

IMPLEMENTATION WITHOUT COMMITMENT IN MORAL HAZARD ENVIRONMENTS*

Bruno Salcedo[†]

The Pennsylvania State University
brunosalcedo.com
bruno@psu.edu

November 2012
Last revised: 3 · 4 · 2013

ABSTRACT I define a solution concept for strategic form games called interdependent-choice equilibrium. It is an extension of correlated equilibrium, in which the mediator is able to choose the timing of his signals and observe the actions taken by the players. The set of interdependent-choice equilibria is nonempty and is given by a finite set of affine inequalities. It characterizes all the outcomes that can be implemented in single shot interactions without the use of binding contracts or any other form of delegation. The results can also be interpreted as robust predictions for environments in which the rules of the game (e.g. order of play and information structure) are unknown.

KEYWORDS Interdependent choices · Solution concepts · Sequential implementation · Mediation · Robust predictions

JEL CLASSIFICATION C72 · D86

This paper is concerned with the possibility of interdependent choices, that is, some agents making choices that are a function of other agents' choices. It is well understood that interdependence is a powerful tool for generating incentives, for instance, in games with contracts, repeated games and sequential games. In such settings, each agent might be willing to choose some action not because of its direct consequences on payoffs, but because of the way other agents will react to his choice. However, trying to explicitly account for every source of interdependence might be fruitless in many instances. Instead, I propose a solution concept that *implicitly* incorporate all such possibilities.

The environment is *partially* described by a finite strategic form game with complete information. The description specifies the payoff-relevant actions and preferences, but says

*Previously circulated with the title “*Implementation through coordination in static environments without commitment*”. The nomenclature was modified to distinguish from coordination games.

[†]I wrote this paper under the invaluable guidance and supervision of Ed Green, and benefited from heated and amenable discussions with Bruno Sultanum. I wish to gratefully acknowledge the useful comments and suggestions from Kalyan Chatterjee, Nail Kashaev, Vijay Krishna and Bulat Gafarov, as well as Wiroy Shin and the attendants of the 1st Prospects in Economic Research Conference at the Pennsylvania State University. Needless to say, all remaining errors are my own.

nothing about the order of play nor the information structure. In particular, choices are not assumed to be independent nor simultaneous. Instead, I only impose two structural assumptions: (i) that the players cannot use binding contracts, delegate their authority over their choices, or use any other form of commitment; and (ii) that the players cannot use monetary transfers or any other payoff-relevant actions which are not already specified in the action sets. The purpose is to characterize the set of outcomes (distributions over payoff-relevant actions) that can be implemented as self-enforcing agreements.

Choice interdependence is particularly relevant in moral hazard environments in which there are no Pareto efficient Nash equilibria. In such situations, agents have incentives to try to extend the set of self-enforceable agreements. With complete information, moral hazard is completely eliminated when agents can write and enforce complete contracts (Coase' theorem), or when they interact repeatedly and are patient enough (folk theorems). This is possible because written contracts or publicly observed histories serve as coordination devices that allow for interdependence. However, my assumptions are specifically tailored to rule out such possibilities. I only consider single-shot interactions without any form of binding agreements, in order to show that the power of choice interdependence is orthogonal to the means by which it is achieved.¹

An illustrative example is provided in Nishihara (1997, 1999), showing that full cooperation in an n -player prisoner's dilemma can be implemented without contracts or repetition. The salient features of Nishihara's model are that players are uncertain about the order of choices, and the information structure allows players to recognize past defections and react to them. Section §1 shows how Nishihara's information structure could be generated via a mediator who helps to coordinate the agents' choices through sequential private recommendations. The rest of the paper extends the analysis to arbitrary environments through a solution concept which I call *interdependent-choice equilibrium*.

Mediated games can be traced back to correlated equilibrium (Aumann, 1974, 1987). A correlated equilibrium can be defined as a Nash equilibrium of a game with a mediator (she) who gives private correlated signals to the players. In Aumann's protocol, the signals are transmitted during a preplay stage, before the actual choices are made. Instead, I consider sequential protocols in which the mediator can choose the timing of choices and observe the actions taken by the agents. This enables the following kind of *sequentially mediated mechanisms*:

1. The game starts with the mediator privately choosing an order for the players and an action profile to be implemented, according to some known distribution.
2. The mediator then visits the players in the chosen order. In each visit, she recommends an action to the visited player and waits to see whether he complies.
3. As long as all players comply, the mediator recommends the actions in the chosen profile. However, if a player deviates, the mediator recommends to the remaining

¹There are different papers that study the effects of choice interdependence in single-shot interactions allowing for different degrees of commitment. Since the cited works exhibit great variety in their purpose and methodology, I defer the comparison with the current paper to section §6. In short, the salient features of my model are that both the timing and the information structure are design variables, and that I do not allow for side payments nor any form of commitment or delegation.

players to punish him.

There are settings with institutional or legal restrictions that only allow for pre-play mediation. My analysis is relevant for settings without such restrictions, in which the mediator may operate sequentially, as in Kissinger’s ‘shuttle diplomacy’ (Hoffman, 2011). One may be concerned that allowing the mediator to control the timing of *choices* might be a strong assumption. I follow this lax approach in order to capture the full extent of the power of choice interdependence in generating incentives. Of course, the model can easily be adapted to settings in which the mediator can only choose the timing of her *signals*, see section §7.1.

A mediated mechanism is incentive compatible if and only if following the mediator’s recommendations constitutes a Nash equilibrium. An interdependent-choice equilibrium is a distribution over action profiles induced by an incentive compatible mediated mechanism. The set of interdependent-choice equilibria is characterized by a finite set of affine inequalities and is thus a convex polytope that can be computed efficiently. Also, it is nonempty because it contains the set of correlated equilibria, and it is always contained in the set of individually rational outcomes.

Section §3 argues that interdependent-choice equilibrium captures the set of all outcomes (distributions over action profiles) that can be implemented without commitment. For that purpose, I define an *extensive form mechanism* to be *any* extensive form game that is consistent with the available information about the environment (payoff-relevant actions and preferences) and with the assumption that players are not able to commit. Theorem 1 states that an outcome can result from a Nash equilibrium of an extensive form mechanism if and only if it is an interdependent-choice equilibrium. In the language of Forges (1986) and Myerson (1986), Theorem 1 is a revelation principle showing that sequentially mediated mechanisms constitute a complete canonical class for implementation without commitment.

The definition of interdependent-choice equilibrium assumes that players will always be willing to punish deviations. However, this might not be ex-post optimal. After a deviation, the player who is supposed to execute the punishment might be better off also deviating. Section §4 analyzes the set of outcomes that can be induced taking this issue into account. I provide a sufficient condition for sequential implementation, as well as complete characterizations of sequential implementation in 2×2 environments and Perfect Bayesian implementation in arbitrary environments. The conditions simply restrict the set of *credible threats* that the mediator can recommend as punishments. The restrictions are rather permissive, since both the sequential structure and the off-path beliefs are design variables, refinements have a relatively small impact.

1 MOTIVATING EXAMPLE: A PRISONER’S DILEMMA

The model is motivated by the following example adapted from Nishihara (1997), which shows that cooperation can be implemented in a single-shot prisoner’s dilemma without repetition nor commitment. Suppose that two suspects of a crime are arrested. The DA has enough evidence to convict them of a misdemeanor but requires a written confession to convict them of the crime they allegedly committed. The DA then offers each prisoner

a sentence reduction in exchange for a confession. Each of the prisoners has to choose whether to behave cooperatively (C) by remaining silent or to defect (D) by confessing. Their preferences are summarized by the payoff matrix in Figure 1, where $B < b < g < G$.

	C	D
C	g, g	B, G
D	G, B	b, b

FIGURE 1 Payoff matrix for the prisoner’s dilemma

In the story told, there is no reason to assume that players will have to make a decision at exactly the same time. Also, even if the prisoners cannot directly communicate with each other, it is by no means clear that their choices need to be independent. Different forms of interdependence could either arise naturally or be artificially constructed. Even so, implementing cooperation remains a non-trivial task. The decision to confess cannot be delegated, and the legal obligation to confess prevents the enforcement of contracts that would bind them to remain silent. However, they could hire a lawyer who would schedule and be present in all the negotiations with the DA and instruct him as follows:

“If the DA offers us a (prisoner’s dilemma) deal you must always recommend that we do not confess, unless one of us has already confessed, in which case you must recommend that we do confess. Other than those recommendations, you must not provide us with any additional information.”




FIGURE 2 A sequential mechanism for the prisoner’s dilemma

The resulting situation can be described by the extensive form game in Figure 2. In the event that the first prisoner to move confesses, the second prisoner will be informed of this choice before making his own. If the first player decides to cooperate, the second prisoner

will remain uninformed about which of the two following events is true: (i) the event in which he is the first prisoner to receive the offer, and (ii) the event in which he is the second one and his accomplice remained silent.

Now consider the strategies represented by the red arrows. They support full cooperation and constitute an equilibrium as long as $G - g \leq g - b$. That is, as long as the benefit that a player can obtain from unilaterally deviating from (C,C) is less or equal to the inefficiency of both players confessing. This is possible because, along the equilibrium path, each prisoner assigns sufficient probability to the event in which: (i) if he cooperates, then his accomplice will remain uninformed and will also cooperate; and (ii) if he confesses, then his accomplice will learn of his defection and will punish him by also confessing.

2 ENVIRONMENT AND INTERDEPENDENT-CHOICE EQUILIBRIUM

The environment is described by $\mathcal{E} = (I, A, u)$. It represents a situation in which players in a finite set $I = \{1, 2, \dots, n\}$ are to make decisions, typical players are denoted by i . The main body of the paper assumes $n = 2$, section §7.2 discusses the general case. Each player i is to choose and perform one and only one action from a finite set $A_i = \{a_i, a'_i, \dots\}$. i 's preferences over action profiles are represented by $u_i : A \rightarrow \mathbb{R}$.

\mathcal{E} is assumed to be common knowledge. At this point, this is the only information assumed to be common knowledge. The description of the strategic environment is only a *partial* characterization of the problem at hand. It says nothing about the order in which choices will be made, nor about the information that each player will have at the moment of making his choice. Also, it does not specify whether players will be able to communicate or which sorts of randomization devices are available.

The salient aspect of choice interdependence is that, at the moment of choosing an action, each player i can believe that some of his opponents might base their decisions on his own. i 's beliefs on his opponents' behavior can be described by a *conjecture* $\lambda_i : A_i \rightarrow \Delta(A)$ which assigns a distribution $\lambda_i(\cdot | a_i) \equiv [\lambda_i(a_i)](\cdot) \in \Delta(A_{-i})$ to each action $a_i \in A_i$. $\lambda_i(a_{-i} | a_i)$ is i 's belief that his opponents will play a_{-i} if he decides to play a_i . Λ_i denotes the set of i 's conjectures. His preferences over his own actions are represented by the expected utility function $U_i : A_i \times \Lambda_i \rightarrow \mathbb{R}$, defined by:

$$(1) \quad U_i(a_i, \lambda_i) = \sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}) \cdot \lambda_i(a_{-i} | a_i)$$

Notice that a player weighs utilities with different weights depending on his choice. This is why it is possible to implement strictly dominated actions.

When a conjecture λ_i is constant with respect to i 's choice, it is isomorphic to a distribution over A_{-i} . Such conjectures are called degenerate and are identified with the corresponding distribution. For degenerate conjectures, equation (1) reduces to the standard definition of expected utility.

Before proceeding, it is convenient to define the notion of action (sub)spaces. An action space is a set of action profiles $A' \subseteq A$ that can be written as a Cartesian product $A' = \times_{i \in I} A'_i$ with $A'_i \subseteq A_i$ for every player i . Let $\mathcal{A} \subsetneq 2^A$ denote the set of action spaces. For each

action space $A' \in \mathcal{A}$, $\Lambda_i(A')$ denotes the set of i 's conjectures that assign full probability to his opponents choosing actions in A'_{-i} , regardless of i 's own choice.

2.1 INTERDEPENDENT-CHOICE EQUILIBRIUM

Interdependent-choice equilibrium is defined in terms of a simple class of extensive form games in which a trusted mediator manages the players through private recommendations. A (sequentially) mediated mechanism is a tuple (α, θ, B) . $\alpha \in \Delta(A)$ is a distribution over action profiles to be implemented. $\theta : A \rightarrow \Delta(I)$ specifies a distribution over the order in which players will move, conditional on the action profile to be implemented. $B \in \mathcal{A}$ is a set of *additional* credible threats that the mediator can recommend as punishments when some player fails to follow a recommendation. The adjective “additional” is important. It means that they can be recommended as punishments regardless of whether they are played along the equilibrium path. The mediator is always allowed to punish deviations by recommending any action which is also recommended along the equilibrium path with positive probability. The effective set of credible threats is thus $B^* = B \cup \left(\times_i \text{supp } \alpha_i \right)$.²

A mediated mechanism represents the following extensive form game. The game begins with the mediator choosing the action profile a^* that she wants to implement, and the player i to move first. She chooses a^* according to α and i according to θ_{a^*} . Then, the mediator “visits” each of the players one by one, visiting i first and $-i$ second. In each visit, she will recommend an action and observe the action actually taken. Let a^r denote the vector of actions recommended by the mediator, and a^p the vector of actions chosen by the players. At the moment of making their choices, the players do not possess any information other than the recommendation they receive. The mediator always recommends the intended action to the first player, i.e. $a_i^r = a_i^*$. She recommends the intended action to the second player if the first player complied, and one of the worst available punishments in B_{-i}^* otherwise. That is:

$$\begin{aligned} a_{-i}^r &= a_{-i}^* && \text{if } a_i^p = a_i^r \\ a_{-i}^r &\in \arg \min_{a_{-i} \in B_{-i}^*} u_i(a_i^p, a_{-i}) && \text{if } a_i^p \neq a_i^r \end{aligned}$$

A mediated mechanism is incentive compatible if and only if following the mediator’s recommendations constitutes a Nash equilibrium. Since only Nash incentive compatibility is required, there are no incentive constraints for the punishments, which occur off the equilibrium path. This is why the mediator can always recommend the worst available punishment without any consideration of the player actually performing it.

Nash incentive compatibility is characterized by a finite set of affine inequalities. Namely, for every action a_i that is recommended with positive probability, i 's expected utility from complying should be greater or equal to the expected utility from deviating to an alternative action a'_i . The expected utility from complying is given by $\sum_{a_{-i}} \alpha(a_{-i}|a_i) u_i(a)$. The

²The only purpose of restricting the set of credible threats, is to ensure that the punishments are incentive compatible, see section §4. Actions that are chosen along the equilibrium can always be used as punishments because the player who performs the punishment will remain uninformed of the deviation and believe that he is along the equilibrium path.

expected utility from deviating equals $\sum_{a_{-i}} \alpha(a_{-i}|a_i) u_i(a'_i, a_{-i})$ if i is the second player to move (in which case it is too late to punish him). Otherwise, it is given by $\underline{u}_i(a'_i, B^*) \equiv \min_{a_{-i} \in B^*} u_i(a'_i, a_{-i})$. The set of interdependent-choice equilibria can thus be defined as follows.

DEFINITION 1 A distribution over action profiles $\alpha \in \Delta(A)$ is an *interdependent-choice equilibrium* with respect to a set of credible threats $B \in \mathcal{A}$, if and only if there exists some $\theta : A \rightarrow \Delta(I)$ such that (α, θ, B) is incentive compatible, i.e. such that for every player $i \in I$ and every pair of actions $a_i, a'_i \in A_i$:

$$\sum_{a_{-i} \in A_{-i}} \alpha(a) \left(u_i(a) - \left(1 - \theta_a(i) \right) u_i(a'_i, a_{-i}) - \theta_a(i) \underline{u}_i(a'_i, B^*) \right) \geq 0$$

$\text{IE}(B)$ denotes the set of interdependent-choice equilibria with respect to B . When $B = A$, I omit the reference to it and simply say that $\alpha \in \text{IE}$ is an interdependent-choice equilibrium.

The inequalities that define interdependent-choice equilibria resemble those which define other solution concepts involving choice interdependence. By setting $\theta_a(i) = \delta \in (0, 1)$ for all i and a , and imposing some restrictions on \underline{u} , one recovers the recursive characterization of payoff corresponding to SPNE of repeated games due to Abreu et al. (1990), hereafter APS. If one sets $\theta_a(i) = 0$ for all i and a , then players cannot readjust their intended plans in order to punish deviations. This case results in the definition of correlated equilibrium. In the opposite extreme, if $\theta_a(i) = 1$ for all i and a , then players are *always* able to punish deviations. This case results in interim individual rationality: each player's expected payoff conditional on his own action should be greater or equal to his minimax value. It also coincides with the notion of conjectural variations equilibrium adapted to finite games and deprived of any consistency restrictions (Figueiras et al., 2004).

From the previous analysis, it follows that the inequalities defining interdependent-choice equilibrium are tighter than those of individual rationality, and weaker than those of correlated equilibrium. Hence the set of interdependent-choice equilibria is always contained in the set of individually rational outcomes, and always contains the set of correlated equilibria (and is thus never empty). Also, it is straightforward to show that IE is monotone with respect to the set of credible threats. The following example, adapted from Aumann (1987), illustrates that the containments between these different solution concepts can be strict.

Example 1 Two partners decide whether to work (W) or shirk (S) in a joint-venture, their payoffs are depicted in Figure 3. The figure shows the sets of payoffs corresponding to individual rationality, Nash equilibrium with public randomization, correlated equilibrium, and interdependent-choice equilibrium. In this particular example, the set of SPNE payoffs of the repeated game with perfect monitoring and public randomization coincides either with the convex hull of the set of Nash equilibrium payoffs or with the set of individually rational payoffs, depending on the discount factor. Also, all the Pareto efficient outcomes correspond to interdependent-choice equilibria, and all but one interdependent-choice equilibria are sequential equilibria of the corresponding mediated game.

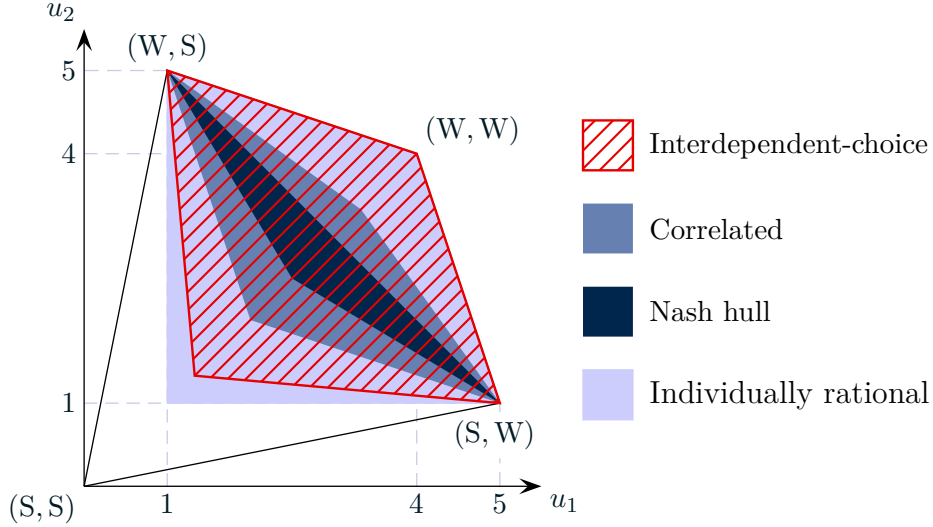


FIGURE 3 Equilibrium payoffs for the chicken game

The existence statement about θ makes Definition 1 appear complicated. However, interdependent-choice equilibrium can also be defined in terms of joint distributions $\gamma \in \Delta(A \times I)$ satisfying the inequalities:

$$\sum_{a_{-i} \in A_{-i}} \left[\gamma(a) u_i(a) - \gamma(a, i) \underline{u}_i(a'_i, B^*) - \gamma(a, -i) u_i(a'_i, a_{-i}) \right] \geq 0$$

This implies that, knowing \underline{u} , the set of interdependent-choice equilibria with respect to B is characterized by a finite set of affine inequalities. This makes it very easy to compute. For small games it can be computed by hand and for moderate games it can be accurately approximated by a computer in fractions of a second. An additional difficulty may arise when $B \neq A$, because \underline{u} depends on B^* which in turn may depend on the support of the marginal of γ over A . This difficulty disappears when this support is guaranteed to be a subset of B .

3 NASH IMPLEMENTATION WITHOUT COMMITMENT

Interdependent-choice equilibria are equilibrium outcomes of mediated games. They are implementable whenever the players can hire a mediator who can control the order of choices, observe the choices made, and generate private recommendations. A natural question is whether there are other mechanisms which can implement additional outcomes and are still consistent with our assumptions. The answer is no. Every equilibrium outcome of a single-shot interaction without commitment or side payments, corresponds to an interdependent-choice equilibrium (the precise meaning of ‘no commitment’ is formalized in the definition of *extensive form mechanisms* in section §3.2).

3.1 EXTENSIVE FORM GAMES

Extensive form games are defined as in Osborne and Rubinstein (1994), with some differences in notation. An extensive form game is a tuple $\mathcal{G} = (M, X, \iota, \mathcal{H}, \sigma_0^*, v)$. M denotes a set of moves. $X \subseteq \cup_{t \in \mathbb{N}} M^t$ denotes a finite set of histories or nodes. \leq denotes precedence among nodes. $M(x) = \{m \in M \mid (x, m) \in X\}$ is the set of moves available at x . $\iota(x) \in I \cup \{0\}$ is the agent moving at x , where 0 represents Nature (or a mediator). Z and Y_i are the sets of terminal nodes and i 's decision nodes respectively. $Z(x) = \{z \in Z \mid x \leq z\}$ is the set of terminal nodes that can be reached after x . \mathcal{H}_i partitions i 's decision nodes into information sets and satisfies perfect recall. σ_0^* specifies the players' common prior beliefs about Nature's choices. Finally, $v_i : Z \rightarrow \mathbb{R}$ represents i 's preferences over terminal nodes. Notice that attention is restricted to *finite* games with *perfect recall*.

A pure strategy for i is a function $s_i : \mathcal{H}_i \rightarrow M$, with $s_i(H_i) \in M(H_i)$ for every $H_i \in \mathcal{H}_i$. A mixed strategy for i is a distribution σ_i over his pure strategies, it is strictly mixed if it has full support. S_i , Σ_i and Σ_i^+ denote the sets of i 's pure, mixed and strictly mixed strategies respectively. Given that Nature chooses according to σ_0^* , a strategy profile σ induces a distribution over nodes $\zeta(\cdot \mid \sigma, \sigma_0^*) \in \Delta(X)$. $\zeta(x \mid \sigma, \sigma_0^*)$ is the probability that the game will reach x if players choose according to σ and Nature chooses according to σ_0^* . When there is no ambiguity, I omit the reference to σ and σ_0^* and simply write $\zeta(x)$. Expected payoffs $V_i : \Sigma \rightarrow \mathbb{R}$ are defined in the obvious way. A Nash equilibrium (NE) is a strategy profile σ^* such that $V_i(\sigma^*) \geq V_i(\sigma'_i, \sigma_{-i}^*)$ for every i and σ' .

3.2 EXTENSIVE FORM MECHANISMS

Loosely speaking, an extensive form mechanism is *any* extensive form game which is consistent with the partial characterization of the environment (players, action spaces and preferences) and with the assumption that players cannot commit nor delegate the authority over their choices. The first requirement for an extensive form game to be an extensive form mechanism is that it must preserve the outcome and preference structure of the environment. That is, there must be a preference-preserving map from terminal nodes (outcomes of the game) to action profiles (outcomes of the environment).

DEFINITION 2 An *outcome homomorphism* from an extensive form game \mathcal{G} to \mathcal{E} is a function τ from terminal nodes **onto** action profiles preserving preferences, i.e. such that $v(z) = u(\tau(z))$ for every terminal node z . \mathcal{G} is *outcome equivalent* to \mathcal{E} if it admits an outcome homomorphism.

Outcome equivalence does not capture the assumption that players cannot commit. For this purpose, each player should *freely* choose his own action at some point in the game. To eliminate commitment, there must exist a relationship that relates moves (choices in the game) with actions (choices in the environment), preserving the choice structure in a precise sense discussed ahead. For the remainder of this section, let \mathcal{G} be outcome equivalent to \mathcal{E} . The subsequent definitions all depend on a fixed outcome homeomorphism τ . To keep the notation simple, I omit explicit references to this dependence. Also, for every player i , $\tau_i : Z \rightarrow A_i$ denotes the projection of τ onto A_i .

For every player i and every corresponding decision node y_i , τ induces a *representation* relationship \approx_{y_i} from the set of moves available at y_i in the game to the set of i 's actions in the environment. A move $m \in M(y_i)$ represents action $a_i \in A_i$ at y_i , if and only if choosing such move at node y_i in the game has the same effect in (payoff-relevant) outcomes as choosing a_i in the environment. This idea is formalized by the following definition.

DEFINITION 3 Given a player $i \in I$ and a decision node $y_i \in Y_i$, a move $m \in M(y_i)$ *represents* an action $a_i \in A_i$ at y_i if and only if:

1. $\tau_i(z) = a_i$ for every $z \in Z(y_i, m)$
2. There exist $m' \in M(y_i)$ and $z \in Z(y_i, m')$ such that $\tau_i(z) \neq a_i$

The representation relationship is denoted by $m \approx_{y_i} a_i$, and $M^{a_i}(y_i)$ denotes the set of moves that represent a_i at y_i . A move is *pivotal* at y_i if it represents some action.

The first requirement is that, if i chooses m at y_i , then the game will end at a terminal node which is equivalent to a_i according to τ_i . This is regardless of any previous or future moves by either i or his opponents. The second requirement is that, after the game reaches y_i , i could still choose a different move m' after which the game remains open to the possibility of ending at a terminal node that is *not* equivalent to a_i .

In the strategic environment, player i must choose one and only one action in A_i . A decision node y_i is pivotal for i according to τ if and only if it represents an analogous choice problem. That is, if and only if for every $a_i \in A_i$ there is a move available at y_i which represents a_i .

DEFINITION 4 A decision node $y_i \in Y_i$ is *pivotal* for player $i \in I$ with respect to τ if and only if $M^{a_i}(y_i) \neq \emptyset$ for every $a_i \in A_i$. Let $D_i \subseteq Y_i$ denote the set of pivotal nodes for i .

It seems natural to require that players should always know whether their moves represent an action. (\mathcal{G}, τ) satisfies full disclosure of consequences if and only if $\approx_{y_i} = \approx_{y'_i}$ whenever y_i and y'_i belong to the same information set.³ When (\mathcal{G}, τ) satisfies full disclosure of consequences, the previous definitions can be extended to talk about information sets instead of nodes. In particular one can say that m represents a_i at H_i , or that H_i is pivotal, and $M^{a_i}(H_i)$ can be defined in the obvious way. A extensive form mechanism for \mathcal{E} can now be defined as an extensive form game \mathcal{G} that admits an outcome homomorphism (outcome equivalence), satisfying full disclosure of consequences, and such that the payoff-relevant components of every terminal node are determined at pivotal nodes (strategic equivalence).

DEFINITION 5 A *extensive form mechanism* for \mathcal{E} is a tuple (\mathcal{G}, τ) consisting of an extensive form game \mathcal{G} and an outcome homomorphism τ that satisfies full disclosure of consequences and such that for every terminal node z and every player i , there exists a pivotal node $y_i \in D_i$ and a pivotal move $m \in M^{\tau_i(z)}$ such that $z \in Z(y_i, m)$.

³ $\approx_{y_i} = \approx_{y'_i}$ means that $m \approx_{y_i} a_i$ if and only if $m \approx_{y'_i} a_i$ for all $m \in M(y_i) = M(y'_i)$ and all $a_i \in A_i$.

The notion of extensive form mechanisms attempts to capture the *most general* class of extensive form games that do not allow for any form of commitment. The definition might appear to be overly complicated or unnatural. For example, one could define a pivotal node for player i simply by requiring that $M(y_i) = A_i$. However, this would exclude possibilities which, ex ante, could have interesting consequences.

Consider for instance a generic 2×2 environment with $A_i = \{a_i, a'_i\}$ for $i \in \{1, 2\}$, and the two mechanisms illustrated in Figure 4. In the first one, player 1 has the option of either making a definitive decision at the beginning of the game, or waiting to see his opponent's choice before making his own (endogenous timing). In the second one, player 1 has the option of partially revealing some information about his choice to player 2 (endogenous signaling). Our main results imply that these features are irrelevant, in that the set of implementable outcomes remains unchanged with or without them. However, this is a result, not an assumption.




FIGURE 4 Valid extensive form mechanisms

3.3 NASH IMPLEMENTATION

A distribution over action profiles $\alpha \in \Delta(A)$ is (Nash, sequentially, ...) implementable if and only if there exist an extensive form mechanism (\mathcal{G}, τ) and a (Nash, sequential, ...) equilibrium $\sigma^* \in \Sigma$ that implements it, i.e. such that for every $a \in A$:

$$\alpha(a) = \zeta^* \left(\tau^{-1}(a) \right) = \sum_{z \in Z} \zeta(z, \sigma^*, \sigma_0^*) \cdot \mathbf{1}(\tau(x) = a)$$

The following theorem asserts that the set of mediated mechanisms constitutes a complete canonical class for Nash implementation.

THEOREM 1 *A distribution over action profiles is Nash implementable if and only if it is an interdependent-choice equilibrium.*

This result is analogous to the central results in Aumann (1987) and other related papers (Bergemann and Morris, 2011a, Forges, 1986, 1993, Myerson, 1986). All of these results follow from similar arguments: given a general mechanism and an equilibrium, it is possible to construct a canonical mechanism that replicates the strategic features but gives players

the minimal amount of information needed to follow the plan of action. By reducing their information, each player's ability to deviate profitably also diminishes.

For interdependent-choice equilibrium, the proof requires two additional considerations. The first one is that in mediated mechanisms players always use the worst available punishments off the equilibrium path. This implies that the distribution of actions played can be replicated only *along the equilibrium path*. However, this is not an issue because using the *worst* available punishments only relaxes the incentive compatibility constraints.

A more complicated issue is that, in order to keep the definition general, I allowed for mechanisms that cannot be transformed into canonical mechanisms in a straightforward manner. For instance, players can delay their choices or partially reveal or acquire information from past play. This is not a significant issue. Intuitively, all the non-pivotal choices can be delegated to the mediator in order to obtain more tractable mechanisms. For instance, to deal with the mechanisms presented in Figure 4, one could consider the simpler mechanisms shown in Figure 5. While this turns out to be insignificant for the result, it does complicate the notation required for the proof.




FIGURE 5 Simple mechanisms corresponding to the general mechanisms from figure 4

4 CREDIBLE THREATS AND THE PROBLEM OF PERFECTION

The goal was to find the set of outcomes that can be implemented without commitment, and yet interdependent-choice equilibrium is defined *as if* the players were able to commit themselves to punish deviations. In particular, one might be concerned with the willingness of a player to carry on a punishment when, ex-post, it might be in his best interest not to do so. This may occur because interdependent-choice equilibrium is defined in terms of Nash incentive compatibility, which only imposes optimality requirements along the equilibrium path. This section addresses this issue by focusing on outcomes that can be implemented as sequential equilibria of extensive form mechanisms.⁴

⁴A very different justification for the commitment to punish can be found in the context of psychological games. For instance, one may assume that players actually derive utility from punishing those who deviated in the past (Dufwenberg and Kirchsteiger, 2004, Vyrastekova and Funaki, 2010).

4.1 A SUFFICIENT CONDITION FOR SEQUENTIAL IMPLEMENTATION

The sufficient condition for sequential implementation provided here is based three observations. First suppose that α is a interdependent-choice equilibrium with respect to $B = \emptyset$ (as a set of additional credible threats). In the corresponding mediated game, all the actions which are recommended as punishments are also recommended along the equilibrium path. This implies that every information set is reached with positive probability. Consequently, following the mediator recommendations constitutes a sequential equilibrium and α is sequentially implementable.

Secondly, given two extensive form mechanisms, one can construct a new mixed mechanism with a first stage in which nature randomizes between them. If the outcome of the randomization is publicly announced, then the incentives remain unchanged. This implies that there is some sequentially implementable distribution α^S with full support, meaning that $\alpha_i^S(a_i) > 0$ for every $a_i \in A_i^S$, where A^S denotes the action space consisting of actions that are played with positive probability in some sequentially implementable distribution. Also, together with the first observation, this implies that every interdependent-choice equilibrium with respect to A^S can be arbitrarily approximated by sequentially implementable distributions.

The two previous observations are true not just for sequential equilibrium, but also for finer refinements which only impose additional restrictions off the equilibrium path, such as perfect and proper equilibria. Specializing the analysis to sequential equilibrium allows for less stringent sufficient conditions. Let A^R be the set of first-order rationalizable actions and recall that sequential equilibrium does not impose any restrictions on the relative likelihood of different deviations. This suggests that any action in A^R can be used as a credible threat. For instance, if a_1 is a best response to a_2 , then one can construct a mediated game in which, whenever 1 is asked to play a_1 , he will believe that it is because 2 deviated by choosing a_2 .⁵

The preceding discussion suggests sufficient conditions for sequential implementation, if $\alpha \in \text{IE}(\emptyset)$, $\alpha \in \text{IE}(A^R)$ or $\alpha \in \text{IE}(A^S)$, then α is (almost) sequentially implementable. However, these conditions are unsatisfactory because they may become intractable in computational terms. On one hand, $\text{IE}(\emptyset)$ and $\alpha \in \text{IE}(A^R)$ can be hard to compute because the set of effective threats may depend on the support of the equilibrium distribution and, consequently, the corresponding inequalities are no longer linear nor continuous. On the other hand, there is no obvious way of finding A^S . Instead, I propose a recursive procedure, much in the spirit of APS, to find a set $A^{\text{IE}} \subseteq A^S$ with the property that every interdependent-choice equilibrium with respect to A^{IE} is sequentially implementable.

Define $T : \mathcal{A} \rightarrow \mathcal{A}$ as $T(A') = \times_{i \in I} T_i(A')$, where:

$$T_i(A') = \left\{ a_i \in A_i \mid \exists \alpha \in \text{IE}(A' \cup A^R) \text{ such that } \alpha_i(a_i) > 0 \right\}$$

⁵In other words, since off-path beliefs are design variables, then can be chosen as to generate ex-post incentives for punishment. Consider for instance the following example from *The Godfather* (1972), dealing with the relationships inside the Italian mafia. After offering a truce agreement, Vito Corleone remains concerned about the wellbeing of his son and makes the following announcement: “*I am a superstitious man, and if some unfortunate accident should befall him... if he’s struck by a bolt of lightning, then I’m going to blame some of the people in this room.*” He is announcing what his interim beliefs would be after observing an hypothetical unexpected event. By doing so, any best response to such beliefs becomes a credible threat.

In words, $T(A')$ is the action space which includes only those actions which can be implemented in some extensive form mechanism with respect to $A' \cup A^R$. Now define the sequence $\{A^n\}_{n \in \mathbb{N}}$ by $A^1 = A$ and $A^{n+1} = T(A^n)$. The following proposition ensures that A^{n+1} converges in finite time to the largest fixed point of T .

PROPOSITION 2 *A^n is a decreasing sequence converging in finite time to a nonempty limit A^{IE} , such that $A^{\text{IE}} = T(A^{\text{IE}})$ and $A' \subseteq A^{\text{IE}}$ whenever $A' \subseteq T(A')$.*

Since $\text{IE}(A^{\text{IE}} \cup A^R)$ is convex and $A^{\text{IE}} = T(A^{\text{IE}})$, it follows that there is a maximum-support interdependent-choice equilibrium $\alpha^{\text{IE}} \in \text{IE}(A^{\text{IE}})$ such that $\text{supp}(\alpha_i^{\text{IE}}) = A_i^{\text{IE}}$ for every player i . All the information sets are reached with positive probability in the mediated game corresponding to α^{IE} . Therefore, following the mediator's recommendations constitutes a sequential equilibrium. This means that α^{IE} is sequentially implementable, and consequently, $A^{\text{IE}} \subseteq A^S$. Hence the following theorem obtains, which provides a tractable sufficient condition for sequential implementation.⁶

THEOREM 3 *If α is an interdependent-choice equilibrium with respect to $A^{\text{IE}} \cup A^R$, then it is sequentially implementable.*

This sufficient condition for sequential implementation is very tractable. The fact that the iterative procedure is monotone is important. It means that in each state, if $\alpha \in \text{IE}(A^n \cup A^R)$ then $\text{supp}(\alpha) \subseteq A^n$. Hence, the set of effective threats coincides with the set of additional threats, and $\text{IE}(A^n \cup A^R)$ is characterized by a finite set of affine inequalities. With this in mind, finding A^{IE} is much simpler in computational terms than the APS algorithm. First, only pure actions are eliminated at each iteration. The elements of the sequence can thus be described with finite information. Also, $T(A^n)$ is defined by a *decreasing* number of affine inequalities. In each iteration, the inequalities corresponding to an eliminated action a_i may be replaced with $\alpha_i(a_i) = 0$. Finally, the convergence of the current procedure occurs in finite time and, since interdependent-choice equilibria is a permissive solution concept, the number of required iterations in most cases should be small (if positive at all).

It is worth to note that the condition is also very permissive. Since both the sequential structure and the off-path beliefs are design variables, restricting attention to sequential equilibria has a relatively small impact. For example, in environments with no strict dominance, $A^{\text{IE}} \cup A^R = A^R = A$, and thus Nash and sequential implementability coincide, hence the following corollary.

COROLLARY 4 *When there are no strictly dominated actions, a distribution is sequentially implementable if and only if it is an interdependent-choice equilibrium.*

⁶For the theorem to be true, there must be at least three players, or it must be feasible for Nature or the mediator to make mistakes (null choices). Otherwise, instead of having that α is sequentially implementable, one can only guarantee that for every $\varepsilon > 0$, there exists a sequentially implementable distribution α' such that $\|\alpha - \alpha'\| < \varepsilon$. The reason for this is made apparent in the last step of the proof.

	L	C	D	R
T	3, 0	0, k	0, 0	0, 3
C	k , 0	6, 6	2, 9	k , 0
D	0, 0	9, 2	5, 5	0, 0
B	0, 3	0, k	0, 0	3, 0

FIGURE 6 Payoff matrix for examples 2 and 3.

Example 2 Consider the environment described by the payoff matrix in Figure 6 with $k = 2$. The central part of the matrix corresponds to a prisoner’s dilemma in which cooperation cannot be achieved through mediation. Adding the additional actions, allows cooperation to be implemented by using T and R as punishments. Furthermore, following recommendations is a sequential equilibrium, as long as the players consider B and L to be the most likely trembles. However, T and R cannot be played with positive probability in any equilibrium. Hence $A^{\text{IE}} = \{(D, D)\} \neq A^{\text{S}}$, and (C, C) is sequentially implementable despite the fact that it is not an interdependent-choice equilibrium with respect to A^{S} .

It is possible to construct more complicated examples which admit distributions that can be properly implementable only if, off the equilibrium path, some agents choose actions outside of $A^{\text{S}} \cup A^{\text{R}}$. This implies that the condition in Theorem 3 is not sufficient. I can show that there exists a set of credible threats B^{S} such that the set of sequentially implementable outcomes coincides with $\text{IE}(B^{\text{S}})$. However, the conditions that define B^{S} are intractable. For proper and perfect implementation it is no longer enough to specify a set of credible threats. The enforceable punishments might actually depend on the equilibrium actions and the specific deviations being punished. In any case, providing conditions that are both sufficient and necessary for arbitrary refinements and arbitrary environments is a complicated task. Instead, the remainder of this section focuses on 2×2 environments.

4.2 C-RATIONALIZABILITY AND 2 BY 2 ENVIRONMENTS

A natural restriction on the set of credible threats is that every punishment should be rationalizable, i.e. it should be a best response to some rational belief of the player performing the punishment (Bernheim, 1984, Pearce, 1984). However, when choice interdependence is possible, the relevant notion of rationalizability is not with respect to simple beliefs but instead with respect to conjectures.

DEFINITION 6 [Conjectural rationalizability]

- An action $a_i^* \in A_i$ is C-rationalizable with respect to $A' \in \mathcal{A}$ if and only if there exists some conjecture $\lambda_i \in \Lambda_i(A')$ for which $a_i^* \in \arg \max_{a_i \in A_i} U_i(a_i, \lambda_i)$. Let $\text{CR}_i(A')$ denote the set of i ’s actions that are C-rationalizable with respect to A'
- An action space $A' \in \mathcal{A}$ is self C-rationalizable if and only if every action profile in A' consists of actions that are C-rationalizable with respect to A' , i.e. if $A' \subseteq \text{CR}(A')$.

- The space of C-rationalizable action profiles $A^{\text{CR}} \in \mathcal{A}$ is the largest self C-rationalizable action space.

A^{CR} is guaranteed to exist because $\text{CR}(\cdot)$ is \subseteq -monotone. Consequently, the union of all self C-rationalizable sets is also self C-rationalizable. It is nonempty because it always contains the set of rationalizable action profiles. C-rationalizability is analogous in many ways to the notion of rationalizability for environments with independent choices. A^{CR} can be found in a tractable way using the notion of absolute dominance, which is analogous to the standard notion of strict dominance.⁷

DEFINITION 7 Given two actions $a_i, a'_i \in A_i$, a_i *absolutely dominates* a'_i with respect to $A' \in \mathcal{A}$ if and only if $\max_{a_{-i} \in A'_{-i}} u_i(a'_i, a_{-i}) < \min_{a_{-i} \in A'_{-i}} u_i(a_i, a_{-i})$.

In other words, a_i is absolutely dominated by a'_i if and only if the best possible payoff from playing a_i is strictly worse than the worst possible payoff from playing a'_i . Absolute dominance is much simpler than strict dominance in computational terms because a player can conjecture different reactions for each alternative action, and hence mixed actions need not be considered. Since ex-post minimization can be achieved, there is no need for interim minimization. In order to find the set of absolutely dominated actions one simply has to find the *pure-action* minimax payoff $\underline{u}_i = \max_{a_i \in A_i} \min_{a_{-i} \in A'_{-i}} u_i(a_i, a_{-i})$, and then eliminate those actions a_i such that $\max_{a_{-i}} u_i(a_i, a_{-i}) < \underline{u}_i$. The following proposition ensures that $\text{CR}(A')$ can be obtained by eliminating absolutely dominated actions, and A^{CR} can be found by repeating this process iteratively.

PROPOSITION 5 *An action is C-rationalizable with respect to A' if and only if it is not absolutely dominated in A' , and the iterated removal of all absolutely dominated actions is order independent and converges to A^{CR} .*

In 2×2 environments without repeated payoffs there are two possibilities. If there are no absolutely dominated actions, then there is an interdependent-choice equilibrium with full support and hence $A^{\text{IE}} = A^{\text{S}} = A$. Otherwise, there is a unique interdependent-choice equilibrium with respect to A^{IE} ; namely, a player chooses his (unique) dominant action and his opponent chooses the (unique) best response to it. In view of Theorem 3, this results in the following characterization of sequential implementation for 2×2 environments.

PROPOSITION 6 *In generic 2×2 environments, a distribution is sequentially implementable if and only if it is a interdependent-choice equilibrium with respect to A^{CR} .*

Example 1 [continued] In the teamwork example, there are no dominated actions and hence $A^{\text{IE}} \cup A^{\text{R}} = A$. This implies that every interdependent-choice equilibrium is sequentially

⁷Absolute dominance is closely related to the payoff-dominant cylinders in Lee (2011). The main difference is that he is interested in dominant action profiles instead of dominated actions.

implementable. Indeed, (S, W) and (W, S) are NE of the simultaneous move game. Every other IE except the one that places full probability on (W, W) has full support, and hence all the punishments are also used along the equilibrium path and following recommendations is a sequential equilibrium of the mediated game. The case of (W, W) is more complicated. Suppose that players rule out mistakes by the mediator. Then if a player deviates in the mediated game, the other player will become informed of this deviation and his best response is to choose S. (W, W) is still sequentially implementable allowing mistakes from the mediator or considering more complicated mechanisms. In any case, (W, W) can be arbitrarily approximated by completely mixed sequentially implementable distributions.

5 PERFECT BAYESIAN IMPLEMENTATION

This section defines a form of perfect Bayesian equilibrium (PBE) which is weaker than sequential equilibrium, and provides sufficient and necessary conditions for PB implementation. Since sequential equilibria are PBE, these conditions are also necessary for sequential implementation. My focus on PBE is partially motivated by the fact that it is the finer refinement for which I can provide a complete characterization. However, PBE might be an interesting solution concept in its own right, see section §5.2.

Sequential equilibrium is defined in terms of sequential rationality and belief consistency. Sequential rationality requires choices to be optimal at the interim stage for every information set in the game. As for belief consistency, Bayes rule no longer applies after receiving unexpected news, but it is still possible to require players to update their beliefs in accordance with some prior assessment of the relative likelihoods of different trembles or mistakes. Belief consistency requires this to be the case, and requires that the prior assessments should be common to all players. My definition of PBE imposes sequential rationality and requires beliefs to be consistent with trembles, but allows players to disagree about which deviations are more likely.⁸

It is useful to allow for mechanisms in which Nature assigns zero probability to some of its available moves. This is because, when faced with a null event, a player can believe that it was Nature who made a mistake instead of necessarily believing that an opponent deviated (intentionally or by accident) from the equilibrium.⁹ In order to define consistent beliefs, it is necessary to introduce new notation to denote players' beliefs about Nature's choices, other than σ_0^* . Let Σ_0 and Σ_0^+ denote the sets of mixed strategies and strictly mixed strategies for Nature.

A conditional belief system for i in \mathcal{G} is a function $\psi_i : \mathcal{H}_i \rightarrow \Delta(Y_i)$ with $[\psi_i(H_i)](H_i) = 1$ for every $H_i \in \mathcal{H}_i$. $\psi_i(y_i|H_i) \equiv [\psi_i(H_i)](y_i)$ is the probability that i assigns to being in y_i

⁸This definition is stronger than other notions of weak PBE in the literature. See for instance chapter §8 in Fudenberg and Tirole (1991), in particular, the examples corresponding to figures 8.1 and 8.9.

⁹It is often assumed that Nature assigns positive probability to all of its available moves, but I am unaware of any good arguments to maintain this assumption. Consider for instance the following quote from Kreps and Wilson (1982): “To keep matters simple, we henceforth assume that the players initial assessments [on Nature's choices] are strictly positive”, page 868. For further discussion see section §7.3.

whenever he is in H_i . Let Ψ_i denote the set of i 's conditional belief systems. An assessment is a tuple $(\psi, \sigma) \in \Psi \times \Sigma$ that specifies both players' (common) prior beliefs on strategies and interim beliefs about the position of the game. An extended assessment is a tuple $(\psi, \sigma, \sigma_0) \in \Psi \times \Sigma \times \Sigma_0$ that also specifies prior beliefs on Nature's choices. Given an assessment (ψ, σ) , an information set H_i for player i and a move $m \in M(H_i)$, let $V_i(m|H_i)$ denote i 's expected payoff from choosing m at H_i . The expectation is taken given his interim beliefs $\psi_i(H_i)$ regarding the current state of the game and assuming that future choices will be made according to σ .

DEFINITION 8 [Perfect Bayesian equilibrium]

- An assessment $(\psi, \sigma) \in \Psi \times \Sigma$ is *weakly consistent* if and only if for every player there exists a sequence of strictly mixed extended assessments $(\psi^n, \sigma^n, \sigma_0^n)$ which satisfy Bayes' rule and converge to $(\psi, \sigma, \sigma_0^*)$.
- An assessment $(\psi, \sigma) \in \Psi \times \Sigma$ is *sequentially rational* if and only if $V_i(s(H_i)|H_i) \geq V_i(m|H_i)$ for every $H_i \in \mathcal{H}_i$, every $m \in M(H_i)$, and every strategy $s_i \in S_i$ such that $\sigma_i(s_i) > 0$
- A *perfect Bayesian equilibrium* (PBE) of \mathcal{G} is an assessment (ψ, σ) that is both sequentially rational and weakly consistent.

Before proceeding to the characterization, it is helpful to understand the requirements implied by PBE. Sequential rationality requires that the choices that occur off the equilibrium path should be optimal. This implies that players must always believe that the *future* choices of their opponents will be rational, and this fact is common knowledge. However, off the equilibrium path, PBE imposes no restrictions on beliefs about *past* choices, nor agreement of beliefs across different players. Hence, the difference between PBE and Nash equilibrium can be thought of as a form of *forward-looking* rationalizability off the equilibrium path.¹⁰

5.1 CREDIBLE THREATS AND PB IMPLEMENTATION

The previous discussion suggests that there are two kind of actions which can always be enforced as credible punishments for PB implementation. Any punishment which is C-rationalizable is admissible because PB implementation does not require that off-path beliefs across agents should agree. Hence, it is always possible to ensure that the player performing the punishment has the conjectures which rationalize it. Additionally, since PBE only imposes belief of rationality for future choices, a player can always hold arbitrary beliefs about *past* deviations by his opponents. This suggests that best responses to arbitrary *degenerate* conjectures are also admissible. Our characterization results from combining these two ideas, using the notion of *forward-looking conjectural rationalizability*.

¹⁰Though related, this is different from forward induction, requiring rationality and strong belief of rationality (Battigalli and Siniscalchi, 2002). Besides forward-looking rationality, forward induction also requires players to try to rationalize past deviations of their opponents as if they had been intended, whenever possible.

DEFINITION 9 [Forward-looking conjectural rationalizability]

- An action $a_i^* \in A_i$ is *FC-rationalizable* with respect to an action sub-space $A' \in \mathcal{A}$ if and only if there exists some conjecture $\lambda_i^1 \in \Lambda_i(A')$, some **degenerate** conjecture $\lambda_i^0 \in \Delta(A_{-i})$, and some $\mu \in [0, 1]$ for which $a_i^* \in \arg \max_{a_i \in A_i} U_i(a_i, \lambda_i)$, where $\lambda_i = \mu \lambda_i^0 + (1 - \mu) \lambda_i^1 \in \Lambda(A)$. Let $\text{FR}_i(A')$ denote the set of i 's actions that are FC-rationalizable with respect to A' .
- An action sub-space $A' \in \mathcal{A}$ is self FC-rationalizable if and only if $A' \subseteq \text{FR}(A')$.
- The space of FC-rationalizable action profiles $A^{\text{FR}} \in \mathcal{A}$ is the largest self FC-rationalizable action space.

As before, A^{FR} is guaranteed to exist because $\text{FR}(\cdot)$ is \subseteq -monotone and thus the union of all self FC-rationalizable sets is also self FC-rationalizable. Also, it is non-empty because it always contains the set of C-rationalizable action profiles.

Intuitively, one can think of λ_i^0 as the arbitrary beliefs (degenerate conjectures) over past deviations, and think of λ_i^1 as the conjectures about future FC-rationalizable choices. With this interpretation, an action a_i is FC-rationalizable with respect to an action space $A' \in \mathcal{A}$ if it is a best response to some conjecture $\lambda_i \in \Lambda_i$ that assigns full probability to actions in A'_{-i} , *only for choices that occur in the future*. λ_i can assign positive probability to any action, provided that this probability is independent from i 's choice. The set of FC-rationalizable actions is exactly the set of credible threats that characterizes PB implementation.

THEOREM 7 *A distribution over action profiles is PB implementable if and only if it is an interdependent-choice equilibrium with respect to A^{FR} .*

There are two interesting corollaries of this result. First, since sequential implementability implies PB implementability, Theorem 7 implies that sequentially implementable distributions belong to $\text{IE}(A^{\text{FR}})$. Hence the result can be interpreted as a necessary condition for sequential implementation in arbitrary environments.

COROLLARY 8 *Every sequentially implementable distribution is an interdependent-choice equilibrium with respect to A^{FR} .*

Second, since C-rationalizable actions are also FC-rationalizable, then, in games with no absolute dominance a distribution is PB implementable if and only if it is a coordinated equilibrium. This means that requiring PB instead of Nash constraints has a small impact, because most games of interest have no absolutely dominated actions.

COROLLARY 9 *When there are no absolutely dominated actions, then a distribution is PB implementable if and only if it is an interdependent-choice equilibrium.*

Example 3 Consider the environment described by the payoff matrix in Figure 6, but now suppose that $k = 5$. Now, player 1 is willing to choose T only if he believes with probability

at least $2/3$ that he is the first player to move and player 2 chooses L. However, player 2 is only willing to choose L if he believes with probability at least $2/3$ that he is the first player to move and player 1 chooses B. Also, player 1 is only willing to choose B if he believes with probability at least $2/3$ that he is the first player to move and player 1 chooses T. Finally, player 2 is only willing to choose T if he believes with probability at least $2/3$ that he is the first player to move and player 1 chooses L. Hence T and L can only be played if the players disagree about the order of play in a way that is not consistent with sequential implementation. Therefore T and L cannot be used for sequential implementation, and the only sequentially implementable outcome is (D, D). However, since there is no absolute dominance, T and L are credible threats for PB implementation and thus (C, C) is PB implementable.

This section concludes with a characterization of the operator FR. Loosely speaking, the following proposition shows that it is equivalent to the elimination of strictly dominated actions in an auxiliary game. Hence finding A^{FR} is no more complicated than finding the set of rationalizable actions of a finite game.

PROPOSITION 10 *An action $a_i \in A_i$ is FC-rationalizable with respect to an action subspace $A' \in \mathcal{A}$ if and only if there is no $\alpha_i \in \Delta(A_i)$ such that:*

1. $\max \{u_i(a_i, a_{-i}) \mid a_{-i} \in A'_{-i}\} < \min \{u_i(\alpha_i, a_{-i}) \mid a_{-i} \in A'_{-i}\}$
2. $u_i(a_i, a_{-i}) < u_i(\alpha_i, a_{-i})$ for every $a_{-i} \in A_{-i} \setminus A'_{-i}$

5.2 A RATIONALE FOR PB EQUILIBRIUM

Loosely speaking, equilibrium refinements such as sequential equilibrium pretend to capture the restriction that choices should be in equilibrium, not only along the equilibrium path, but also in every ‘subgame’.¹¹ In contrast, the notion of PB equilibrium here proposed requires equilibrium only along the equilibrium path, and only imposes rationalizability in every ‘subgame’. I believe that there are situations in which this former approach is more adequate.

Equilibrium is not a straightforward consequence of rational behavior. Assuming rationality and common certainty of rationality only guarantees that choices are rationalizable, in order to guarantee equilibrium one must also assume mutual knowledge of beliefs or strategies (Brandenburger, 1992). Mutual knowledge of beliefs is usually justified in terms of: focal points, communication or repetition. For instance Lewis (1969), provides an eloquent analysis of these justifications in the context of coordination games. These justifications appear less appealing when it comes down to strategies and beliefs off the equilibrium path.

Determining whether an equilibrium is sequential or not is a difficult task. Hence, it seems hard to come up with a universal argument to defend sequential equilibria as focal points. Furthermore, there are complexity issues related to finding equilibria. Planning for

¹¹I am using the word ‘subgame’ informally to refer to different decision points even if they do not correspond to proper subgames.

all possible contingencies or agreeing on their likelihood seems an arduous task, specially when null events are concerned. Finally, there may be natural dynamic processes which converge to equilibrium, Players beliefs might become concordant not because of planning but because of experience. However, repetition provides no experience about null events which only happen off the equilibrium path (Fudenberg and Levine, 1993). Therefore there might be situations in which (i) it makes sense to assume mutual knowledge of beliefs *exclusively* along the equilibrium path; and yet (ii) rationality and common certainty of rationality may also be defended in every subgame. In such situations, PBE is a more natural solution concept than sequential equilibrium.

6 RELATED LITERATURE

The incentives arising from choice interdependence have been studied in a variety of ways, that differ from the current paper both in methodology and purpose. This section compares the present work with a sample of representative papers within the related literature. The discussion is divided into three sections. Section §6.1 considers papers in which interdependence is modeled explicitly as the result of sequential mechanisms. Section §6.2 considers papers that allow for tacit choice interdependence without being too explicit about the mechanisms behind it. Section §6.3 discusses papers on robust predictions that are independent from structural assumptions.

6.1 EXPLICIT INTERDEPENDENCE

Revision games Perhaps the model closest to ours is that of Kamada and Kandori (2009, 2012). Kamada and Kandori consider an environment in which players choose their actions during a continuous time interval and, before the actual play of the game, they might learn the intended actions of their opponents and have a chance to revise their own. The revision opportunities are stochastic and exogenous. This model is particularly appealing in settings in which actions are not instantaneous, for instance because a considerable amount of planning or preparation is involved. In such settings there is a time lapse from the moment of choosing an action to the moment of performing it. During this time lapse, an agent could receive unexpected information and change his mind.

I contrast, choices in my model are instantaneous. This is captured by the fact that, once a player makes reaches a pivotal decision point and makes a move representing an action, the decision is final. Instantaneous choices is an adequate assumption for settings in which the amount of time between choosing and performing an action is short, and hence the possibility of revision during such interval is negligible. The results of the current paper show that a phenomenon similar to the revision effects of Kamada and Kandori may prevail in such settings, provided that choices do not happen exactly at the same time.

Another difference is in the nature of the results. Kamada and Kandori take both the timing and the sequential structure as given, and ask which outcomes are possible in equilibria of the corresponding game. In contrast, I take the timing and information as design variables and ask which outcomes are possible in equilibria of *some* extensive form mechanism.

Mediators There are different extensions of Aumann (1974, 1987) which also increase the power of the mediator. Perhaps the closest example can be found in Myerson (1986), who allows for sequential mediation and dependence on past play, but takes the timing of choices and recommendations as exogenous. Hence my mechanisms extend those in Myerson's model. On the other hand, several papers have considered mediated games in which the players can either commit ex-ante to follow the mediators' recommendations (Forgó, 2010, Moulin and Vial, 1978), or delegate their choice to the mediator who can then commit to follow pre-specified instructions (Ashlagi et al., 2009, Monderer and Tennenholtz, 2009). In both cases, the mediator serves as a commitment mechanism drastically reducing or completely eliminating moral hazard. In contrast, I do not allow for any form of commitment or delegation from part of the players.

Pre-play negotiations There are a number of papers that allow pre-play negotiation phase in which players can make binding agreements. This literature can be traced back to Kalai (1981), for more recent analyses see Bade et al. (2009), Kalai et al. (2010), Renou (2009). In other papers, actions are not contractible, but players can sign binding agreements on action-contingent monetary transfers, e.g. Jackson and Wilkie (2005). In contrast, my model refers to situations without enforceable contracts nor monetary transfers. Environments without commitment can be easily thought of. Monetary transfers or other form of side punishments or compensations are much harder to rule out. One could think of intermediate situations in which transfers are possible but actions are not verifiable. In such settings, transfers would have to occur freely without the use of contracts. This extension is left as an open problem.

Distributed games The work of Monderer and Tennenholtz (1999), introduces a model of distributed games mimicking software interaction in networks. The sequential and incentive structure in their model is similar to mine. There are two important differences. First, the signals are generated by players and not by a mediator, which is only possible because they consider information structures which do not satisfy perfect recall. Also, each player plays the game a number of times on different locations, while my model considers single-shot interactions.

Sequential choices The role of choice interdependence is also present in some models related to Stackelberg equilibrium, resulting from games in which players move sequentially and actions may be observed at each stage. The literature on endogenous timing and price leadership (e.g. Van Damme and Hurkens (1999)), allows oligopolistic firms to choose whether to publicly announce their prices early on, or to wait to see other firm's prices before choosing their own. Also related is Solan and Yariv (2004), which considers sequential games in which the second player to move can acquire signals about the choices of the first player. In both cases, choice interdependence results from the timing and information structure, but only a restricted class of mechanisms is considered. The literature on endogenous leadership takes the information structure is given, while the literature on espionage takes the sequential structure as given. In contrast, the current work extends both models by taking both the sequential and the information structures as design variables.

Quantum games The literature on quantum games endows the players with quantum randomization devices. Players' choices can thus be contingent on quantum states that are entangled in a functional relationship. This enables for different forms of choice interdependence (Eisert et al., 1999). However, the precise interpretation of quantum randomization for human decision making remains an open problem.

Extensive form mechanisms There are different notions of extensive form mechanisms elsewhere in the literature differ from ours in that they allow for commitment. For example, the literature on exact implementation with incomplete information uses what I call outcome equivalent extensive form games (Moore and Repullo, 1988, Serrano and Vohra, 1997). Another example can be found in the work of Kalai (2004, 2006) on robust predictions. Kalai's definition requires that: (i) every player gets to move along each history of play and (ii) each player i has a strategy that guarantees $\tau(z) = a_i$ independently of his opponent's choices for each $a_i \in A_i$. The following example illustrates the forms of commitment implicit in Kalai's definition.

Consider a generic 2×3 environment with $A_1 = \{a_1, a'_1\}$ and $A_2 = \{a_2, a'_2, a''_2\}$. Figure 7 illustrates extensive forms that fit the definition of Kalai but allow for total and partial commitment respectively. The left panel shows a mechanism in which player 1 can choose an action but, in the event that player 2 chooses a_2 he can then review his original choice. This means that player 1 can commit to choosing an action *conditional* on player 2 not choosing a_2 . The definition of representing moves rules out this possibility. The right panel shows a mechanism in which player 2 can commit to not using a_2 before making a final choice between a'_2 and a''_2 . I rule out this possibility by requiring that payoff-relevant choices are made in pivotal nodes.




FIGURE 7 Mechanisms with commitment

6.2 TACIT INTERDEPENDENCE

The results of this paper are driven not by actual interdependence but by conjectural beliefs. It is not important whether players choices depend on each other, it only matters that each player *believes* that the choices of his opponents might depend on his own. There are papers that study tacit choice interdependence without being too specific about the mechanism which generates it. In conjectural variations equilibrium (CVE), see for instance Figuières et al. (2004), all the players act as Stackelberg leaders, *as if* they were the first

player to move and other players could react to their choices. A similar idea results in the notion of cooperative equilibrium from Halpern and Rong (2010) and in Rapoport (1965)'s account of the prisoner's dilemma. In what follows I argue that, such behavior is hard to justify in any extensive form game corresponding to a single interaction.¹² For simplicity, the exposition is phrased in terms of Rapoport's framework.

Rapoport argues that the two prisoners are identical ex-ante and the situation is symmetrical. Consequently, each prisoners believes that his opponent's choice will match his own. In practical terms, this means that he is not really choosing between cooperating or defecting, but between the outcomes (D, D) and (C, C) the former one being strictly better. This argument is considered to be equivalent to an individual choice problem known as Newcomb's paradox, and it has been questioned for conflicting with the notion that players have independent free wills (Gibbard and Harper, 1980, Lewis, 1979).

Rapoport's solution is not consistent with the sequential nature of my model. In an extensive form mechanism, the choice of i can only influence the choice of j if it precedes it. Since precedence is asymmetric, in a state of the world in which i influences j , j cannot influence i . In contrast, Rapoport requires both choices to be entangled in a functional relationship, that is, they have to influence each other simultaneously. In particular, recall that in the prisoner's dilemma example it is only possible to sustain cooperation as an interdependent-choice equilibrium when $G - g \leq g - b$. This is because an extensive form mechanism can generate incentives for a player to cooperate only when he is the first mover (otherwise he will take his accomplice's action as given and the domination argument holds). Hence, interdependent-choice equilibrium can be thought of as a refinement of models with tacit choice interdependence, which requires that the players' conjectures should be consistent with the acyclic nature of sequential choices.

Unrestricted tacit cooperation can still arise if one is willing to assume that: (i) players move at the same time and their choices are affected by a common force, such as a quantum randomization devices; or (ii) the order in which choices are made depends on the chosen actions as in revision games;¹³ or (iii) players disagree on their prior assessments about the order of choices, for instance, consider the following parable showing that unconditional cooperation could arise in the mechanism from Figure 2, if the common prior assumption was relaxed.

Suppose that the prisoners agree in believing that the DA will always visit the prisoner with the higher profile first, but also agree to disagree about their reputations. Each prisoner is certain that he is more famous than his opponent, that his opponent is certain of the opposite fact and these beliefs are also common certainty. With this subjective prior assessments, the strategy to remain silent unless you receive evidence that your accomplice has confessed is always sequentially rational and cooperation can be sustained without any additional restrictions on payoffs.

Along the equilibrium path, each prisoner will cooperate because he will be certain about the following events: (i) that he is the first player to move, (ii) that if he confesses then

¹²It it should be noted that CVE are often thought of as reduced forms of a repeated game, see for instance Kalai and Stanford (1985).

¹³For a more concrete example, consider a two firm Stackelberg duopoly in which the first firm to set a price also offers a price matching guarantee.

his accomplice will learn about his defection before making his choice and will also confess, and (iii) that if he cooperates then his opponent will (mistakenly) think that he is the first mover and will remain silent. Although the justification is different, unrestricted cooperation arises because players hold exactly the same conjectures postulated by Rapoport. Notice that, although one of the prisoners is necessarily mistaken, the prisoners' prior assessments of the outcome coincide and thus this 'equilibrium' is self-confirming in the sense of Fudenberg and Levine (1993), as long as the realized order is not revealed.

6.3 ROBUST PREDICTIONS

Thus far I have analyzed interdependent choices from an implementation perspective. However, the analysis also admits a positive perspective. Game theoretic predictions are very sensitive to the rules of the game (e.g. order of play and information structure) which are often unknown to the Economist. This point has been stressed out for instance in the seminal work of Sutton (1991) who looks for detail-free robust predictions in IO. Given Theorem 1, the set of interdependent-choice equilibria can be interpreted as a weak but robust predictions indicating anything that is possible in equilibrium independently of any structural or informational assumptions other than no commitment.

Going back to the prisoner's dilemma story, suppose for instance that the prisoners are taken to different cells. The DA (she) visits each one of them sequentially to offer them the deal. The prisoners believe that they are both equally likely to be the first one to receive the offer. Also, assume that the (non-strategic) DA will always try to convince each prisoner that his accomplice has already confessed. This norm of behavior is common knowledge to the prisoners who will only believe the words of the DA when they are true, e.g. because only then will she be able to show a written confession as evidence of her claim.¹⁴ This would result in the mechanism from Figure 2, which could very well arise naturally even though an outside observer might be tempted to assume that choices are independent.

The empirical results from Rapoport (1997) and Muller and Sadanand (2003) suggest that choice interdependence has significant effects in actual decision making. In an experimental setting, they find significant differences in behavior across different extensive form games with a common reduced form. For instance, they consider bargaining games with different orders of play and private actions. Since the actions of the first mover are not observable, they have no effect on the set of equilibria. However, they observe significant differences between the first and second movers. They rationalize this phenomenon by introducing uncertainty about the independence of choices. See also Vyrastekova and Funaki (2010).

Thinking of robust predictions, the current work is similar to Bergemann and Morris (2011a,b), with the difference that they deal with adverse selection instead of moral hazard. They study environments with incomplete information and are interested in robustness with

¹⁴The choices of the DA are in fact enabling coordination. Alternatively, one could assume that the DA is young and unexperienced. That this is her first big case and that obtaining a confession might provide a big boost to her career. Hence, she will remain nervous unless she obtains a confession, in which case she will become confident. The prisoners will not believe anything the DA says because they know that talking is cheap. However, they can observe whether she is nervous or confident and use this signal to coordinate their actions.

respect to hierarchies of beliefs (Wilson (1987) doctrine). They define a notion of Bayesian correlated equilibrium which characterizes the distributions over outcomes that can result as an equilibrium of the Bayesian game corresponding to *some* belief structure.

With this interpretation, this paper is also closely related to Kalai (2004). However, the nature of the results is significantly different. Using my language, Kalai’s main result is that, under some anonymity and continuity assumptions, the set of equilibria of the *simultaneous move game* are robust in that they remain to be approximate Nash equilibria of *any* extensive form mechanisms as long as the number of players is large enough. In contrast, I find the set of outcomes that can arise in *some* extensive form mechanism. Interdependent-choice equilibria might not be robust to changes in the rules of the game, but in some cases they can Pareto dominate all the Nash equilibria of the simultaneous move game. Kalai’s work shows the *limits* of coordination in anonymous games with *many* players, while the present work shows the *possibilities* of coordination in games with *few* players.

7 SUMMARY AND DISCUSSION

The current paper proposes a model to analyze the role of choice-interdependence as a mechanism to generate incentives in moral hazard environments. It proposes a class of mediated mechanisms in which a mediator manages the game through private recommendations. Two salient aspects of the model are that the recommendations are sequential and occur during the actual play of the game, and that they can depend on previous choices. The interdependence between choices and signals may generate powerful incentives that, for instance, allow for cooperation in the prisoner’s dilemma.

Interdependent-choice equilibrium is defined as a Nash equilibrium of a mediated mechanisms. The set of interdependent-choice equilibria admits a canonical characterization consisting of a finite set of affine inequalities. The paper also provides conditions for implementation according to different equilibrium refinements that take into account the problem of perfection. The conditions simply restrict the set of credible threats that the mediator can recommend as credible threats off the equilibrium path. Different sets of credible threats offer necessary and/or sufficient conditions for different solution concepts. The implications are summarized in Figure 8.

Interdependent-choice equilibrium can also be interpreted as a refinement of other solution concepts involving tacit choice-interdependence, or as a robust solution concept for environments in which the information and sequential structure are unknown. The rest of this section discusses some limitations and extensions of the model.

7.1 ALTERNATIVE RESTRICTIONS

In order to characterize the possibilities of choice interdependence, I imposed minimal restrictions on our set of extensive form mechanisms. Essentially, I only require that payoff-relevant choices have to be made at pivotal information sets. This of course would not be useful in environments in which further restrictions apply. Additional restrictions can be incorporated by adjusting the worst punishment functions \underline{w}_i .

For instance, mediated mechanisms attribute a great amount of power to the mediator

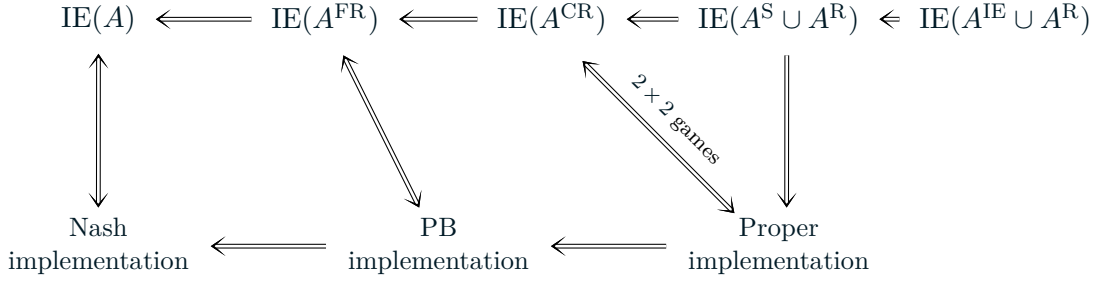


FIGURE 8 Summary of results

who can control the order in which players make their choices and perfectly observe the choices made. Instead of assuming that the mediator controls the order of choices, it is much more plausible to assume that she can control the order of her recommendations but that players remain open to the possibility of acting before or after they encounter the mediator. In such cases, the mediator could not recommend action-specific punishments. A player who intended to deviate would make his choice after the mediator has left and thus the mediator could no longer observe the specific deviation. I claim that the set of implementable outcomes under these conditions can be characterized by replacing \underline{w}_i with the constant minimax punishment $\underline{w}'_i(a'_i) = \min_{\alpha_{-i} \in \Delta(B^*)} \max_{a_i \in A_i} u_i(a, \alpha_{-i})$. Notice that this modification makes no difference when each player has at most two actions, including the prisoner's dilemma and the team work example considered previously.

To give another example, suppose that all deviations from equilibrium are publicly observed. In that case, one could replace the worst punishments function \underline{w}_i with the weaker version $\underline{w}'_i(a'_i) = \min\{u_i(a'_i, a_{-i}) \mid a_{-i} \in \text{BR}_{-i}(a'_i)\}$ where BR_{-i} is $-i$'s best response correspondence. In any case, ones would still obtain an interesting solution concept that would lie between the set of correlated equilibria and the set of interdependent-choice equilibria.

7.2 MANY PLAYERS

The results of the paper extend naturally to n -player games. However the required notation becomes significantly cumbersome. For one thing, when the mediator chooses an ordering of the players she is no longer choosing the player who moves first but an entire enumeration n of i . Hence θ_a must be a distribution over such enumerations. The incentive constraints then become:

$$\sum_{a_{-i} \in A_{-i}} \alpha(a) u_i(a) \geq \sum_{a_{-i} \in A_{-i}} \sum_n \alpha(a) \theta_a(n) \cdot \min \left\{ u_i \left(a'_i, a'_{n^+(i)}, a_{n^-(i)} \right) \mid a_{n^+(i)} \in B_{n^+(i)}^* \right\}$$

where $n^+(i) = \{j \in I \mid n(j) > n(i)\}$ and $n^-(i) = \{j \in I \mid n(j) < n(i)\}$ are the set of players that move before and after i according to n respectively. While the solution concepts can be extended to environments with many players, it is not entirely clear whether their interpretation remains valid. I have assumed that coordination is costless and monitoring

is perfect. These assumptions might not carry on to environments with large number of players, in which small frictions might add up or propagate. In some cases, the effectiveness of choice interdependence might dissipate as the number of players increases.

7.3 MISTAKES BY NATURE

I have allowed players to attribute deviations from the equilibrium path to unexpected moves by Nature (or a mediator). In equilibrium, when an agent finds himself in the equilibrium path he may believe that is because Nature made a mistake and not because his opponent decided to deviate. This feature because it makes the analysis tractable and because I am unaware of any solid argument against it. Since this is not a common feature in other models, it requires some justification.

Consider for instance the hypothetical situation of a loving marriage after the wife finds unfamiliar lingerie mixed in the laundry. A plausible explanation is that the husband deviated from the marital arrangement by involving in an extramarital relationship, and made the mistake of bringing home evidence of his deviation. However, more often than not, a *trusting* wife is likely to ignore this story and instead recur to intricate explanations involving unexpected chance events. Back to our abstract environment, each player i knows that his opponents cannot gain from deviating, as long as he sticks to his equilibrium strategy. Hence he has no reason to be suspicious about them, and attributing deviations to Nature seems reasonable.

A key element in this previous example is the trusting nature of the relationship. This line of thought might find less favor in situations in which the agents have reasons to be suspicious about each other. A extramarital affair is bound to be the favored explanation by a suspicious wife who expects to be cheated. The sense in which allowing for null chance moves is sensible might depend on the level of trust or suspicion among the agents. However, the example here provided describes a common situation in which I judge this feature to be both appropriate and necessary for explaining observed behavior. In any case, one could strict attention to refinements that exclude this possibility. In particular, the sufficient condition for sequential implementation in Theorem 3 can be easily adapted to exclude of null chance moves, see footnote 6.

7.4 INCOMPLETE INFORMATION

An important part of the implementation literature focuses on the important problem of eliciting private information. In contrast, I consider a complete information environment in order to emphasize the role of interdependent choices in generating incentives to prevent moral hazard. Moral hazard is often ruled out by assuming that a principal has all the bargaining power and can commit to enforce any individually rational and incentive compatible mechanism. An extension of interdependent-choice equilibrium to incomplete information environments could account for the set of outcomes that are implementable when commitment is not possible. It could also be used as a set of admissible coalition deviations for collusion-proof implementation as in Lamy (2008). The resulting definition would be adequate for settings in which legal restrictions prevent the agents from enforcing side contracts. Preliminary analysis shows that in the simplest single object allocation

problems, the second price auction is collusion-proof while the first price auction is not.

7.5 MEDIATION TECHNOLOGY

The definition of interdependent-choice equilibrium relies on games with a non-strategic trusted mediator who manages the play and can observe actions perfectly. Changing these assumptions may affect the set of implementable outcomes. With imperfect monitoring, the need to punish deviations might necessarily introduce inefficiency as it does in repeated games (Green and Porter, 1984). When non-strategic mediators are out of the question, an interesting question is whether the mediator can be replaced by cheap-talk (Vida and Forges, 2013), or whether it is possible to generate the required information structures with transparent mediators (Izmalkov et al., 2005).

REFERENCES

- Abreu, D., Pearce, D., and Stacchetti, E. (1990). Towards a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):1041–1063.
- Ashlagi, I., Monderer, D., and Tennenholtz, M. (2009). Mediators in position auctions. *Games and Economic Behavior*, 67(1):2–21.
- Aumann, R. J. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96.
- Aumann, R. J. (1987). Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55(12):1–18.
- Bade, S., Guillaume, H., and Renou, L. (2009). Bilateral commitment. *Journal of Economic Theory*, 144(4):1817–1831.
- Battigalli, P. and Siniscalchi, M. (2002). Strong belief and forward induction reasoning. *Journal of Economic Theory*, 106(2):356–391.
- Bergemann, D. and Morris, S. (2011a). Correlated equilibrium in games with incomplete information. Cowles Foundation, mimeo.
- Bergemann, D. and Morris, S. (2011b). Robust predictions in games with incomplete information. Cowles Foundation, mimeo.
- Bernheim, D. (1984). Rationalizable strategic behavior. *Econometrica*, 52(4):1007–1028.
- Brandenburger, A. (1992). Knowledge and equilibrium in games. *The Journal of Economic Perspectives*, 6(4):83–101.
- Dufwenberg, M. and Kirchsteiger, G. (2004). A theory of sequential reciprocity. *Games and Economic Behavior*, 47(2):268–298.
- Eisert, J., Wilkens, M., and Lewenstein, M. (1999). Quantum games and quantum strategies. *Physical Review Letters*, 83(15):3077–3080.
- Figuières, C., Jean-Marie, A., Quérrou, N., and Tidball, M. (2004). *Theory of conjectural variations*, volume 2 of *Series on Mathematical Economics and Game Theory*. World Scientific.

- Forges, F. (1986). An approach to communication equilibria. *Econometrica*, 54(6):1375–1385.
- Forges, F. (1993). Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision*, 35(3):277–310.
- Forgó, F. (2010). A generalization of correlated equilibrium: a new protocol. *Mathematical Social Sciences*, 60(1):186–190.
- Fudenberg, D. and Levine, D. K. (1993). Self-confirming equilibrium. *Econometrica*, 61(3):523–545.
- Fudenberg, D. and Tirole, J. (1991). *Game theory*. MIT Press.
- Gibbard, A. and Harper, W. L. (1980). Counterfactuals and two kinds of expected utility. In Harper, W. L., Stalnaker, R., and Pearce, G., editors, *Ifs, conditionals, beliefs, decision, chance and time*. Springer.
- Green, E. J. and Porter, R. H. (1984). Noncooperative collusion under imperfect price information. *Econometrica*, 52(1):87–100.
- Halpern, J. Y. and Rong, N. (2010). Cooperative equilibrium. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*.
- Hoffman, D. A. (2011). Mediation and the art of shuttle diplomacy. *Negotiation Journal*, 27(3):263–309.
- Izmailkov, S., Micali, S., and Lepinski, M. (2005). Rational secure computation and ideal mechanism design. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 585–594.
- Jackson, M. O. and Wilkie, S. (2005). Endogenous games and mechanisms: side payments among players. *The Review of Economic Studies*, 72(2):543–566.
- Kalai, A. T., Kalai, E., Lehrer, E., and Samet, D. (2010). A commitment folk theorem. *Games and Economic Behavior*, 69(1):123–137.
- Kalai, E. (1981). Preplay negotiations and the prisoner’s dilemma. *Mathematical Social Sciences*, 1(4):375–379.
- Kalai, E. (2004). Large robust games. *Econometrica*, 72(6):1631–1665.
- Kalai, E. (2006). Structural robustness of large games. Northwestern University, mimeo.
- Kalai, E. and Stanford, W. G. (1985). Conjectural variations strategies in accelerated cournot games. *International Journal of Industrial Organization*, 3(2):133–152.
- Kamada, Y. and Kandori, M. (2009). Revision games. Harvard University, mimeo.
- Kamada, Y. and Kandori, M. (2012). Asynchronous revision games. Harvard University, mimeo.
- Kreps, D. M. and Wilson, R. (1982). Sequential equilibria. *Econometrica*, 50(4):863–894.
- Lamy, L. (2008). Mechanism design with partially-specified participation games. Technical report. Paris School of Economics, mimeo.
- Lee, B. (2011). *Essays in Game Theory*. PhD thesis, The Pennsylvania State University.
- Lewis, D. (1969). *Convention*. Harvard University Press.

- Lewis, D. (1979). Prisoner's dilemma is a Newcomb problem. *Philosophy and Public Affairs*, 8(3):235–240.
- Monderer, D. and Tennenholtz, M. (1999). Distributed games. *Games and Economic Behavior*, 28(1):55–72.
- Monderer, D. and Tennenholtz, M. (2009). Strong mediated equilibrium. *Artificial Intelligence*, 173(1):180–195.
- Moore, J. and Repullo, R. (1988). Subgame perfect implementation. *Econometrica*, 56(5):1191–1220.
- Moulin, H. and Vial, J. (1978). Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3-4):201–221.
- Muller, A. and Sadanand, A. (2003). Order of play, forward induction, and presentation effects in two-person games. *Experimental Economics*, 6(1):5–25.
- Myerson, R. B. (1986). Multistage games with communication. *Econometrica*, 54(2):323–358.
- Nishihara, K. (1997). A resolution of n-person prisoners' dilemma. *Economic Theory*, 10(3):531–540.
- Nishihara, K. (1999). Stability of the cooperative equilibrium in n-person prisoners' dilemma with sequential moves. *Economic Theory*, 13(2):483–494.
- Osborne, M. J. and Rubinstein, A. (1994). *A course in game theory*. MIT Press.
- Pearce, D. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050.
- Rapoport, A. (1965). *Prisoner's dilemma*. University of Michigan Press.
- Rapoport, A. (1997). Order of play in strategically equivalent extensive form games. *International Journal of Game Theory*, 26(1):113–136.
- Renou, L. (2009). Commitment games. *Games and Economic Behavior*, 66(1):488–505.
- Serrano, R. and Vohra, R. (1997). Non-cooperative implementation of the core. *Social Choice and Welfare*, 14(4):513–525.
- Solan, E. and Yariy, L. (2004). Games with espionage. *Games and Economic Behavior*, 47(1):172–199.
- Sutton, J. (1991). *Sunk costs and market structure: price competition, advertisement and the evolution of concentration*. MIT Press.
- Van Damme, E. and Hurkens, S. (1999). Endogenous stackelberg leadership. *Games and Economic Behavior*, 28(1):105–129.
- Vida, P. and Forges, F. (2013). Implementation of communication equilibria by correlated cheap talk: the two-player case. *Theoretical Economics*, 8(1).
- Vyrastekova, J. and Funaki, Y. (2010). Cooperation in a sequential n-person prisoner's dilemma: the role of information and reciprocity. mimeo.
- Wilson, R. (1987). Game theoretic analysis of trading processes. In Bewley, T. F., editor, *Advances in Economic Theory*. Cambridge University Press.

A NOTATION

For every player i , let $-i = I \setminus \{i\}$. For every collection of sets $\{X_i \mid i \in I\}$, for every function $f : Y \rightarrow \mathbb{R}^I$ and for every set of players $J \subseteq I$, let $X_J = \{x_J, x'_J, \dots\} = \times_{i \in J} X_i$, and let $f_J : B \rightarrow \mathbb{R}^J$ denote the projection of f on \mathbb{R}^J . In particular, $X = \{x, x', \dots\} \equiv X_I = \times_{i \in I} X_i$, $X_{-i} = \{x_{-i}, x'_{-i}, \dots\} = \times_{j \in I \setminus \{i\}} X_j$.

For every finite set X , $\Delta(X) = \{\xi \in \mathbb{R}_+^X \mid \|\xi\|_1 = 1\}$ denotes the set of distributions on X . $\xi(Y) = \sum_{x \in Y} \xi(x)$ denotes the probability of the event $Y \subseteq X$ according to ξ . When $X = \times_{i \in I} X_i$ is a Cartesian product, $\xi_i(x_i)$ and $\xi(x_{-i} \mid x_i)$ denote marginal and conditional distributions respectively.

B PROOFS

B.1 NASH IMPLEMENTATION

Proof of Theorem 1. The sufficiency follows straight from the definitions: a mediated game is an extensive form mechanism and interdependent-choice equilibria result from NE of mediated games. To establish necessity, fix a mechanism (\mathcal{G}, τ) , a NE σ^* and let α be the induced distribution. We will show that α is an interdependent-choice equilibrium.

For that purpose consider any pair of actions $a_i^*, a'_i \in A_i$ with $\alpha_i(a_i^*) > 0$ and $a'_i \neq a_i^*$. For each information set $H_i \in \mathcal{H}_i$, let $M^*(H_i)$ be the moves that represent a_i^* at H_i and are chosen with positive probability. Also let \mathcal{H}_i^* be the set of information sets *along the equilibrium path* in which i chooses a move representing a_i^* with positive probability, i.e.:

$$M^*(H_i) = \left\{ m \in M^{a_i^*}(H_i) \mid (\exists s_i \in S_i)(\sigma_i^*(s_i) > 0 \wedge s_i(H_i) = m) \right\}$$

$$\mathcal{H}_i^* = \{ H_i \in \mathcal{H}_i \mid \zeta^*(H_i) > 0 \wedge M^*(H_i) \neq \emptyset \}$$

where ζ^* is the distribution over nodes induced by σ^* . All expectations and conditional distributions used in the proof are with respect to ζ^* .

Now consider some $H_i \in \mathcal{H}_i^*$. Since H_i must be a pivotal information set, there exists some $m' \in M^{a'_i}(H_i)$. Since σ^* is a NE we know that for each $m^* \in M^*(H_i)$ we have:

$$(2) \quad \mathbb{E}[u_i(a_i^*, a_{-i}) \mid H_i, m^*] \geq \mathbb{E}[u_i(a'_i, a_{-i}) \mid H_i, m']$$

where H_i, m denotes the set of nodes $H_i \times \{m\}$ for $m \in \{m^*, m'\}$.

Now let $\Phi^{H_i} \subseteq H_i$ denote the event that τ_{-i} is already determined at H_i , i.e.:

$$(3) \quad \Phi^{H_i} = \left\{ y_i \in H_i \mid (\forall z, z' \in Z(y_i))(\tau_{-i}(z) = \tau_{-i}(z')) \right\}$$

and let $\bar{\Phi}^{H_i} = H_i \setminus \Phi^{H_i}$. Notice that the probability of Φ^{H_i} and the probability of $\tau_{-i}^{-1}(a_{-i})$ conditional on $\bar{\Phi}^{H_i}$ are independent from i 's choice at H_i . Hence:

$$(4) \quad \begin{aligned} \mathbb{E}[u_i(a'_i, a_{-i}) \mid H_i, m'] &= \zeta^* \left(\Phi^{H_i} \mid H_i, m' \right) \mathbb{E} \left[u_i(a'_i, a_{-i}) \mid H_i, m', \Phi^{H_i} \right] \dots \\ &\quad \dots + \zeta^* \left(\bar{\Phi}^{H_i} \mid H_i, m' \right) \mathbb{E} \left[u_i(a'_i, a_{-i}) \mid H_i, m', \bar{\Phi}^{H_i} \right] \\ &= \zeta^* \left(\Phi^{H_i} \mid H_i, m^* \right) \mathbb{E} \left[u_i(a'_i, a_{-i}) \mid H_i, m^* \right] \dots \end{aligned}$$

$$\begin{aligned} & \dots + \zeta^* \left(\bar{\Phi}^{H_i} | H_i, m^* \right) \mathbb{E} \left[u_i(a'_i, a_{-i}) | H_i, m', \bar{\Phi}^{H_i} \right] \\ & \geq \zeta^* \left(\Phi^{H_i} | H_i, m^* \right) \mathbb{E} \left[u_i(a'_i, a_{-i}) | H_i, m^* \right] + \zeta^* \left(\bar{\Phi}^{H_i} | H_i, m^* \right) \underline{w}_i(a'_i) \end{aligned}$$

Equations (2) and (4) together yield the following inequality *which does not depend on m'* :

$$\mathbb{E} \left[u_i(a_i^*, a_{-i}) | H_i, m^* \right] \geq \zeta^* \left(\Phi^{H_i} | H_i, m^* \right) \mathbb{E} \left[u_i(a'_i, a_{-i}) | H_i, m^* \right] + \zeta^* \left(\bar{\Phi}^{H_i} | H_i, m^* \right) \underline{w}_i(a'_i)$$

We can obtain one of these inequalities for each point in the game in which i chooses a_i^* with positive probability according to σ^* . To obtain the desired inequality we simply have to “integrate” over them in order to obtain:

$$\sum_{a_{-i} \in A_{-i}} \zeta^*(a_i^*, a_{-i}) u_i(a_i^*, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} \left[\zeta^*(-i, a_i^*, a_{-i}) u_i(a'_i, a_{-i}) + \zeta^*(i, a_i^*, a_{-i}) \underline{w}_i(a'_i) \right]$$

since this was for arbitrary i , a_i^* and a'_i , it follows that α is a CE. ■

B.2 C-RATIONALIZABILITY AND FC-RATIONALIZABILITY

Proof of Proposition 5. C-rationalizable actions are clearly not absolutely dominated. For the opposite direction, fix an action $a_i^* \in A'_i$ that is not absolutely dominated in A' , we will show that $a_i^* \in \text{CR}(A')$. Let $a_{-i}^* \in \arg \max_{a_{-i} \in A'_{-i}} u_i(a_i^*, a_{-i})$. For each $a'_i \in A'_i$ there exists some $a_{-i}(a'_i) \in A'_{-i}$ such that $u_i(a_i^*, a_{-i}^*) \geq u_i(a'_i, a_{-i}(a'_i))$ (otherwise we would have $a'_i \succ_{A'} a_i^*$). Hence a_i^* is a best response to $\lambda_i \in \Lambda_i(A')$ with $\lambda_i(a_{-i}^* | a_i^*) = 1$ and $\lambda_i(a_{-i}(a'_i) | a'_i) = 1$ for all $a'_i \in A'_i \setminus \{a_i^*\}$.

Now let A^n be the sequence of action subspaces generated by iteratively removing *all* absolutely dominated actions. For each $n \in \mathbb{N}$ fix some arbitrary action profile $a^0 \in A^n$ and for every i let a_i^* be a best response to a_{-i}^0 (it exists because the game is finite). Clearly, a^* is in A^{n+1} . Thus, by induction, A^n is a weakly decreasing sequence of *nonempty* action subspaces. Since \mathcal{A} is finite, this implies that A^n must converge in finite time to some nonempty limit A^∞ . Since undominated actions are C-rationalizable, this implies that A^∞ is self C-rationalizable and thus $A^\infty \subseteq A^{\text{CR}}$. Now recall that A^{CR} is C-rationalizable with respect to itself. This implies that for every n , if $A^{\text{CR}} \subseteq A^n$ then $A^{\text{CR}} \subseteq A^{n+1}$. Since $A^{\text{CR}} \subseteq A = A^1$, by induction we have that $A^{\text{CR}} \subseteq A^n$ for all $n \in \mathbb{N}$ and therefore $A^{\text{CR}} \subseteq A^\infty$.

We still have to show order independence. For that purpose, define an absolute-dominance elimination operator to be any function $U : \mathcal{A} \rightarrow \mathcal{A}$ ($U(A')$ is the space of actions that are not eliminated) such that for every $A' \in \mathcal{A}$: (i) never adds new actions, i.e. $U(A') \subseteq A'$; (ii) never eliminates undominated actions, i.e. $\text{CR}(A') \subseteq U(A')$; and (iii) if there are dominated actions then it always eliminates at least one, i.e. $\text{CR}(A') \neq A'$ implies $U(A') \neq A'$. Now define $A^m = U^m(A)$. As before, the monotonicity of U and the finiteness of \mathcal{A} imply that A^m converges to a limit A^∞ in finite time. By (ii) we know that $A^{\text{CR}} \subseteq A^\infty$ and by (iii) we know that $A^\infty \subseteq \text{CR}(A^\infty)$ and thus $A^\infty \subseteq A^{\text{CR}}$. ■

Proof of Proposition 10. We know that $a_i \in \text{FR}_i(A')$ if and only if it is a best response to some conjecture $\lambda_i = \mu \lambda_i^0 + (1 - \mu) \lambda_i^1$ with $\lambda_i^0 \in \Delta(A_{-i} \setminus A'_{-i})$, $\lambda_i^1 \in \Lambda_i(A'_{-i})$ and $\mu \in [0, 1]$. Without loss of generality we can choose a conjecture λ_i^1 that makes a_i^* more attractive, i.e.:

$$(5) \quad \lambda_i^1(a_{-i} | a_i) = \mathbb{1}(a_i = a_i^*, a_{-i} = \bar{a}_{-i}) + \mathbb{1}(a_i \neq a_i^*, a_{-i} = \underline{a}_{-i}(a'_i))$$

with $\bar{a}_{-i} \in \arg \max_{a_{-i} \in A'_{-i}} u_i(a_i^*, a_{-i})$ and $\underline{a}_{-i}(a_i) \in \arg \min_{a_{-i} \in A'_{-i}} u_i(a_i, a_{-i})$. Hence we have $a_i^* \in \text{FR}_i(A')$ if and only if:

$$(6) \quad (1 - \mu) \max_{a_{-i} \in A'_{-i}} \left\{ u_i(a_i^*, a_{-i}) \right\} + \sum_{a_{-i} \in A_{-i} \setminus A'_{-i}} \mu \lambda_i^0(a_{-i}) u_i(a_i^*, a_{-i}) \\ \geq (1 - \mu) \min_{a_{-i} \in A'_{-i}} \left\{ u_i(a'_i, a_{-i}) \right\} + \sum_{a_{-i} \in A_{-i} \setminus A'_{-i}} \mu \lambda_i^0(a_{-i}) u_i(a'_i, a_{-i})$$

for every $a'_i \in A_{-i}$. That is, if and only if it is a best response to some independent belief in the game $(I, \tilde{A}, \tilde{u})$ with $\tilde{A}_i = A_i$, $\tilde{A}_{-i} = (A_{-i} \setminus A'_{-i}) \cup \{a_{-i}^0\}$ and $\tilde{u} : \tilde{A} \rightarrow \mathbb{R}$ given by:

$$(7) \quad \tilde{u}_i(a_i, a_{-i}) = \begin{cases} u_i(a_i, a_{-i}) & \text{if } a_{-i} \in A_{-i} \setminus A'_{-i} \\ \max_{a_{-i} \in A'_{-i}} u_i(a_i^*, a_{-i}) & \text{if } a_i = a_i^* \wedge a_{-i} = a_{-i}^0 \\ \min_{a_{-i} \in A'_{-i}} u_i(a_i, a_{-i}) & \text{if } a_i \neq a_i^* \wedge a_{-i} = a_{-i}^0 \end{cases}$$

The result then follows from the well known equivalence never between best responses and dominated actions, cf. Lemma 3 in Pearce (1984). \blacksquare

B.3 SEQUENTIAL IMPLEMENTATION

Proof of Proposition 2. Monotonicity of A^n is shown by induction. By definition $A^2 \subseteq A^1 = A$. Now suppose that $A^{n+1} \subseteq A^n$. Since IE is a monotone correspondence, then so is T. This implies that $A^{n+2} = T(A^{n+1}) \subseteq T(A^n) = A^{n+1}$. Hence by the induction principle it follows that $A^{n+1} \subseteq A^n$ for all $n \in \mathbb{N}$. Since \mathcal{A} is finite and $\{A^n\}$ is a monotone sequence in \mathcal{A} , we know that it converges in finite time to a limit, hence A^{IE} is well defined.

Since the convergence occurs in finite time, there exists some m such that $A^{\text{IE}} = A^m = T(A^m)$, which implies that A^{IE} is nonempty and is indeed a fixed point of T. Also, since IE is nonempty-valued, this implies that A^{IE} is nonempty. Finally, let $A' \in \mathcal{A}$ be such that $A' \subseteq T(A')$. By definition $A' \subseteq A = A^1$. Now suppose that $A' \subseteq A^n$ for some $n \in \mathbb{N}$. Then, by monotonicity of T, $A' \subseteq T(A') \subseteq T(A^n) = A^{n+1}$. Hence, by the induction principle, $A' \subseteq A^n$ for all $n \in \mathbb{N}$. Consequently, $A' \subseteq A^{\text{IE}} = \bigcap_{n \in \mathbb{N}} A^n$. \blacksquare

Proof of Theorem 3. By Proposition 2, we know that $A^{\text{IE}} = T(A^{\text{IE}})$. The fact that $A^{\text{IE}} \subseteq T(A^{\text{IE}})$ implies that $\text{supp}(\alpha) \subseteq A^{\text{IE}} \subseteq A^{\text{IE}} \cup A^{\text{R}}$ for every $\alpha \in \text{IE}(A^{\text{IE}} \cup A^{\text{R}})$. Therefore $\text{IE}(A^{\text{IE}} \cup A^{\text{R}})$ is characterized by a finite set of affine inequalities, and is thus a convex set. The fact that $A^{\text{IE}} \supseteq T(A^{\text{IE}})$ implies that there exists some $\alpha^* \in \text{IE}(A^{\text{IE}} \cup A^{\text{R}})$ such that $\alpha_i^*(a_i) > 0$ for every i and every $a_i \in A_i^{\text{IE}}$.

Consider the mediated game \mathcal{G}^* which implements α^* using $A^{\text{IE}} \cup A^{\text{R}}$ as the set of additional credible threats. For every $a_i \in A_i^{\text{R}} \setminus A_i^{\text{IE}}$, let $\lambda_i \in \Delta(A_{-i})$ be a degenerate conjecture for which a_i is a best response. Every information set in which player i is asked to use a_i occurs off the equilibrium path, and i may believe that the most likely tremble leading to it corresponds to $-i$ choosing according to λ_i . Hence choosing a_i is indeed sequentially rational. For $a_i \in A_i^{\text{IE}} \setminus A_i^{\text{R}}$, the information set in which i is asked to choose a_i is along the equilibrium path. Therefore, by the definition of IE, choosing a_i is optimal. Hence following recommendations constitutes a sequential equilibrium.

Now consider any $\alpha \in \text{IE}(A^{\text{IE}} \cup A^{\text{R}})$ and the corresponding mediated game \mathcal{G} . Let $\hat{\mathcal{G}}$ be the extensive form mechanism in which (i) the mediator randomizes between \mathcal{G} and \mathcal{G}^* with

probabilities $(1 - \varepsilon)$ and ε respectively; and (ii) players are only informed about recommendations, in particular they are never told whether they are in \mathcal{G} or \mathcal{G}^* . Now suppose that all players agree that trembles in \mathcal{G}^* are more likely than trembles in \mathcal{G} . This means that, whenever they are asked to perform an action in $A^{\text{IE}} \cup A^{\text{R}}$, they will believe either that they are along the equilibrium path or that they are in \mathcal{G}^* . From the previous analysis it follows that complying remains to be sequentially optimal. And hence the distribution $m\hat{\alpha}c = (1 - \varepsilon)\alpha + \varepsilon\alpha^*$ is sequentially implementable. Notice that $\hat{\alpha}$ approximates α as ε approaches 0. Finally, if one allows for $\varepsilon = 0$ (meaning that the mediator can make mistakes), then $\hat{\alpha} = \alpha$ and thus α is sequentially implementable. ■

Proof of Proposition 6. Let $A_i = \{a_i, b_i\}$ for $i \in I$. If the game has no dominated strategies, then it has a completely mixed NE, which is also a proper equilibrium of the simultaneous move game. Hence every action is properly implementable and the result follows from Theorem 3. Now suppose that some player i has an absolutely dominated strategy, say b_i , and let a_{-i} be the unique best response to a_i . Then (a_i, a_{-i}) is the unique CE with respect to A^{CR} , and it is a proper equilibrium of the simultaneous move game. In the remaining cases no player has absolutely dominated strategies, but at least one player has a dominated strategy.

First suppose that player 2 has no dominated strategies but b_1 is dominated by a_1 . Without loss of generality assume that a_2 is a best response to a_1 . Since the game has no repeated payoffs and player 2 has no dominated strategies, this implies that (a_1, a_2) is a strict NE and that b_2 is the unique best response to b_1 . Since b_1 is dominated but not absolutely dominated, then it is a best response either to λ_1 or λ'_1 with:

$$\begin{aligned} \lambda_1(b_2|b_1) = 1 \quad \wedge \quad \lambda_1(a_2|a_1) = 1 \\ \lambda'_1(a_2|b_1) = 1 \quad \wedge \quad \lambda'_1(b_2|a_1) = 1 \end{aligned}$$

In the first case it suffices to have player 1 move first. By backward induction he knows that player 2 will choose a_2 if he chooses a_1 and b_2 if he chooses b_1 . Hence, choosing b_1 is the best response. Since the equilibrium is proper and the outcome is (b_1, b_2) , it follows that every every action is properly implementable and the result follows from Theorem 3.




FIGURE 9 Implementation of b_1 when it is the only dominated action.

If b_1 is a best response to λ'_1 , then it can be implemented as an equilibrium of the mechanism in Figure (9), when $\varepsilon > 0$ is small enough. The equilibrium strategies are represented with red arrows. Player 1 is willing to choose a_1 because it is a best response to a_2 . Player 1 is willing to choose b_1 because his conjectures at that moment are close enough to λ'_1 and, since there are no repeated payoffs, b_1 is a strict best response to λ'_1 . Player 2 is willing to choose b_2 because it is a best response to b_1 . He is willing to choose a_2 when ε is sufficiently close to 0, because (a_1, a_2) is a *strict* NE. Since all the information sets are on the equilibrium path, the equilibrium is proper. Hence every every action is properly implementable and, the result follow from Theorem 3.

The only remaining case to consider is when both players have dominated strategies, say b_1 and b_2 . In this case (a_1, a_2) is a strict NE and there are two possibilities. Define λ_1 and λ'_1 as before and the corresponding conjectures for player 2:

$$\begin{aligned} \lambda_2(b_1|b_2) = 1 & \quad \wedge \quad \lambda_2(a_1|a_2) = 1 \\ \lambda'_2(a_1|b_2) = 1 & \quad \wedge \quad \lambda'_2(b_1|a_2) = 1 \end{aligned}$$




FIGURE 10 Implementation of b_1 when b_2 is also dominated.

If b_i is a best response to λ_i , then it can be implemented as a NE of the game where i moves first and $-i$ chooses b_{-i} along the equilibrium path and punishes deviations with a_{-i} . If b_i is a best response to λ'_i , then it can be implemented as a NE of the mechanism depicted in Figure 10. These NE fail to be subgame perfect, however, consider an extended mechanism in which Nature randomizes between the mechanism that implements b_1 and b_2 . In this mechanism, all actions are played with positive probability along the equilibrium path. Hence we can connect the punishment nodes with the equilibrium path as to include them in the same information sets, and the equilibrium becomes sequential. Hence every every action is properly implementable and, as before, the result follow from Theorem 3. ■

B.4 PB IMPLEMENTATION

In this appendix we provide a proof for Proposition 7 regarding the necessary and sufficient conditions for PB implementation. The proof is divided in two parts regarding necessity and sufficiency respectively.

To establish necessity it suffices to show that given a PBE of an extensive form mechanism, every action played with positive probability, either on or off the equilibrium path, is in A^{FR} .

Then the proof of Theorem 1 applies simply replacing $\underline{w}_i(a'_i, A)$ with $\underline{w}_i(a'_i, A^{\text{FR}})$. This fact is established by the following lemma. Given an extensive form mechanism and an equilibrium σ^* , let $A_i^* \subseteq A_i$ denote the set of actions that i plays with positive probability in some information set, i.e.:

$$A_i^* = \left\{ a_i \in A_i \mid \left(\exists H_i \in \mathcal{H}_i \right) \left(\exists s_i \in S_i \right) \left(\sigma_i^*(s_i) > 0 \wedge s_i(H_i) \in M^{a_i}(H_i) \right) \right\}$$

LEMMA 11 *For every PBE of an extensive form mechanism we have $A^* \subseteq A^{\text{FR}}$.*

Proof. Fix some action $a_i^* \in A_i^*$ that is chosen with positive probability in some information set $H_i \in \mathcal{H}_i$. Fix some move $m^{a_i^*} \in M^{a_i^*}(H_i)$ that represents a_i^* and is chosen with positive probability. For every possible deviation $a'_i \in A_i \setminus \{a_i^*\}$ pick any move $m^{a'_i} \in M^{a'_i}(H_i)$ representing a'_i . Now let $\mu = \psi_i(\Phi^{H_i} \mid H_i, m_i^*) \in [0, 1]$, where Φ^{H_i} is the event that τ_{-i} is already determined at H_i , as defined in equation (3). Finally, let $\lambda_i^0 \in \Delta(A_{-i})$ and $\lambda_i^1 \in \Delta(A^*)$ be the conjecture given by:

$$\lambda_i^0(a_{-i}) = \zeta_i(\tau_{-i}^{-1}(a_{-i}) \mid H_i, \Phi^{H_i}) \quad \lambda_i^1(a_{-i} \mid a_i) = \zeta_i(\tau_{-i}^{-1}(a_{-i}) \mid H_i, m^{a_i}, \bar{\Phi}^{H_i})$$

and let $\lambda_i = \mu \lambda_i^0 + (1 - \mu) \lambda_i^1$. Sequential rationality together with the fact that $\zeta_i(\Phi^{H_i} \mid H_i, m)$ and $\zeta_i(\tau_{-i}^{-1}(a_{-i}) \mid H_i, m, \Phi^{H_i})$ are independent from m , imply that for every deviation a'_i :

$$\begin{aligned} \sum_{a_{-i} \in A_{-i}} \lambda_i(a_{-i} \mid a_i^*) u_i(a_i^*, a_{-i}) &= \sum_{a_{-i} \in A_{-i}} \zeta_i(\tau_{-i}^{-1}(a_{-i}) \mid H_i, m^{a_i^*}) u_i(a_i^*, a_{-i}) \\ &\geq \sum_{a_{-i} \in A_{-i}} \zeta_i(\tau_{-i}^{-1}(a_{-i}) \mid H_i, m^{a'_i}) u_i(a'_i, a_{-i}) \\ &= \sum_{a_{-i} \in A_{-i}} \lambda_i(a_{-i} \mid a'_i) u_i(a_{-i}, a_{-i}) \end{aligned}$$

This means that a_i^* is a best response to the conjecture λ_i , and therefore $a_i^* \in \text{FR}(A^*)$. Since this is true for all i and all $a_i^* \in A_i^*$, it follows that $A^* \subseteq \text{FR}(A^*)$ and thus $A^* \subseteq A^{\text{FR}}$. \blacksquare

Proof of sufficiency for Proposition 7. Let α be a CE with respect to A^{FR} . We will construct an extensive form mechanism (\mathcal{G}, τ) and a PBE (σ, ψ) that implement α . We start with the mediated game \mathcal{G}^0 that implements α as a NE, and players choose only actions in A^{FR} . Then we will add new off-path histories that will make the off path-punishments be incentive compatible. Since we are only adding nodes off the equilibrium path, it is straightforward that the strategies constitute a NE that induces α . We only have to ensure that the choices that are made off the equilibrium path are sequentially rational and that the beliefs is off-equilibrium information sets are consistent.

In our construction, all the players' information sets are pivotal, have a unique pivotal move representing each move in the environment, and all the moves in each pivotal information set are pivotal, i.e. $M(H_i) = \cup_{a_i \in A_i} M^{\tau, a_i}(H_i)$ and $\#M(H_i) = \#A_i$ for every $H_i \in \mathcal{H}_i$. We specify the equilibrium strategies by labelling each information set with the distribution of actions that the corresponding player is supposed to follow. For instance $H_i^{a_i}$ represents a pivotal information set in which i is supposed to choose the only available pivotal move in $M^{a_i}(H_i^{a_i})$. We then "join" those information sets in which the same player chooses the same action, so that there are no different information sets with the same label.¹⁵ This implies that the only information that a player has at the moment of making his choice is the action that he is supposed to choose.

¹⁵This is possible because each player only makes one pivotal decision along each path, and all available moves are pivotal. This implies that each player has one and only one decision node along each path and perfect recall is trivially satisfied.

Fix a player i and some action $a_i \in A_i^{\text{FR}} \setminus \text{supp}(\alpha_i)$. Since A^{FR} is self FC-rationalizable, we know that there exist some conjecture $\lambda_i \in \Lambda_i$ for which a_i^0 is a best response, and such that can be written as $\lambda_i = (1 - \theta)\lambda_i^0 + \theta\lambda_i^3$ for some $\theta \in [0, 1]$, some $\lambda_i^0 \in \Delta(A_{-i})$ and some $\lambda_i^3 \in \Lambda_i(A^{\text{FR}})$. We can further decompose $(1 - \theta)\lambda_i^0$ as $(1 - \theta)\lambda_i^0 = \mu\lambda_i^1 + \eta\lambda_i^2$ with $\mu, \eta \in [0, 1]$, $\lambda_i^1 \in \Delta(A_{-i} \setminus A_{-i}^{\text{FR}})$ and $\lambda_i^2 \in \Delta(A_{-i}^{\text{FR}})$. Without loss of generality we can assume that $\lambda_i^3(\bar{a}_{-i}|a_i^0) = 1$ and $\lambda_i^3(\underline{a}_{-i}|a_i')$ for every $a_i' \neq a_i^0$ with $\bar{a}_{-i} \in \arg \max_{a_{-i} \in A_{-i}^{\text{FR}}} \{u_i(a_i^0, a_{-i})\}$ and $\underline{a}_{-i}(a_i') \in \arg \min_{a_{-i} \in A_{-i}^{\text{FR}}} \{u_i(a_i', a_{-i})\}$.

The entire mechanism starts from an initial node where Nature chooses between the equilibrium path or other paths. For each action $a_i \in \text{FR}_i^\infty \setminus \text{supp}(\alpha_i)$ we will construct a set of paths on which player i is willing to choose such action and believe that the future choices of his opponents will be restricted to FR_{-i}^∞ . These sets are all generic and they correspond to the ones depicted in Figure 11.

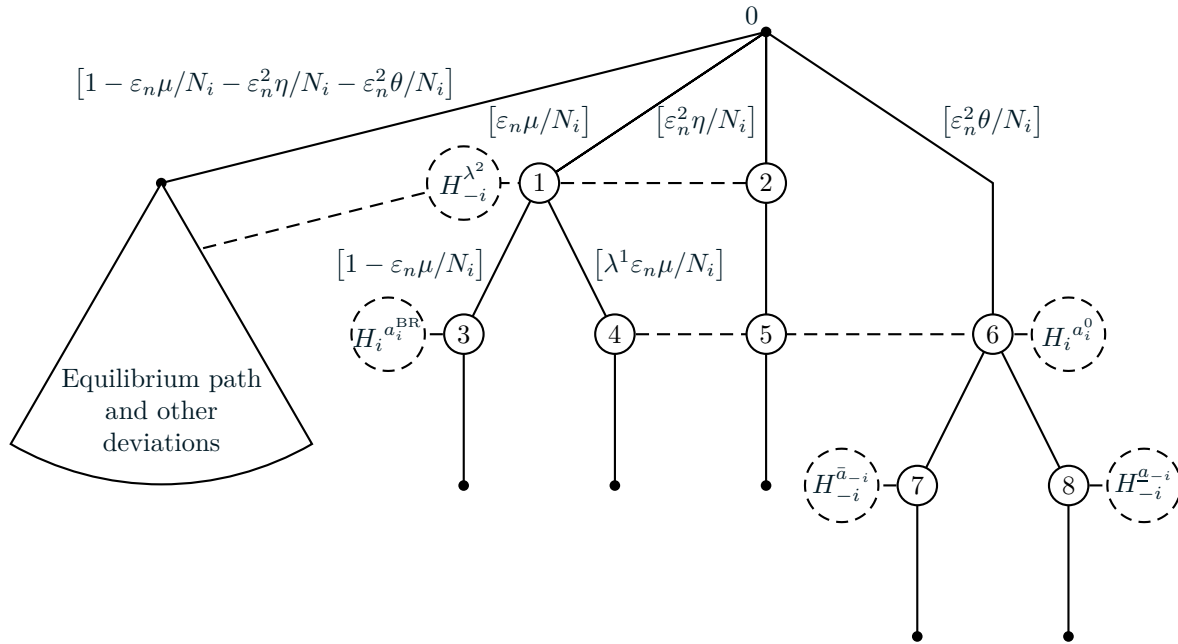


FIGURE 11 Incentives for $a_i^0 \in A_i^{\text{FR}} \setminus A_i^*$

Within brackets, we have specified the sequence of mixed strategies that converges to the equilibrium beliefs, given a sequence of numbers $\epsilon_n \in (0, 1)$ converging to 0 and the number of FC-rationalizable actions $N_i = \#A_i^{\text{FR}}$.¹⁶ The limit of this sequence will generate weakly consistent beliefs. Hence it only remains to verify the incentive constraints:

- At nodes (1) and (2), player $-i$ is willing to make choices according to λ_i^2 because he believes that he is on the equilibrium path (where all the recommendations are incentive compatible by construction).
- At nodes (7) and (8) it might be the case that \bar{a}_{-i} and \underline{a}_{-i} are not best responses to a_i^0 or a_i' . However, they are in FR_{-i}^∞ and thus, either they are chosen along the equilibrium path or they are chosen in other off-the equilibrium paths that are designed to generate

¹⁶This sequence is not strictly mixed, there are still a number of actions that have null probability, but we are not interested in them. We could obtain strictly mixed sequences by assigning sufficiently small probabilities to them (of order ϵ^3 or lower) but this would only complicate the exposition unnecessarily.

incentives for them. Since $-i$ will consider the deviations in this path to be unlikely (of order at most ε^3), the incentives for these actions are independent from what happens in this figure.

- c) First suppose that the information sets for i are fully contained in the figure:
- At (3) player i is supposed to choose an action which is a best response to λ_i^2 (which always exists by finiteness of the game. Hence this choice is trivially incentive compatible.
 - It is straight forward to see that equilibrium beliefs for player i would generate a conjecture λ_i at H_i^0 and thus he would be willing to choose a_i^0 .
- d) Now suppose that either H_i^0 or H_i^{BR} appear in other parts of the game. There are only two possibilities:
- They appear as punishments to deviations from the equilibrium path or in some analogous block corresponding to some a_{-i}^0 in the position analogous to (7) or (8). Since all these possibilities have probability of order ε^3 or lower, whatever happens there is irrelevant for player i .
 - They can appear in the equilibrium path, or in another block for some other a'_i in the positions of (3) - (8). In any such case, the corresponding action will also be a best response to the conditional beliefs and thus to the average beliefs.

The mechanics behind the argument are as follows. Every action $a_i^0 \in \text{FR}_i^\infty$ can be rationalized by some combination of equilibrium or arbitrary previous choices and future choices in FR_{-i}^∞ . Whenever i is asked to choose a_i^0 we will be naive and believe that he is in exactly that situation in which choosing a_i^0 is in his best interest. Since weak consistency does not imply any consistency requirements *across* players, this can always be done even if it implies that *i must be certain that his opponent is or will be mistaken.* ■

0///