote Γ's IOD. Also, $\forall \mathbf{\Theta} \in \mathbf{\Sigma}$, $g(\mathbf{\Theta}) \in \mathbf{\Sigma}'$ and $\mu_{\mathbf{\Theta}} = \mu_{\mathbf{\Theta}'}$. By Proposition 1,

$$\begin{aligned} \mu_{\mathbf{\Theta}} &= \prod_{k=1}^{n} \mu_{\mathbf{\Theta}}\left(\tilde{a}_k | \widetilde{pa}_k^{\mathcal{G}}\right) \\ \mu_{g(\mathbf{\Theta})} &= \prod_{j=1}^{n} \mu_{\mathbf{\Theta}}\left(\tilde{a}'_j | \widetilde{pa}_j^{\mathcal{F}}\right). \end{aligned}$$

For each $j \in N$, outcome compatibility allows us to substitute $\tilde{a}'_j$ with $\tilde{a}_k$ where $j = f(k)$ in the second equation. If necessary, the order of the product of conditional probabilities in $\mu_{g(\mathbf{\Theta})}$ can be altered to match that of $\mu_{\mathbf{\Theta}}$. Then the result holds where $\widetilde{pa}_k^{\mathcal{F}}$ is as explained in Footnote 12.

B.7 Proposition 4

Lemma 5 *Given a game of semi-perfect information* Γ *and a probability distribution* μ *on* $(Z, 2^Z)$ *such that*

$$\forall z \in Z, \quad \mu\left(z\right) = \prod_{k \in N} \mu\left(\tilde{a}_k | \widetilde{pa}_k\right)\left(z\right),$$

there exists a strategy $\mathbf{\Theta} \in \mathbf{\Sigma}$ *such that* $\mu_{\mathbf{\Theta}} = \mu$, *$\mu$-a.s.*

Proof. Given $z \in Z$, $k \in N$ such that $\mu\left(\widetilde{pa}_k^{-1}(z)\right) > 0$, assume $a_{k,j} = \tilde{a}_k(z)$, let $\pi_{k,i} \equiv \widetilde{pa}_k(z)$, and then define $\theta_{i,j}^{k,z} \equiv \mu\left(\tilde{a}_k|\widetilde{pa}_k\right)(z)$. By the definition of conditional independence, for all z' such that $\widetilde{pa}_k(z') = \pi_{k,i}$, $\theta_{i,j}^{k,z'} = \theta_{i,j}^{k,z}$. Hence, we can write $\theta_{i,j}^k$ without ambiguity. If $\mu\left(\widetilde{pa}_k^{-1}(z)\right) = 0$, choose $\theta_{i,j}^k \in [0,1]$ on the paths corresponding to each $z' \in \widetilde{pa}_k^{-1}(z)$ such that $\sum_{j=1}^{l_k} \theta_{i,j}^k = 1$. Then,

$$\mu(z) = \prod_{k \in N} \mu\left(\tilde{a}_k|\widetilde{pa}_k\right)(z) = \prod_{k=1}^{n} \theta_{i,j}^k.$$

Define the matrix $\Theta_k = (\theta_{i,j}^k, i = 1, \ldots, p_k; j = 1, \ldots, l_k)$, where l_k is the number of actions in A_k and p_k is the number of values taken by $\widetilde{pa}_k$. We claim $\mathbf{\Theta} = (\Theta_k)_{k \in N}$ is a well-defined behavior strategy and that $\mu_{\mathbf{\Theta}} = \mu$, μ-a.s.

To show that Θ_k is a valid behavior strategy for the k-th move in Γ, recall that a valid behavior strategy at move k is a matrix Θ_k in which each row is a probability distribution on A_k corresponding to an information set for that move. By the definition of the μ-conditional probability of $\tilde{a}_k$ given $\widetilde{pa}_k$ when $\mu\left(\widetilde{pa}_k^{-1}(z)\right) > 0$ and by construction otherwise, for all $k \in N$, $\theta_{i,j}^k \in [0,1]$ and $\sum_{a_{k,j} \in A_k} \theta_{i,j}^k = 1$. Hence, each row of Θ_k constitutes a probability distribution on A_k. As Γ is a game of *semi-perfect information*, the number of information sets at move k is equal to the number of distinct action profiles $\{(a_j)_{j \in \Pi_k}\} = \{\pi_i^k, i = 1, \ldots, p_k\}$. Hence the matrix Θ_k implies a complete behavior strategy for each player. Finally, recall that μ_Θ is defined (in the text) as: $\mu_\Theta(z) = \prod_{k \in N} \tilde{\theta}_k(z)$, which obtains by defining $\tilde{\theta}_k$ in the obvious way. ■

Proof of Proposition 4

Step 1 (constructing the game form): Given $\mathbf{A}$ and $\mathcal{G}$, the corresponding game form $\Gamma_{\mathcal{G}} = (N, \mathcal{X}(X, \mathbf{E}), \mathbf{A})$ is constructed as follows. $\mathbf{A}$ and N are given. The rest of the game form is constructed by defining a valid order of play and defining the game tree and the information sets using that order.

By definition, the IOD $\mathcal{G}$ is a directed acyclic graph (DAG). Every DAG can be well-ordered, that is numbers can be assigned to nodes such that $a \to b$ then the number assigned to a is less than the one assigned to b. An algorithm that does this is the Topological sort algorithm (see Cowell et al

(1999)). Any assignment of numbers to nodes that makes $\mathcal{G}$ well-ordered defines an order of play that can be used to define a game form. Take one such assignment and assume wlog that the order of play corresponds to labels in $\mathcal{G}$.

Recall that X_k denotes the set of nodes assigned to player/move k. The game tree $(X, \mathbf{E})$ is then constructed as follows: create a node, x_1, as the root of the tree and assign player 1 to that node: $X_1 = \{x_1\}$. For each of the actions in A_1 create a new node x_2'. Assign player 2 to each of the new nodes, $x_2' \in X_2$. Add an edge from the root node to each of the new nodes. For player 2 do the following: take a node assigned to 2 and create a node, x_3', for each of the actions in A_2. Assign the new nodes to player 3, $x_3' \in X_3$ and add an edge from the old node to the new nodes. Repeat this for every node assigned to player 2. Then, proceed recursively forward (from j to $j+1$) through all players until you reach player n. Then proceed as before only instead of assigning the new nodes, x', to (the non-existent) player $n+1$, label them as terminal nodes, $x' \in Z$, where Z denotes the set of terminal nodes.

Finally, construct the information sets, $\mathcal{X} = \{\mathcal{X}_1, \ldots, \mathcal{X}_m\}$. Let $\mathcal{X}_1 = \{X_1\}$. For each player $k = 2, \ldots, n$ let Pa_k be the set of players j such that $j \to k$. Let A_k^{Pa} be the set of all possible action sequences that can be taken by the parents of k, i.e. $(a_{r1}, \ldots, a_{ri}) \in A_k^{Pa}$ if there exists z such that $\widetilde{pa}_k(z) = (a_{r1}, \ldots, a_{ri})$. Then, define player k's information sets $\mathcal{X}_k = \{X_{k,1}, \ldots, X_{k,m_k}\}$ as follows: for each distinct $(a_{r1}, \ldots, a_{ri}) \in A_k^{Pa}$ define a distinct information set $X_{k,r}$ such that for all x, x' assigned to player k, for all z, z' such that $\tilde{x}_k(z) = x$ and $\tilde{x}_k(z') = x'$, $\widetilde{pa}_k(z) = \widetilde{pa}_k(z')$ iff $\{x\} \cup \{x'\} \subseteq X_{k,r}$. That is, two nodes are in the same information set for player k iff the actions taken by the parents of player k along the path to those nodes are exactly the same.

Step 2. (Semi-perfect information game form) It is clear that $\Gamma_{\mathcal{G}}$ is a game of semi-perfect information by construction (see the construction of the information sets).

Step 3. (only if) By construction, $\mathcal{G}$ is implied by $\Gamma_{\mathcal{G}}$. Then, it follows from Proposition 1 that if $\mu = \mu_{\mathbf{\Theta}}$ for some $\mathbf{\Theta} \in \mathbf{\Sigma}_{\mathcal{G}}$ then μ will admit recursive factorization according to $\mathcal{G}$.

Step 4. (if) Suppose μ admits recursive factorization according to $\mathcal{G}$. Then, define the following strategy for player k:

$$\tilde{\theta}_k(z) = \mu(\tilde{a}_k(z) | \tilde{h}_k(z))$$

By hypothesis, for each z,

$$\mu(z) = \prod_{k=1}^{n} \mu(\tilde{a}_k(z)|\widetilde{pa}_k(z))$$

so that for all k

$$\tilde{\theta}_k(z) = \mu(\tilde{a}_k(z)|\widetilde{pa}_k(z))$$

All that is left to check is that $\tilde{\theta}_k(z)$ is measurable with respect to the information sets of $\mathcal{X}_k$. Consider an arbitrary k and information set $X_{k,r}$. For z, z' such that $\tilde{x}_k(z), \tilde{x}_k(z') \in X_{k,r}$, let $\theta^z = \mu(\tilde{a}_k(z)|\widetilde{pa}_k(z))$ and $\theta^{z'} = \mu(\tilde{a}_k(z')|\widetilde{pa}_k(z'))$. Because $\tilde{x}_k(z), \tilde{x}_k(z') \in X_{k,r}$ by construction of $\Gamma_{\mathcal{G}}$, $\widetilde{pa}_k(z) = \widetilde{pa}_k(z') = \pi$. If $\tilde{a}_k(z) = \tilde{a}_k(z') = a^*$ then $\mu(\tilde{a}_k(z)|\widetilde{pa}_k(z)) = \mu(\tilde{a}_k(z)|\widetilde{pa}_k(z)) = \mu(a^*|\pi) = \theta^k_{r,*}$—the proposed strategies are measurable.

B.8 Proposition 5

We use the notion and vocabulary of dependency models presented in Appendix A. First, we note the following.

Lemma 6 *Every IOD is a directed, acyclic graph.*

Proof. By definition, the IOD is directed. Because of the move order imposed by the order of play in the game form, the definition of the IOD includes the requirement: $i \to k$ only if $i < k$, which implies the IOD is acyclic. ■

Proof of Proposition 5

By Theorem 7, $\mathcal{G} = (N, \to)$ and $\mathcal{F} = (N, \to')$ satisfy $\mathcal{E} = \mathcal{E}'$ and $\mathbf{S} = \mathbf{S}'$ iff every distribution that can be factored according to the recursive basis $B_{\mathcal{G}}$ can also be factored according to $B_{\mathcal{F}}$ and viceversa. This result does not prove the result directly because the set of probability distributions that can be factored according to $B_{\mathcal{G}}$ is greater than the set of probabilities that can be generated by strategies in the game form generating $\mathcal{G}$ (consider for example the game forms in Figure 4).

($\Leftarrow$) By Theorem 7, $\mathcal{G}$ and $\mathcal{F}$ satisfy $\mathcal{E} = \mathcal{E}'$ and $\mathbf{S} = \mathbf{S}'$ implies that all $\mathbf{\Theta} \in \mathbf{\Sigma}$, $\mu_{\mathbf{\Theta}}$ admits factorization according to $\mathcal{F}$ and for every $\mathbf{\Theta}' \in \mathbf{\Sigma}'$, $\mu_{\mathbf{\Theta}'}$ admits factorization according to $\mathcal{G}$, i.e. $\Gamma \doteqdot \Gamma'$

($\Rightarrow$) Take a maximally revealing $\mathbf{\Theta} \in \mathbf{\Sigma}$. Then $\mathcal{G}$ is by definition a perfect map of the dependency model implied by $\mu_{\mathbf{\Theta}}$. By hypothesis, $\mu_{\mathbf{\Theta}}$ can be factored according to $\mathcal{F}$ so that the dependency model implied by $\mu_{\mathbf{\Theta}}$ is an I-map of $M_{\mathcal{F}}$, that is $M_{\mathcal{G}} \subseteq M_{\mathcal{F}}$. Similarly, take a maximally revealing $\mathbf{\Theta} \in \mathbf{\Sigma}'$.

By the same argument, $M_{\mathcal{F}} \subseteq M_{\mathcal{G}}$, so that $\mathcal{G}$ and $\mathcal{F}$ are observationally equivalent and, by Theorem 7, satisfy $\mathcal{E} = \mathcal{E}'$ and $\mathbf{S} = \mathbf{S}'$.

References

[1] 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.

[2] Cowell, R. G., Dawid, A. P., Lauritzen, S. L., Spiegelhalter, D. J. 1999. Probabilistic Networks and Expert Systems. Springer, New York.

[3] Demuynck, T., and Lauwers, L. 2005. Nash rationalizability of collective choice over lotteries. Working Papers of Faculty of Economics and Business Administration, Ghent University, Belgium: #05/301.

[4] Elmes, S., Reny, P. 1994. On the strategic equivalence of extensive form games. Journal of Economic Theory, 62(1), 1-23.

[5] Geiger, D., Verma, T. S., J. Pearl, J. 1990. Identifying independence in Bayesian networks. Networks, 20, 507-34.

[6] Kakade, S., Kearns, M., Langford, J., Ortiz, L. 2003. Correlated Equilibria in Graphical Games. Proceedings of the 4th ACM conference on Electronic commerce. San Diego, 42 - 47.

[7] Kalai, E. 2004. Large Robust Games. Econometrica 72 (6), 1631–1665.

[8] Glazer, J., Rubinstein, A. 1996. An Extensive Game as a Guide for Solving a Normal Game. Journal of Economic Theory, 70(1), 32-42.

[9] Kearns, M., Littman, M., Singh, S. 2001. Graphical Models for Game Theory. Proceedings of UAI.

[10] Kohlberg, E., Mertens J. 1986. On the Strategic Stability of Equilibria. Econometrica, 54(5), 1003-1038.

[11] Koller, D. Milch, B. 2002. Multi-agent influence diagrams for representing and solving games. Seventeenth International Joint Conference on Artificial Intelligence, Seattle, WA.

[12] Kuhn, H. W. 1953. Extensive Games and the Problem of Information. Contributions to the Theory of Games II, 193-216.

[13] La Mura, P. 2002. Game Networks. Unpublished manuscript.

[14] Meek, C. 1997. Strong completeness and faithfulness in Bayesian networks, Proceedings of Eleventh Conference on Uncertainty in Artificial Intelligence, Montreal, QU, 411-418.

[15] Paz, A., Pearl, J., Ur, S. 1993. A Net Characterization of Graphs Based on Interception Relations. Mimeo.

[16] Pearl, J. 1986. Fusion, propagation and structuring in belief networks. Artificial Intelligence 29, 241-88.

[17] —— 1988. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. North Holland, Amsterdam.

[18] —— 2000. Causality: Models, Reasoning, and Inference. Cambridge University Press.

[19] ——, Paz, A. 1985. Graphoids: A graph based logic for reasoning about relevance relations. Tech rep 850083 (R-53-L), Cognitive Systems Laboratory, UCLA.

[20] Ray, I., Snyder, S. 2004. Observable Implications of Nash and Subgame-Perfect Behavior in Extensive Games. Unpublished manuscript.

[21] ——, Zhou, L. 2001. Game Theory via Revealed Preferences. Games and Economic Behavior 37, 415–24.

[22] Ryall, M. D. 1997. Causal learning in extensive-form games. Ph.D. dissertation. UCLA.

[23] Rubinstein, A. 1991. Comments on the Interpretation of Game Theory. Econometrica, 59(4), 909-924.

[24] Spirtes, P., Glymour, C. and Scheines, R. 2000. Causation, Prediction, and Search. The MIT Press, Cambridge.

[25] Spohn, W. 1980. Stochastic independence, causal independence and shieldability. J. Phil. Logic, 9, 73-99.

[26] Sprumont, Y. 2000. On the Testable Implications of Collective Choice Theories. Journal of Economic Theory 93, 205-32.

[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.