3

GAMES WITH OBSERVABLE DEVIATORS

Pierpaolo BATTIGALLI

Abstract

An n-person dynamic game has *observable deviators* if each player at each move is given n separate pieces of information respectively concerning her own past behavior and that of her $n-1$ opponents. This paper defines and characterizes the observable deviators property for extensive form games with incomplete information and illustrates how this property can be used to simplify game theoretic analysis. In particular, it is proved that sequential equilibria can be characterized by simple and transparent restrictions on players' assessments.

Introduction

Consider a dynamic Cournot oligopoly game such that each firm has (possibly) private information about its own technology and in each period chooses output and effort in research and development (R&D). Suppose that the output choices are observed at the beginning of the following period. We may also assume that some firms have "spies" reporting information about the technology and R&D efforts of some opponents, but in general such reports are either incomplete or available for only a subset of competitors. This is an extensive form game with incomplete and imperfect (not even almost perfect) information. Yet the information structure is simple enough that different solution concepts turn out to be equivalent. What makes the information structure relatively simple is the following feature: *at any point of the game the total information of a player about her opponents' types and strategic behavior is derived from separate pieces of "marginal" information about each one of them.* This would not be the case if, for example, only the past market prices were observed, as this would provide information about the competitors' total output which cannot be decomposed into information about each single competitor's output.

When the information structure has the above stated property we say that the game exhibits "observable deviators". More formally, a game has *observable deviators* if for each information set h the set of profiles of type–strategy pairs (one pair for each player) jointly consistent with h is a Cartesian product of the sets of each player's type–strategy pairs individually consistent with h. The reason for the "observable deviators" terminology is more easily understood in the complete information case. Suppose that some strategy profile s is expected, say by player i, but this player gets the move at an information set h which is inconsistent with s. Then, if the stated property holds, player i can use the information contained in h to identify which opponents must have deviated from s. Note that both characterizations of the property make clear that every two–person game with perfect recall has observable deviators.

The above definition of the observable deviators property uses a mixture of primitive extensive form concepts (information sets) and derived concepts (strategies). Although this definition turns out to be useful for the analysis of some solution concepts, a more direct characterization in terms of primitive extensive form concepts is desirable. This paper provides such a characterization for general extensive form games with incomplete information and then proceeds to give sharper characterizations of a stronger property: observable deviators in the agent form.

One advantage of considering such somewhat simplified information structures is that it is possible to obtain equivalence results about game theoretic concepts which differ in more general environments, thus sharpening our understanding of the theory and simplifying the analysis.

Battigalli (1996b) shows that in (finite) complete information games with observable deviators an intuitive and transparent stochastic independence property completely characterizes Kreps and Wilson's (1982) notion of "full consistency" of assessments.

Even simpler and weaker properties characterize full consistency in multistage games with observable deviators in the agent form. These results concern the properties of players' conditional beliefs in games and can be used to characterize both solution concepts assuming common beliefs, (such as Kreps and Wilson's (1982) sequential equilibrium) and solution concepts that allow for heterogeneous beliefs (such as some notions of extensive form rationalizability[1]).

In the present paper these equivalence results are extended to games with incomplete information. Furthermore, it is shown that, in multistage games with observable deviators in the agent form, a rather weak property of conditional beliefs called "general reasonableness" (Fudenberg and Tirole (1991)) is sufficient (if coupled with sequential rationality) to characterize sequential equilibrium outcomes.

Unlike Kreps and Wilson's (1982) notion of full consistency, the properties of conditional beliefs used in this paper can be easily formulated for infinite games. Thus the above mentioned equivalence results suggest possible extensions of solution concepts for finite games, such as sequential equilibrium, to infinite games. For this reason we consider infinite games in general and then restrict our attention to finite games when we characterize full consistency.

Related Literature. The observable deviators property has been defined independently by Fudenberg and Levine (1993) and Battigalli (1994). The latter paper compares different ways to model players' conditional expectations in extensive games, viz. Pearce's (1984) systems of conjectures and Myerson's (1986) conditional probability systems. Fudenberg and Levine (1993) use the observable deviators property to characterize self confirming equilibria. In particular, their results imply that, in games with observable deviators, self confirming equilibria with independent, unitary beliefs[2] are outcome equivalent to Nash equilibria.

Fudenberg and Tirole (1991) provide a quite simple and intuitive characterization of sequential equilibria of multistage games of incomplete information with observed actions, which are a special case of the incomplete information games with observable deviators analyzed here. By considering a less general case, they are able to state a weaker and simpler sufficient condition for consistency

[1] See e.g. Pearce (1984) and the modified solution proposed by Battigalli (1996a). Reny (1992) proposes an iterative removal procedure to define the notion of "explicable equilibrium." Using the results of the present paper one can show that in games with observable deviators the strategies that survive this deletion procedure are exactly those which are rationalizable in the sense of Battigalli (1996a).

[2] A self confirming equilibrium is a vector σ of mixed strategies such that, for each player i, each pure strategy s_i with positive probability is a best response to some probabilistic belief $b^i(s_i)$ about the opponents, and the probability distribution on terminal nodes induced by s_i and $b^i(s_i)$ coincides with the distribution induced by s_i and σ_{-i}. Beliefs are independent if they are given by a product measure and are unitary if $b^i(s_i) = b^i$ is independent of s_i, so that also σ_i is a best response to b^i.

than the one given here. They also consider general (finite) extensive games and put forward the "general reasonableness" notion mention above. But they do not provide a characterization result for games with (partially) unobserved actions.[3]

Battigalli (1996b) analyzes sequential equilibria and weaker notions of perfect Bayesian equilibrium in finite games of complete information without chance moves. As we mentioned above, some equivalence results are derived for games with observable deviators. The present paper is a continuation and extension of Battigalli (1996b).

Kohlberg and Reny (1992) give a complete characterization of Kreps and Wilson's full consistency and show that the stochastic independence condition put forward by Battigalli (1994,1996b) is only necessary, not sufficient, for full consistency in more general games. Swinkels (1993) provides additional insight on Kohlberg and Reny's characterization.

1 Observable Deviators in Extensive Form Games with Incomplete Information

In Section 1.1, we introduce a formal description of extensive form games with incomplete information, which extends the formulation given in Kreps and Wilson (1982). The readers already familiar with the latter can skip Section 1.1 and refer to the Appendix, Section A.1, which contains a list with the notation and corresponding terminology. Note that some symbols are new and some differ from Kreps and Wilson (1982) because we consider a more general and more structured class of games.[4] In Section 1.2, we characterize the observable deviators property and the stronger property "observable deviators in the agent form.".

[3] They claim that general reasonableness is necessary and sufficient for consistency, but Battigalli (1996b) shows that this is claim is incorrect even in the restricted class of multistage games with observable deviators in the agent form. For more on this see Battigalli (1996b, Section 5) and Section 2.2 of the present paper.

[4] Kreps and Wilson consider finite games of imperfect information with a random move selecting the initial node, which may be interpreted as a state of the world in a game with incomplete information, or a complete strategy of the chance player in a game with complete information. Here we consider a slightly more general formulation. We do not assume in general that the game is finite. We distinguish between ordinary chance moves, which may take place at any point of the game, and the initial move by Nature choosing the state of the world. Furthermore, in our framework the state of the world is explicitly represented as a vector of types.

1.1 Extensive Form Games with Incomplete Information

An *n-person extensive game with incomplete information* is a tuple

$$\Gamma = (\Theta, X, \prec; \iota; H; A, \alpha; \rho, \pi_0; u)$$

with the following elements:

(i) $\Theta = \Theta_1 \times \ldots \times \Theta_n$ and $(X, \prec)$ is an arborescence with initial nodes $\theta = (\theta_1, \ldots, \theta_n) \in \Theta_1 \times \ldots \times \Theta_n \subset X$ and precedence relation $\prec \subset X \times X$. $\preceq$ denotes the reflexive closure of $\prec$, that is, $x \preceq y$ if and only if $x \prec y$ or $x = y$. $f(x)$ denotes the set of *immediate successors* of node x. Z denotes the set of *complete paths* and Z^* the set of *terminal nodes*. Whenever every complete path is finite and no confusion may arise, we write $Z^* = Z$, identifying a complete path with its terminal node. $X \backslash Z^*$ is the set of *decision (or chance) nodes*. $Z(x)$ $(Z^*(x))$ denotes the set of complete paths (terminal nodes) containing (following) a given node x. For any given set of nodes $Y \subseteq X$, $Z(Y) := \bigcup_{x \in Y} Z(x)$ $(Z^*(Y) := \bigcup_{x \in Y} Z^*(x))$.

An initial node $\theta \in \Theta$ represents a state of the world and $\theta_i \in \Theta_i$ is the *type* (private information) of player i (see (iii) below). A (complete) path corresponds to a pair given by a state of the world and a (complete) history of choices in the game. We say that a game Γ has *complete information* if Θ is a singleton. For any node $y \in X$, $\vartheta(y)$ denotes the unique initial node $\theta \in \Theta$ such that $\theta \preceq y$ and $\vartheta_j(y)$ is the j^{th} component of $\vartheta(y)$.

(ii) Let $I := \{0, 1, \ldots, n\}$ and $N = \{1, \ldots, n\}$ denote respectively the *players set* and the *personal players set*. Index 0 denotes the *chance* player. For many purposes there will be no need to distinguish between chance and personal players. $\iota : X \backslash Z^* \to I$ is the *players function* and $X_i := \iota^{-1}(i)$ is the set of i's decision nodes (chance nodes if $i = 0$). We will often refer to groups of players $J \subseteq I$. For example, $X_J := \bigcup_{j \in J} X_j$ denotes the set of decision (or chance) nodes of players in J. For any vector of elements $(v_0, v_1, ..., v_n)$ we write $v_J = (v_i)_{i \in J}$. Singletons are identified with their unique element whenever convenient. The complement of J in I is denoted $-J$.

(iii) H is the *information partition*, i.e. a partition of $X \backslash Z^*$ such that the players function ι is H-measurable. H_i denotes the induced partition of X_i, that is, the collection of player i's information sets. $H(x)$ denotes the information set containing node x.[5]

Since θ_i represents i's private information, we assume that at each information set $h \in H_i$ $(i \in N)$, player i knows the element θ_i of θ, i.e. if $x', x'' \in h \in H_i$, then $\vartheta_i(x') = \vartheta_i(x")$.

(iv) A is the set of *actions*; the *actions function* $\alpha : X \backslash \Theta \rightarrow A$ specifies the last action $a = \alpha(t)$ taken to reach node t.[6] For each $h \in H$, $A(h)$ $((A(x))$ denotes the set of feasible action at information set h (node x). The *immediate successor of node xfollowing action* a is denoted by $f(x, a)$. Thus for all $x \in X \backslash Z^{*}M$, $\alpha(f(x, a)) = a$.

We assume that the tuple $(\Theta, X, \prec; \iota; H; A, \alpha)$ satisfies the *perfect recall* property.[7]

It will be convenient to have a specific notation for the initial action (or set of actions) leading from a node (or set of nodes) to a successor (or a set of successors) of that node. For all $h \in H$, $y \in X$, $a \in A(h)$, we write $a = \alpha(h \rightarrow y)$ if and only if there is one (and, by perfect recall, only one) node $x \in h$ such that $f(x, a) \preceq y$. For all $x \in X$, $Y \subseteq X$, $A(x \rightarrow Y) := \{a \in A(x) \mid \exists y \in Y : f(x, a) \preceq y\}$ denotes the set of initial actions which may lead from node x to some node in Y.

(v) $\rho \in \Delta(\Theta)$ is a probability measure on Θ. We assume that there are *independent types* – that is, ρ is the product of n probability measures $\rho_i \in \Delta(\Theta_i)$, $i = 1, \ldots, n$.[8] π_0 is a function which assigns a probability measure $\pi_0(\cdot \mid h) \in \Delta(A(h))$ to each information set $h \in H_0$ of the chance player. We assume that ρ and each $\pi_0(\cdot \mid h)$, $h \in H_0$, have full support.

(vi) $u = (u_1, \ldots, u_n)$ is n-tuple of payoff functions $u_i : Z \rightarrow \mathbb{R}$.

[5] H is such that if $H(x') = H(x")$, then there is a bijection between $f(x')$ and $f(x")$.

[6] and is such that for each $x \in X \backslash Z^{*}$, α is one to one on $f(x)$ and $\alpha(f(x')) = \alpha(f(x"))$ if $H(x') = H(x")$.

[7] See Kreps and Wilson (1982) for a precise definition using this notation. Note that they assume for notational convenience that $A(h') \cap A(h") = \emptyset$ whenever $h' \neq h''$.

[8] Myerson (1991, Chapter 1) shows that, as long as no property of the payoff functions is assumed (as is the case in this paper), there is no loss of generality in assuming independent types.

Types, Strategies, and Probabilities

A *strategy* for player i is a function $s_i : H_i \to A$ such that $s_i(h) \in A(h)$ for all $h \in H_i$. A *complete strategy profile* s is an element of the Cartesian set $S := S_0 \times S_1 \times \ldots \times S_n$. A complete strategy profile induces a complete history (feasible sequence of actions) of the game; a pair $(\theta, s) \in \Theta \times S$ induces a complete path $\zeta(\theta, s)$.[9]

For the purposes of this paper it is often more convenient to represent (θ, s) as the $(2n+2)$-tuple

$$(\theta_0, s_0; \theta_1, s_1; \ldots; \theta_n, s_n) \in (\Theta_0 \times S_0) \times (\Theta_1 \times S_1) \times \ldots \times (\Theta_n \times S_n),$$

where θ_0 is a fictitious type which belongs to the singleton $\Theta_0 := \{\theta_0\}$. For notational convenience we write $\mathcal{S}_i := (\Theta_i \times S_i)$ and $\mathcal{S} := \times_{i \in I} \mathcal{S}_i$. In general, a given player i's payoff is affected by the private information and the behavior of each opponent j. Thus the set of type–strategy pairs $\mathcal{S}_j$ captures well the relevant uncertainty about opponent j. When considering a group of players J, we sometime abuse notation writing $(\theta_J, s_J) \in \mathcal{S}_J$, $(\theta, s) \in \mathcal{S}_J \times \mathcal{S}_{-J}$. An element of $\mathcal{S}_J$ is called a J-profile.

A *strategy* s_i is *consistent with node* y, written $s_i \in S_i(y)$ if for all nodes $x \in X_i$ before y, s_i selects the action on the path to y. Formally, $S_i(y) = \{s_i \in S_i \mid \forall x \in X_i, (x \prec y) \Leftrightarrow [s_i(H(x)) = \alpha(H(x) \to y)]\}$. A *$J$-profile* $(\theta_J, s_J) \in \mathcal{S}_J$ is *consistent with a node* y, written $(\theta_J, s_J) \in \mathcal{S}_J(y)$, if for all $j \in J$, $\theta_j = \vartheta_j(y)$ and $s_j \in S_j(y)$ (where by convention $\vartheta_0(y) = \theta_0$ for all y). A J-profile (θ_J, s_J) is *consistent with a set of nodes* Y, if it is consistent with some node in Y, i.e. if $(\theta_J, s_J) \in \mathcal{S}_J(Y) := \bigcup_{y \in Y} \mathcal{S}_J(y)$. The subscript J is omitted if $J = I$. Clearly $\mathcal{S}(y)$ is the set of type–strategy profiles inducing a path through y. It follows immediately from the definitions that, for each node y and each $J \subseteq I$, $\mathcal{S}_J(y) = \times_{j \in J} \mathcal{S}_j(y)$. But in general the set of profiles consistent with a *set* of nodes is *not* Cartesian: $\mathcal{S}_J(Y) \subseteq \times_{j \in J} \mathcal{S}_j(Y)$.

An important implication of the perfect recall assumption is that for all players $i \in I$, information sets $h \in H_i$ and nodes $x \in h$, $\mathcal{S}_i(x) = \mathcal{S}_i(h)$. In fact, at each node $x \in h$ player i knows her type $\theta_i = \vartheta_i(x)$ and the sequence of player i's information sets and actions on the path to x is the same for every $x \in h$.

[9] $\zeta(\theta, s)$ is defined by induction: the first element is θ; given the k^{th} element x, the $(k+1)^{th}$ element is $x' = f(x, s_{\iota(x)}(H(x)))$.

This implies that $\mathcal{S}(h) = \bigcup_{x \in h}[\mathcal{S}_i(x) \times \mathcal{S}_{-i}(x)] = \mathcal{S}_i(h) \times \mathcal{S}_{-i}(h)$. Therefore, we obtain the following:

Remark 1. In two–person games with perfect recall and without chance moves (and in one–person games with chance moves) each set of type–strategy profiles consistent with an information set h is Cartesian.

A *mixed strategy* for player i is a probability measure on the set S_i of player i's pure strategies. But for the purposes of this paper it is more convenient to consider a different concept. A *distributional strategy* for player i is a probability measure $\sigma_i \in \Delta(\mathcal{S}_i)$ on the set $\mathcal{S}_i = \Theta_i \times S_i$ of player i's type–strategy pairs.[10]

A distributional strategy σ_i is *feasible* given the probability measure on types ρ_i, written $\sigma_i \in \Delta(\mathcal{S}_i; \rho_i)$, if $\sigma_i(\hat{\Theta}_i \times S_i) = \rho_i(\hat{\Theta}_i)$ for all (measurable) subsets $\hat{\Theta}_i \subseteq \Theta_i$. Clearly, distributional strategies coincide with mixed strategies in games of complete information. A *behavioral strategy* for player i is a collection of probability measures $\pi_i = (\pi_i(\cdot \mid h))_{h \in H_i} \in \Pi_i := \times_{h \in H_i} \Delta(A(h))$. Note that the behavioral strategy of the chance player is part of the description of the game. The set of behavioral strategy profiles containing the exogenously given behavioral strategy π_0 of the chance player is denoted by $\Pi := \{\pi_0\} \times \Pi_1 \times \cdots \times \Pi_n$.

A *system of beliefs* is a collection of probability measures $\mu = (\mu(\cdot \mid h))_{h \in H} \in M := \times_{h \in H} \Delta(h)$. An *assessment* is a pair $(\mu, \pi) \in M \times \Pi$.

Consider a vector λ of distributional (or behavioral) strategies. Assuming that the type–strategy pairs of different players (or the actions at different information sets) are chosen independently, one can compute the probability $P^\lambda(y)$ of reaching a node y given λ.[11]

Consider, for simplicity, a finite game. There is a canonical way to derive a distributional strategy $\delta_i(\pi_i)$ from a behavioral strategy π_i and to derive a behavioral strategy $\beta_i(\sigma_i)$ from a distributional strategy σ_i:

$$\delta_i(\pi_i)(\theta_i, s_i) = \rho_i(\theta_i) \prod_{h \in H_i} \pi_i(s_i(h) \mid h);$$

[10] This is not exactly Milgrom and Weber's (1985) definition, because in their framework pure strategies in S_i are not type–contingent.

[11] These probabilities are computed in a straightforward way: $P^\sigma(y) = \prod_{i \in I} \sigma_i(\mathcal{S}_i(y))$; $P^\pi(\theta) = \rho(\theta)$, if $y = f(x, a)$, $P^\pi(y) = P^\pi(x) \cdot \pi_{\iota(x)}(a \mid H(x))$.

$$\beta_i(\sigma_i)(a \mid h) = \sigma_i(\mathcal{S}_i(f(x,a)))/\sigma_i(\mathcal{S}_i(x)), \text{ if } \sigma_i(\mathcal{S}_i(x)) > 0, x \in h \in H_i, a \in A(h)$$

(by perfect recall and knowledge of i's type $\mathcal{S}_i(x)$ and $\mathcal{S}_i(f(x,a))$ are independent of $x \in h$; $\beta_i(\sigma_i)$ is arbitrary if $\sigma_i(\mathcal{S}_i(x)) = 0$). Given $\pi = (\pi_i)_{i \in I}$ and $\sigma = (\sigma_i)_{i \in I}$, let $\delta(\pi) = (\delta_i(\pi_i))_{i \in I}$ and $\beta(\sigma) = (\beta_i(\sigma_i))_{i \in I}$. It can be checked that these transformations preserve the probability of reaching any given node:[12]

$$\forall y \in X, P^{\pi}(y) = P^{\delta(\pi)}(y) \text{ and } P^{\sigma}(y) = P^{\beta(\pi)}(y).$$

1.2 Games with Observable Deviators

Intuitively, a game has observable deviators if each player at each information set is given n separate pieces of information concerning the type and behavior of each opponent (including the chance player). Since the uncertainty of player i about player j is represented by the set $\mathcal{S}_j$, it is natural to consider the "strategic information" about j conveyed by information set $h \in H_i$, that is $\mathcal{S}_j(h)$. If the information that i gets at h is given by (information about herself and) n separate pieces of information about the opponents, it must be the case that $\mathcal{S}_{-i}(h) = \times_{j \neq i}\mathcal{S}_j(h)$. Otherwise, $\mathcal{S}_{-i}(h) \subset \times_{j \neq i}\mathcal{S}_j(h)$ and i has joint information about her opponents which cannot be obtained from the information concerning each one of them separately. We know that by perfect recall (and knowledge of type) $\mathcal{S}(h) = \mathcal{S}_i(h) \times \mathcal{S}_{-i}(h)$, $h \in H_i$. Therefore we have the following definition:

Definition 1. *A game of incomplete information has* observable deviators *iff*[13]

$$\forall h \in H, \ \mathcal{S}(h) = \mathcal{S}_0(h) \times \mathcal{S}_1(h) \times \ldots \times \mathcal{S}_n(h).$$

[12] This is a version Kuhn's (1953) theorem on the equivalence between mixed and behavior strategies in perfect recall games. Since every player knows her type, we can interpret a distributional strategy as a mixed strategy of a fictitious player *with perfect recall* choosing $(\theta_i, s_i) \in \mathcal{S}_i$.

[13] This property was called "strategic decomposability of information" in the working paper version of Battigalli (1994).

Note that every game with two players and no chance moves has observable deviators (see Remark 1).

Remark 2. Let $\hat{\mathcal{S}}$ be a subset of any Cartesian product $\mathcal{S} = \mathcal{S}_0 \times \ldots \times \mathcal{S}_n$. Then $\hat{\mathcal{S}} = \hat{\mathcal{S}}_0 \times \ldots \times \hat{\mathcal{S}}_n$ if and only if for all $i = 0, 1, \ldots, n$, $\hat{\mathcal{S}} = \hat{\mathcal{S}}_i \times \hat{\mathcal{S}}_{-i}$ (where $\hat{\mathcal{S}}_i$ and $\hat{\mathcal{S}}_{-i}$ are the projections of $\hat{\mathcal{S}}$ on $\mathcal{S}_i$ and $\mathcal{S}_{-i}$ respectively). Therefore a game has observable deviators if and only if

$$\forall i \in I, \forall h \in H, \mathcal{S}(h) = \mathcal{S}_i(h) \times \mathcal{S}_{-i}(h).$$

The terminology "observable deviators" is due to the following characterization (cf. Fudenberg and Levine (1993)). Consider, just for simplicity, a complete information game and suppose that the players are expected to play according to a strategy profile s, but then player i deviates to a different strategy s'_i thus inducing a path through an information set h, which would not have been reached otherwise. Whoever gets the information corresponding to h realizes that *someone* has deviated from s, but does not necessarily know who. However, if the game has observable deviators, it is possible to identify player i as the deviator. Consider the following property:[14]

(FL) $\forall i \in I, \forall(\theta, s) \in \mathcal{S}, \forall(\theta, s) \in \mathcal{S}_i, \forall(\theta'_i, s'_i) \in \mathcal{S}_{-i}, \forall h \in H$,
if $(\theta'_i, s'_i; \theta_{-i}, s_{-i}) \in \mathcal{S}(h)$ and $(\theta, s) \notin \mathcal{S}(h)$,
then $(\theta_i, s_i; \theta'_{-i}, s'_{-i}) \notin \mathcal{S}(h)$.

Proposition 1. *A game of incomplete information has observable deviators if and only if property (FL) holds.*

Proof. (Only if) Suppose that $(\theta'_i, s'_i; \theta_{-i}, s_{-i}) \in \mathcal{S}(h)$ and $(\theta, s) \notin \mathcal{S}(h)$. Then $(\theta_{-i}, s_{-i}) \in \mathcal{S}_{-i}(h)$. If the game has observable deviators, $\mathcal{S}(h) = \mathcal{S}_i(h) \times \mathcal{S}_{-i}(h)$ and $(\theta_i, s_i; \theta_{-i}, s_{-i}) \notin \mathcal{S}(h)$ implies that $(\theta_i, s_i) \notin \mathcal{S}_i(h)$. This in turn implies that $(\theta_i, s_i; \theta'_{-i}, s'_{-i}) \notin \mathcal{S}_i(h) \times \mathcal{S}_{-i}(h) = \mathcal{S}(h)$.

(If) Suppose that the game does *not* have observable deviators. By Remark 2, there are some $i \in I$ and $h \in H$ such that $\mathcal{S}(h) \subset \mathcal{S}_i(h) \times \mathcal{S}_{-i}(h)$. Pick $(\theta, s) \in (\mathcal{S}_i(h) \times \mathcal{S}_{-i}(h)) \backslash \mathcal{S}(h)$. Since $(\theta, s) \in \mathcal{S}_i(h) \times \mathcal{S}_{-i}(h)$, there are (θ'_i, s'_i)

[14] FL stands for Fudenberg and Levine, since this is (the extension to incomplete information games of) the "observable deviators" property in their (1993) paper.

and (θ'_{-i}, s'_{-i}) such that $(\theta_i, s_i; \theta'_{-i}, s'_{-i}) \in \mathcal{S}(h)$ and $(\theta'_i, s'_i; \theta_{-i}, s_{-i}) \in \mathcal{S}(h)$. Thus property (FL) does not hold. □

Definition 1 and property (FL) provide a characterization of the observable deviators property in terms of strategies and information sets. These "strategic form/extensive form" characterizations turn out to be useful in game theoretic analysis, but "observable deviators" clearly is a property of the information structure. Therefore it is conceptually desirable to characterize the property exclusively in terms the primitive elements of the extensive form.[15]

Consider the following story. At the beginning of the game each player i says that her type is θ_i. Given the state θ resulting from these reports a "planner" or "mediator" recommends that everybody acts so as to reach some node t coming after θ. Then someone (the "planner" or one of the players) gets hard evidence that at least one player has not been honest, because the play has reached a set of nodes h (for example an information set) which is not intersected by any path through t (formally $Z(y) \cap Z(h) = \emptyset$). Who is to blame? The following definition provides a sufficient condition to conclude that at least one player in a given group J has either lied or deviated from the proposed course of action.

Definition 2. *Let $t \in X$, $h \subseteq X$, $J \subseteq I$. Group J is an* observed liar/deviator *with respect to t given evidence h, if for all $y \in h$, either $\vartheta_J(y) \neq \vartheta_J(t)$ (someone in J must have lied about her type) or there are $t', y' \in X_J$ such that $t' \prec t$, $y' \prec y$, $H(t') = H(y')$ and $\alpha(H(y') \to y) \neq \alpha(H(t') \to t)$ (someone in J must have deviated from the the action she was supposed to choose in order to reach t).*

Continuing with the above scenario, suppose that group $-i$ is not an observed liar/deviator with respect to t given h. Then it may be the case that player i with a unilaterally dishonest behavior caused h to be reached instead of t. So she is a *suspect unilateral liar/deviator* with respect to t given h. The observable deviators property can be characterized as follows: whenever player i is a suspect unilateral liar/deviator, then she actually is an observed liar/deviator.

(OD) $\forall t \in X$, $\forall h \in H$, $\forall i \in I$, if $Z(t) \cap Z(h) = \emptyset$ and group $-i$ is not an observed liar/deviator with respect to t given evidence h, then player i is an observed liar/deviator with respect to t given evidence h.

[15] There are game theoretic solution concepts which are stated in terms of paths and do not mention strategies at all (see e.g. Greenberg (1990) and Bonanno (1992)).

Note that we could have stated a similar, stronger property which holds for all groups $J \subset I$ and not only for single players $i \in I$, but it turns out that such a property is equivalent to (OD).

Proposition 2. *A game of incomplete information has observable deviators if and only if property (OD) holds.*

Proof. (If) Suppose that the game does *not* have observable deviators. Then, by Remark 2, there are there are $i \in I$, $h \in H$, such that $\mathcal{S}(h) \neq \mathcal{S}_i(h) \times \mathcal{S}_{-i}(h)$. Since it is always the case that $\mathcal{S}(h) \subseteq \mathcal{S}_i(h) \times \mathcal{S}_{-i}(h)$, there must be some $(\theta, s) \in (\mathcal{S}_i(h) \times \mathcal{S}_{-i}(h)) \backslash \mathcal{S}(h)$. Since $(\theta, s) \notin \mathcal{S}(h)$, there is some t on the path $\zeta(\theta, s)$ such that $Z(t) \cap Z(h) = \emptyset$. We show that neither $-i$ nor i are observed liar/deviators with respect to t given h and therefore (OD) fails. By choice of $t, (\theta_J, s_J) \in \mathcal{S}_J(t)$ for all $J \subset I$ (in particular $J = i, -i$). Since $(\theta_{-i}, s_{-i}) \in \mathcal{S}_{-i}(h)$, there is some $y \in h$ such that $(\theta_{-i}, s_{-i}) \in \mathcal{S}_{-i}(y)$. Group $-i$ cannot be an observed liar because $(\theta_{-i}, s_{-i}) \in \mathcal{S}_{-i}(y) \cap \mathcal{S}_{-i}(t)$ implies $\vartheta_{-i}(t) = \theta_{-i} = \vartheta_{-i}(y)$. Now suppose that there are $t', y' \in X_{-i}$ such that $t' \prec t$, $y' \prec x$ and $H(t') = H(y')$. Then $(\theta_{-i}, s_{-i}) \in \mathcal{S}_{-i}(y) \cap \mathcal{S}_{-i}(t)$ implies $\alpha(H(y') \to y) = s_{\iota(y')}(H(y')) = s_{\iota(t')}(H(t')) = \alpha(H(t') \to y)$. Thus $-i$ cannot be an observed deviator.

Since $(\theta_i, s_i) \in \mathcal{S}_i(h)$, there is some $w \in h$ such that $(\theta_i, s_i) \in \mathcal{S}_i(w)$. Thus $(\theta_i, s_i) \in \mathcal{S}_i(w) \cap \mathcal{S}_i(t)$. An argument similar to the one above shows that i cannot be an observed liar/deviator.

(Only if) Suppose that (OD) does not hold, viz. there are $i \in I$, $h \in H$, $t \in X$ such that $Z(t) \cap Z(h) = \emptyset$ and neither group $-i$ nor player i are observed deviators with respect to t given evidence h. Since $-i$ is not an observed deviator, there is $y \in h$ such that

(a) $\vartheta_{-i}(y) = \vartheta_{-i}(t)$ and whenever $y' \prec y$, $t' \prec t$ and $H(t') = H(y')$, then $\alpha(H(y') \to y) = \alpha(H(t') \to t)$.

Take a profile (θ_{-i}, s_{-i}) such that

(b) $\theta_{-i} = \vartheta_{-i}(y)$,

(c) for all $y' \in X_{-i}$, if $y' \prec y$, then $s_{\iota(y')}(H(y')) = \alpha(H(y') \to y)$,

(d) for all $t' \in X_{-i}$, if $t' \prec t$ and there is no $y' \prec y$ with $H(y') = H(t')$, then $s_{\iota(t')}(H(t')) = \alpha(H(t') \rightarrow t)$.

(clearly such a profile exists). Then (a), (b), (c) and (d) imply that $(\theta_{-i}, s_{-i}) \in \mathcal{S}_{-i}(y) \cap \mathcal{S}_{-i}(t)$.

With a similar argument one can show that, since i is not an observed deviator, then there must be some $w \in h$ and $(\theta_i, s_i) \in \mathcal{S}_i(w) \cap \mathcal{S}_i(t)$. Thus $(\theta_i, s_i) \in \mathcal{S}_i(w) \subseteq \mathcal{S}_i(h)$, $(\theta_{-i}, s_{-i}) \in \mathcal{S}_{-i}(y) \subseteq \mathcal{S}_{-i}(h)$ and $(\theta, s) \in \mathcal{S}(t)$. Since $(\theta, s) \in \mathcal{S}(t)$ and $Z(t) \cap Z(h) = \emptyset$, $(\theta, s) \notin \mathcal{S}(h)$. Thus $(\theta, s) \in (\mathcal{S}_i(h) \times \mathcal{S}_{-i}(h)) \backslash \mathcal{S}(h)$ and the game does not have observable deviators. □

An interesting class of games with observable deviators is given by multistage games with possibly incomplete information and observed actions (modulo essentially simultaneous moves). More generally it is sufficient that at each stage τ each moving player i has a separate information partition H^{τ}_{ij} over the previous stage actions of each opponent j. Figure 1 shows some examples of extensive forms with and without observable deviators. Nodes in the same information set are joined by a dashed line. The indexes i, j, k represent players. As long as perfect recall is satisfied, different indexes do not necessarily represent different players. Extensive form (a) has observable deviators. Also extensive form (b) - the only one with incomplete information - has observable deviators. But note that the corresponding extensive form with imperfect (not incomplete) information in which there is a unique initial node where the chance player chooses θ' or θ'' does not have observable deviators. Finally, extensive form (c) has observable deviators if and only if $i = j$.

Extensive forms (b) and (c) do not satisfy the following strengthening of the observable deviator property. Given an extensive form game Γ, let $\mathcal{A}(\Gamma)$ denote the *agent form* of Γ, viz. the extensive game obtained by assigning to each information set $h \in H_i$ of each player i in Γ a separate player (i, h) - also called *agent* - with payoff function u_i.[16]

For each $h \in H$, let $S^{\mathcal{A}}_h$ denote the set of strategies of player/agent (i, h) in $\mathcal{A}(\Gamma)$. A game Γ has *observable deviators in the agent form* if $\mathcal{A}(\Gamma)$ has observable deviators, viz. for all $h' \in H, S^{\mathcal{A}}(h') = \times_{h \in H} S^{\mathcal{A}}_h(h')$. Clearly neither (b) nor (c) in Figure 1 have observable deviators in the agent form. In general, no game of incomplete information where the actions of a privately informed player are (partially) observed can have observable deviators in the agent form.

[16]The normal form of $\mathcal{A}(\Gamma)$ is Selten's (1975) *agent normal form*.

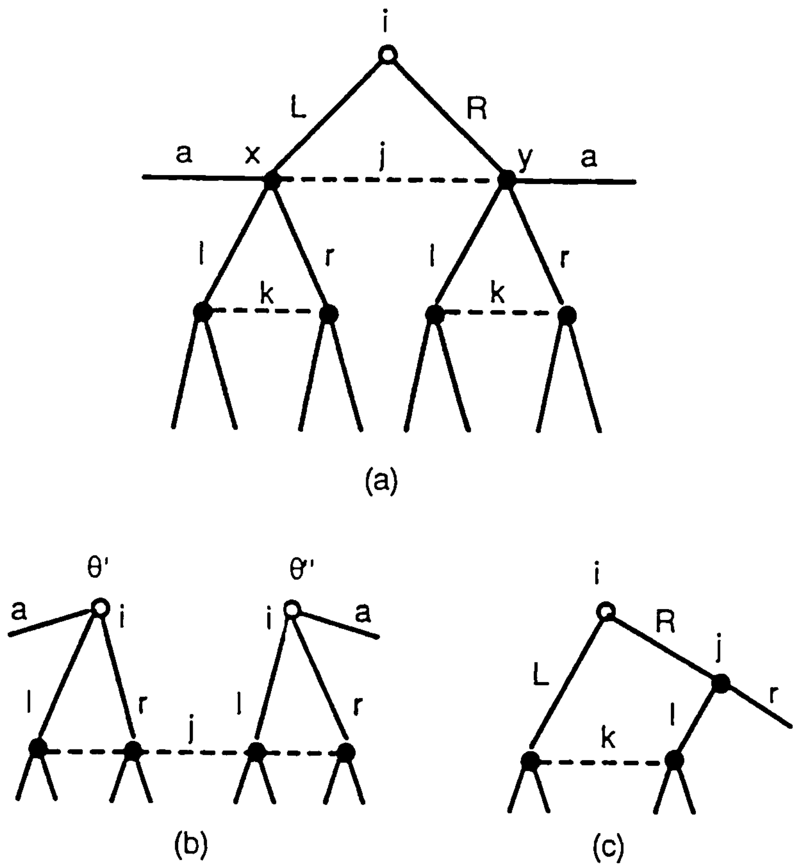

Figure 1 Three extensive forms.

This is a consequence of the following proposition, which characterizes games with observable deviators in the agent form. Recall that $A(x)$ is the set of feasible actions at x and $A(x \rightarrow h) \subseteq A(x)$ is the subset of actions which may lead to h from x.

Proposition 3. *A game has observable deviators in the agent form if and only if the following property holds:*
(ODAF) for all $h \in H, x, y \in h, x' \prec x$, if $A(x' \rightarrow h) \subset A(x')$ ($\subset$ means

"strict inclusion"), then there is one (and only one) node $y' \prec y$ such that $H(x') = H(y')$ and $A(x' \to h) = A(y' \to h)$.

Proof. See the Appendix. □

The characterization of games with observable deviators in the agent form is simplified for multistage games. Formally, an extensive form game has a *multistage* structure if for every information set $h \in H$ and every pair of nodes $x, y \in h$, the number of predecessors of x and y is the same (this is the class of extensive form games analyzed by Von Neumann and Morgenstern (1944)). A *stage* is a maximal set of nodes with the same number of predecessors. A superscript is used to denote the number of predecessors of a given node or information set, i.e. x^k (h^k) is a node (information set) with k predecessors. Note that all the extensive forms in Figure 1 except (c) have a multistage structure.[17]

The following corollary is an immediate consequence of Proposition 3.

Corollary 1. *A multistage game has observable deviators in the agent form if and only if for all stages $k \geq 0, \ell > k$, information sets h^ℓ and nodes $x^\ell, y^\ell \in h^\ell$, $x^k \prec x^\ell$, $y^k \prec y^\ell$, either*

$$H(x^k) = H(y^k) \text{ and } A(x^k \to h^\ell) = A(y^k \to h^\ell),$$

or

$$A(x^k \to h^\ell) = A(x^k) \text{ and } A(y^k \to h^\ell) = A(y^k).$$

[17] Extensive form (c) in Figure 1 can be transformed into a multistage one by adding a dummy node with only one feasible action for player i or j between the initial node and the left node of player k. Often such "irrelevant transformations" can turn a non multistage game into a multistage one.

2 Sequential Equilibria in Games with Observable Deviators

In this section we consider different properties for systems of conditional probabilities and equilibrium assessments in extensive games and we show that they are equivalent for the class games with observable deviators. Section 2.1 considers general systems of conditional probabilities on possibly uncountable spaces and introduces a stochastic independence property. In Section 2.2 we focus on finite extensive games with observable deviators and incomplete information. We show that for such games stochastic independence and full consistency are equivalent. Then we provide a simpler characterization of full consistency and sequential equilibria for the subset of games with observable deviators in the agent form. While it is not clear how to define the full consistency property for general infinite games, the characterizing properties discussed here are easily defined in the more general case.

2.1 Stochastic Independence for Systems of Conditional Probabilities

Consider a set $\mathcal{S}$, a sigma-algebra $\mathcal{A}$ of subsets of $\mathcal{S}$ and a collection of subsets $\mathcal{B} \subseteq \mathcal{A}$ which does not contain the empty set and is closed under finite union. We interpret $\mathcal{S}$ as a set of elementary events or states, $\mathcal{A}$ as the collection of relevant events and $\mathcal{B}$ as a collection of "hypotheses" or possible observations about the state. For example, if $\mathcal{S}$ is finite or countable, we may have $\mathcal{A} = 2^{\mathcal{S}}$ the power set of $\mathcal{S}$ and $\mathcal{B} = 2^{\mathcal{S}} \setminus \{\emptyset\}$. In game theoretic analysis the set of states may be the set of J-profiles of types and strategies $\mathcal{S}_J$ ($J \subseteq I$), the collection of hypotheses may be the algebra generated by $\{\mathcal{S}_J(h), h \in H\}$ (possible observations for some players), or by $\{\mathcal{S}_J(x), x \in X\}$ (the possible observations of a completely and perfectly informed "referee"), or the power set of $\mathcal{S}_J$ (minus $\{\emptyset\}$). A perhaps more appealing modeling choice is to have Z (the set of complete paths) as the set of elementary events. Whenever a sigma-algebra $\mathcal{A}$ on $\mathcal{S}$ is understood, $\Delta(\mathcal{S})$ denotes the set of probability measures on $\mathcal{S}$ with domain $\mathcal{A}$.

A *conditional probability system* (or simply "conditional system") on $(\mathcal{S}, \mathcal{A}, \mathcal{B})$ is a collection of conditional probability measures $\{\sigma(\cdot \mid B)\}_{B \in \mathcal{B}}$ jointly satisfying Bayes rule, or – more precisely – a function $\sigma(\cdot \mid \cdot) : \mathcal{A} \times \mathcal{B} \to \mathbb{R}$ such that

(i) $\forall (E, B) \in \mathcal{A} \times \mathcal{B}$, if $B \subseteq E$ then $\sigma(E \mid B) = 1$,

(ii) $\forall B \in \mathcal{B}$, $\sigma(\cdot \mid B) \in \Delta(\mathcal{S})$,

(iii) $\forall (E, F, B) \in \mathcal{A} \times \mathcal{A} \times \mathcal{B}$, if $F \cap B \in \mathcal{B}$ then $\sigma(E \cap F \mid B) = \sigma(E \mid F \cap B)\sigma(F \mid B)$.

(Note that (i) and (iii) imply $\sigma(E \mid B) = \sigma(E \cap B \mid B)$) A conditional probability system is *complete* if it is defined for the triple $(\mathcal{S}, 2^{\mathcal{S}}, 2^{\mathcal{S}} \backslash \{\emptyset\})$. In this case condition (iii) can be replaced by

(iii') $\forall E, F, B \in 2^{\mathcal{S}} \times (2^{\mathcal{S}} \backslash \{\emptyset\}) \times (2^{\mathcal{S}} \backslash \{\emptyset\})$, if $E \subseteq F \subseteq B$ then $\sigma(E \mid B) = \sigma(E \mid F)\sigma(F \mid B)$.

The set of complete conditional probability systems on $\mathcal{S}$ is denoted by $\Delta^{*}(\mathcal{S})$.[18]

In Section 2.2 we will focus on complete conditional probability systems over finite spaces. In this case $\Delta^{*}(\mathcal{S})$ can be regarded as a subset of a Euclidean space endowed with the usual metric of pointwise convergence, and it is easy to see that $\Delta^{*}(\mathcal{S})$ is a compact set.

For $\mathcal{S}$ *finite*, let $\Delta^{\circ}(\mathcal{S})$ denote the set of strictly positive probability measures on $\mathcal{S}$ and similarly let $\Delta^{*\circ}(\mathcal{S})$ denote the set of strictly positive conditional systems on $\mathcal{S}$. There is an obvious bijection between $\Delta^{*\circ}(\mathcal{S})$ and $\Delta^{\circ}(\mathcal{S}) : \Delta^{*\circ}(\mathcal{S}) \ni \sigma \mapsto \sigma(\cdot \mid \mathcal{S}) \in \Delta^{\circ}(\mathcal{S})$ and $\Delta^{\circ}(\mathcal{S}) \ni \sigma \mapsto (\sigma(E \mid B))_{E,B \subseteq \mathcal{S}, B \neq \emptyset} \in \Delta^{*\circ}(\mathcal{S})$, where $\sigma(E \mid B) = \sigma(E \cap B)/\sigma(B)$.

Result 1. (Myerson (1986)) Let $\mathcal{S}$ be finite. Then $\Delta^{*}(\mathcal{S})$ is the closure of $\Delta^{*\circ}(\mathcal{S})$.

Now assume that $(\mathcal{S}, \mathcal{A}, \mathcal{B})$ (finite or infinite) is a product of $(n+1)$ spaces $(\mathcal{S}_i, \mathcal{A}_i, \mathcal{B}_i)$ in the sense that: $\mathcal{S} = \mathcal{S}_0 \times \mathcal{S}_1 \times \ldots \times \mathcal{S}_n$, $\mathcal{A}$ is the product sigma-algebra obtained from $(\mathcal{A}_i)_{i \in I}$ and $\mathcal{B}$ contains all the products $B_0 \times B_1 \times \ldots \times B_n$, $(B_i \in \mathcal{B}_i)$. We also assume that $\mathcal{S}_i \in \mathcal{B}_i$ for all $i \in I = \{0, 1, \ldots, n\}$. For any $J \subset I$ (with at least two elements), let $\mathcal{A}_J$ denote the product sigma algebra obtained from $(\mathcal{A}_j)_{j \in J}$ and let $\mathcal{B}_J := \{B_J \in \mathcal{A}_J \mid B_J \times \mathcal{S}_{J} \in \mathcal{B}\}$. For any conditional system σ on $(\mathcal{S}, \mathcal{A}, \mathcal{B})$ we can define a *marginal conditional system* σ_J on $(\mathcal{S}_J, \mathcal{A}_J, \mathcal{B}_J)$ $(\emptyset \neq J \subseteq I)$ as follows:

[18] Rênyi (1955) analyzes general *conditional probability spaces*, viz. four-tuples $(\mathcal{S}, \mathcal{A}, \mathcal{B}, \sigma)$ whereby σ satisfies (i),(ii) and (iii). Myerson (1986) independently used the concept of a complete conditional probability system (on a finite space) in the context of game theoretic analysis.

$$\forall E_J \in \mathcal{A}_J, \forall B_J \in \mathcal{B}_J, \sigma_J(E_J \mid B_J) = \sigma(E_J \times S_{-J} \mid B_J \times S_{-J})$$

We are interested in conditional systems satisfying a stochastic independence condition which extends the usual product property of probability measures. In the game theoretic analysis this property reflects the assumption that the types and strategic choices of different players are stochastically independent.

Definition 3. *A conditional system σ satisfies the* stochastic independence property *if for all partitions $\{J_0, J_1, \ldots, J_K\}$ of I, all $(E_{J_0}, E_{J_1}, \ldots, E_{J_K}) \in \times_k \mathcal{A}_{J_k}$ and all $(B_{J_0}, B_{J_1}, \ldots, B_{J_K}) \in \times_k \mathcal{B}_{J_k}$,*[19]

$$\sigma(\times_k E_{J_k} \mid \times_k B_{J_k}) = \prod_k \sigma_{J_k}(E_{J_k} \mid B_{J_k}).$$

Remark 3. If σ satisfies the stochastic independence property then the following property holds:
(SI) for every bipartition $J \subset I$, all $(A_J, B_J) \in \mathcal{A}_J \times \mathcal{B}_J$, $C_{-J}, D_{-J} \in \mathcal{B}_{-J} \times \mathcal{B}_{-J}$,

$$\sigma(A_J \times C_{-J} \mid B_J \times C_{-J}) = \sigma(A_J \times D_{-J} \mid B_J \times D_{-J}).$$

If σ is a complete conditional system, then (SI) is not only necessary, but also sufficient for stochastic independence, as one can verify applying Bayes rule (iii') and (SI) iteratively.[20]

The meaning of (SI) is that the marginal conditional probabilities about group J are independent of information exclusively concerning the complementary group $-J$.

Consider again complete conditional systems on *finite* spaces. Given $S = S_0 \times S_1 \times \ldots \times S_n$ let $I\Delta^*(S)$ denote the set of conditional systems with the independence property. Clearly $I\Delta^*(S)$ is a compact subset of $\Delta^*(S)$.

[19] This is a strengthening of Renyi's notion of product of conditional probability spaces.
[20] See Battigalli (1996b). Battigalli and Veronesi (1996) provide a decision theoretic axiomatization of (SI) for complete conditional systems on finite spaces.

It is easily checked that $I\Delta^{*}(\mathcal{S}_0 \times \mathcal{S}_1 \times \ldots \times \mathcal{S}_n) \cap \Delta^{*\circ}(\mathcal{S}_0 \times \mathcal{S}_1 \times \ldots \times \mathcal{S}_n)$ is exactly the set of conditional systems derived from strictly positive product distributions. Let $\Psi(\mathcal{S}_0 \times \mathcal{S}_1 \times \ldots \times \mathcal{S}_n)$ denote the closure of this set. Note that also $\Psi(\mathcal{S}_0 \times \mathcal{S}_1 \times \ldots \times \mathcal{S}_n)$ is a compact set. The conditional systems in $\Psi(\mathcal{S}_0 \times \mathcal{S}_1 \times \ldots \times \mathcal{S}_n)$ are called *fully consistent* (cf. Kreps and Wilson (1982), McLennan (1989a,b), Myerson (1991)). Since $I\Delta^{*}(\mathcal{S}_0 \times \mathcal{S}_1 \times \ldots \times \mathcal{S}_n)$ is closed, we have: $\Psi(\mathcal{S}_0 \times \mathcal{S}_1 \times \ldots \times \mathcal{S}_n) \subseteq I\Delta^{*}(\mathcal{S}_0 \times \mathcal{S}_1 \times \ldots \times \mathcal{S}_n)$, i.e. every fully consistent conditional system satisfies the stochastic independence property. It can be shown that the converse does not hold.

2.2 Sequential Equilibria in Finite Games with Observable Deviators

In this section we provide characterizations of the sequential equilibrium concept for *finite* extensive games of incomplete information with observable deviators.

Recall that an *assessment* is a pair $(\mu, \pi) \in M \times \Pi$, where μ is a system of beliefs and π is a profile of behavioral strategies. An assessment (μ, π) is *sequentially rational* if for every personal player $i \in N$ and every information set $h \in H_i$, π_i maximizes i's conditional expected payoff given $\mu(\cdot \mid h)$ and π_{-i}. An assessment (μ, π) is *fully consistent* if there exists a sequence of strictly positive assessments $\{(\mu^k, \pi^k)\} \subseteq M \times \Pi$ converging to (μ, π)[21] and such that μ^k is derived from π^k (and ρ) via Bayes rule.[22]

A fully consistent and sequentially rational assessment is called *sequential equilibrium* (Kreps and Wilson (1982)).

Since assessments are lists of conditional probabilities of actions and nodes, they can be derived from richer conditional systems on $\mathcal{S}$ or Z. An assessment is a parsimonious description of players' expectations, since it lists just the necessary elements appearing in the sequential rationality condition. But the analysis of richer systems of conditional probabilities allows the formulation of "reasonable" restrictions on assessments.

[21] Recall that the behavioral strategy of the chance player π_0 is given ($\Pi = \{\pi_0\} \times \Pi_1 \times \ldots \times \Pi_n$) and strictly positive. Thus $\pi_0^k = \pi_0$, all k.

[22] Since π^k is strictly positive, for every information set h, we can set $\mu^k(x \mid h) = P^{\pi^k}(x)/(\sum_{y \in h} P^{\pi^k}(y))$.

An extended assessment is a triple (λ, μ, π) whereby λ is a conditional system and (μ, π) is derived from λ. In the spirit of traditional equilibrium analysis we assume here that an extended assessment describes the players' common conditional expectations or reflects the point of view of an outside observer. But most of our analysis can also be applied to solution concepts with heterogeneous conditional expectations. We analyze extended assessments (λ, μ, π) whereby λ is either a conditional system on $\mathcal{S}$ or a conditional system on Z.[23]

In both cases we require that λ is consistent with the prior ρ on Θ. Let $\Delta^*(\mathcal{S};\rho)$ $(\Delta^*(Z;\rho))$ be the set of conditional systems σ on $\mathcal{S} = \Theta \times S$ (ν on Z) such that for all $\theta \in \Theta$, $\sigma(\{\theta\} \times S \mid \Theta \times S) = \rho(\theta)$ $(\nu(Z(\theta) \mid Z) = \rho(\theta))$.

Definition 4. *A* strategic extended assessment *is a triple* $(\sigma, \mu, \pi) \in \Delta^*(\mathcal{S};\rho) \times M \times \Pi$ *such that*

$$\forall h \in H, \forall x \in h, \mu(x \mid h) = \sigma(\mathcal{S}(x) \mid \mathcal{S}(h)) \quad (2.1)$$
$$\forall h \in H, \forall x \in h, \forall a \in A(h), \pi(a \mid h) = \sigma(\mathcal{S}(f(x,a) \mid \mathcal{S}(x)) \quad (2.2)$$

A tree–extended assessment *is a triple* $(\nu, \mu, \pi) \in \Delta^*(Z;\rho) \times M \times \Pi$ *such that*

$$\forall h \in H, \forall x \in h, \mu(x \mid h) = \nu(Z(x) \mid Z(h)) \quad (2.3)$$
$$\forall h \in H, \forall x \in h, \forall a \in A(h), \pi(a \mid h) = \nu(Z(f(x,a)) \mid Z(x)). \quad (2.4)$$

It can be shown that if (σ, μ, π) is a strategic extended assessment and $\nu(Z(Y) \mid Z(W)) = \sigma(\mathcal{S}(Y) \mid \mathcal{S}(W))$ for all $Y, W \subseteq X$, then (ν, μ, π) is a tree–extended assessment. Equations (2.2) and (2.4) say that the given conditional system σ or ν is consistent with a unique profile of behavioral strategies, because the conditional probability of an action a at a node x in an information set h is independent of the particular node x selected in h. In games of imperfect information this condition can be violated by some conditional systems $\sigma \in \Delta^*(\mathcal{S};\rho)$ and $\nu \in \Delta^*(Z;\rho)$. Thus we are actually considering a restricted set of conditional systems. However, every conditional system $\sigma \in \Delta^*(\mathcal{S};\rho)$ with the

[23] Since we are now analyzing finite games, with an abuse of notation we can identify complete paths and terminal nodes.

stochastic independence property (and consistent with the given probabilities of chance moves π_0) satisfies (2.2) for some $\pi \in \Pi$.[24]

Note that if (σ, μ, π) is strictly positive, (2.1) and (2.2) are satisfied if and only if the profile of distributional strategies $(\sigma_i(\cdot \mid \mathcal{S}_i))_{i \in I}$ is derived from $\pi((\sigma_i(\cdot \mid \mathcal{S}_i))_{i \in I} = \delta(\pi))$ if and only if π is derived from the profile of distributional strategies $(\pi = \beta((\sigma_i(\cdot \mid \mathcal{S}_i))_{i \in I}))$ (note that we are implicitly using the independent types assumption). This fact can be used to prove the following result.

Result 2. An assessment (μ, π) is fully consistent if and only if there exists a fully consistent conditional system $\sigma \in \Psi(\mathcal{S}_0 \times \mathcal{S}_1 \times \ldots \times \mathcal{S}_n)$ such that (σ, μ, π) is a strategic extended assessment.

This result and the fact that fully consistent conditional systems necessarily satisfy the stochastic independence property yield the following corollary:

Corollary 2. *An assessment (μ, π) is fully consistent only if it is part of a strategic extended assessment (σ, μ, π) where σ has the independence property.*

The converse of Corollary 2 does not hold for general extensive games (see Kohlberg and Reny, (1992)), but it holds for games with observable deviators, as the following proposition shows.

Proposition 4. *In every game of incomplete information with independent types and observable deviators an assessment (μ, π) is fully consistent if and only if it is part of a strategic extended assessment (σ, μ, π) such that σ has the stochastic independence property (cf. Battigalli (1996b, Proposition 3.1)).*

Proof. (Only if) The "only if" part is Corollary 2.

(If) Assume that (σ, μ, π) is a strategic extended assessment such that $\sigma \in I\Delta^{\circ}(\mathcal{S})$. Each marginal conditional system σ_i is the limit of a sequence of

[24] $\forall i \in I, \forall h \in H_i, \forall x \in h, \forall a \in A(h), \mathcal{S}_{-i}(f(x,a)) = \mathcal{S}_{-i}(x)$ and, by perfect recall, $\mathcal{S}_i(x) = \mathcal{S}_i(h), \mathcal{S}_i(f(x,a)) = \mathcal{S}_i(h,a)$, where $\mathcal{S}_i(h,a) := \{(\theta_i, s_i) \in \mathcal{S}_i(h) \mid s_i(h) = a\}$. By stochastic independence, $\sigma(\mathcal{S}(f(x,a)) \mid \mathcal{S}(x)) = \sigma(\mathcal{S}_i(f(x,a)) \times \mathcal{S}_{-i}(x) \mid \mathcal{S}_i(x) \times \mathcal{S}_{-i}(x)) = \sigma_i(\mathcal{S}_i(h,a) \mid \mathcal{S}_i(h))$.

In the context of complete information games, Battigalli (1996b) gives a less restrictive definition of strategic extended assessments, but the difference is immaterial when the stochastic independence property is satisfied.

conditional systems $\{\sigma_i^k\}_0^\infty$ where each σ_i^k is generated by a strictly positive distributional strategy $\sigma_i^k(\cdot \mid \Theta_i \times S_i)$. Let $\hat{\sigma}^k$ be the conditional system derived from the strictly positive product measure $\times_{i \in I}\sigma_i^k(\cdot \mid S_i)$. Note that $\hat{\sigma}^k \in \Psi(\mathcal{S})$. By compactness, $\{\hat{\sigma}^k\}_0^\infty$ has a cluster point $\hat{\sigma} \in \Psi(\mathcal{S})$. By construction, σ and $\hat{\sigma}$ have the same marginals σ_i and induce the same π. Since σ is consistent with the prior ρ on Θ so is $\hat{\sigma}$. By the assumption of observable deviators, for all $h \in H, x \in h, \mathcal{S}(h) = \mathcal{S}_1(h) \times \ldots \times \mathcal{S}_n(h)$. Hence stochastic independence yields

$$\mu(x \mid h) = \sigma(\mathcal{S}(x) \mid \mathcal{S}(h)) = \prod_{j=0}^{j=n} \sigma_j(\mathcal{S}_j(x) \mid \mathcal{S}_j(h)) = \hat{\sigma}(\mathcal{S}(x) \mid \mathcal{S}(h)).$$

Thus $(\hat{\sigma}, \mu, \pi)$ is a strategic extended assessment and full consistency of $\hat{\sigma}$ implies full consistency of (μ, π). □

We now consider conditional systems on the set Z of terminal nodes (or complete paths) instead of the set $\mathcal{S}$ of strategy profiles. The advantage of this approach is that the involved events are observable (at least in principle), while subsets of $\mathcal{S}$ do not always correspond to observable events.

Definition 5. *(cf. Fudenberg and Tirole (1991)). A tree-extended assessment (ν, π, μ) is* generally reasonable *if*

$$\begin{aligned} &\forall h \in H, \forall a \in A(h), \forall x, y \in h : \\ &\quad \nu(Z(f(x,a))|Z(\{f(x,a), f(y,a)\})) = \nu(Z(x)|Z(\{x,y\})). \end{aligned} \tag{2.5}$$

Equation (2.5) says that, according to ν, the choice of player $\iota(h)$ at $h \in H$ cannot affect the relative probabilities of the events "x is reached" and "y is reached", because $\iota(h)$ cannot distinguish x from y. (Equality (2.5) is implied by (2.3) and (2.4) if $\pi(a \mid h) > 0$ and $\mu(x \mid h) + \mu(y \mid h) > 0$.) It can be shown that this pattern of conditional probabilities on nodes and actions is implied by the independence property of conditional systems on types–strategies (cf. Battigalli (1996b, Proposition 5.2)).

It has been conjectured that a generally reasonable extended assessment contains a fully consistent assessment.[25] This would imply that the sequential

[25] See Fudenberg and Levine (1991, Proposition 6.1).

equilibrium concept is characterized by sequential rationality and general reasonableness. But Battigalli (1996b) shows that general reasonableness is strictly weaker than stochastic independence even in the restricted class of multistage games with observable deviators in the agent form, where stochastic independence is equivalent to full consistency (by Proposition 4 full consistency is equivalent to stochastic independence in all games with observable deviators). To see that general reasonableness is weaker than stochastic independence we can consider the extensive form depicted in Figure 1 (a). It is possible to construct a generally reasonable extended assessment such that the probability of actions L and a is one and player k's conditional belief about player j's choice (ℓ or r) depends on player i's observed choice (L or R). To rule out this possibility Battigalli (1996b) considers the following additional property:

$$\forall h \in H, \forall x, y \in h, \forall a, b \in A(h):$$
$$\nu(Z(f(x,a)) \mid Z(\{f(x,a), f(x,b)\})) = \nu(Z(f(y,a)) \mid Z(\{f(y,a), f(y,b)\})). \tag{2.4'}$$

(2.4') is a completion of (2.4) in the following sense: if (ν, μ, π) is a tree–extended assessment satisfying (2.4'), for each h, ν uniquely induces a whole conditional system $\pi(\cdot \mid \cdot; h) \in \Delta^*(A(h))$ with prior $\pi(\cdot \mid h)$ instead of the simple probability measure $\pi(\cdot \mid h)$. The conditional system $\pi(\cdot \mid \cdot; h)$ is derived from the "binary" conditional probabilities

$$\pi(a \mid \{a,b\}; h) = \nu(Z(f(x,a) \mid Z(\{f(x,a), f(x,b)\})$$

where $x \in h$ can be arbitrarily chosen. A (generally reasonable) tree–extended assessment (ν, μ, π) can fail to satisfy (2.4') if $\pi(a \mid h) = 0 = \pi(b \mid h)$ as in the example mentioned above. It can be shown that if (ν, μ, π) is derived from a strategic extended assessment satisfying stochastic independence, then also (2.4') is satisfied (see Battigalli (1996b, Proposition 5.2')). The following proposition yields a characterization of sequential equilibria in terms of tree–extended assessments:

Proposition 5. *(Battigalli (1996b, Proposition 5.3)) In every multi-stage game with observable deviators in the agent form the following statements are equivalent:*

(i) (μ, π) is part of a tree-extended assessment (ν, μ, π) satisfying (2.4') and (2.5),

(ii) (μ, π) is part of a strategic extended assessment (σ, μ, π) such that σ satisfies the stochastic independence property,

(iii) (μ, π) is fully consistent.

We have already mentioned that (2.4') is necessary to obtain the equivalence of (i) with (ii) and (iii). Building on the example mentioned above, one can also show that there are generally reasonable and sequentially rational extended assessments (ν, μ, π) which do not satisfy (2.4') and are such that the behavioral strategies profile π is not part of any sequential equilibrium. However, it turns out that in multi–stage games with observable deviators in the agent form, general reasonableness and sequential rationality are sufficient to characterize sequential equilibrium *outcomes*. That is:

Proposition 6. *For every multistage game with observable deviators in the agent form, every generally reasonable and sequentially rational extended assessment $(\hat{\nu}, \hat{\mu}, \hat{\pi})$ induces the same probability measure over terminal nodes as some sequential equilibrium (μ^*, π^*).*

Here we prove this claim for games with the extensive form depicted in Figure 1 (a) (independently of payoffs). This gives a hint about the general proof contained in the appendix. Fix a sequentially rational and generally reasonable extended assessment $(\hat{\nu}, \hat{\mu}, \hat{\pi})$. First note that we may assume without loss of generality that condition (2.4') holds at every information set except $\{x, y\}$, because $(\hat{\nu}, \hat{\mu}, \hat{\pi})$ satisfies (2.4). We consider two cases:

(i) either $\hat{\pi}(a) < 1$ or $0 < \hat{\nu}(Z(x) \mid Z(\{x, y\})) < 1$,

(ii) $\hat{\pi}(a) = 1$ and $\hat{\nu}(Z(x) \mid Z(\{x, y\})) \in \{0, 1\}$.

In case (i), (2.4) implies (2.4') and, by Proposition 5 $(\hat{\mu}, \hat{\pi})$ is a sequential equilibrium. This is easily checked if $\hat{\pi}(A'') < 1$. If $0 < \hat{\nu}(Z(x) \mid Z(\{x, y\})) < 1$, apply the following result:

Remark 4. For every $\nu \in \Delta^*(Z)$, if (2.5) holds, then for all $h \in H, x, y \in h, a, b \in A(h)$, if $0 < \nu(Z(x) \mid Z(\{x, y\})) < 1$ then

$$\nu(Z(f(x,a)) \mid Z(\{f(x,a), f(x,b)\})) = \nu(Z(f(y,a)) \mid Z(\{f(y,a), f(y,b)\})).$$

In case (ii), without loss of generality assume $\hat{\nu}(Z(x) \mid Z(\{x,y\})) = 1$. Then for every action a by player i or j the probability of reaching player k's right information set given $(\hat{\mu}, \hat{\pi}/a)$ is zero. Thus we can modify the belief and the local strategy of k after R' without affecting the incentives of i and j. The outcome–equivalent sequential equilibrium (μ^*, π^*) is obtained assigning to the right information set of player k the same belief (and corresponding best reply) of the left information set.

3 Conclusions

In this paper we have provided some characterizations of the observable deviators property for extensive form games with incomplete information. By considering this somewhat restricted class of games it is possible to simplify the game theoretic analysis and improve our understanding of some solution concepts.

We focused our attention on Kreps and Wilson's (1982) sequential equilibrium concept, providing alternative characterizations which emphasize the structural properties of beliefs. In particular, we have shown that every generally reasonable and sequentially rational assessment is outcome equivalent to a sequential equilibrium assessment.

While the original definition of sequential equilibrium can only be applied to finite games and a few special infinite games, the alternative characterizing properties can be easily extended to general, infinite extensive form games.

Appendix

A.1 Notation on Extensive Form Games with Incomplete Information

Notation	*Terminology*
$(X, \prec)$	arborescence
$x \preceq y$	either $x \prec y$ or $x = y$
$\theta = (\theta_1, \ldots, \theta_n) \in \Theta = \Theta_1 \times \ldots \times \Theta_n$	types vectors,initial nodes
$\vartheta(y)$, $(\vartheta : X \rightarrow \Theta)$	initial predec. of node y
$f(x)$	immed.successors of node x
$p \in Z$	complete paths
$Z(Y)$, $(Y \subseteq X)$	complete paths through Y
$z \in Z^*$	terminal nodes
$\iota : X \backslash Z^* \rightarrow \{0, 1, \ldots, n\}$	player function (0 =chance)
$h \in H$, $(h \in H_i)$	information sets (for i)
$H(x)$	inf. set containing node x
$a \in A(x)$ (or $A(h)$)	actions at x (or h)
$f(x, a)$	immed.succ. of x after a
$\alpha(h \rightarrow y)$ $(h \in H)$	first action from h to y
$A(x \rightarrow Y) = \{a \in A(x) \mid \exists y \in Y : f(x,a) \preceq y\}$	actions from $x \in X$ to $Y \subset X$
$Z(x, \hat{A}) = \{z \in Z(x) \mid \exists a \in \hat{A}\text{: } f(x,a) \preceq z\}$	term. succ.of actions $a \in \hat{A}$ from x
$\rho \in \Delta(\Theta)$, $\rho = \rho_1 \times \ldots \times \rho_n$	product prior on types
$\pi_0 \in (\times_{h \in H_0} \Delta(A(h)))$	prob. of chance moves
$\mathcal{S}_i = \Theta_i \times S_i$	i's type–strategy pairs
$S_h^A = \{(a \text{ if } h) : a \in A(h)\}$ $(h \in H_i)$	strategies of agent (i, h)
$\pi \in \Pi = \times_{h \in H} \Delta(A(h))$	behavioral profiles
$P^\pi(y \mid x)$, $(\pi \in \Pi)$	prob. of reaching y from x given π
$\mu \in M := \times_{h \in H} \Delta(h)$, $(\mu, \pi) \in M \times \Pi$	beliefs, assessments
$\sigma(\cdot \mid \cdot) \in \Delta^*(\mathcal{S}), \nu(\cdot \mid \cdot) \in \Delta^*(Z)$	cond.prob.syst. on $\mathcal{S}, Z$
x^k, h^k	stage k nodes, inform. sets

A.2 Proof of Proposition 3

We must show that a game has observable deviators in the agent form if and only if the following property holds:

(ODAF) for all $h \in H, x, y \in h, x' \prec x$, if $A(x' \to h) \subset A(x')$ ($\subset$ means "strict inclusion"), then there is one (and only one) node $y' \prec y$ such that $H(x') = H(y')$ and $A(x' \to h) = A(y' \to h)$.

(*If*) Assume that (ODAF) holds. Take an information set h. Then each $S_g^{\mathcal{A}}(h)$ ($h \neq g \in H$) is given by one of the following expressions:

(i) If we can find $x' \in g$ and $x \in h$ such that $x' \prec x$ and $A(x' \to h) \subset A(x')$, then $S_g^{\mathcal{A}}(h) = \{a \text{ if } g \mid a \in A(x' \to h)\}$, because in this case (ODAF) implies that each path reaching h intersects g at some node y' and and agent $(\iota(g), g)$ can prevent h from being reached if and only if he chooses $a \in A(g) \backslash A(y' \to h) = A(g) \backslash A(x' \to g)$;

(ii) Otherwise, either information set g is not intersected by any path reaching h, or it is intersected but no action at g can prevent h from being reached, because - by (ODAF) - for all $x' \in g$, $A(x') = A(x' \to h)$.

Consider a strategy profile $s = (a_g \text{ if } g)_{g \in H} \in (\times_{g \in H} S_g^{\mathcal{A}}(h))$. We have to show that for each initial predecessor θ of h, the induced path $\zeta(\theta, s)$ intersects h. Let $\zeta(\theta, s) = (\theta = x_0, x_1, \ldots)$ be the induced path (thus $x_{k+1} = f(x_k, a_{H(x_k)})$). We show that for each k, if there is an $x \in h$ such that $x_k \prec x$, then there is $y \in h$ such that $x_{k+1} \preceq y$. Since either $h \subseteq \Theta$ or $\theta = x_0$ is a predecessor of h, this implies the thesis.

Let $x_k \prec x \in h$. Then - by (i) or (ii) - $(a_{H(x_k)} \text{ if } H(x_k)) \in S_{H(x_k)}(h)$ implies $a_{H(x_k)} \in A(x_k \to h)$. Hence there is $y \in h$ such that $x_{k+1} = f(x_k, a_{H(x_k)}) \preceq y$.

(*Only if*) Assume that (ODAF) does not hold. Then there are two possible (non mutually exclusive) cases.

(i) There are an information set h with two distinct nodes $x, y \in h$ and a node $x' \prec x$ such that $A(x' \to h) \subset A(x')$ and for every $y' \prec y$, $H(y') \neq H(x')$. Consider an agent form strategy profile $(s_g^{\mathcal{A}})_{g \in H}$ inducing a path through node y. Since the induced path does not intersect $H(x')$, we may assume without loss of generality that $s_{H(x')}^{\mathcal{A}}$ prescribes an action $a^* \in A(x') \backslash A(x' \to h)$ if $H(x')$ is reached. Now consider a profile $(t_g)_{g \in H}$ inducing a path through

x. Clearly $(s_g^{\mathcal{A}})_{g\in H} \in S(h)$ and $(t_g^{\mathcal{A}})_{g\in H} \in S(h)$, thus $(s^{\mathcal{A}}_{H(x')}, t^{\mathcal{A}}_{-H(x')}) \in S^{\mathcal{A}}_{H(x')}(h) \times S^{\mathcal{A}}_{-H(x')}(h)$. But $(s^{\mathcal{A}}_{H(x')}, t^{\mathcal{A}}_{-H(x')}) \notin S^{\mathcal{A}}(h)$, because the induced path reaches x' where choice a^* prevents h from being reached.

(ii) There are an information set h with two distinct nodes $x, y \in h$ and nodes $x' < x, y' < y$ such that $H(y') = H(x')$, $A(x' \to h) \subset A(x')$, and $A(x' \to h) \neq A(y' \to h)$. We may assume without loss of generality that there is an action $a^* \in A(y' \to h)\backslash A(x' \to h)$ (if, for every choice of x, x', y, y', $A(y' \to h)\backslash A(x' \to h)$ were empty, then h would satisfy condition (ODAF)). Let $w \in h$ be a node satisfying $f(y', a^*) \preceq w$. Consider an agent form strategy profile $(s_g^{\mathcal{A}})_{g\in H}$ inducing a path through $w \in h$ and a profile $(t_g^{\mathcal{A}})_{g\in H}$ inducing a path through x. Clearly $(s_g^{\mathcal{A}})_{g\in H} \in S^{\mathcal{A}}(h)$ and $(t_g^{\mathcal{A}})_{g\in H} \in S^{\mathcal{A}}(h)$, thus $(s^{\mathcal{A}}_{H(x')}, t^{\mathcal{A}}_{-H(x')}) \in S^{\mathcal{A}}_{H(x')}(h) \times S^{\mathcal{A}}_{-H(x')}(h)$. But $(s^{\mathcal{A}}_{H(x')}, t^{\mathcal{A}}_{-H(x')}) \notin S^{\mathcal{A}}(h)$, because the induced path reaches x' where choice a^* prevents h from being reached.

We have shown that in both cases we can find two information sets h and g such that $S^{\mathcal{A}}(h) \subset S_g^{\mathcal{A}}(h) \times S_{-g}^{\mathcal{A}}(h)$. Therefore deviators are not observable in the agent form if (ODAF) does not hold.

A.3 Proof of Proposition 6

We have to prove that in every multistage game with observable deviators in the agent form every sequentially rational and generally reasonable extended assessment $(\hat{\nu}, \hat{\mu}, \hat{\pi})$ induces the same probability distribution on terminal nodes as some sequential equilibrium (μ^*, π^*).

The proof is in three parts. In part I we define and characterize the class $H(\mu, \pi)$ of information sets which are relevant in order to check the equilibrium outcome conditions for a given assessment (μ, π). In part II we construct a fully consistent assessment (μ^*, π^*) such that $\hat{\pi}$ and π^* coincide on $H(\hat{\mu}, \hat{\pi})$. Since $H(\hat{\mu}, \hat{\pi})$ is a superset of the information sets reached with positive probability by $\hat{\pi}$ and $\pi^*, \hat{\pi}$ and π^* induce the same probability distribution on the set of terminal nodes Z. In part III we show that $\hat{\mu}$ and μ^* coincide on $H(\hat{\mu}, \hat{\pi}) = H(\mu^*, \pi^*)$ and (μ^*, π^*) is a sequential equilibrium.

I. (*Definition and characterization of (μ, π)-relevant information sets.*) For any given assessment (μ, π) let $P^{\mu,\pi}(h \mid h')$ denote the conditional probability of reaching h from an information set $h' \in H$. From the point of view of an agent h' who is trying to maximize his conditional expected utility given (μ, π), the only relevant components of (μ, π) are $\mu(\cdot \mid h')$ and the local (randomized) strategies $\pi(\cdot \mid h)$ at those information sets h which can be reached with positive conditional probability choosing some action $a \in A(h')$. This motivates the following definition.

Definition A.1. *An information set h is a (μ, π)-*follower *of h' if there is an action $a \in A(h')$ such that $P^{\mu,\pi/a}(h \mid h') > 0$, where π/a is the behavioral profile obtained from π replacing $\pi(\cdot \mid h')$ with a. An information set h is (μ, π)-*relevant, *written $h \in H(\mu, \pi)$, if there are information sets $h_0, \ldots, h_\ell, \ldots, h_L$ such that $h_0 \subseteq \Theta, h_L = h$ and $h_{\ell+1}$ is a (μ, π)-follower of h_ℓ, $\ell = 0, \ldots, L-1$.*

For example, consider an assessment (μ, π) for the extensive form depicted in Figure 1 (a) whereby the probability of L and a is one: the information sets of players i, j and the left information set of player k are (μ, π)-relevant, while the right information set of j is not (μ, π)-relevant. This means that we can modify (μ, π) on the latter information set leaving the best response conditions at the other information sets unaltered.

Clearly, every information set h reached from an initial node with positive π-probability is (μ, π)-relevant.

We use a superscript such as ϑ, t, T to denote the stage containing a node, action, or information set. A path from h^ϑ to h^{T+1} is a sequence $(x^\vartheta, a^\vartheta, \ldots, x^t, a^t, \ldots, x^T, a^T, x^{T+1})$ of nodes and actions such that $x^\vartheta \in h^\vartheta$, $x^{T+1} \in h^{T+1}$, $x^{t+1} = f(x^t, a^t)$, for all $t = \vartheta, \ldots, T$. Equivalently, the sequence is such that $x^\vartheta \in h^\vartheta$, $x^{t+1} = f(x^t, a^t)$, $a^t \in A(x^t \to h^{T+1})$ for all $t = \vartheta, \ldots, T$.

Remark A.1. h^{T+1} is a (μ, π)-follower of h^ϑ, $\vartheta \leq T$, if and only if there is a path $(x^\vartheta, a^\vartheta, \ldots, x^t, a^t, \ldots, x^T, a^T, x^{T+1})$ from h^ϑ to h^{T+1} such that $\mu(x^\vartheta \mid h^\vartheta) > 0, \pi(a^t \mid H(x^t)) > 0$ and $H(x^t)$ is a (μ, π)-follower of h^ϑ for all $t = \vartheta + 1, \ldots, T$. This is made apparent by the following equation

$$P^{\mu,\pi/a^\vartheta}(h^{T+1} \mid h^\vartheta) = \\ = \sum_{x^\vartheta \in h^\vartheta} \mu(x^\vartheta \mid h^\vartheta) \left[\sum_{x^{T+1} \in h^{T+1}, f(x^\vartheta, a^\vartheta) \leq x^{T+1}} P^\pi(x^{T+1} \mid f(x^\vartheta, a^\vartheta)) \right].$$

All the following claims hold for multi–stage games with observable deviators in the agent form.

Claim A.1. *Let* $(x^\vartheta, a^\vartheta, \ldots, x^t, a^t, \ldots, x^T, a^T, x^{T+1})$ *be a path from* $h^\vartheta \in H$ *to* $h^{T+1} \in H (\vartheta \leq T)$ *and consider a new path* $(w^\vartheta, c^\vartheta, \ldots, w^t, c^t, \ldots, w^T, c^T, w^{T+1})$ *with the following properties:*

(i) $w^\vartheta \in h^\vartheta, c^\vartheta \in A(w^\vartheta \to h^{T+1})$,

(ii) $\forall t \in \{\vartheta+1, \ldots, T\}$, $H(w^t) = H(x^t) \Rightarrow c^t = a^t$.

Then $w^{T+1} \in h^{T+1}$.

Proof. We show by induction that $c^t \in A(w^t \to h^{T+1})$ for $t = \vartheta, \ldots, T$. This implies that $w^{T+1} \in h^{T+1}$and $c^\vartheta \in A(w^\vartheta \to h^{T+1})$ by (i). Assume that $c^\tau \in A(w^\tau \to h^{T+1})$ for $\tau = \vartheta, \ldots, t-1$. Then w^t and x^t are predecessor of h^{T+1}, and Corollary 1 implies that either

(a) $H(w^t) = H(x^t)$ and $A(w^t \to h^{T+1}) = A(x^t \to h^{T+1})$, or

(b) $H(w^t) \neq H(x^t)$ and $A(w^t \to h^{T+1}) = A(w^t)$.

In case (a), (ii) implies $c^t = a^t$ and we obtain $c^t = a^t \in A(x^t \to h^{T+1}) = A(w^t \to h^{T+1})$. In case (b) it is trivially true that $c^t \in A(w^t \to h^{T+1})$. This concludes the proof of Claim A.1. □

For any non terminal node x and subset of actions $\hat{A} \subseteq A(x)$, let $Z(x, \hat{A})$ be the set of terminal nodes following x after some action $a \in \hat{A}$. (Clearly $Z(x, A(x)) = Z(x)$.) Thus $\nu(Z(f(x,a)) \mid Z(x, \hat{A}))$ gives the conditional probability of a given $\hat{A}$ evaluated at node x.

Claim A.2. *If* h^{T+1} *is a*(μ, π)*-follower of* h^ϑ $(\vartheta \leq T)$, *then there is a path* $(w^\vartheta, c^\vartheta, \ldots, w^t, c^t, \ldots, w^T, c^T, w^{T+1})$ *from* h^ϑ *to* h^{T+1} *such that* $\mu(w^\vartheta \mid h^\vartheta) > 0$, $\nu(Z(f(w^\vartheta, c^\vartheta)) \mid Z(w^\vartheta, A(w^\vartheta \to h^{T+1}))) > 0$ *and* $\pi(c^t \mid H(w^t)) > 0$ *for all* $t = \vartheta+1, \ldots, T$.

Proof. Take a path $(x^\vartheta, a^\vartheta, \ldots, x^t, a^t, \ldots, x^T, a^T, x^{T+1})$ from h^ϑ to h^{T+1} as in Remark A.1 and construct a new path $(w^\vartheta, c^\vartheta, \ldots, w^t, c^t, \ldots, w^T, c^T, w^{T+1})$ as follows:

$$w^{\vartheta} = x^{\vartheta},$$
$$c^{\vartheta} \in A(x^{\vartheta} \to h^{T+1}),$$
$$\nu(Z(f(x^{\vartheta}, c^{\vartheta})) \mid Z(x^{\vartheta}, A(x^{\vartheta} \to h^{T+1}))) > 0;$$

for all $t = \vartheta + 1, \ldots, T$,

$$\begin{aligned} H(w^t) = H(x^t) &\Rightarrow c^t = a^t \\ (\text{thus } \pi(c^t \mid H(w^t)) &= \pi(a^t \mid H(x^t)) > 0), \end{aligned}$$

otherwise choose any c^tsuch that $\pi(c^t \mid H(w^t)) > 0$.

Note that $(w^{\vartheta}, c^{\vartheta}, \ldots, w^t, c^t, \ldots, w^T, c^T, w^{T+1})$satisfies (i) and (ii) of Claim A.1. Therefore $w^{T+1} \in h^{T+1}$ and all the desired properties are satisfied. This concludes the proof of Claim A.2. □

Claim A.3. *Assume that (ν, μ, π) is generally reasonable and h^{T+1} is a (μ, π)-follower of h^{ϑ} $(\vartheta \leq T)$. Let $(y^{\vartheta}, b^{\vartheta}, \ldots, y^t, b^t, \ldots, y^T, b^T, y^{T+1})$ be a path from h^{ϑ} to h^{T+1}. Then $\mu(y^{T+1} \mid h^{T+1}) > 0$ only if $\mu(y^{\vartheta} \mid h^{\vartheta}) > 0, P^{\pi}(y^{T+1} \mid f(y^{\vartheta}, b^{\vartheta})) > 0$ and $\nu(Z(f(y^{\vartheta}, b^{\vartheta})) \mid Z(y^{\vartheta}, A(y^{\vartheta} \to h^{T+1}))) > 0$. Furthermore, if $Z(h^{T+1}) \subseteq Z(h^{\vartheta})$, then $\mu(y^{T+1} \mid h^{T+1}) > 0$ if and only if $\mu(y^{\vartheta} \mid h^{\vartheta}) > 0$, $P^{\pi}(y^{T+1} \mid f(y^{\vartheta}, b^{\vartheta})) > 0$ and $\nu(Z(f(y^{\vartheta}, b^{\vartheta})) \mid Z(y^{\vartheta}, A(y^{\vartheta} \to h^{T+1}))) > 0$.*

Proof. Take a path $(w^{\vartheta}, c^{\vartheta}, \ldots, w^T, c^T, w^{T+1})$ from h^{ϑ} to h^{T+1} as in Claim A.2. We first show that $\mu(y^{T+1} \mid h^{T+1}) > 0$ implies $\mu(y^{\vartheta} \mid h^{\vartheta}) > 0$ and $P^{\pi}(y^{T+1} \mid f(y^{\vartheta}, b^{\vartheta})) > 0$. Construct a new path ($x^{\vartheta}$, a^{ϑ}, ..., x^t, a^t, .., x^T, a^T, x^{T+1}) as follows (see Figure 2):

$$\begin{aligned} x^{\vartheta} &= w^{\vartheta}, \\ a^{\vartheta} &= b^{\vartheta}, \end{aligned}$$

for all $t = \vartheta + 1, \ldots, T$,

$$\begin{aligned} H(x^t) = H(w^t) &\Rightarrow a^t = c^t \\ (\text{thus } \pi(a^t \mid H(x^t)) &= \pi(c^t \mid H(w^t)) > 0), \end{aligned}$$

otherwise choose any a^t such that $\pi(a^t \mid H(w^t)) > 0$.

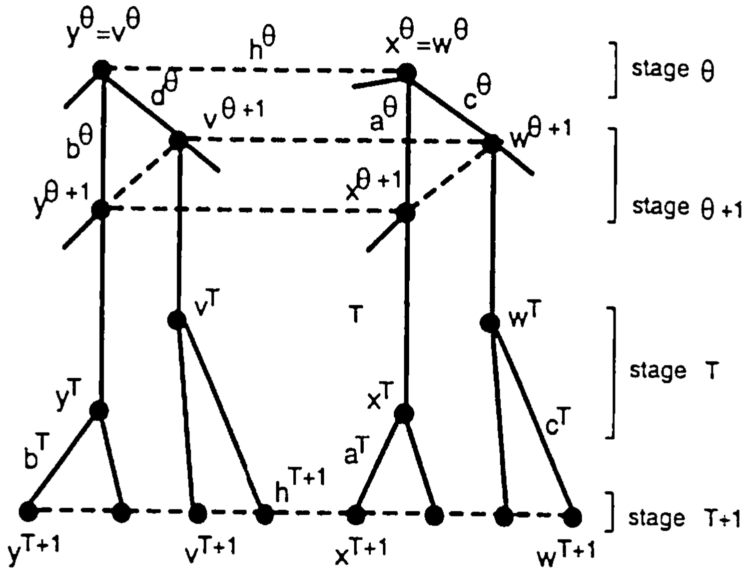

Figure 2 A fragment of a multistage extensive form with observable deviators in the agent form.

Since $y^\vartheta, w^\vartheta \in h^\vartheta$ and $y^{T+1}, w^{T+1} \in h^{T+1}$, the construction and Corollary 1 yield

$$a^\vartheta = b^\vartheta \in A(y^\vartheta \to h^{T+1}) = A(w^\vartheta \to h^{T+1}) = A(x^\vartheta \to h^{T+1}).$$

Thus Claim A.1 implies $x^{T+1} \in h^{T+1}$. By construction, $\mu(x^\vartheta \mid h^\vartheta) > 0$ and $P^\pi(x^{T+1} \mid f(x^\vartheta, b^\vartheta)) > 0$. (2.5) and $\mu(x^\vartheta \mid h^\vartheta) > 0$ yield

$$\begin{aligned}&\nu(Z(f(x^\vartheta, b^\vartheta)) \mid Z(\{f(x^\vartheta, b^\vartheta), f(y^\vartheta, b^\vartheta)\})) = \\ &= \nu(Z(x^\vartheta) \mid Z(\{x^\vartheta, y^\vartheta\})) = \frac{\mu(x^\vartheta \mid h^\vartheta)}{\mu(\{x^\vartheta, y^\vartheta \mid h^\vartheta)}.\end{aligned}$$

Therefore we obtain

$$\frac{\nu(Z(y^{T+1}) \mid Z(\{x^{T+1}, y^{T+1}\}))}{\nu(Z(x^{T+1}) \mid Z(\{x^{T+1}, y^{T+1}\}))} =$$
$$= \frac{P^{\pi}(y^{T+1} \mid f(y^{\vartheta}, b^{\vartheta}))\nu(Z(f(y^{\vartheta}, b^{\vartheta})) \mid Z(\{f(x^{\vartheta}, b^{\vartheta}), f(y^{\vartheta}, b^{\vartheta})\}))}{P^{\pi}(x^{T+1} \mid f(x^{\vartheta}, b^{\vartheta}))\nu(Z(f(x^{\vartheta}, b^{\vartheta})) \mid Z(\{f(x^{\vartheta}, b^{\vartheta}), f(y^{\vartheta}, b^{\vartheta})\}))} =$$
$$= \frac{P^{\pi}(y^{T+1} \mid f(y^{\vartheta}, b^{\vartheta}))\mu(y^{\vartheta} \mid h^{\vartheta})}{P^{\pi}(x^{T+1} \mid f(x^{\vartheta}, b^{\vartheta}))\mu(x^{\vartheta} \mid h^{\vartheta})}. \tag{A.1}$$

As $x \in h^{T+1}$and $\mu(x^{\vartheta} \mid h^{\vartheta}) > 0$, (A.1) implies that $\mu(y^{T+1} \mid h^{T+1}) > 0$ only if $\mu(y^{\vartheta} \mid h^{\vartheta}) > 0$ and $P^{\pi}(y^{T+1} \mid f(y^{\vartheta}, b^{\vartheta})) > 0$.

Now we show that $\mu(y^{T+1} \mid h^{T+1}) > 0$ implies $\nu(Z(f(y^{\vartheta}, b^{\vartheta})) \mid Z(y^{\vartheta}, A(y^{\vartheta} \to h^{T+1}))) > 0$. Construct a path $(v^{\vartheta}, d^{\vartheta}, \ldots, v^{T}, d^{T}, v^{T+1})$ as follows (see Figure 2):

$$v^{\vartheta} = y^{\vartheta},$$
$$d^{\vartheta} = c^{\vartheta},$$

for all $t = \vartheta + 1, \ldots, T$,

$$H(v^t) = H(w^t) \Rightarrow d^t = c^t$$
$$(\text{thus } \pi(d^t \mid H(v^t)) = \pi(c^t \mid H(w^t)) > 0),$$

otherwise choose any d^t such that $\pi(d^t \mid H(v^t)) > 0$.

By construction $P^{\pi}(v^{T+1} \mid f(v^{\vartheta}, d^{\vartheta})) = P^{\pi}(v^{T+1} \mid f(y^{\vartheta}, c^{\vartheta})) > 0$. By Claim A.1, $v^{T+1} \in h^{T+1}$. Since $\mu(y^{\vartheta} \mid h^{\vartheta}) > 0$ and $\mu(w^{\vartheta} \mid h^{\vartheta}) > 0$, $0 < \nu(Z(y^{\vartheta}) \mid Z(\{y^{\vartheta}, w^{\vartheta}\})) < 1$, Thus Remark 4 yields

$$\nu(Z(f(y^{\vartheta}, b^{\vartheta})) \mid Z(y^{\vartheta}, A(h^{\vartheta} \to h^{T+1}))) =$$
$$= \nu(Z(f(w^{\vartheta}, b^{\vartheta})) \mid Z(w^{\vartheta}, A(h^{\vartheta} \to h^{T+1})))$$

and

$$\begin{aligned}\nu(Z(f(y^{\vartheta},c^{\vartheta})) \mid Z(y^{\vartheta},A(h^{\vartheta} \to h^{T+1}))) &= \\ = \nu(Z(f(w^{\vartheta},c^{\vartheta})) \mid Z(w^{\vartheta},A(h^{\vartheta} \to h^{T+1}))) &> 0.\end{aligned}$$

(Corollary 1 allows us to write $A(y^{\vartheta} \to h^{T+1}) = A(h^{\vartheta} \to h^{T+1}) = A(w^{\vartheta} \to h^{T+1})$; for the inequality see Claim A.2) and we obtain

$$\begin{aligned}&\frac{\nu(Z(y^{T+1}) \mid Z(\{w^{T+1},y^{T+1}\}))}{\nu(Z(v^{T+1}) \mid Z(\{v^{T+1},y^{T+1}\}))} = \\ &= \frac{P^{\pi}(y^{T+1} \mid f(y^{\vartheta},b^{\vartheta}))\nu(Z(f(y^{\vartheta},b^{\vartheta})) \mid Z(y^{\vartheta},A(h^{\vartheta} \to h^{T+1})))}{P^{\pi}(v^{T+1} \mid f(y^{\vartheta},c^{\vartheta}))\nu(Z(f(y^{\vartheta},c^{\vartheta})) \mid Z(y^{\vartheta},A(h^{\vartheta} \to h^{T+1})))} = \\ &= \frac{P^{\pi}(y^{T+1} \mid f(y^{\vartheta},b^{\vartheta}))\nu(Z(f(w^{\vartheta},b^{\vartheta})) \mid Z(y^{\vartheta},A(h^{\vartheta} \to h^{T+1})))}{P^{\pi}(v^{T+1} \mid f(y^{\vartheta},c^{\vartheta}))\nu(Z(f(w^{\vartheta},c^{\vartheta})) \mid Z(y^{\vartheta},A(h^{\vartheta} \to h^{T+1})))}.\end{aligned} \tag{A.1'}$$

Since $v^{T+1} \in h^{T+1}$, (A.1') implies that $\mu(y^{T+1} \mid h^{T+1}) > 0$ only if $\nu\,(Z(f\,(y^{\vartheta}, b^{\vartheta})) \mid Z(y^{\vartheta},\, A(y^{\vartheta} \to h^{T+1}))) > 0$.

In order to prove that the "if" part of the claim holds when $Z(h^{T+1}) \subseteq Z(h^{\vartheta})$, we first show that $\mu(y^{\vartheta} \mid h^{\vartheta}) > 0$ implies $\nu(Z(v^{T+1}) \mid Z(\{w^{T+1},v^{T+1}\})) > 0$. (2.5) and $\mu(w^{\vartheta} \mid h^{\vartheta}) > 0$ yield

$$\frac{\nu(Z(f(y^{\vartheta},c^{\vartheta})) \mid Z(\{f(y^{\vartheta},c^{\vartheta}),f(w^{\vartheta},c^{\vartheta})\}))}{\nu(Z(f(w^{\vartheta},c^{\vartheta})) \mid Z(\{f(y^{\vartheta},c^{\vartheta}),f(w^{\vartheta},c^{\vartheta})\}))} = \frac{\mu(y^{\vartheta} \mid h^{\vartheta})}{\mu(w^{\vartheta} \mid h^{\vartheta})}$$

Since $\nu(Z(f(w^{\vartheta},c^{\vartheta})) \mid Z(w^{\vartheta},A(w^{\vartheta} \to h^{T+1}))) > 0$ and $P^{\pi}(w^{T+1} \mid f(w^{\vartheta},c^{\vartheta})) > 0$, we obtain

$$\begin{aligned}&\frac{\nu(Z(v^{T+1}) \mid Z(\{v^{T+1},w^{T+1}\}))}{\nu(Z(w^{T+1}) \mid Z(\{v^{T+1},w^{T+1}\}))} = \\ &= \frac{P^{\pi}(v^{T+1} \mid f(y^{\vartheta},c^{\vartheta}))\nu(Z(f(y^{\vartheta},c^{\vartheta})) \mid Z(\{f(y^{\vartheta},c^{\vartheta}),f(w^{\vartheta},c^{\vartheta})\}))}{P^{\pi}(w^{T+1} \mid f(w^{\vartheta},c^{\vartheta}))\nu(Z(f(w^{\vartheta},c^{\vartheta})) \mid Z(\{f(y^{\vartheta},c^{\vartheta}),f(w^{\vartheta},c^{\vartheta})\}))} = \\ &= \frac{P^{\pi}(v^{T+1} \mid f(y^{\vartheta},c^{\vartheta}))\mu(y^{\vartheta} \mid h^{\vartheta})}{P^{\pi}(w^{T+1} \mid f(w^{\vartheta},c^{\vartheta}))\mu(w^{\vartheta} \mid h^{\vartheta})}.\end{aligned} \tag{A.1''}$$

Since $P^{\pi}(v^{T+1} \mid f(y^{\vartheta}, c^{\vartheta})) > 0$, (A.1") shows that $\mu(y^{\vartheta} \mid h^{\vartheta}) > 0$ implies $\nu(Z(v^{T+1}) \mid Z(\{w^{T+1}, v^{T+1}\})) > 0$.

Assume that $Z(h^{T+1}) \subseteq Z(h^{\vartheta})$, i.e. every node in h^{T+1} is preceded by some node in h^{ϑ}. Let $\hat{w}^{T+1}$ be any node such that $\mu(\hat{w}^{T+1} \mid h^{T+1}) > 0$. By the "only if" part of the proof we know that the path $(\hat{w}^{\vartheta}, \hat{c}^{\vartheta}, \ldots, \hat{w}^{T}, \hat{c}^{T}, \hat{w}^{T+1})$ from h^{ϑ} to h^{T+1} has the same properties of path $(w^{\vartheta}, c^{\vartheta}, \ldots, w^{T}, c^{T}, w^{T+1})$ stated in Claim A.2 and we may assume without loss of generality that they are the same path. Thus $\mu(w^{T+1} \mid h^{T+1}) > 0$. Now assume that $\mu(y^{\vartheta} \mid h^{\vartheta}) > 0$, $P^{\pi}(y^{T+1} \mid f(y^{\vartheta}, b^{\vartheta})) > 0$ and $\nu(Z(f(y^{\vartheta}, b^{\vartheta})) \mid Z(y^{\vartheta}, A(y^{\vartheta} \rightarrow h^{T+1}))) > 0$. We have shown that $\mu(y^{\vartheta} \mid h^{\vartheta}) > 0$ implies $\nu(Z(v^{T+1}) \mid Z(\{w^{T+1}, v^{T+1}\})) > 0$, therefore $\mu(v^{T+1} \mid h^{T+1}) > 0$. Thus $\mu(y^{T+1} \mid h^{T+1}) > 0$ if and only if $\nu(Z(y^{T+1}) \mid Z(\{y^{T+1}, v^{T+1}\})) > 0$. But (A.1') shows that $\nu(Z(y^{T+1}) \mid Z(\{y^{T+1}, v^{T+1}\})) > 0$ if and only if $\nu(Z(f(y^{\vartheta}, b^{\vartheta})) \mid Z(y^{\vartheta}, A(h^{\vartheta} \rightarrow h^{T+1}))) > 0$ and $P^{\pi}(y^{T+1} \mid f(y^{\vartheta}, b^{\vartheta})) > 0$. Therefore, if $Z(h^{T+1}) \subseteq Z(h^{\vartheta})$, then $[\mu(y^{\vartheta} \mid h^{\vartheta}) > 0, \nu(Z(f(y^{\vartheta}, b^{\vartheta})) \mid Z(y^{\vartheta}, A(h^{\vartheta} \rightarrow h^{T+1}))) > 0$ and $P^{\pi}(y^{T+1} \mid f(y^{\vartheta}, b^{\vartheta})) > 0]$ implies $\mu(y^{T+1} \mid h^{T+1}) > 0$. This concludes the proof of Claim A.3. □

Let h be a (μ, π)-relevant information set ($h \in H(\mu, \pi)$). By definition there is some chain $(h_0, \ldots, h_\ell, \ldots, h_L)$ of information sets such that $h_0 \subseteq \Theta$, $h_{\ell+1}$ is a (μ, π)-follower of h_ℓ ($\ell = 0, \ldots, L-1$) and $h_L = h$. We say $(h_0, \ldots, h_\ell, \ldots, h_L)$ is a *(μ, π)-chain for h*. A (μ, π)-chain for h is *minimal* if there is no strictly shorter subchain which is also a (μ, π)-chain for h. In other words, $(h_0, \ldots, h_\ell, \ldots, h_L)$ is a minimal (μ, π)-chain for $h = h_L$ if for all $k = 0, \ldots, L-2$, all $\ell = k+2, \ldots, L$, h_ℓ is *not* a (μ, π)-follower of h_k.

The followings are obvious consequences of these definitions:

- if $(h_0, \ldots, h_\ell, \ldots, h_L)$ is a (minimal) (μ, π)-chain for h_L, then $(h_0, \ldots, h_\ell)$ is a (minimal) (μ, π)-chain for h_ℓ
- for every (μ, π)-relevant information set there is at least one minimal (μ, π)-chain.

Claim A.4. *Let h be a (μ, π)-relevant information set and let $(h_0, \ldots, h_\ell, \ldots, h_L)$ be a minimal (μ, π)-chain for h, then $Z(h_{\ell+1}) \subseteq Z(h_\ell)$ for $\ell = 0, \ldots, L-1$. Therefore for every (μ, π)-relevant h there is a (μ, π)-chain $(h_0, \ldots, h_{L-1}, h_L = h)$ such that $Z(h) \subseteq Z(h_{L-1})$ (i.e. every node in h is a follower of some node in h_{L-1}).*

Proof. . We prove the claim by induction on the length $L+1$ of the minimal (μ,π)–chains $(h_0,\ldots,h_\ell,\ldots,h_L)$. The claim is trivially true for the minimal chain of length 1. Assume it is true for all minimal chains of length $L \geq 1$ or less. Let $(h_0,\ldots,h_\ell,\ldots,h_{L-1},h_L)$ be a (μ,π)–chain for $h_L = h^{T+k}$ such that $(h_0,\ldots,h_\ell,\ldots,h_{L-1})$ is minimal for $h_{L-1} = h^{T+1}$. We show that if $Z(h^{T+k})$ is not included in $Z(h^{T+1})$ then $(h_0,\ldots,h_\ell,\ldots,h_L)$ is not minimal for h^{T+k}. This implies the thesis.

If $Z(h^{T+k})$ is not included in $Z(h^{T+1})$, there are $T+1$–stage predecessors of h^{T+k} in different information sets and Corollary 1 implies that $A(x^{T+1} \to h^{T+k}) = A(x^{T+1})$ for each x^{T+1} before h^{T+k}. Then Claim A.2 implies that there is a path $(y^{T+1},\ldots,y^{T+k})$ from h^{T+1} to h^{T+k} such that $\mu(y^{T+1} \mid h^{T+1}) > 0$ and $P^{\pi}(y^{T+k} \mid y^{T+1}) > 0$. Let $h_{L-2} = h^{\vartheta}$. Since $(h_0,\ldots,h_\ell,\ldots,h_{L-1})$ is minimal for $h_{L-1} = h^{T+1}$, the induction hypothesis implies that $Z(h^{T+1}) \subseteq Z(h^{\vartheta})$. Therefore there is a path $(y^{\vartheta}, b^{\vartheta}, \ldots, y^{T}, b^{T}, y^{T+1})$ from some node $y^{\vartheta} \in h^{\vartheta}$ to $y^{T+1} \in h^{T+1}$. Since $\mu(y^{T+1} \mid h^{T+1}) > 0$, Claim A.3 implies that $\mu(y^{\vartheta} \mid h^{\vartheta}) > 0$ and $P^{\pi}(y^{T+1} \mid f(y^{\vartheta}, b^{\vartheta})) > 0$. Therefore $(y^{\vartheta}, b^{\vartheta}, \ldots, y^{T+1}, b^{T+1}, \ldots, y^{T+k})$ is a path from $h^{\vartheta} = h_{L-2}$ to $h^{T+k} = h_L$ such that $\mu(y^{\vartheta} \mid h^{\vartheta}) > 0$and $P^{\pi}(y^{T+k} \mid f(y^{\vartheta}, b^{\vartheta})) > 0$, which implies that h_L is a (μ,π)–follower of h_{L-2} and $(h_0,\ldots,h_\ell,\ldots,h_L)$ is not minimal for $h_L = h^{T+k}$. This concludes the proof of Claim A.4. □

II. (*Construction of an equivalent assessment.*) Let $(\hat{\nu},\hat{\mu},\hat{\pi})$ be a sequentially rational and generally reasonable extended assessment of the extensive game Γ and consider the agent form of Γ. For each $h \in H(\hat{\mu},\hat{\pi})$ modify $u_{i,h}$ on $Z(h)$ $(u_{i,h} = u_i, h \in H_i)$ to obtain a payoff function $\hat{u}_{i,h}$ such that all actions (local strategies) a with $\hat{\pi}(a \mid h) > 0$ are payoff equivalent for agent (i,h) and all actions b with $\hat{\pi}(b \mid h) = 0$ are strictly dominated conditional on h. Let $\hat{\Gamma}$ denote the resulting modified game.

Now we construct a sequential equilibrium (μ^*,π^*) of $\hat{\Gamma}$. Let $\{\hat{\nu}^k\}$ be a sequence of strictly positive probability measures (corresponding to conditional systems) converging to $\hat{\nu}$. Define a corresponding sequence of perturbed games $\hat{\Gamma}^k$ with mixed strategy spaces restricted as follows ($\pi(\cdot \mid h)$ denotes a generic randomized strategy of agent $(\iota(h),h)$ in $\hat{\Gamma}^k$):

For all $h \in H(\hat{\mu},\hat{\pi})$, fix some $\hat{x}_h \in h$ with $\hat{\mu}(\hat{x}_h \mid h) > 0$,. Then, for all $a,b \in A(h)$, $\hat{\pi}(a \mid h) > 0$ implies $[\pi(a \mid h) \geq \hat{\pi}(a \mid h) - 1/k]$, and $\hat{\pi}(b \mid h) = 0$ implies $[\pi(b \mid h) \geq \hat{\nu}^k(Z(f(\hat{x}_h,b)))/\hat{\nu}^k(Z(\hat{x}_h))]$ (note that $\hat{\pi}(b|h) = 0$ implies $\lim_{k\to\infty} \hat{\nu}^k(Z(f(\hat{x}_h,b)))/\hat{\nu}^k(Z(f(\hat{x}_h))) = 0$, because $\hat{\pi}(b \mid h) = \hat{\nu}(Z(f(\hat{x}_h,b)) \mid \hat{\nu}(Z(\hat{x}_h))))$

For all $h \notin H(\hat{\mu}, \hat{\pi})$, for all $a \in A(h)$, $\pi(a \mid h) \geq 1/k$.

Let (ν^*, μ^*, π^*) be a cluster point of a sequence $\{(\nu^k, \mu^k, \pi^k)\}$, whereby π^k is a constrained equilibrium of $\hat{\Gamma}^k$ and (ν^k, μ^k) is derived from the strictly positive profile π^k. By construction, (μ^*, π^*) is fully consistent and (ν^*, μ^*, π^*) is generally reasonable. It is trivially true by construction that $\hat{\pi}$ and π^* coincide on the $(\hat{\mu}, \hat{\pi})$-relevant information sets:

$$\forall h \in H(\hat{\mu}, \hat{\pi}), \forall a' \in A(h), \hat{\pi}(a' \mid h) = \pi^*(a' \mid h). \tag{A.2}$$

The construction also implies that

$$\begin{aligned} &\forall h \in H(\hat{\mu}, \hat{\pi}), \forall a', a'' \in A(h): \\ &\quad \lim_{k \to \infty} [\pi^k(a' \mid h)/\pi^k(a'' \mid h)] = \lim_{k \to \infty} \frac{\hat{\nu}^k(Z(f(\hat{x}_h, a')))}{\hat{\nu}^k(Z(f(\hat{x}_h, a'')))}. \end{aligned} \tag{A.3}$$

((A.3) is obvious if either $\hat{\pi}(a' \mid h) > 0$ or $\hat{\pi}(a'' \mid h) > 0$; if for some $b \in h, \hat{\pi}(b \mid h) = 0$, then b is dominated in $\hat{\Gamma}$ and, in every constrained equilibrium π^k, $\pi^k(b \mid h) = [\hat{\nu}^k(Z(f(\hat{x}_h, b)))/\hat{\nu}^k(Z(\hat{x}_h))]$).

Remark 4 and (A.3) yield

$$\begin{aligned} &\forall h \in H(\hat{\mu}, \hat{\pi}), \forall \hat{A} \in 2^{A(h)} \backslash \{\emptyset\}, \forall a \in \hat{A}, \forall x \in h: \\ &\quad \hat{\mu}(x \mid h) > 0 \Rightarrow \nu^*(Z(f(x,a)) \mid Z(x, \hat{A})) = \\ &\quad = \lim_{k \to \infty} \frac{\hat{\nu}^k(Z(f(\hat{x}_h, a)))}{\hat{\nu}^k(Z(f(\hat{x}_h, \hat{A})))} = \hat{\nu}(Z(f(x,a)) \mid Z(x, \hat{A})). \end{aligned} \tag{A.4}$$

III. ((μ^*, π^*) *is a sequential equilibrium.*) It is easy to see that (μ^*, π^*) is a sequential equilibrium of the modified game $\hat{\Gamma}$. We must show that (μ^*, π^*) is also a sequential equilibrium of the original game Γ. For all non relevant h ($h \notin H(\hat{\mu}, \hat{\pi})$), the local mixed strategy $\pi^*(\cdot \mid h)$ is a best reply to (μ^*, π^*) conditional on h even in the original game Γ, because agent h has the same payoffs in both games Γ and $\hat{\Gamma}$. We must show that for all $(\hat{\mu}, \hat{\pi})$-relevant information sets h ($h \in H(\hat{\mu}, \hat{\pi})$), $\pi^*(\cdot \mid h)$ ($= \hat{\pi}(\cdot \mid h)$, see (A.2)) is a local best reply to (μ^*, π^*) even in the original game Γ. Since the best reply conditions for agent h depend only on $\mu^*(\cdot \mid h)$ and $\pi^*(\cdot \mid g)$ with g (μ^*, π^*)–follower of h, it is sufficient to show that $H(\hat{\mu}, \hat{\pi}) = H(\mu^*, \pi^*)$ and $\mu^*(\cdot \mid h) = \hat{\mu}(\cdot \mid h)$.

We show by induction on the stages that for all t–stage information sets $h^t, h^t \in H(\hat{\mu}, \hat{\pi})$ if and only if $h^t \in H(\mu^*, \pi^*)$, and $h^t \in H(\hat{\mu}, \hat{\pi})$ implies $\mu^*(\cdot \mid h^t) = \hat{\mu}(\cdot \mid h^t)$.

The statement is trivially true for the 0–stage information sets $h^0 \subseteq \Theta$. Assume that the statement is true for all the stages $t = 0, \ldots T$ and consider an information set h^{T+1}.

Let $h^{T+1} \in H(\hat{\mu}, \hat{\pi})$ and let $(h^0, \ldots, h^{\vartheta}, h^{T+1})$ be a $(\hat{\mu}, \hat{\pi})$–chain for h^{T+1}. By Remark A.1 there is a path $(x^{\vartheta}, a^{\vartheta}, \ldots, x^t, a^t, \ldots, x^T, a^T, x^{T+1})$ from h^{ϑ} to h^{T+1} such that $\hat{\mu}(x^{\vartheta} \mid h^{\vartheta}) > 0$, $\hat{\pi}(a^t \mid H(x^t)) > 0$ and $H(x^t) \in H(\hat{\mu}, \hat{\pi})$ for $t = \vartheta + 1, \ldots, T$ ($H(x^t)$ is a $(\hat{\mu}, \hat{\pi})$–follower of $h^{\vartheta} \in H(\hat{\mu}, \hat{\pi})$). The induction hypothesis and (A.2) imply that $\hat{\mu}(\cdot \mid h^{\vartheta}) = \mu^*(\cdot \mid h^{\vartheta})$, $\hat{\pi}(\cdot \mid H(x^t)) = \pi^*(\cdot \mid H(x^t))$ and $H(x^t) \in H(\mu^*, \pi^*)$ for $t = \vartheta, \ldots, T$. Therefore Remark A.1 yields $h^{T+1} \in H(\mu^*, \pi^*)$. If $h^{T+1} \in H(\mu^*, \pi^*)$ a symmetric argument yields $h^{T+1} \in H(\hat{\mu}, \hat{\pi})$. Therefore $h^{T+1} \in H(\hat{\mu}, \hat{\pi})$ if and only if $h^{T+1} \in H(\mu^*, \pi^*)$.

We must show that $h^{T+1} \in H(\hat{\mu}, \hat{\pi})$ implies $\mu^*(\cdot \mid h^{T+1}) = \hat{\mu}(\cdot \mid h^{T+1})$. Let $h^{T+1} \in H(\hat{\mu}, \hat{\pi})$. By Claim A.4 there is a $(\hat{\mu}, \hat{\pi})$–chain $(h^0, \ldots, h^{\vartheta}, h^{T+1})$ for h^{T+1} such that every node of h^{T+1} follows some node of h^{ϑ} (i.e. $Z(h^{T+1}) \subseteq Z(h^{\vartheta})$). Then Corollary 1 implies that for all $y^{\vartheta} \in h^{\vartheta}$, $A(y^{\vartheta} \to h^{T+1}) = A(h^{\vartheta} \to h^{T+1})$. Furthermore, Claim A.3 holds in the "if and only if" version. We first show that $\mu^*(\cdot \mid h^{T+1})$ and $\hat{\mu}(\cdot \mid h^{T+1})$ have the same support. Let $(y^{\vartheta}, b^{\vartheta}, \ldots, y^t, b^t, \ldots, y^T, b^T, y^{T+1})$ be a path from h^{ϑ} to h^{T+1}. By Claim A.3, $\mu^*(y^{T+1} \mid h^{T+1}) > 0$ if and only if $[\mu^*(y^{\vartheta} \mid h^{\vartheta}) > 0, \nu^*(Z(f(y^{\vartheta}, b^{\vartheta})) \mid Z(y^{\vartheta}, A(h^{\vartheta} \to h^{T+1}))) > 0$ and $\pi^*(b^t \mid H(y^t)) > 0$ for $t = \vartheta + 1, \ldots, T]$. Note that $H(y^t) \in H(\mu^*, \pi^*)$ for $t = \vartheta, \ldots, T$. By the induction hypothesis, (A.2), (A.4) and Claim A.3, $\mu^*(y^{T+1} \mid h^{T+1}) > 0$ if and only if $[\hat{\mu}(y^{\vartheta} \mid h^{\vartheta}) = \mu^*(y^{\vartheta} \mid h^{\vartheta}) > 0$, $\hat{\nu}(Z(f(y^{\vartheta}, b^{\vartheta})) \mid Z(y^{\vartheta}, A(h^{\vartheta} \to h^{T+1}))) = \nu^*(Z(f(y^{\vartheta}, b^{\vartheta})) \mid Z(y^{\vartheta}, A(h^{\vartheta} \to h^{T+1}))) > 0$ and $\hat{\pi}(b^t \mid H(y^t)) = \pi^*(b^t \mid H(y^t)) > 0$ for $t = \vartheta + 1, \ldots, T]$ if and only if $\hat{\mu}(y^{T+1} \mid h^{T+1}) > 0$.

Now it is sufficient to show that, for all paths $(x^{\vartheta}, a^{\vartheta}, \ldots, x^t, a^t, \ldots, x^T, a^T, x^{T+1})$, $(y^{\vartheta}, b^{\vartheta}, \ldots, y^t, b^t, \ldots, y^T, b^T, y^{T+1})$ from h^{ϑ} to h^{T+1} such that $\hat{\mu}(y^{T+1} \mid h^{T+1}) > 0$ and $\hat{\mu}(x^{T+1} \mid h^{T+1}) > 0$,

$$\frac{\mu^*(y^{T+1} \mid h^{T+1})}{\mu^*(x^{T+1} \mid h^{T+1})} = \frac{\hat{\mu}(y^{T+1} \mid h^{T+1})}{\hat{\mu}(x^{T+1} \mid h^{T+1})}$$

But this follows immediately from the induction hypothesis, (A.2), (A.4) and the following equations, which hold for $(\nu, \mu, \pi) \in \{(\hat{\nu}, \hat{\mu}, \hat{\pi}), (\nu^*, \mu^*, \pi^*)\}$:

$$\mu(y^{T+1} \mid h^{T+1})/\mu(x^{T+1} \mid h^{T+1}) = \frac{\nu(Z(y^{T+1}) \mid Z(\{x^{T+1}, y^{T+1}\}))}{\nu(Z(y^{T+1}) \mid Z(\{x^{T+1}, y^{T+1}\}))} =$$
$$= \frac{\mu(y^{\vartheta} \mid h^{\vartheta})\nu(Z(f(y^{\vartheta}, b^{\vartheta})) \mid Z(y^{\vartheta}, A(h^{\vartheta} \to h^{T+1})))P^{\pi}(y^{T+1} \mid f(y^{\vartheta}, b^{\vartheta}))}{\mu(x^{\vartheta} \mid h^{\vartheta})\nu(Z(f(y^{\vartheta}, c^{\vartheta})) \mid Z(y^{\vartheta}, A(h^{\vartheta} \to h^{T+1})))P^{\pi}(x^{T+1} \mid f(y^{\vartheta}, c^{\vartheta}))}.$$

References

[1] Battigalli, P.(1994), Structural Consistency and Strategic Independence in Extensive Games, *Ricerche Economiche*, 48, pp. 357–376.

[2] Battigalli, P. (1996a), Strategic Rationality Orderings and the Best Rationalization Principle, *Games and Economic Behavior*, in press.

[3] Battigalli, P.(1996b), Strategic Independence and Perfect Bayesian Equilibria, *Journal of Economic Theory*, in press.

[4] Battigalli, P. and Veronesi, P. (1996), A Note on Stochastic Independence without Savage–Null Events, *Journal of Economic Theory*, in press.

[5] Bonanno, G. (1992), Rational Beliefs in Extensive Games, *Theory and Decision*, 33, pp. 153–176.

[6] Fudenberg, D. and Levine, D. (1993), Self-Confirming Equilibrium, *Econometrica*, 61, pp. 523–45.

[7] Fudenberg, D. and Tirole, J. (1991), Perfect Bayesian equilibrium and Sequential Equilibrium, *Journal of Economic Theory* 53, pp. 236–260.

[8] Greenberg, J. (1990), *The Theory of Social Situations*. Cambridge U.K.: Cambridge University Press.

[9] Kohlberg, E. and Reny, P. (1992), "On the Rationale of Perfect Equilibrium", Mimeo.

[10] Kreps, D. and Wilson, R. (1982), Sequential Equilibria, *Econometrica*, 50, pp. 863–894.

[11] Kuhn, H.W. (1953), Extensive Games and the Problem of Information, in: Kuhn, H.W. and Tucker , A.W. (eds), *Contributions to the Theory of Games*, Vol. II. Princeton: Princeton University Press.

[12] McLennan, A. (1989a), The Space of Conditional Systems is a Ball, *International Journal of Game Theory*, **18**, pp. 125–139.

[13] McLennan, A. (1989b), Consistent Conditional Systems in Noncooperative Game Theory," *International Journal of Game Theory*, **18**, pp. 141-174.

[14] Milgrom, P. and Weber, R. (1985), Distributional Strategies for Games with Incomplete Information, *Mathematic of Operations Research*, 10, pp. 619–632.

[15] Myerson, R. (1986), Multistage Games with Communication, *Econometrica*, **54**, pp. 323–358.

[16] Myerson, R. (1991), *Game Theory. Analysis of Conflict.* Cambridge Ma: Harvard University Press.

[17] Pearce, D. (1984), Rationalizable Strategic Behavior and the Problem of Perfection, *Econometrica*, 52, pp. 1029–1050.

[18] Rênyi, A. (1955), On a New Axiomatic Theory of Probability, *Acta Mathematica Academiae Scientiarum Hungaricae*, 6, pp. 285–335.

[19] Reny, P. (1992), Backward Induction, Normal Form Perfection and Explicable Equilibria, *Econometrica*, 60, pp. 627–649.

[20] Selten, R. (1975), Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games, *International Journal of Game Theory*, 4, pp. 25–55.

[21] Swinkels, J. (1993), "Independence for Conditional Probability Systems", Mimeo, Kellogg Graduate School of Management, Northwestern University.

[22] Von Neumann, J. and Morgenstern, O. (1944), *Theory of Games and Economic Behavior.* Princeton: Princeton University Press.