Sequential Information Design*

Laura Doval † Jeffrey C. Ely ‡

April 29, 2020

Abstract

We study games of incomplete information as both the information structure and the extensive form vary. An analyst may know the payoff-relevant data but not the players' private information, nor the extensive form that governs their play. Alternatively, a designer may be able to build a mechanism from these ingredients. We characterize all outcomes that can arise in an equilibrium of some extensive form with some information structure. We show how to specialize our main concept to capture the additional restrictions implied by extensive-form refinements.

*The latest version of this paper can be found here. This research is supported by grants from the National Science Foundation. The second author thanks the Einaudi Institute of Economics and Finance (EIEF) for its hospitality. We thank the Editor Joel Sobel and four anonmyous referees for feedback that has greatly improved this paper. We also thank Ricardo Alonso, Andreas Blume, Gabriel Carroll, Marciano Siniscalchi, George Mailath, Ludovic Renou, Jesse Shapiro, and Ali Hortaçsu. Wade Hann-Caruthers was an outstanding research assistant. When revising this paper in December 2019, we discovered that the name of our concept, *coordinated equilibrium* was also previously used in Salcedo (2012), an early version of Salcedo (2017).

†Division of the Humanities and Social Sciences, California Institute of Technology. ldoval@caltech.edu.

‡Department of Economics, Northwestern University. jeffely@northwestern.edu.

1 Introduction

To summarize the possible outcomes of a strategic interaction, after having specified the players' actions and payoffs, an analyst makes assumptions about their information and the extensive form that ties all of these elements together. Conventional equilibrium concepts begin by fixing both and proceed to provide a set of solutions for that fully specified game. We are interested instead in fixing only the *base game* (the actions and payoffs) and summarizing all of the equilibrium outcomes consistent with some specification of the information structure and extensive form. For example in applied modeling the analyst may be unable to observe these details and/or be agnostic about them. She may not like her solutions to rely on any particular assumption.

We develop a framework based on variations of an umbrella solution concept we term *coordinated equilibrium* which, for any given base game identifies the full range of equilibrium outcomes associated with all such assumptions. We show further how to impose restrictions on coordinated equilibria in order to capture outcomes associated with stronger and stronger extensive-form refinements.

For complete-information environments, Aumann (1974) proposed the concept of correlated equilibrium. The set of correlated equilibrium outcomes of a base game characterizes the Nash equilibrium outcomes of all specifications satisfying two restrictions: the payoffs are common knowledge and the players move simultaneously. Bergemann and Morris (2016) found the appropriate generalization, Bayes' correlated equilibrium, that does the same for games with incomplete information. That is, across all specifications in which the restriction to simultaneous moves is maintained but the information structure concerning payoffs is unrestricted, an outcome can arise as a Bayesian Nash Equilibrium if and only if it belongs to the set of Bayes' correlated equilibria of the base game.

Returning to the domain of complete-information environments, the path-breaking work of Salcedo (2017) explores variations of the extensive form beyond simultaneous moves. Salcedo (2017) proposes the Interdependent Choice Equilibrium (ICE) concept to delineate Nash equilibrium outcomes of complete-information extensive-form elaborations of a given base game.

We consider both incomplete information and non-trivial extensive forms, bringing forth new challenges as well as applications. Solution concepts for extensive-form games span a range of refinements embodying various notions of sequential rationality and consistency of beliefs. The gap between unrefined Nash equilibrium and the strongest versions of perfection are particularly significant in the presence of incomplete information; indeed, this is the domain in which most of the theoretical discussion of refinements has taken place. We show how to progressively specialize our baseline coordinated equilibrium concept in order to characterize the equilibrium sets across a range of solution concepts including Bayesian Nash Equilibrium, Perfect Bayesian Equilibrium and a strong refinement of Sequential Equilibrium.

Our work opens the door to the study of the *design* of fully dynamic extensive-form asymmetric information structures. As with all of the aforementioned solution concepts, an equivalent dual perspective identifies coordinated equilibrium and its variations as the feasible set of outcomes that can be implemented by a designer who, taking the base game as given, builds an extensive-form mechanism to coordinate play toward some objective. For example, to raise funds for a public project, a principal can decide on the sequence in which she approaches investors as well as the information disclosed to each about the progress so far. An agenda setter for a deliberative committee who structures the pattern of discussion and voting controls both the way in which preferences are aggregated and the information about preferences that gets revealed along the way. A mediator engaging in "shuttle diplomacy" chooses which parties to negotiate with first and what details about past agreements to disclose in subsequent negotiations.

Finally, coordinated equilibrium is expressed as a family of linear incentive constraints derived completely from the base game. That is, like correlated equilibrium, Bayes' correlated equilibrium and ICE, coordinated equilibria can be computed as the solution to a linear program. We compute the solution sets for a variety of illustrative examples.

Related Literature As in the literature on communication equilibria (e.g., Aumann (1974, 1987), Myerson (1982), Forges (1986), Myerson (1986)) we study a solution concept, coordinated equilibrium, and the canonical device that implements it. The literature on information design, including Bayesian Persuasion (Kamenica and Gentzkow (2011)), Bayes' correlated equilibrium (Bergemann and Morris (2016)) and sequential Bayes' correlated equilibrium (Makris et al. (2018)), takes a similar approach. These papers provide the appropriate extensions of correlated equilibrium and sequential communication equilibrium to the case in which the analyst is agnostic about the players' information structure. de Oliveira and Lamba (2019) apply similar ideas to a single-agent dynamic choice problem to study rationalizable decision paths. All these papers take the extensive form as given and study equilibrium outcomes as the information structure varies.

Salcedo (2017) on the other hand takes as given a complete information base game and considers the set of equilibria of extensive-form games consistent with it. We discuss Salcedo (2017) in further detail in Section 7. Nishihara (1997) shows that by designing the extensive form it is possible to achieve cooperation in the prisoners' dilemma even if it is played only once. Gershkov and Szentes (2009) consider communication protocols whose outcomes can be achieved as Bayes' Nash equilibria in a canonical extensive form. They do this in the context of a sequential information acquisition game, where players need to disclose the information acquired to the social planner.[1] In a similar vein, Kremer et al. (2014) use position uncertainty to motivate agents to experiment with an alternative of unknown value. More recently, Best and Quigley (2017) use the uncertainty over the extensive form to motivate a sender who cannot commit to an information structure, while Gallice and Monzon (2018) consider the benefits of using a sequential move protocol to motivate agents to contribute to a public good. Finally, Peters (2015) studies games of competing mechanisms from a similar perspective. To capture competing mechanisms, the extensive forms in Peters (2015) allow players to deviate without being detected by their opponents.

Sutton (1991, 2001, 2003) discusses the consequences of the assumptions the analyst makes on game forms for predictions of interest to industrial organization. He proposes what he calls the *bounds approach* (see Shaked and Sutton (1987)): instead of making a point prediction based on a particular game, the analyst should strive to make predictions that are robust to a range of model specifications. Our work echoes the latter approach. Indeed, as the main example in Section 2 shows, an analyst's prediction in a standard quantity competition game are sensitive to the assumptions she makes about the extensive form: while Cournot competition has a unique correlated equilibrium, the set of coordinated equilibria is quite large.

Organization The rest of the paper is organized as follows. The next section contains an informal overview of our results through a series of examples. In Section 3 we introduce our basic coordinated equilibrium concept. Section 4 and Section 5 show how to specialize the concept in order to characterize extensive-form refinements from Perfect Bayesian Equilibrium in Section 5 and a strong refinement of sequential equilibrium in Section 4. In Section 6 we present our converse results. Further examples and proofs are in the Appendices.

[1] Bognar et al. (2015) also highlight the benefit of having agents move sequentially, as opposed to simultaneously, in a voting game.

2 Overview

This section is a high-level overview of the results in the paper and their relationship to the existing literature on information design. Readers wishing to go straight to the formal developments may want to skip this section.

Setup A base game G is described by the set of players $i \in \{1, \ldots, N\}$, each player's finite set of available actions A_i, a finite set of payoff-relevant states of the world Θ, and each player's payoff function $u_i : A \times \Theta \to \mathbf{R}$ where $A = \prod_i A_i$ is the set of action profiles. In some applications we may also specify a (common) prior probability distribution over states $\rho \in \Delta(\Theta)$.

A base game with two players and two states of the world is depicted below in Figure 1a. In this base game player 1 chooses from two actions U and D, player 2 chooses from three actions L, M, and R and there are two states of the world θ and θ'.

	L	M	R
U	$2,2$	$-1,-4$	$-1,3$
D	$3,0$	$0,-1$	$4,-1$

θ

	L	M	R
U	$2,2$	$-1,3$	$-1,-4$
D	$3,0$	$4,-1$	$0,-1$

θ'

Figure (a) Base Game

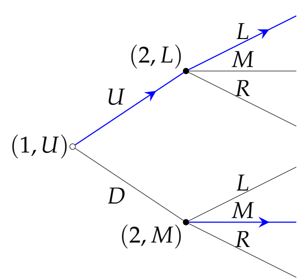

Figure (b) Example of a Plan

Figure 1: A base game and a plan

Plans, Obedience, Coordinated Equilibrium We develop a framework to derive the set of equilibrium outcomes under various solution concepts. The basic element of our framework is the concept of a *plan*. A plan is a tree that represents a fully contingent path of action recommendations for the players. An example of a plan is depicted in Figure 1b. In this plan player 1 moves first and is recommended to play U. Player 2 is recommended to play L if 1 played U but to play M if 1 instead played D.

The set of all possible plans is P. The *obedient path* of a plan is the path that results when all players play according to the plan. The obedient path in Figure 1b leads to the action profile (U, L).

A *coordinated equilibrium* is a joint probability distribution $\pi \in \Delta(\Theta \times P)$ over states and plans which is obedient: conditional on having selected from π a plan whose obedient path calls on Player i to play action a_i, the choice of a_i in fact maximizes i's conditional expected payoff under the assumption that all other players will also be obedient. To express these obedience constraints, let $\langle a_i \rangle$ be the subset of plans whose obedient path has i play action a_i. We denote by $u_i(b_i, p, \theta)$ the payoff to player i from (possibly *dis*-obediently) choosing b_i within plan p under the assumption that all other players are obedient. Then

$$\sum_{\theta \in \Theta} \pi(\theta) \sum_{p \in \langle a_i \rangle} \pi(p \mid \theta) \left[u_i(a_i, p, \theta) - u_i(b_i, p, \theta) \right] \geq 0 \qquad (1)$$

is the constraint that player i finds it in her interest to obediently select a_i rather than b_i. If these obedience constraints hold for all players i and actions a_i, b_i, then π is a coordinated equilibrium.

Specializing further, if we are also given a prior ρ for the base game we may add the additional constraint that the coordinated equilibrium π be consistent with ρ, i.e., the marginal of π on states agrees with ρ.

A basic result is that the *outcome* of a coordinated equilibrium, i.e. its induced joint distribution over states and action profiles, is a (Bayesian) Nash equilibrium of an appropriately chosen extensive-form game.[2] Nash equilibrium is not a suitable solution concept for extensive-form games, especially under incomplete information, so instead our main results show how to specialize the definition of coordinated equilibrium in order to characterize various extensive-form refinements.

Cournot We use the example of Cournot duopoly to illustrate. Suppose that price is a function $P(Q \mid \theta)$ of total output $Q = q_1 + q_2$ (downward sloping) and a demand shock $\theta \in \Theta$. There is a given probability distribution $\rho(\theta)$ over realizations of the shock. The firms have constant marginal costs $c > 0$ and supply quantities q_i (taken from some large finite set) to maximize expected profits $q_i (P(Q \mid \theta) - c)$. Let $\bar{Q}^\theta$ be the zero-profit industry output in state θ, i.e.

[2] Indeed if we restrict attention to complete-information games, where Θ is a singleton, the set of coordinated equilibrium outcomes coincides with the set of ICE outcomes from Salcedo (2017).

$P\left(\bar{Q}^{\theta} \mid \theta\right) = c$.

For any value of the demand shock θ, consider the plan in which firm 1 moves first and produces zero output $q_1 = 0$, firm 2 responds to $q_1 = 0$ by producing the monopoly output, call it $q_2 = Q^{M,\theta}$, and firm 2 responds to any other choice $\hat{q}_1$ of firm 1 by flooding the market, i.e. $q_2 = \max\left\{\bar{Q}^{\theta} - \hat{q}_1, 0\right\}$. In this plan, call it $p(\theta)$, any choice of Firm 1 would earn zero profits and therefore Firm 1 is willing to be obedient and produce zero. Along the obedient path Firm 2 earns the monopoly profit, a best-response and therefore also obedient. Consider then the coordinated equilibrium π whose marginal $\pi(\theta)$ agrees with ρ and for which $\pi(p(\theta) \mid \theta) = 1$ for every demand shock θ.

It is indeed a coordinated equilibrium. It assigns probability 1 to plans in which Firm 1 optimally chooses $q_1 = 0$ and Firm 2 plays a best-response. However it fails to capture the requirements that would come from sequential rationality in an extensive form. Indeed it is merely an elaborate version of the textbook entry deterrence equilibrium used to illustrate the shortcomings of Nash equilibrium. Firm 1 is obedient only because she anticipates that Firm 2 would impose zero profits on herself in order to punish a deviation.

Sequential Rationality and Extensive-Form Refinements Therefore, in order to capture outcomes that are consistent with sequential rationality, we refine the coordinated equilibrium concept. In particular we consider coordinated equilibria that use only certain subsets of plans. Consider a family B of non-empty action subsets $B_1, \ldots, B_N$ and let P^B be the subset of plans that only ever recommend actions in some subset B.

We are interested in results of the following form: If a coordinated equilibrium π assigns probability 1 to plans in P^B then there exists an extensive form in which its outcome can arise as an equilibrium satisfying solution concept Y.

For solution concepts Y requiring sequential rationality, such as Perfect Bayesian Equilibrium or Sequential Equilibrium, we must restrict the set of plans, i.e., $P^B \subsetneq P$. The idea is that some action choices cannot be made sequentially rational and we want to restrict attention to plans that never make use of such actions.

Self-Contained Coordinated Equilibria We demonstrate one such result here. Consider the set of actions C defined by the following iterative procedure. First, we eliminate from consideration actions that can never be played in a coordinated equilibrium. For each player i, let C_i^1 be the set of i's actions that can arise with positive probability in some coordinated equilibrium. Next, let C_i^2 be the set of actions that can arise with positive probability in some coordinated equilibrium assigning probability 1 to plans in P^{C^1}, i.e. those plans that only use actions that survived the previous step of elimination. Continuing in this way we arrive at a fixed point consisting of nonempty subsets of actions C_i for each player.

Consider a coordinated equilibrium that assigns probability 1 to plans in P^C. It only involves actions that can arise in coordinated equilibria involving actions that can arise in coordinated equilibria ... etc. For that reason, we call these *self-contained* coordinated equilibria. We show in Theorem 2 that for every self-contained coordinated equilibrium there is an extensive form with a sequential equilibrium which yields (essentially) the same outcome.

We return to Cournot duopoly to illustrate. For any value of the demand shock θ pick any quantity $q_1 \in [0, \bar{Q}^\theta]$ and consider the following plan. Firm 1 produces q_1. If firm 1 is obedient firm 2 produces the best-response to q_1, and if firm 1 produces any quantity $\hat{q}_1 \neq q_1$ then firm 2 produces $q_2 = \max\left\{\bar{Q}^\theta - \hat{q}_1, 0\right\}$. This plan, together with the state θ is a *degenerate* coordinated equilibrium, i.e. the measure assigning probability 1 to θ and the plan just described satisfies the obedience constraints in Equation 1.

Thus, for firm 1 every quantity $q_1 \in [0, \bar{Q}]$ where $\bar{Q} = \max_\theta \bar{Q}^\theta$ survives the first round of elimination and likewise for firm 2. Any other quantity results in negative profits and therefore cannot be played in a coordinated equilibrium. Those quantities are eliminated. It follows that $C_i^1 = [0, \bar{Q}]$.

Moreover the degenerate coordinated equilibrium plans only used actions that themselves belong to $[0, \bar{Q}]$. Therefore the elimination process reaches its fixed point after the first step, $C_i = [0, \bar{Q}]$. The self-contained coordinated equilibria are those π assigning probability 1 to plans using these quantities. We show how to construct an extensive-form game possessing a sequential equilibrium whose outcome is arbitrarily close to that of any such π.

A Canonical Extensive Form For Cournot Duopoly Consider the following extensive form. Nature moves first and chooses an element of $\Theta \times P^C$. Following each possible move by Nature (θ, p), we complete the tree by appending the tree from p. Finally, for each player i and action a_i we construct an information set which joins all of the nodes at which, in the plan to which the node belongs, i is called upon to play a_i.

Nature's choice probabilities are given by the following mixture. With probability $(1 - \varepsilon)$ Nature's choice is a draw from the target self-contained coordinated equilibrium distribution π. The remaining probability $\varepsilon > 0$ is spread uniformly over the degenerate coordinated equilibria. Note that this distribution, call it ψ is a mixture of coordinated equilibria and is therefore itself a coordinated equilibrium, due to the linearity of the obedience constraints.

This extensive form can be interpreted as a random, contingent recommendation system. Each information set for player i is associated with a quantity $q_i \in C_i$ and we can interpret it as the event that player i has been recommended to produce quantity q_i. The strategy for player i in which she produces the recommended quantity at each of her information sets is what we call the *obedient* strategy.

When both firms play their obedient strategies, then at firm i's information set associated with quantity q_i, firm i knows the following (and nothing more): Nature selected a state and plan from the coordinated equilibrium ψ, and the selected plan together with obedient play has led to a node where i is expected to play q_i. This information is exactly what is captured in the obedience constraint Equation 1 for coordinated equilibrium and therefore it is a sequential best-reply for firm i to play her obedient strategy. The obedient strategy profile is therefore sequentially rational and results in an outcome that can be made arbitrarily close (by taking ε to zero) to that of the target coordinated equilibrium π.

Self-Contained Coordinated Equilibria Satisfy All Refinements Indeed the obedient strategy profile is a sequential equilibrium. To see why notice that every information set arises with positive probability. This is because every quantity in C occurs along the obedient path of one of the degenerate self-contained coordinated equilibria receiving positive probability from Nature's move. Because there are no off-path information sets, sequential equilibrium imposes no restrictions beyond (on-path) sequential rationality. The same is true of all belief-based refinements of sequential equilibrium (for example the Intuitive Criterion). We conclude that (essentially)[3] any self-contained coordinated equilibrium outcome can arise in an extensive-form game under all conventional solution concepts.[4]

[3]The construction above assigned ε probability to arbitrary coordinated equilibria to ensure that all information sets arise with positive probability. This is necessary only for boundary coordinated equilibria, thus the "essential" qualifier.

[4]In the example of Cournot duopoly, we actually obtain that any individually rational and feasible payoff can be sustained in a coordinated equilibrium. We have already argued that there is a coordinated equilibrium where a firm obtains monopoly profits. Note that there is also a

For example, here is a self-contained coordinated equilibrium in which the firms perfectly "collude." Fix any demand shock θ and consider the 50-50 lottery over the following two plans. In the first plan firm 1 is recommended to produce half of the monopoly output. If she is obedient then firm 2 is recommended to do the same. If firm 1 produces any other output, firm 2 is recommended to flood the market. The second plan is identical to the first but with the roles of the two firms reversed.

In this lottery over plans the only recommendations that occur with positive probability are the recommendations to produce the collusive quantity. Let us check the obedience constraint. If, e.g., firm 1 were to deviate to any quantity q_1, she will earn zero with probability $1/2$ (in the first plan) and some positive profit Π with the remaining probability (in the second plan.) But Π cannot be larger than the monopoly profit and therefore the expected payoff from the deviation cannot be larger than half of the monopoly profit. Since the latter is the expected payoff from obedience, we see that the obedience constraint is satisfied. This is therefore a coordinated equilibrium and it is self-contained because all of the recommended quantities (the collusive quantities and the market-flooding quantities) belong to $C = [0, \bar{Q}]$.

Perfect Bayesian Equilibrium Self-contained coordinated equilibrium outcomes can be generated by extensive-form equilibria that do not rely on any specification of out-of-equilibrium beliefs. In order to obtain that strong property, the set C iteratively removes from consideration actions that cannot be made sequentially rational on the path of play. Nevertheless, in many games such actions can be made sequentially rational *off* the path of play. Therefore, in order to characterize less stringent refinements such as Perfect Bayesian Equilibrium, we will need to identify this larger set of actions, which we call D, and consider coordinated equilibria using the larger set of plans P^D.

We illustrate via the example in Figure 2:

coordinated equilibrium where each firm makes 0 profits. It is achieved by using the plan where the first firm produces $\overline{Q}^{\theta}$ and the second firm best replies by producing 0. Noting that the set of coordinated equilibrium payoffs is convex implies the claim.

	c_3	d_3
c_2	$0, 1, 1$	$0, 1, 0$
d_2	$0, 0, 1$	$0, 0, 0$

d_1

	c_3	d_3
c_2	$1, 1, 1$	$0, 1, 0$
d_2	$0, 0, 1$	$-1, 0, 0$

c_1

	c_3	d_3
c_2	$-2, -1, -1$	$-2, \Delta, -2$
d_2	$-2, -2, \Delta$	$-2, 0, 0$

X_1

Figure 2: A complete information base game

This is a three-player base game with complete information. Player 1 chooses the matrix while Players 2 and 3 choose the rows and columns respectively. In this game the sets C_i are the singletons $\{c_i\}$. In the language of Salcedo (2017) the action X_1 is *absolutely dominated*: the maximum payoff it can earn Player 1 is smaller than the minimum payoff from d_1. Absolutely dominated actions can never be played with positive probability in a coordinated equilibrium. Indeed iterative removal of absolutely dominated actions leads to the unique self-contained equilibrium outcome (c_1, c_2, c_3).

However, when $\Delta \leq 1$ there exists a sequential equilibrium of a suitably chosen extensive form in which Player 1 plays the action d_1, supported by the threat that a deviation to c_1 would be punished by 2 and 3 playing (d_2, d_3).

An example of such an extensive-form is presented in Appendix A. Our main result Theorem 3 ensures the existence of such an extensive-form and a Perfect Bayesian Equilibrium by identifying sets of actions D_i that survive an elimination procedure that is weaker than the procedure that defines C_i. In the game of Figure 2, when $\Delta \leq 1$ the sets D_i include the actions d_i for each player. Theorem 3 establishes that for any coordinated equilibrium using only plans in P^D, there is an extensive-form with a Perfect Bayesian Equilibrium which yields the same outcome.

The sets D_i are constructed by identifying actions that can be rationalized following deviations by one or more other players. For example in the game of Figure 2, the profile (d_2, d_3) can arise in a restricted coordinated equilibrium in which we check only the incentive constraints of players 2 and 3, taking as given that player 1 has "deviated" to X_1. This fact ensures that d_2 and d_3 belong to D. Moreover, it reveals how to specify beliefs at off-path information sets that are mutually consistent among players, satisfy Bayes' rule where possible, and which ensure the sequential rationality of off-path threats.

On the other hand, when $\Delta > 1$ the actions (d_2, d_3) can no longer be rationalized in this way and they no longer belong to our set D. We discuss this case further in Appendix A.

Sequential Information Design We have motivated our analysis by considering an external analyst who is agnostic about the extensive form and information structure. Our results can also be applied to study the *design* of extensive-form information structures to implement outcomes in a base game with incomplete information. Our last example will emphasize this perspective.

Consider the base game from Figure 1, reproduced below in Figure 3a. For illustrative purposes we will assume a prior ρ which assigns equal probability to the two states.

	L	M	R
U	$2,2$	$-1,-4$	$-1,3$
D	$3,0$	$0,-1$	$4,-1$

θ

	L	M	R
U	$2,2$	$-1,3$	$-1,-4$
D	$3,0$	$4,-1$	$0,-1$

θ'

Figure (a) Base Game

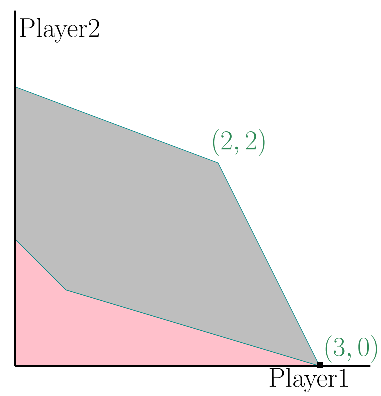

Figure (b) Payoffs

Figure 3: Self-Contained Coordinated Equilibrium

In Figure 3b, the set of all payoff profiles achievable in self-contained coordinated equilibria is displayed in grey. That is, for every payoff pair in the interior of the grey region it is possible to design an extensive-form information structure consistent with the prior ρ such that the resulting game has a sequential equilibrium achieving the target expected payoffs.

The mechanisms that implement this range of payoffs control the incentives of the players by selectively hiding and revealing information to different players at different points in time and after different contingencies. To illustrate the incentive power of sequential information design, the comparison with static information design is instructive.

For example, we may consider all possible games of incomplete information obtained from the base game by designing an interim, possibly asymmetric, information structure but restricting the players to move simultaneously. Then the set of all implementable outcomes is summarized by the concept of Bayes' Correlated Equilibrium (BCE) due to Bergemann and Morris (2016). In the example of Figure 3a there is only one Bayes' Correlated Equilibrium outcome: the players play the pure action profile (D, L).

To see why note that the action U is *ex post* strictly dominated by D. Thus, regardless of the information structure, Player 1 must play D with probability 1 in any Bayesian Nash Equilibrium. Given this, the unique best-reply for Player 2, independent of the state, is to play L.

Without the flexibility afforded by the extensive form, information design has no power in this example. Similarly, designing the extensive form without also designing the information structure leads to the same conclusion. Salcedo (2017) was the first to study this problem and he introduced the solution concept of Interdependent Choice Equilibrium (ICE) to summarize the set of Nash equilibria of complete-information extensive-form games.

In the example there is only one ICE outcome, again the pure profile (D, L). The reasons are different because, as Salcedo (2017) observed, although U is strictly dominated, in an extensive-form mechanism Player 1 may be incentivized to play U if he expects that a deviation would be punished by Player 2.[5]

But when the two states have equal probability, both M and R earn player 2 a negative expected payoff regardless of the action of player 1. Player 2 earns at least 0 from L and therefore regardless of the extensive form, Player 2 must play L and therefore Player 1 must play D.

Implementing an outcome such as (U, L) requires designing the information structure *within* the extensive form. Our results not only delineate the implementable outcomes but also show how to construct these extensive-forms. Intuitively, in the example information must be disclosed to player 2 in a way that is contingent on the behavior of player 1. When player 1 plays the recommended action U, player 2 should be kept uninformed about θ so that she is dissuaded from playing M and

[5]Indeed Salcedo (2017) shows that for certain specifications of the payoffs in the Prisoner's Dilemma there are ICE outcomes in which both players cooperate. See also Nishihara (1997).

R. But when player 1 deviates to D, some information must be disclosed to player 2 in order for her punishment to be tailored to the state (L in state θ and R in state θ'.) Making these punishments incentive-compatible for Player 2 requires further obfuscation. Below is an example of an extensive-form game and a self-contained equilibrium in which (U, L) is played. The analysis is in Appendix A.

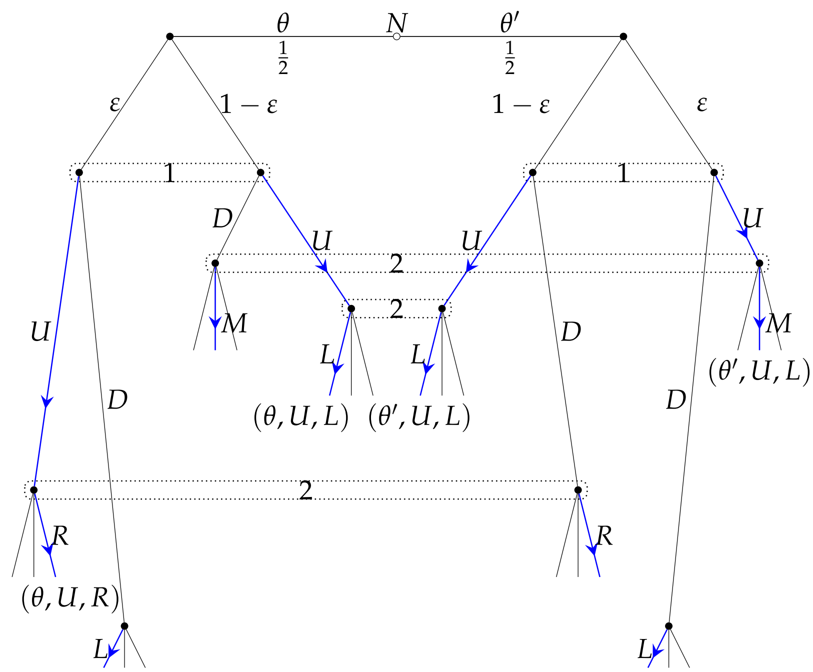

Figure 4: A sequential equilibrium

Converses So far we have described our results in terms of sufficient conditions. For example, in order to guarantee that an outcome can arise from a self-contained equilibrium of an extensive-form game it is sufficient that there is a self-contained coordinated equilibrium with the same outcome. These results are proven by construction. We use a simple extensive form in which along every path each player i has exactly one opportunity to move and at that opportunity she chooses from her set of actions A_i. The game in Figure 4 is one such extensive form.

However, we may go a step further and consider abstract extensive forms analogous to indirect mechanisms in mechanism design. In an abstract extensive form players may move multiple times along a path and moves are not necessarily equated with directly choosing actions. Instead we simply associate each terminal node with an action profile. The players "choose" an action profile by following the path that leads to an associated terminal node. The question arises whether a larger set of outcomes than those we identify can arise as equilibria of the broader class of abstract extensive forms.

Our goal is to establish exact converses of our results, i.e. coordinated equilibrium as a *necessary* condition. We prove converses under two sets of restrictions on extensive forms. These restrictions slightly generalize the class of extensive-form mechanisms first introduced by Salcedo (2017). First, we impose three properties that preserve the autonomy of the players and the non-cooperative nature of the game. Every tuple consisting of a state of the world and action profile must be associated with at least one terminal node. That is, the extensive-form must not restrict the set of feasible outcomes. When a player (implicitly) chooses an action a_i by following some path, she must know that her action will be a_i. That is, there cannot be an information set bundling that path with one leading to a different action b_i. Finally, player i cannot "delegate" the choice from her action set A_i to another player j. That is, a move by player j in the extensive-form cannot reduce the set of actions in A_i that i can feasibly choose through subsequent moves.

The simple extensive forms we use to prove our sufficient conditions satisfy all of the above properties and therefore the sufficient conditions apply when we allow for this broader class of abstract extensive forms. To establish necessity we impose one further and more substantive restriction (and one that is also satisfied by our simple extensive forms). Players cannot make partial commitments: at any stage of the extensive-form the set of actions that player i can "choose" via subsequent moves is either a singleton (i.e., she has already made her choice) or the entire set A_i.

In many environments it would be quite natural to suppose that players can choose their actions by a progressive sequence of partial choices. For example, when contributing to a public good it may be possible to make an initial non-refundable contribution and maintain the option of making additional contributions later, perhaps contingent on the contributions of others. And indeed such possibilities will expand the set of equilibrium outcomes, as can be seen from the work of Renou (2009), Bade et al. (2009), and more recently Dutta and Ishii (2016), who specifically study the power of partial commitments.

3 Coordinated Equilibrium

We consider base games of incomplete information defined as follows. There are N players, each player i chooses an action from the finite set A_i. Payoffs $u_i(a, \theta)$ depend on action profiles a and the state of the world θ, an element of the finite set Θ. In some cases we might also fix a common prior ρ over Θ.

The basic unit of our analysis is a plan.

Definition 1. *A plan is a tree where each node belongs to a player i and is labeled with an action $a_i \in A_i$. At every node belonging to i there is exactly one branch for every action in A_i and every path through the tree passes through exactly one node for each player.*

We let P be the set of all plans. A plan can be thought of as a contingent sequence of recommendations to the players. We interpret a node labeled a_i as player i being called upon to act and being recommended to play action a_i. Every plan has an *obedient path*, the path that results from always following the branch associated with the recommended action. We define the subset of plans $\langle a_i \rangle \subset P$ to be the set of all plans whose obedient path has player i play the action a_i.

Similarly, we can consider the path that results from a possible deviation by player i to an action b_i. Consider following the obedient path until we reach the node belonging to i, then following the branch associated with b_i. After that we continue following the recommended branches for the subsequent players.[6] This will result in a unique action profile, and for a given state of the world θ, this yields a payoff to player i which we denote by $u_i(b_i, p, \theta)$.

Definition 2. *A coordinated equilibrium is a distribution $\pi \in \Delta(\Theta \times P)$ such that for each player i the obedience constraints hold i.e.*

$$\sum_{\theta \in \Theta, p \in \langle a_i \rangle} \pi(\theta, p) \left[u_i(a_i, p, \theta) - u_i(b_i, p, \theta) \right] \geq 0 \tag{2}$$

for all $a_i, b_i \in A_i$.

In some applications the analyst may fix a common prior ρ over states. A coordinated equilibrium π is consistent with ρ if the marginal of π on Θ coincides with

[6] Note that b_i may be the action recommended to player i, in which case we are simply following the obedient path again.

ρ. When we associate each plan with the action profile that results from its obedient path, a coordinated equilibrium induces a distribution $\alpha \in \Delta(\Theta \times A)$ over states and action profiles. We call α a coordinated equilibrium *outcome*.

Imagine that player i knows that a state and a plan will be drawn from the distribution π, and she receives the recommendation to play a_i. She learns that a plan from the set $\langle a_i \rangle$ has been selected, and she learns nothing else. If she assumes that all players (if any) who have previously moved have followed their recommendations, and all players who will subsequently move will follow their recommendations, then the left-hand side of Equation 2 is proportional to the conditional expected payoff gain from being obedient herself. If the inequality holds then player i will obey the recommendation to play a_i.

Indeed we can directly translate this intuition into the construction of an extensive-form game in which "obedience" is a Bayesian Nash Equilibrium and yields the same outcome. The root node belongs to Nature. Nature has exactly one branch associated with each element of $\Theta \times P$. Nature's choice probabilities are given by π. We then complete the construction of the game tree by appending after every one of Nature's branches (θ, p), the "continuation" tree given by the plan p. Finally, we build information sets. For every player i and action $a_i \in A_i$ there is a single information set containing every node belonging to i labeled with action a_i.

Refer to this extensive-form game with imperfect and incomplete information as $\Gamma(\pi)$. It is easy to see that the information sets in $\Gamma(\pi)$ model exactly the inference underlying the obedience constraint. At her information set associated with action a_i, player i knows (nothing more than) that Nature has selected a plan whose obedient path eventually recommends a_i to player i.

For a given base game G we refer to the class of extensive forms $\Gamma(\pi)$, as π varies, as the *canonical* extensive form(s) for G. The *obedient strategy* for player i is the strategy for $\Gamma(\pi)$ which, at each information set, selects the recommended action. By virtue of the obedience constraints in Equation 2, if π is a coordinated equilibrium[7], then the obedient strategy profile is a Bayesian Nash Equilibrium of $\Gamma(\pi)$.

Theorem 1. *If π is a coordinated equilibrium with outcome α then the obedient strategy profile is a Bayesian Nash Equilibrium of $\Gamma(\pi)$ and yields outcome α.*

[7]In some applications, the analyst may fix a prior ρ over states, and confine attention to games with common prior ρ. Then Theorem 1 implies that the outcomes of coordinated equilibria consistent with ρ are generated by Bayesian Nash equilibria of $\Gamma(\pi)$ where π has marginal distribution over Θ equal to ρ. The same observation applies to each of the refinements we consider below.

Theorem 1 extends the result of Salcedo (2017) to incomplete information environments.[8]

4 Self-Contained and Sequential Equilibrium

The obedience constraint in Equation 2 is vacuous for actions a_i that never arise along the obedient path of any plan p in the support of a coordinated equilibrium. Nevertheless, such actions may very well arise following a deviation from p by some other player. This is the sense in which coordinated equilibrium captures Nash equilibrium, but not sequentially rational outcomes of extensive forms. Indeed any π assigning probability 1 to the entry deterrence plan in the Cournot example from Section 2 is a (degenerate) coordinated equilibrium.

In this section we consider the opposite extreme: the strongest belief-based refinement of extensive-form equilibrium. To introduce it, recall that in an extensive-form game an information set h is on the path of play of a strategy profile σ if some node in h is reached with positive probability when play follows σ. An information set h is *reachable* given a strategy profile σ if there is some strategy σ_i' for i such that h is on the path of play of (σ_i', σ_{-i}).

Refinements of Bayesian Nash Equilibrium are based on scrutinizing the beliefs at information sets that are reachable but not on the path of play. An equilibrium survives all refinements when there are no such information sets.

Definition 3. *Consider a finite extensive-form game with perfect recall. A Bayesian Nash Equilibrium σ is self-contained if for each player i every information set that is reachable by i given σ_{-i} is on the path of play.*

Proposition 1. *A self-contained Bayesian Nash Equilibrium is a sequential equilibrium strategy profile.*

We derive a refinement of coordinated equilibrium that identifies outcomes of a base game that could arise as self-contained Bayesian Nash equilibria of some extensive form. It is based on restricting the set of plans under consideration. Intuitively we want to rule out plans that use actions which cannot be played along the equilibrium path of some equilibrium of some extensive form. In fact

[8] Theorem 1 is reminiscent of the result in Myerson (1982) which ties correlated equilibrium to the Nash equilibrium of the base game when the players communicate with a mediator.

we also want to avoid using actions which *can* be played but only under the threat of actions which are ruled out by the previous consideration. Indeed we want to iteratively remove such actions from the available plans.

It turns out that coordinated equilibrium itself provides the criterion for elimination. For each player i let C_i^1 be the set of actions a_i such that there exists a coordinated equilibrium outcome $\alpha \in \Delta(\Theta \times A)$ assigning positive probability to a_i. Next for each $k > 1$, consider the set of plans $P^{C^{k-1}}$ such that all actions recommended to any player j, on or off the obedient path, belong to C_j^{k-1}.

We inductively define C_i^k to be the set of actions a_i such that there exists a coordinated equilibrium $\pi \in \Delta(\Theta \times P^{C^{k-1}})$ assigning positive probability to some plan with a_i on its obedient path. Since the set of actions is finite, this elimination procedure terminates at a fixed point. We use C_i to refer to the set of actions that survive[9] for player i, and P^C the surviving set of plans.

Definition 4. *A coordinated equilibrium is self-contained if it assigns probability 1 to plans in P^C.*

Note that, since C is a fixed point, for every $a_i \in C_i$ there exists a self-contained coordinated equilibrium whose outcome assigns positive probability to a_i. Say that two outcomes α, α' are within ε distance if $|\alpha'(\theta, a) - \alpha(\theta, a)| \leq \varepsilon$ for every pair (θ, a). In the following result we consider canonical extensive forms in which Nature selects from the subset $\Theta \times P^C$, we call these $\Gamma^C(\pi)$.

Theorem 2. *For any self-contained coordinated equilibrium outcome α and for any $\varepsilon > 0$, there is a (canonical) extensive-form with a self-contained Bayesian Nash equilibrium whose outcome is within ε distance of α.*

Proof. Let π be a self-contained coordinated equilibrium with outcome α. For all i, for all $a_i \in C_i$ there is a self-contained coordinated equilibrium, call it π_{a_i} whose outcome assigns positive probability to a_i. In particular, π_{a_i} assigns positive probability to a plan with a_i on its obedient path. Let M be the sum of the cardinality of the sets C_i for $i = 1, \ldots, N$. Consider the following mixture for any $\varepsilon \in (0, 1)$:

$$\psi(\theta, p) = (1 - \varepsilon)\pi(\theta, p) + \frac{\varepsilon}{M} \sum_{i=1}^{N} \sum_{a_i \in A_i} \pi_{a_i}(\theta, p)$$

[9] This set is non-empty and includes all actions that can be played with positive probability in a Bayes' correlated equilibrium of the base game.

Since the obedience constraints are satisfied for each of the π_{a_i} individually, they are satisfied for ψ. Moreover ψ assigns probability 1 to plans in P^C. Thus ψ is a self-contained coordinated equilibrium. Let α' be its outcome.

For any action profile a, let $\langle a \rangle$ be the set of plans whose path yields action profile a. Then $\alpha(\theta, a) = \sum_{\Theta \times \langle a \rangle} \pi(\theta, p)$ and $\alpha'(\theta, a) = \sum_{\Theta \times \langle a \rangle} \psi(\theta, p)$. By construction for any (a, θ) we have

$$(1 - \varepsilon)\alpha(\theta, a) \leq \alpha'(\theta, a) \leq (1 - \varepsilon)\alpha(\theta, a) + \varepsilon$$

and thus $|\alpha'(\theta, a) - \alpha(\theta, a)| \leq \varepsilon$.

By Theorem 1 the obedient strategy profile is a Bayesian Nash Equilibrium of $\Gamma^C(\psi)$ and yields the same outcome as ψ. By construction, for every $a_i \in C_i$ the mixture ψ assigns at least probability $\varepsilon/M > 0$ to a plan in which a_i is on the obedient path. Thus, all of the (reachable) information sets are on the path and the obedient strategy profile is a self-contained Bayesian Nash equilibrium of $\Gamma(\psi)$.

Its outcome is the same as that of ψ and we have shown that the latter is within ε distance of π. $\square$

It should be clear from the proof that the outcomes of all *interior* self-contained coordinated equilibria can be *exactly* realized by a self-contained Bayesian Nash equilibrium. The approximation is needed only for boundary points.

In cases in which we are given a prior ρ, we can further specialize Theorem 2. In Doval and Ely (2019) we show that, as long as ρ has full support, the ψ used in the proof can be chosen to be consistent with ρ so that the resulting self-contained Bayesian Nash Equilibrium outcome exactly coincides with α in its marginal over states and differs by ε only in terms of the distribution of action profiles.

5 Perfect Bayesian Equilibrium

From the point of view of a planner, who *designs* the game and information structure to implement some solution, it may make sense to stop at self-contained coordinated equilibrium. The solutions obtained will not rely on any assumption about how beliefs at off-path information sets are formed. However, for an external analyst studying a game outside of her control, and who is agnostic about the extensive form, self-contained coordinated equilibrium is too restrictive. We next study the weaker solution concept of Perfect Bayesian Equilibrium, allowing for off-path information sets but requiring that off-path beliefs satisfy Bayes' rule where possible.

We begin with the definition of Perfect Bayesian Equilibrium. In a finite extensive-form game with perfect recall let Σ_i denote the set of pure strategies for player i. Nature is denoted $i = 0$, and analogously to the players a pure strategy for Nature specifies a move at each chance node. The set of pure strategy profiles (including a strategy for Nature) is Σ, and the set of pure strategy profiles of the opponents of i (including Nature) is Σ_{-i}.

A *behavioral strategy* for player i is a rule β_i which specifies for each information set h_i belonging to player i, a probability distribution $\beta_i(\cdot \mid h_i)$ over the moves $M_i(h_i)$ available at h_i. A *system of beliefs* ν specifies a probability distribution $\nu(\cdot \mid h_i)$ over the nodes within each information set h_i. An assessment is a pair (β, ν) where β is a profile of behavioral strategies and ν is a system of beliefs.

Bayes' rule where possible is the condition that the updating of beliefs should be consistent with the strategy profile even at information sets that are off the path. We adopt a strong version of this concept which employs conditional probability systems (see, e.g., Myerson (1986); Battigalli (1996); Watson (2017)).

A conditional probability system over Σ is a collection of probability measures $\mu(\cdot \mid S) \in \Delta\Sigma$, one for each non-empty subset S of strategy profiles satisfying the following two conditions:

1. $\mu(S \mid S) = 1$ for all $S \subset \Sigma$,
2. For all $A \subset B \subset C \subset \Sigma$, the chain rule of conditional probabilities holds

$$\mu(A \mid C) = \mu(A \mid B) \cdot \mu(B \mid C).$$

Given any information set h_i belonging to player i we can identify the subset of strategy profiles σ_{-i} for the opponents of i such that h_i is reachable when σ_{-i} is played. We will conserve notation and use h_i to refer both to the information set in the extensive form and the associated set of strategy profiles. Note that by perfect recall, every strategy profile in $h_i \subset \Sigma_{-i}$ can reach one and only one node within the information set h_i.

Similarly, to any node y, we can associate the set of pure strategy profiles on whose path lies y. We will use y to refer both to the node in the extensive form and to the set of strategy profiles that reach it.

By Kuhn's Theorem, each profile of behavioral strategies can be associated with a unique equivalent probability distribution over pure strategy profiles. Similarly, for each node y, a profile of behavioral strategies β defines a unique equivalent *conditional* probability distribution $\beta^{\Sigma}(y) \in \Delta\Sigma$ over pure strategy profiles. Letting $\prec$ denote the predecessor relation for nodes y in the extensive form, the probability assigned to profile σ is

$$\beta^{\Sigma}(y)\,[\sigma] = \prod_{\{h:(\nexists y' \in h):y' \prec y\}} \beta(\sigma(h) \mid h)$$

if $\sigma \in y$ and zero otherwise.

The chain rule in item 2 for conditional probability systems is the formal statement of Bayes' rule "where possible". Whenever $\mu(B \mid C) > 0$ we may divide by it and obtain Bayes' formula for $\mu(A \mid B)$. Otherwise Bayes' rule is not possible. This is our motivation for the following definition.[10]

Definition 5. *An assessment (β, ν) satisfies Bayes' rule where possible if there is a conditional probability system μ such that for each information set h_i and node y,*

1. $\mu(y \mid h_i) = \nu(y \mid h_i)$ *and*
2. *If* $\mu(y \mid h_i) > 0$ *then* $\mu(\cdot \mid y) = \beta^{\Sigma}(y)$.

A conditional probability system satisfies Bayes' rule where possible, by definition. The first condition above says that μ yields the same beliefs at information sets as the system of beliefs ν, therefore ν satisfies Bayes' rule where possible. The second condition says that the "transition" probabilities used to apply Bayes' rule (where possible) are the correct ones, namely the transition probabilities from the behavioral strategy profile β.

For any profile of behavioral strategies β and any node y belonging to player i, let $\Upsilon(\beta \mid y)$ denote player i's expected continuation payoff beginning from y when the players follow β.

Definition 6. *An assessment (β, ν) is sequentially rational if for all players i, for every information set h_i*

$$\sum_{y \in h_i} \nu(y \mid h_i)\left[\Upsilon(\beta \mid y) - \Upsilon(\beta_i', \beta_{-i} \mid y)\right] \geq 0$$

[10] Throughout, when we consider cylinder sets like $h_i \subset \Sigma_{-i}$, the conditional probability $\mu(\cdot \mid h_i)$ will be understood to mean $\mu(\cdot \mid \Sigma_i \times h_i)$. Likewise if $C \subset \Sigma_{-i}$ then $\mu(C \mid A)$ means $\mu(\Sigma_i \times C \mid A)$.

for all behavioral strategies β_i' .

Definition 7. *A Perfect Bayesian Equilibrium is an assessment that is sequentially rational and satisfies Bayes' rule where possible.*

Deviant Coordinated Equilibrium For the concepts of Bayesian Nash and Self-Contained equilbrium we used coordinated equilibria to identify actions that can be rationalized on the path of play of an extensive form. In order to identify additional actions that can also be played off-path we will need an extended concept which we call Deviant Coordinated Equilibrium. A *deviant plan* is a pair $q = (p, x)$ where p is a plan and x is a node of p. The interpretation of a deviant plan (p, x) is that p was the plan, but possibly due to a sequence of deviations from the path of the plan we have arrived at node x. The path of a deviant plan (p, x) consists of all the nodes reached by obedient play in p *starting* from x. In particular, x itself is on the path of (p, x) but none of its predecessors are. Also, if x is the initial node of p then the path of the deviant plan (p, x) is just the obedient path of p itself. Let Q be the set of all deviant plans and for any family of subsets B, let Q^B be the set of all deviant plans (p, x) that only recommend actions in B, i.e. $p \in P^B$.

Let $\langle a_i \rangle$ be the set of all deviant plans q such that the action a_i is recommended to player i at some node on its path. For any deviant plan q and state θ, if there is a node belonging to i on the path of q=(p, x), then we write $u_i(a_i, q, \theta)$ for the payoff to i from the action profile that results when, starting from x, all players are obedient except possibly player i who chooses action a_i.

Definition 8. *A deviant coordinated equilibrium relative to a family of action subsets $B = (B_1, \dots, B_N)$ is a probability $\pi \in \Delta(\Theta \times Q)$ such that for each i, the obedience constraints*

$$\sum_{\theta \in \Theta, q \in \langle a_i \rangle} \pi(\theta, q) \left[u_i(a_i, q, \theta) - u_i(b_i, q, \theta) \right] \geq 0 \tag{3}$$

hold for all actions $a_i \notin B_i$ and for all $b_i \in A_i$.

Notice that in a deviant coordinated equilibrium the obedience constraint for an action a_i does not include plans in which a_i precedes the deviant node (such plans will not belong to $\langle a_i \rangle$). Thus, a deviant coordinated equilibrium takes as given any deviations and checks incentives only after the deviations have occurred. A deviant coordinated equilibrium relative to B also ignores the obedience constraints for the pre-specified set of actions B_i. The role of the family B will be clear when we use Definition 8 in the iterative procedure below. Finally notice that any coordinated equilibrium is a deviant coordinated equilibrium relative to any family B, by simply specifying that the deviant node in each plan is the initial node.

We will use deviant coordinated equilibrium to define an iterative procedure leading to a family of action subsets $D = (D_1, \ldots, D_N)$ which we will then prove represents all of the actions that can be played on or off the path of a Perfect Bayesian Equilibrium.

The procedure is a nested iterative application of two operators. The first operator removes actions that (we will show) cannot be played on the equilibrium path. Let J be a given family of non-empty action sets, $J_1, \ldots, J_N$. Let $\mathcal{C}_i(J)$ be the set of actions for player i that can be played with positive probability in a coordinated equilibrium $\pi \in \Delta(\Theta \times P^J)$ and $\mathcal{C}(J)$ be the family $\mathcal{C}_1(J), \ldots \mathcal{C}_N(J)$.

Notice that if we were to iteratively apply the operator $\mathcal{C}$, the sequence $\mathcal{C}(A)$, $\mathcal{C}(\mathcal{C}(A))\ldots$ would be precisely the decreasing sequence of non-empty subsets $C^1, C^2 \ldots$ terminating at the set of self-contained action sets C. To obtain instead the set of actions that can be played in a Perfect Bayesian equilibrium we will first *add* some of the removed actions back to the set $\mathcal{C}(A)$ before applying $\mathcal{C}$ again.

Adding back actions that cannot be played *on* the equilibrium path but can be played *off* the equilibrium path is the key to constructing the set D. Just as the removal of actions involves iteratively applying the coordinated equilibrium concept, to determine which actions should be *restored* we will iteratively apply *deviant* coordinated equilibrium.

Define $D_i^{1,1}$ to be the the set of actions for i that are on the path of a deviant coordinated equilibrium $\pi \in \Delta(\Theta \times P^A)$ relative to $\mathcal{C}(A)$.

Continuing in this fashion we define for each $k \geq 2$, the sets $D_i^{1,k}$ to be the set of actions for i that are on the path of a deviant coordinated equilibrium $\pi \in \Delta(\Theta \times Q^A)$ relative to $D^{1,k-1}$.

This is an increasing sequence of sets because the set of obedience constraints is shrinking (see Lemma 1 below). The sequence therefore terminates and we will call the terminal family $D^{1,\infty}$. The remainder of the procedure continues in the same way, nesting the iterative addition of actions within the iterative application of $\mathcal{C}$. Formally we set $D^{0,\infty} = A$ and for each $l \geq 1$, $k \geq 1$ we define the families $D^{l,k}$ as follows

- $D_i^{l,0} = \mathcal{C}_i\left(D^{l-1,\infty}\right)$.

- $D_i^{l,k}$ is the set of actions for i that are on the path of a deviant coordinated equilibrium $\pi \in \Delta(\Theta \times P^{D^{l-1,\infty}})$ relative to $D^{l,k-1}$,
- $D^{l,\infty} = \bigcup_k D^{l,k}$,

We will illustrate the workings of the procedure using the three player complete-information game depicted in Figure 5a. Player 1 chooses the matrix while players 2 and 3 choose the rows and columns respectively. Iterative removal of absolutely dominated actions yields the unique profile (d_1, d_2, d_3). That is, repeated application of the operator $\mathcal{C}$ would eliminate X, c_2, and c_3. However, as we will show the set D includes in addition the actions c_2 and c_3 as the c_2 will be restored by the inner nest and as a result c_3 will never be removed.

	c_3	d_3
c_2	$-1,-1,1$	$-1,1,0$
d_2	$-1,-1,1$	$-1,2,0$

X

	c_3	d_3
c_2	$1,0,0$	$1,0,-5$
d_2	$1,3,0$	$1,3,3$

d_1

Figure (a) Base Game

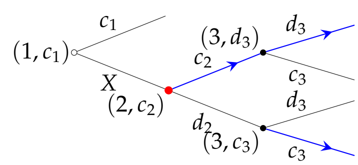

Figure (b) Deviant Plan

Figure 5: Illustrating the Set D

Let us trace through the nested iterative procedure to derive the set D for this game. First, the action X is absolutely dominated and will not be included in $D^{1,0} = \mathcal{C}(A)$. In fact X is the only action that can be removed in the first step and therefore $D^{1,0} = [\{X\}, A_2, A_3]$. Next we apply the inner nest to possibly restore removed actions. Given that X is absolutely dominated it cannot be incentivized even after a deviation by another player and hence X will not be added back. Thus $D^{1,\infty} = D^{1,0}$.

We then apply the outer operator $\mathcal{C}$ to $D^{1,\infty} = [\{X\}, A_2, A_3]$ to obtain $D^{2,0}$. Within this smaller set of actions, the action c_2 is absolutely dominated and thus removed. This is the only action that can be removed so we have $D^{2,0} = [\{X\}, \{d_2\}, A_3]$. At this stage the inner nest considers deviant coordinated equilibria relative to the set $D^{2,0}$. The plans in these involve some sequence of possible deviations (for which incentive constraints are not checked) followed by obedient play where we check the incentives only of those actions not belonging to $D^{2,0}$. Figure 5b is a degenerate (applying probability 1 to a single plan) deviant coordinated equilibrium of this form.

Notice that player 2 obeys the recommendation to play c_2 because it leads to the profile (X, c_2, d_3) earning a payoff of 1. The alternative action d_2 would be met with the response c_3 and a payoff of -1. Thus 2 is obedient and the path of this deviant coordinated equilibrium passes through the recommendation that player 3 play d_3. It is *not* optimal for player 3 to obey this recommendation because given the preceding choices of players 1 and 2, the best-response for 3 is the action c_3. However, d_3 belongs to the set $D^{2,0}$ and therefore this incentive constraint is ignored and the plan above is a (degenerate) deviant coordinated equilibrium relative to $D^{2,0}$. Since c_2 is on its path we have $D^{2,1} = [\{X\}, A_2, A_3]$.

Finally, since $D^{2,1}$ is identical to $D^{1,1}$ the remainder of the procedure will simply repeat previous steps and thus $D^{2,\infty} = D^{1,\infty} = D^{\infty,\infty} = D$.

The following lemma establishes that this procedure has a well-defined terminal family.

Lemma 1. *For each $l \geq 1$, $k \geq 0$,*

1. $D^{l,k} \subset D^{l,k+1}$,
2. $\emptyset \neq D^{l,\infty} \subset D^{l-1,\infty}$

Thus, we define $D = \bigcap_l D^{l,\infty}$. We are going to show that the outcomes of coordinated equilibria that use plans in P^D can arise as Perfect Bayesian equilibria. Actions outside the set D cannot be made sequentially rational even off the path. Therefore an obedient strategy profile cannot be part of a Perfect Bayesian Equilibrium of a canonical extensive form in which Nature's moves include the choice of plans that recommend such actions. In the following result we consider canonical extensive forms in which Nature selects from the subset $\Theta \times P^D$, we call these $\Gamma^D(\pi)$.

Theorem 3. *If π is a coordinated equilibrium which assigns probability 1 to plans in P^D and has outcome α, then the obedient strategy profile is a Perfect Bayesian Equilibrium of $\Gamma^D(\pi)$ and yields outcome α.*

The proof of Theorem 3 is in Appendix B. The idea is to use deviant coordinated equilibria to build a conditional probability system over strategy profiles.[11] A conditional probability system can be thought of as a lexicographic hierarchy of probability measures. At the top level we assign probability 1 to the obedient (i.e. equilibrium) strategy profile. All information sets in $\Gamma^D(\pi)$ corresponding to action recommendations with positive probability under α will be on the path and beliefs at such information sets are derived from Bayes' rule. Progressively lower level probability measures are built from the deviant coordinated equilibria that rationalize the remaining actions in D.

6 Converses

The protocol through which players jointly determine the outcome of a game may be less simple and direct than the canonical extensive forms used in our sufficient conditions. At the most abstract level, an extensive form is an arbitrary finite game tree each of whose terminal nodes is associated with an outcome, i.e., a tuple consisting of a state of the world and action profile in the base game.

Limits on the set of outcomes that can arise in equilibria will depend on the class of abstract extensive forms admitted. We consider a class of extensive forms which satisfy two categories of restrictions. Restrictions in the first category are meant to impose the decentralized and autonomous nature of non-cooperative play. To introduce them we will need some notation for a finite extensive-form game.

First, we use the terminology of *moves* for the branches in the extensive form to distinguish from *actions* in the base game. The players make (sequences of) moves in the extensive form to determine the profile of actions in the base game. Let Z be the set of terminal nodes. The *outcome function* $\gamma : Z \to \Theta \times A$ assigns an outcome to each terminal node.

To represent the incomplete information from the base game, we have Nature move at the initial node and select an element of Θ. Formally, if y is any non-initial node, then $\gamma_\Theta(z) = \gamma_\Theta(z')$ for any two terminal nodes z and z' that follow y, where γ_Θ is the projection of γ on Θ.

[11] Theorem 3 is reminiscent of Theorem 2 in Myerson (1986), where he shows that the set of communication equilibria which do not recommend *codominated* actions coincides with the set of *sequential* communication equilibria. Indeed, the set D in our construction plays a similar role to the set of actions that are not codominated in Myerson (1986). See also Sugaya and Wolitzky (2017).

If y is a non-terminal node, then for each player i we define $\gamma_i(y)$ to be the set of i's actions in the base game which are available at node y. In particular $a_i \in \gamma_i(y)$ if and only if there is a terminal node z succeeding y such that i's action in the profile $\gamma(z)$ is a_i.

Next, say that player i's action is *determined* at a node y if $\gamma_i(y)$ is a singleton. Say that a move m determines action a_i at node y' if (i) player i's action is not determined at y', and (ii) $\gamma_i(y) = \{a_i\}$ at the node y immediately following the move m starting from y'.

First, we require that every outcome in the base game is associated with at least one terminal node of the extensive form, i.e. γ is onto. Second, player i should *know her own action*. If m is a move for i at a node y that determines action a_i, then m determines the same action a_i at all nodes in the information set containing y. Thirdly, there should be *no delegation*: player i cannot delegate her choice of action to another player j. If i's action is not determined at a node y at which j moves, then i's action remains undetermined after any of j's moves at y.

The preceding conditions we view as natural restrictions preserving non-cooperative play. We impose a final condition, which we view as a more substantive restriction. We assume there are *no partial commitments*: if a player's action is undetermined at a node y, then all of her actions remain available, i.e.$\gamma_i(y) = A_i$.[12]

We call any finite extensive form with perfect recall that satisfies all of these conditions admissible.

Definition 9. *Fix a base game. A finite extensive form with perfect recall is admissible if it satisfies know-your-own-action, no-delegation, and no-partial-commitments and if γ is onto.*

We establish converses of our three main results within the class of admissible extensive forms.[13]

Theorem 4. *Let Γ be an admissible extensive form for a given base game.*

[12]No partial commitments is a vacuous restriction for games with only two actions. Moreover, note that no delegation together with no partial commitments implies that a player cannot affect via her moves what actions in the base game another player can choose.

[13]In Doval and Ely (2019) we show examples in which each of the admissibility conditions are violated and equilibrium outcomes exist outside of the set of coordinated equilibria.

1. *If Γ has a Bayesian Nash Equilibrium whose outcome is α then α is a coordinated equilibrium outcome.*
2. *If Γ has a self-contained equilibrium whose outcome is α then α is a self-contained coordinated equilibrium outcome.*
3. *If Γ has a Perfect Bayesian Equilibrium whose outcome is α then α is the outcome of a coordinated equilibrium assigning probability 1 to P^D.*

We conclude by observing that the combination of Theorem 4 and for example Theorem 3 implies that Γ^D is in a certain sense the *canonical* extensive form for the solution concept of Perfect Bayesian Equilibrium.

Corollary 1. *For any base game, an outcome can result from a Perfect Bayesian Equilibrium of an admissible extensive form if and only if it can arise from an obedient Perfect Bayesian Equilibrium of Γ^D.*

7 Relation to Salcedo (2017)

In this section we discuss the relation between our coordinated equilibrium concept and the path-breaking work of Salcedo (2017). We build upon Salcedo (2017) by considering incomplete information and by focusing on a range of refinements.

Salcedo (2017) studied complete information games and was the first to propose a solution concept, ICE, based on linear obedience constraints to capture the set of Nash equilibria of extensive-form games derived from a complete-information base game. Salcedo (2017) also obtained a converse result, proposing restrictions on the class of complete-information extensive forms whose Nash equilibria yield ICE outcomes.

Our Theorem 1 generalizes the basic result in Salcedo (2017) to incomplete information environments. As we demonstrate in Figure 3, incomplete information, together with the design of the extensive adds as an incentive device the partial and selective disclosure of information through the course of play. Our class of admissible extensive forms also generalize and slightly broaden the class identified by Salcedo (2017).

A major focus of our work is understanding how variations of the coordinated equilibrium concept can capture various extensive-form refinements. Salcedo (2017) also considers a refinement of Nash equilibrium. In particular, for 2 player games Salcedo (2017) specializes the ICE concept to capture a novel extensive-form refinement based on player-specific trembles. Building on this work we consider general n-player games with incomplete information and study how to modify the coordinated equilibrium concept to capture Perfect Bayesian equilibrium and also self-contained equilibrium.

8 Conclusion and Further Directions

We have presented methods to summarize equilibrium outcome distributions of games of incomplete information as both the information structure and the extensive form vary.

We see several avenues worth exploring and left for future work. In empirical work, similar to Syrgkanis et al. (2017), Magnolfi and Roncoroni (2017) and Bergemann et al. (2019), who use Bayes'correlated equilibrium to make inferences about structural parameters in auctions, our concepts could be used as a framework for inference that is robust to extensive-form information structures. For instance, starting from the work of Bresnahan and Reiss (1991), industrial organization models of entry make assumptions about the order of moves (see Ciliberto and Tamer (2009) and Grieco (2014)) that could be dispensed with under our framework.

A second goal would be to characterize additional solution concepts. We have characterized self-contained Bayesian Nash equilibrium, a strict refinement of sequential equilibrium. Obtaining an exact characterization of sequential equilibrium outcomes remains an open problem.

Our characterization is based on the view that the specified action sets summarize completely the range of strategic choices made in the extensive form. This is formalized by our admissibility conditions. In some settings however it is natural to suppose that players can reach their final choices in stages, for example a firm may progressively increase its output along the path of play. This would represent a departure from the no-partial-commitments condition. We conjecture that a variation of coordinated equilibrium could summarize the outcomes in that larger set of extensive forms.

Finally, an analyst may have prior knowledge of some partial structure of the extensive form. For example, in an entry game she may impose the restriction that firms first make entry decisions and then make output choices after commonly observing the identity of entrants. The timing and information structure within those separate stages, however may vary. Deriving the restrictions on equilibrium outcomes that would be implied by such prior knowledge is another interesting direction.

References

AUMANN, R. J. (1974): "Subjectivity and correlation in randomized strategies," *Journal of mathematical Economics*, 1, 67–96.

——— (1987): "Correlated equilibrium as an expression of Bayesian rationality," *Econometrica: Journal of the Econometric Society*, 1–18.

BADE, S., G. HAERINGER, AND L. RENOU (2009): "Bilateral commitment," *Journal of Economic Theory*, 144, 1817–1831.

BATTIGALLI, P. (1996): "Strategic independence and perfect Bayesian equilibria," *Journal of Economic Theory*, 70, 201–234.

BERGEMANN, D., B. A. BROOKS, AND S. MORRIS (2019): "Counterfactuals with Latent Information," *Available at SSRN*.

BERGEMANN, D. AND S. MORRIS (2016): "Bayes correlated equilibrium and the comparison of information structures in games," *Theoretical Economics*, 11, 487–522.

BEST, J. AND D. QUIGLEY (2017): "Persuasion for the long run," *Available at SSRN*.

BOGNAR, K., T. BÖRGERS, AND M. MEYER-TER VEHN (2015): "An optimal voting procedure when voting is costly," *Journal of Economic Theory*, 159, 1056–1073.

BRESNAHAN, T. F. AND P. C. REISS (1991): "Entry and competition in concentrated markets," *Journal of political economy*, 99, 977–1009.

CILIBERTO, F. AND E. TAMER (2009): "Market structure and multiple equilibria in airline markets," *Econometrica*, 77, 1791–1828.

DE OLIVEIRA, H. AND R. LAMBA (2019): "Rationalizing dynamic choices," *Available at SSRN*.

DOVAL, L. AND J. C. ELY (2019): "Sequential Information Design: Additional Material," Click here.

DUTTA, R. AND R. ISHII (2016): "Dynamic commitment games, efficiency and coordination," *Journal of Economic Theory*, 163, 699–727.

FORGES, F. (1986): "An approach to communication equilibria," *Econometrica: Journal of the Econometric Society*, 1375–1385.

GALLICE, A. AND I. MONZON (2018): "Cooperation in Social Dilemmas through Position Uncertainty," *The Economic Journal*.

GERSHKOV, A. AND B. SZENTES (2009): "Optimal voting schemes with costly information acquisition," *Journal of Economic Theory*, 144, 36–68.

GRIECO, P. L. (2014): "Discrete games with flexible information structures: An application to local grocery markets," *The RAND Journal of Economics*, 45, 303–340.

KAMENICA, E. AND M. GENTZKOW (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

KREMER, I., Y. MANSOUR, AND M. PERRY (2014): "Implementing the "wisdom of the crowd"," *Journal of Political Economy*, 122, 988–1012.

MAGNOLFI, L. AND C. RONCORONI (2017): "Estimation of discrete games with weak assumptions on information," *manuscript. University of Wisconsin-Madison*.

MAKRIS, M., L. RENOU, ET AL. (2018): "Information design in multi-stage games," Tech. rep., Queen Mary University.

MYERSON, R. B. (1982): "Optimal coordination mechanisms in generalized principal–agent problems," *Journal of mathematical economics*, 10, 67–81.

——— (1986): "Multistage games with communication," *Econometrica: Journal of the Econometric Society*, 323–358.

NISHIHARA, K. (1997): "A resolution of N-person prisoners' dilemma," *Economic theory*, 10, 531–540.

PETERS, M. (2015): "Reciprocal contracting," *Journal of Economic Theory*, 158, 102–126.

RENOU, L. (2009): "Commitment games," *Games and Economic Behavior*, 66, 488–505.

SALCEDO, B. (2012): "Implementation Through Coordination in Static Environments Without Commitment," academia.edu.

——— (2017): "Interdependent choices," Tech. rep., University of Western Ontario.

SHAKED, A. AND J. SUTTON (1987): "Product differentiation and industrial structure," *The Journal of Industrial Economics*, 131–146.

SUGAYA, T. AND A. WOLITZKY (2017): "Revelation Principles in Multistage Games," Tech. rep., Technical Report, Discussion paper, Stanford University and MIT.

SUTTON, J. (1991): *Sunk costs and market structure: Price competition, advertising, and the evolution of concentration*, MIT press.

——— (2001): *Technology and market structure: theory and history*, MIT press.

——— (2003): "Market structure: The bounds approach," *Handbook of Industrial Organization*, 3.

SYRGKANIS, V., E. TAMER, AND J. ZIANI (2017): "Inference on auctions with weak assumptions on information," *arXiv preprint arXiv:1710.03830*.

WATSON, J. (2017): "A General, Practicable Definition of Perfect Bayesian Equilibrium," UC San Diego.

A Examples

A.1 The game in Figure 4

In the game in Figure 4 Nature first chooses the state (each with probability $1/2$) and then selects one of two possible payoff-irrelevant signals with probabilities $1 - \varepsilon$ and ε, where $0 < \varepsilon \leq 2/5$. Then the two players select their actions sequentially with Player 1 choosing first. The arrows in the figure represent the equilibrium strategy profile.

At the information set where Player 2 plays R, she does so because Bayes' rule assigns probability 1 to the node at which the state is θ and Player 1 has played U. Similarly at the information set where Player 2 plays M, she does so because Bayes' rule assigns probability 1 to the node at which the state is θ' and Player 1 has played U. At the non-singleton information set where Player 2 plays L she is uncertain of the state, assigning probability $1/2$ to each state. With such beliefs her expected payoff from either M or R would be $3/2$, smaller than the 2 she earns. At her singleton information set she knows that Player 1 has played D and she plays her unique best-response L.

At both of his information sets, Player 1 anticipates the very likely outcome (U, L) earning a payoff of 2. He expects that a deviation to D would induce a different response from Player 2 depending on which node has been reached in the information set. In expectation Player 1's payoff from the deviation would be $5\varepsilon \leq 2$.

All non-singleton information sets arise with positive probability. Therefore the strategy profile, which we have just shown to be sequentially rational, is a sequential equilibrium. Indeed because there are no off-path beliefs to be determined, this profile would survive all conventional refinements.

Finally note that as ε approaches zero, the probability of the outcome (U, L) approaches 1.

A.2 The Game in Figure 2

Consider the extensive form below in Figure 6 where $q = 1/2$. The blue branches indicate the pure strategies of the players. Following the move d_1, Players 2 and 3 will play their best-responses (c_2, c_3), we have truncated that portion of the tree to reduce clutter.

The strategy profile is a Perfect Bayesian Equilibrium when paired with the following system of beliefs. At Player 2's information set, she assigns probabilities $(0, q, 1-q)$ to nodes e, f, and g, and at Player 3's information set he assigns probabilities $(0, q, 1-q)$ to nodes ϵ, ϕ, and γ.

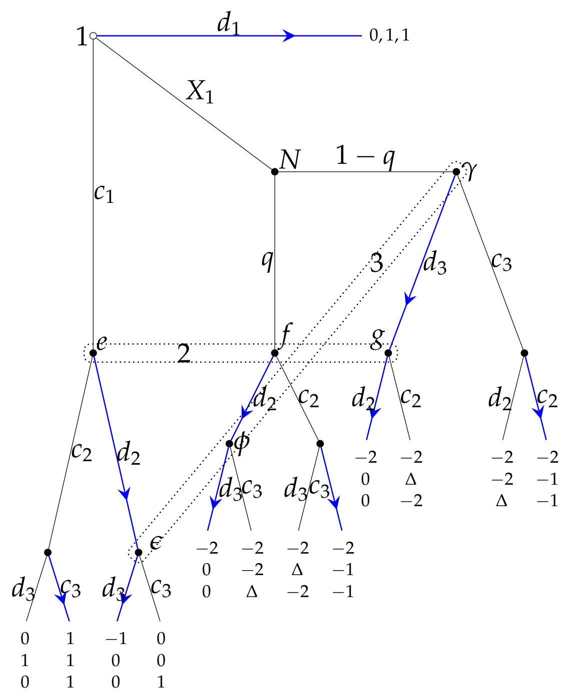

Figure 6: The game from Figure 2

This system of beliefs satisfies Bayes' rule where possible, given the strategy profile and Nature's move probabilities $(q, 1-q)$: conditional on reaching Nature's move, the conditional probabilities of arriving at nodes f and g are q and $1-q$ respectively. Similarly for ϕ and γ.

The choice of d_1, earning a payoff of 0, is sequentially rational for Player 1 because a deviation to c_1 would earn -1 and a deviation to X_1 would earn -2.

Consider now Player 2. Conditional on her non-singleton information set, and given the beliefs specified above, her expected payoff from the equilibrium action d_2 is 0, while the action c_2 would yield $-q + \Delta(1-q) \leq 1 - 2q$, since $\Delta \leq 1$. Similarly, Player 3 earns 0 at her non-singleton information set but her expected payoff from deviating would be $\Delta q - (1-q) \leq 1 - 2q$. Since $q = 1/2$, sequential rationality is satisfied for both players.

We have therefore exhibited a Perfect Bayesian Equilibrium in which Player 1 plays d_1 with probability 1.

We can also use this example to illustrate our necessary conditions. Recall that when $\Delta > 1$ the action d_1 is excluded from the set D. It follows then from our converse result Theorem 4 that there is no Perfect Bayesian equilibrium in which d_1 can be played. Let us now show that the strategy profile exhibited in Figure 6 is no longer a Perfect Bayesian equilibrium when $\Delta > 1$.

In order for the strategy profile to be sequentially rational, a necessary condition is that the choice of d_2 maximize the expected payoff of Player 2 at her information set. Let $(\nu(e), \nu(f), \nu(g))$ be the probabilities assigned to nodes e, f, and g. At that information set the action d_2 earns Player 2 a payoff of 0 and it is therefore sequentially rational iff

$$\nu(e) \cdot 1 + \nu(f) \cdot (-1) + \nu(g) \cdot \Delta \leq 0.$$

This inequality holds only if

$$\frac{\nu(f)}{\nu(g)} \geq \Delta$$

By the analogous calculation, at Player 3's information set the action d_3 is sequentially rational iff

$$\nu(\epsilon) \cdot 1 + \nu(\phi) \cdot \Delta + \nu(\gamma) \cdot (-1) \leq 0$$

and only if

$$\frac{\nu(\gamma)}{\nu(\phi)} \geq \Delta.$$

In a Perfect Bayesian Equilibrium the system of beliefs satisfies Bayes' rule where possible. This places restrictions on the beliefs at the information sets of Players 2 and 3. In particular, given the probabilities of Nature's moves, $(q, 1-q)$, and given the pure strategies of the players, we must have

$$\frac{\nu(f)}{\nu(g)} = \frac{q}{1-q} \qquad \text{and} \qquad \frac{\nu(\gamma)}{\nu(\phi)} = \frac{1-q}{q}$$

and thus

$$\Delta \leq \frac{q}{1-q} \leq \frac{1}{\Delta}$$

is a necessary condition for the profile to be a Perfect Bayesian Equilibrium. When $\Delta > 1$ there is no q that can satisfy this condition. Indeed there is no Perfect Bayesian Equilibrium of this or any admissible extensive form in which Player 3 plays d_3.

B Proof of Theorem 3

Preliminaries Suppose $\sigma_{-i} \in h_i$, and σ_i is any pure strategy for player i. There is a unique node y within h that is reachable given σ_{-i}, and the profile (σ_i, σ_{-i}) results in a unique terminal node beginning from y. We define $U_i(\sigma \mid h_i)$ to be player i's payoff at that terminal node. Given a conditional probability system μ, we then write

$$U_i(\sigma_i, \mu \mid h_i) = \sum_{\tilde{\sigma}_{-i} \in \Sigma_{-i}} \mu(\tilde{\sigma}_{-i} \mid h_i) \cdot U_i(\sigma_i, \tilde{\sigma}_{-i} \mid h_i).$$

The following lemma is a characterization of sequential rationality in terms of pure strategies and conditional probability systems. Its proof is a simple re-arranging of definitions.

Lemma 2. *If an assessment (β, ν) satisfies Bayes' rule where possible with conditional probability system μ then (β, ν) is sequentially rational if and only if for every information set h_i and σ_i such that $\mu(\sigma_i \mid \Sigma) > 0$,*

$$U_i(\sigma_i, \mu \mid h_i) \geq U_i(\sigma_i', \mu \mid h_i),$$

for all σ_i'.

Next we present the proof of Lemma 1.

Proof of Lemma 1. We show that $D^{l,k} \subset D^{l,k+1}$ by induction on k for fixed l. Let $k = 0$. To see that $D^{l,0} \subset D^{l,1}$, note that a coordinated equilibrium $\pi \in \Delta(\Theta \times P^{D^{l-1,\infty}})$ is a deviant coordinated equilibrium $\pi \in \Delta(\Theta \times P^{D^{l-1,\infty}})$ relative to $D^{l,0}$. Having established that $D^{l,k'} \subset D^{l,k'+1}$ for all $0 \leq k' < k$, we show that $D^{l,k} \subset D^{l,k+1}$. To see this, note that $D^{l,k-1} \subset D^{l,k}$ implies that a deviant coordinated equilibrium $\pi \in \Delta(\Theta \times P^{D^{l-1,\infty}})$ relative to $D^{l,k-1}$ is a deviant coordinated equilibrium $\pi \in \Delta(\Theta \times P^{D^{l-1,\infty}})$ relative to $D^{l,k}$.

We show that $D^{l,\infty} \subset D^{l-1,\infty}$ for all $l \geq 1$ by induction on l. For $l = 1$, we have that $D^{1,\infty} \subset D^{0,\infty} \equiv (A_1, \ldots, A_N)$, by definition. Assume then that we have established that $D^{l',\infty} \subset D^{l'-1,\infty}$ for $1 \leq l' < l$. To show that $D^{l,\infty} \subset D^{l-1,\infty}$, first note that $D^{l,0} \subset D^{l-1,0}$ since $D^{l-1,\infty} \subset D^{l-2,\infty}$ implies that any coordinated equilibrium $\pi \in \Delta(\Theta \times P^{D^{l-1,\infty}})$ is a coordinated equilibrium $\pi \in \Delta(\Theta \times P^{D^{l-2,\infty}})$. Inductively, it follows that $D^{l,k+1} \subset D^{l-1,k+1}$. After all, any deviant coordinated equilibrium $p \in \Delta(\Theta \times P^{D^{l-1,\infty}})$ relative to $D^{l,k}$ is a deviant coordinated equilibrium

$p \in \Delta(\Theta \times P^{D^{l-2,\infty}})$ relative to $D^{l-1,k}$ since $D^{l,k} \subset D^{l-1,k}$, and $P^{D^{l-1,\infty}} \subset P^{D^{l-2,\infty}}$ because $D^{l-1,\infty} \subset D^{l-2,\infty}$.

Since $D^{l,\infty} = \cup_k D^{l,k}$ and $\cup_k D^{l,k} \subset \cup_k D^{l-1,k}$, it follows that $D^{l,\infty} \subset D^{l-1,\infty}$. $\square$

Before proceeding with the proof of Theorem 3, we provide one last piece of notation that is useful for the construction. For $l \geq 1, k \geq 1$, let $Q^{l,k}$ denote the set of deviant plans (p, x) such that $p \in P^{D^{l-1,\infty}}$ and such that the recommended action at each node preceding $x$belong to $D^{l,k}$. Since the labels of the nodes leading to the deviant node are payoff irrelevant, $D_i^{l,k+1}$ coincides with the set of actions for i that are on the path of a deviant coordinated equilibrium $\pi \in \Delta(\Theta \times Q^{l,k})$ relative to $D^{l,k}$. In what follows, when we define deviant coordinated equilibrium we do so using the sets $Q^{l,k}$ as opposed to the whole set $P^{D^{l-1,\infty}}$.

Proof of Theorem 3. The bulk of the proof is the construction of a conditional probability system μ which we then translate into a system of beliefs which, when coupled with the obedient strategy profile, forms a Perfect Bayesian Equilibrium assessment.

Consider the canonical extensive form $\Gamma^D(\pi)$, where Nature's mixed/behavioral strategy is the given coordinated equilibrium π. Let σ^*_{-0} denote the obedient strategy profile of the players (excluding Nature) in $\Gamma^D(\pi)$. Because σ^*_{-0} is a pure strategy profile and because Nature moves only at the beginning in Γ^D, the probability distribution $\beta^\Sigma(y)$ is degenerate for any non-initial node y of Γ^D. It assigns probability 1 to the strategy profile which leads to y and is elsewhere obedient. Call this strategy profile $\sigma^*|_y$.

To begin with we set $\mu(\cdot \mid \Sigma) = \left(\pi, \sigma^*_{-0}\right)$. That is, the unconditional probability distribution $\mu(\cdot \mid \Sigma)$ over pure strategy profiles assigns probability 1 to the players choosing σ^*_{-0} and has a marginal distribution over Nature's pure strategies that coincides with π. Let F^0 be the support of $\mu(\cdot \mid \Sigma)$.

To define $\mu(\cdot \mid \Sigma \setminus F^0)$ we follow the inductive procedure that defined D. Let E_i^1 be the set of actions for i that have positive probability in the outcome of the coordinated equilibrium π. By the fixed-point property of D, we have $E_i^1 \subset D_i^{\infty,1}$.

Define inductively R^k to be the set of deviant plans (p, x) such that $p \in P^D$ and such that the recommended action at each node preceding x belongs to E^k. Then

E_i^{k+1} is the set of a_i on the path of some deviant coordinated equilibrium relative to E^k which assigns probability 1 to R^k.

Since $E^1 \subset D^{\infty,1}$ it follows that $E^k \subset D^{\infty,k}$. Furthermore $D^{\infty,1} \subset E^2$ because any coordinated equilibrium is a deviant coordinated equilibrium. Similarly $D^{\infty,k} \subset E^{k+1}$. The sandwiching $D^{\infty,k} \subset E^{k+1} \subset D^{\infty,k+1}$ implies a common fixed point: $E^\infty = D^{\infty,\infty} = D$.

For each player i, for each $a_i \in E_i^k$ there is a deviant coordinated equilibrium $\pi_{a_i} \in \Delta(\Theta \times R^{k-1})$ relative to E^{k-1} with a_i on its path. Define

$$\psi^k = \frac{1}{|E^k|} \sum_{i \in N, a_i \in E_i^k} \pi_{a_i}$$

By the linearity of the obedience constraints, the mixture $\psi^k \in \Delta(\Theta \times R^{k-1})$ is a deviant coordinated equilibrium relative to E^{k-1}. Moreover

$$a_i \in E_i^k \text{ if and only if } \psi^k(\langle a_i \rangle) > 0. \tag{4}$$

We will translate ψ^k into a probability $\tilde{\mu}^k$ over strategy profiles by translating deviant plans into pure strategy profiles. Let $q = (p, x)$ be a deviant plan and θ a state. Construct the associated strategy profile $\phi(\theta, q) = (\sigma_0, \sigma_1, \dots, \sigma_N)$ as follows. Nature's pure strategy σ_0 is (θ, p). The pure strategy σ_i for player i is obedient at all nodes that do not strictly precede x in Γ^D. If y strictly precedes (i.e. is on the path to) x then $\sigma_i(y) = b_i$ where b_i is the branch from y along the path to x. Note that $\phi(\theta, q) = \sigma^*|_y$ for any y that does not strictly precede x.

In particular, $\phi(\theta, q)$ replicates the path of q. To elaborate, let $\langle i, a_i \rangle$ denote the information set in Γ^D at which player i is recommended a_i. A node labeled (i, a_i) is on the path of q only if $\phi(\theta, q)$ reaches the information set $\langle i, a_i \rangle$ and is obedient thereafter. We record here one implication of this that will be used later. Suppose $\tilde{\sigma} = \phi(\theta, q)$ and $\tilde{\sigma}_{-i} \in \langle i, a_i \rangle$. Consider the payoff to i from playing a_i in state θ when the plan is q, i.e. $u_i(a_i, q, \theta)$. This payoff is the same as $U_i(\tilde{\sigma} \mid \langle i, a_i \rangle)$. Indeed, since $\tilde{\sigma}_i$ is obedient at and after $\langle i, a_i \rangle$, this payoff is also the same as $U_i(\sigma_i^*, \tilde{\sigma}_{-i} \mid \langle i, a_i \rangle)$.

Now for any $\sigma \in \Sigma$, set

$$\tilde{\mu}^k(\sigma) = \psi^k(\phi^{-1}(\sigma)).$$

From Equation 4 we have

$$a_i \in E_i^k \text{ if and only if } \tilde{\mu}^k(\langle i, a_i \rangle) > 0. \tag{5}$$

Finally inductively define[14]

$$\mu(\cdot \mid \Sigma \setminus F^{k-1}) = \tilde{\mu}^k(\cdot \mid \Sigma \setminus F^{k-1})$$

and

$$F^k = F^{k-1} \bigcup \operatorname{supp} \mu(\cdot \mid \Sigma \setminus F^{k-1}).$$

for each $k = 1, \ldots, \bar{k}$, where $\bar{k}$ is defined by $E^\infty = E^{\bar{k}} \neq E^{\bar{k}-1}$. From Equation 5 we have

$$\text{If } a_i \in E_i^k \setminus E_i^{k-1} \text{ then } \langle i, a_i \rangle \subset \Sigma \setminus F^{k-1} \text{ and } \mu(\langle i, a_i \rangle \mid \Sigma \setminus F^{k-1}) > 0. \tag{6}$$

Let $F = F^{\bar{k}}$. If $F \neq \Sigma$ then extend μ by defining $\mu(\cdot \mid \Sigma \setminus F)$ arbitrarily.[15]

Now we verify the conditions for Perfect Bayesian Equilibrium. We define the system of beliefs ν as follows. For each information set $\langle i, a_i \rangle$ and node $y \in \langle i, a_i \rangle$, set $\nu(y \mid \langle i, a_i \rangle) = \mu(y \mid \langle i, a_i \rangle)$. By construction the first condition in the definition of Bayes' rule where possible is satisfied.

For the second condition consider any node y and information set $\langle i, a_i \rangle$ of Γ^D such that $\mu(y \mid \langle i, a_i \rangle) > 0$. It is enough to show that $\mu(\sigma^*|_y \mid y) = 1$.

Let k be such that a_i belongs to E_i^k but not E_i^{k-1}. Then by Equation 6, $\langle i, a_i \rangle \subset \Sigma \setminus F^{k-1}$ and $\mu(\langle i, a_i \rangle \mid \Sigma \setminus F^{k-1}) > 0$. Suppose σ is such that $\mu(\sigma \mid y) > 0$. Then by the chain rule, $\mu(\sigma \mid \Sigma \setminus F^{k-1}) = \mu(\sigma \mid y) \cdot \mu(y \mid \langle i, a_i \rangle) \cdot \mu(\langle i, a_i \rangle \mid \Sigma \setminus F^{k-1})$ and since all of the factors on the right-hand side are positive we have $\mu(\sigma \mid \Sigma \setminus F^{k-1}) > 0$, in particular $\tilde{\mu}^k(\sigma) > 0$. By construction of $\tilde{\mu}^k$ this means that $\sigma = \phi(\theta, q)$ for some deviant plan $q = (p, x)$ where $q \in R^{k-1}$. By the definition of ϕ, we have $\sigma = \sigma^*|_y$ for all nodes y' that do not strictly precede x. But since $q \in R^{k-1}$, the only nodes that can strictly precede x are those that recommend actions in E^{k-1}. Since $a_i \notin E^{k-1}$, the node y does not strictly precede x and hence $\sigma = \sigma^*|_y$ and we have proven that $\mu(\sigma^*|_y \mid y) = 1$.

[14] The right-hand side is the conventional conditional probability, i.e. $\tilde{\mu}^k(\sigma|\Sigma \setminus F^{k-1}) = \tilde{\mu}^k(\sigma)/\tilde{\mu}^k(\Sigma \setminus F^{k-1})$.

[15] In Doval and Ely (2019), we show that given a conditional probability system, μ, on a set, X_1. and a measure over a set X_2 such that $X_1 \cap X_2 = \emptyset$, one can construct a new conditional probability system, $\tilde{\mu}$ on $X_1 \cup X_2$, that coincides with μ on $2^{X_1} \setminus \{\emptyset\}$.

We finally turn to sequential rationality. Consider any information set $\langle i, a_i \rangle$ such that $a_i \in D_i$. By definition

$$U_i(\sigma_i^*, \mu \mid \langle i, a_i \rangle) = \sum_{\tilde{\sigma}_{-i} \in \Sigma_{-i}} \mu(\tilde{\sigma}_{-i} \mid \langle i, a_i \rangle) \cdot U_i(\sigma_i^*, \tilde{\sigma}_{-i} \mid \langle i, a_i \rangle).$$

Let k be such that a_i belongs to E_i^k but not E_i^{k-1}. Then by Equation 6, $\langle i, a_i \rangle \subset \Sigma \setminus F^{k-1}$ so that

$$\mu(\tilde{\sigma}_{-i} \mid \langle i, a_i \rangle) = \mu(\tilde{\sigma}_{-i} \mid \langle i, a_i \rangle \cap \Sigma \setminus F^{k-1})$$

and by the chain rule $\mu(\tilde{\sigma}_{-i} \mid \Sigma \setminus F^{k-1}) = \mu(\tilde{\sigma}_{-i} \mid \langle i, a_i \rangle \cap \Sigma \setminus F^{k-1}) \cdot \mu(\langle i, a_i \rangle \mid \Sigma \setminus F^{k-1})$. Since $a_i \in E_i^k \setminus E_i^{k-1}$, by Equation 6 it follows that $\mu(\langle i, a_i \rangle \mid \Sigma \setminus F^{k-1}) > 0$, hence

$$\mu(\tilde{\sigma}_{-i} \mid \langle i, a_i \rangle) = \frac{\mu(\tilde{\sigma}_{-i} \mid \Sigma \setminus F^{k-1})}{\mu(\langle i, a_i \rangle \mid \Sigma \setminus F^{k-1})}.$$

We can therefore write

$$\begin{aligned}
U_i(\sigma_i^*, \mu \mid \langle i, a_i \rangle) &= \frac{1}{\mu(\langle i, a_i \rangle \mid \Sigma \setminus F^{k-1})} \sum_{\tilde{\sigma}_{-i} \in \langle i, a_i \rangle} \mu(\tilde{\sigma}_{-i} \mid \Sigma \setminus F^{k-1}) \cdot U_i(\sigma_i^*, \tilde{\sigma}_{-i} \mid \langle i, a_i \rangle) \\
&= \frac{1}{\tilde{\mu}^k(\langle i, a_i \rangle \mid \Sigma \setminus F^{k-1})} \sum_{\tilde{\sigma}_{-i} \in \langle i, a_i \rangle} \tilde{\mu}^k(\tilde{\sigma}_{-i} \mid \Sigma \setminus F^{k-1}) \cdot U_i(\sigma_i^*, \tilde{\sigma}_{-i} \mid \langle i, a_i \rangle) \\
&= \frac{1}{\tilde{\mu}^k(\langle i, a_i \rangle \mid \Sigma \setminus F^{k-1})} \sum_{\tilde{\sigma}_{-i} \in \langle i, a_i \rangle} \frac{\tilde{\mu}^k(\tilde{\sigma}_{-i})}{\tilde{\mu}^k(\Sigma \setminus F^{k-1})} \cdot U_i(\sigma_i^*, \tilde{\sigma}_{-i} \mid \langle i, a_i \rangle) \\
&= \frac{1}{\tilde{\mu}^k(\langle i, a_i \rangle)} \sum_{\tilde{\sigma}_{-i} \in \langle i, a_i \rangle} \tilde{\mu}^k(\tilde{\sigma}_{-i}) \cdot U_i(\sigma_i^*, \tilde{\sigma}_{-i} \mid \langle i, a_i \rangle) \\
&= \frac{1}{\tilde{\mu}^k(\langle i, a_i \rangle)} \sum_{\tilde{\sigma}_{-i} \in \langle i, a_i \rangle} \sum_{\tilde{\sigma}_i \in \Sigma_i} \tilde{\mu}^k(\tilde{\sigma}) \cdot U_i(\sigma_i^*, \tilde{\sigma}_{-i} \mid \langle i, a_i \rangle) \\
&= \frac{1}{\psi^k(\langle a_i \rangle)} \sum_{(\theta, q): q \in \langle a_i \rangle} \psi^k(\theta, q) \cdot u_i(a_i, q, \theta)
\end{aligned}$$

And similarly, if σ_i' were an alternative strategy for player i which plays b_i at information set $\langle i, a_i \rangle$, then the same derivation leads to

$$U_i(\sigma_i', \mu \mid \langle i, a_i \rangle) = \frac{1}{\psi^k(\langle a_i \rangle)} \sum_{(\theta, q): q \in \langle a_i \rangle} \psi^k(\theta, q) \cdot u_i(b_i, q, \theta)$$

and since ψ^k is a deviant coordinated equilibrium relative to E^{k-1} and $a_i \notin E_i^{k-1}$ the obedience constraint holds and implies that

$$U_i(\sigma_i^*, \mu \mid \langle i, a_i \rangle) \geq U_i(\sigma_i', \mu \mid \langle i, a_i \rangle)$$

i.e., by Lemma 2 sequential rationality is satisfied. $\square$

C Proof of Theorem 4

We will prove item 3 of the Theorem, the other two parts are similar (but simpler). We present the proof for the case of pure-strategy Perfect Bayesian equilibria of *simple* extensive forms. These are admissible extensive forms in which Nature moves only once and at the beginning of the game. We show in Doval and Ely (2019) that any Perfect Bayesian Equilibrium of any admissible extensive form is outcome equivalent to a pure-strategy equilibrium of a simple extensive form. The idea is that any mixing by the players can be replicated by random moves by Nature at the beginning of the game.

Extensive form notation: We denote finite extensive forms with perfect recall by Γ with nodes V. We use y to denote generic nodes of the tree. The precedence relation in the tree is denoted by $\prec$. The subset $Z \subset V$ denotes the terminal nodes, with typical element z.

The outcome function $\gamma : Z \mapsto A \times \Theta$ associates each terminal node to an outcome of the base game. We assume that γ is onto. For a non-terminal node, y, we let $\gamma(y) = \cup_{z \in Z : y \prec z} \gamma(z)$, denote the set of outcomes consistent with terminal nodes that can be reached from y. For any $y \in V$, we let $\gamma_\Theta(y)$, $\gamma_A(y)$, and $\gamma_i(y)$ denote the projection of $\gamma(y)$ on Θ, A, and A_i, respectively.

The initial node of the extensive form belongs to nature, who selects the state according to the prior distribution, ρ. Thus, we assume that $\gamma_\Theta(y)$ is a singleton if y is not the initial node.

Finally, we denote by H_i player i's information partition, that is, a partition of the set of nodes, y, such that player i moves at node y. We denote by h_i, h_i' generic elements of H_i and we let $M(h_i)$ denote the set of moves available to player i at h_i.

Preliminaries Throughout we fix an admissible extensive form Γ, and a pure strategy Perfect Bayesian Equilibrium which we refer to by its pure strategy profile σ^* and its associated conditional probability system μ. We construct the following subsets of strategy profiles. First

$$F^1 = \operatorname{supp} \mu(\cdot \mid \Sigma),$$

and inductively define $\overline{F}^k = \cup_{l=1}^{k-1} F^l$ and

$$F^k = \operatorname{supp} \mu(\cdot \mid \Sigma \setminus \overline{F}^k).$$

Define $E^k = \gamma_A(F^k)$ and $\overline{E}^k = \gamma_A(\overline{F}^k)$ to be the set of action profiles that are reached by strategies in F^k and $\overline{F}^k$ respectively. Let $B = \bigcup_{k=1}^{\infty} E^k$.

The following lemma is the main idea of the proof.

Lemma 3. *For each k there exists a map $\phi_k : F^k \to \Theta \times Q$ such that $\tilde{\sigma}$ and $\phi_k(\tilde{\sigma})$ have the same outcome*[16] *and the measure $\pi \in \Delta(\Theta \times Q)$ defined by*

$$\pi(\theta, q) = \mu(\phi_k^{-1}(\theta, q) \mid \Sigma \setminus \overline{F}^k). \tag{7}$$

is a deviant coordinated equilibrium relative to $\overline{E}^k$ assigning probability 1 to plans in P^B whose recommended action prior to the deviant node belongs to $\overline{E}^k$.

Theorem 4 is proven by a repeated application of Lemma 3

Proof of Theorem 4. First note that in Lemma 3 for the case of $k = 1$, we have $\overline{E}^k = \emptyset$, and therefore π is a coordinated equilibrium. The distribution over terminal nodes implied by $\mu(\mid \Sigma)$ coincides with that of the outcome of π. In item 2 of Definition 5, taking y to be the initial node, we see that this equals the outcome of the equilibrium strategy profile. All that remains to prove the Theorem is to show that $B \subset D$.

Set $D^{0,\infty} = A$, we will show inductively that $B \subset D^{l,\infty}$ for all l, hence $B \subset D$. By definition, $B \subset A$, hence the inductive hypothesis holds for $l = 0$.

Assume now that $B \subset D^{l,\infty}$; we show that $B \subset D^{l+1,\infty}$. Lemma 3 provides a coordinated equilibrium $\pi \in \Delta(\Theta \times P^B) \subset \Delta(\Theta \times P^{D^{l,\infty}})$ in which the actions that occur with positive probability are $E^1 = \gamma_A(F^1)$. Thus $E^1 = \overline{E}^2 \subset D^{l+1,1}$.

[16] The outcome of a strategy $\tilde{\sigma}$ is the outcome associated with the terminal node reached by $\tilde{\sigma}$. In a slight abuse of notation, we sometimes denote it by $\gamma(\tilde{\sigma})$.

Assume now $\overline{E}^k \subset D^{l+1,k-1}$ for $k \geq 2$. We show next that $\overline{E}^{k+1} \subset D^{l+1,k}$. This will conclude the proof because $B = \cup_k \overline{E}^k \subset \cup_k D^{l+1,k} = D^{l+1,\infty}$.

Lemma 3 for k provides a deviant coordinated equilibrium $\pi \in \Delta(\Theta \times P^B) \subset \Delta(\Theta \times P^{D^{l,\infty}})$ relative to $\overline{E}^k \subset D^{l+1,k-1}$ such that all the actions recommended along the path to the deviant nodes belong to $\overline{E}^k \subset D^{l+1,k-1}$ and the set of actions played along the path equaling $\gamma_A(F^k) = E^k$. Thus $E^k \subset D^{l+1,k}$, and $\overline{E}^{k+1} = \overline{E}^k \cup E^k \subset D^{l+1,k-1} \cup D^{l+1,k} = D^{l+1,k}$. □

In the remainder of this section we will construct the mappings ϕ_k and prove Lemma 3. The following lemma distills the essential property of an admissible extensive form. It shows that as long as a player has not yet determined her action, she has a continuation strategy which will guarantee that she plays any action b_i in the base game. The proof is mechanical and deferred to Appendix D.

Say that a node y is *decisive for a_i* if player i's action is not determined at y, but i has a move m_i at y which determines action a_i. In that case we say that m_i is decisive for a_i. If h_i is the information set that contains y, then by the property of know-your-own-action, m_i is decisive for a_i at every node in h_i, so we say that h_i is decisive for a_i.

Lemma 4. *For any decisive information set h_i belonging to i, for any action $b_i \in A_i$ there is a strategy $\sigma_i \in \Sigma_i$ such that for all $\sigma_{-i} \in h_i$ the path of the profile (σ_i, σ_{-i}) passes through h_i and leads to an outcome in which player i plays b_i.*

Proof of Lemma 3. Let $\sigma \in F^k$, and $(a, \theta) = \gamma(\sigma)$. We will construct a deviant plan $q = (p, x)$ having the same outcome and incentives and we will define $\phi(\sigma) = (\theta, q)$.

We begin by building a single path which will be the obedient path of q. For each individual action a_j in the profile a let $y^1, \ldots, y^N$ be the corresponding decisive nodes that are on the path of σ, arranged in their order along the path.We let $\iota(j)$ be the permutation of $1, \ldots, N$ that denotes the player who moves at node y^j. (So in particular the node y^j is a decisive node for action $a_{\iota(j)}$.)

Construct a path in plan p consisting of $N+1$ consecutive nodes where the jth node along the path has $\iota(j)$ moving, and the $N+1$st node is terminal. The branch that connects the jth node to the $j+1$st node is labeled with the action $a_{\iota(j)}$ in the profile a.)

The deviant node x will be one of the nodes along this path. Let h^j be the information set that contains y^j. Pick the j such that y^j is the earliest decisive node along the path of σ such that $h^j \subset \Sigma \setminus \overline{F}^k$. We set x to be the jth node along the path just constructed. If there is no such decisive node along the path of σ we set x to be the terminal node of the path just constructed.

The recommended actions along the path are defined differently for nodes before and after x. For $j' \geq j$ the action recommended at the j'th node is $a_{\iota(j')}$. For $j' < j$ choose any action from $\overline{E}^k$ as a recommendation at the j'th node. Thus, the path we have constructed is the obedient path.

Next we build the paths that fork from the obedient path. Consider $j' \geq j$, write $i = \iota(j')$ and consider any action b_i different from the recommended action at the j'th node. Lemma 4 provides a strategy σ_i' such that the profile (σ_i', σ_{-i}) passes through $h^{j'}$ and leads to an outcome in which the action profile is (b_i, z_{-i}) for some $z_{-i} \in A_{-i}$. Add the branch labeled b_i leading from the j'th node and append the obedient continuation path that leads to the action profile (b_i, z_{-i}).

Complete the remainder of plan p arbitrarily. We now establish properties of the resulting mapping $\phi_k : F^k \to \Theta \times Q$. First, by construction if $(\theta, q) \in \phi_k(F^k)$ then q belongs to P^B and all of the action recommendations preceding the deviant node belong to $\overline{E}^k$.

Next say that a decisive information set h_i for an action a_i is σ^*-decisive if at information set h_i, player i's equilibrium pure strategy σ_i^* selects the move that determines action a_i. Denote by (i, a_i) the family of all information sets that are σ^*-decisive for a_i. And define $\langle i, a_i \rangle$ to be the set of all pure-strategy profiles leading to information sets in (i, a_i):

$$\langle i, a_i \rangle = \{\sigma \in \Sigma : \text{There exist } y, h_i \text{ such that } \sigma \in y \in h_i \in (i, a_i)\}$$

We will now demonstrate that if $a_i \in E_i^k \setminus \overline{E}_i^k$ then

$$\phi_k^{-1}\left(\{(\theta, q) : q \in \langle a_i \rangle\}\right) = \langle i, a_i \rangle . \tag{8}$$

First, if h_i is a decisive information set for a_i, then we claim $\mu(h_i \mid F^{k'}) = 0$ for all $k' < k$. If not, then let k' be the smallest for which $\mu(h_i \mid F^{k'}) > 0$. There exists $\sigma \in h_i$ such that $\mu(\sigma \mid F^{k'}) > 0$ and by the chain rule $\mu(\sigma \mid h_i) \cdot \mu(h_i \mid F^{k'}) = \mu(\sigma \mid F^{k'})$

implying also $\mu(\sigma \mid h_i) > 0$. By the definition of Bayes' rule where possible this requires $\sigma(h_i) = \sigma^*(h_i)$ and thus $a_i \in \gamma_i(F^{k'})$, i.e. $a_i \in \overline{E}_i^{k'}$, a contradiction.

To show Equation 8, let $\sigma \in F^k$ with $\phi(\sigma) = (\theta, q) \in \langle a_i \rangle$. Then the obedient path of q has a node at which player i is recommended to play a_i. Since $a_i \notin \overline{E}^k$ this node weakly succeeds the deviant node of q. That means that the path of σ has a decisive node y for a_i at which σ selects the move that determines a_i. All that remains to show $\sigma \in \langle i, a_i \rangle$ is that $\sigma^*(y) = a_i$. Letting h_i be the information set that contains y we have that h_i is a decisive information set for a_i and therefore $\mu(h_i \mid F^{k'}) = 0$ for all $k' < k$. Thus $h_i \subset \Sigma \setminus \overline{F}^k$ and by the chain rule

$$\begin{aligned}\mu(y \mid \Sigma \setminus \overline{F}^k) &= \mu\left(y \mid h_i \cap [\Sigma \setminus \overline{F}^k]\right) \cdot \mu\left(h_i \cap [\Sigma \setminus \overline{F}^k] \mid \Sigma \setminus \overline{F}^k\right) \\ &= \mu(y \mid h_i) \cdot \mu(h_i \mid \Sigma \setminus \overline{F}^k).\end{aligned}$$

Since $\mu(h_i \mid \Sigma \setminus \overline{F}^k) \geq \mu(y \mid \Sigma \setminus \overline{F}^k) \geq \mu(\sigma \mid \Sigma \setminus \overline{F}^k) > 0$ this implies $\mu(y \mid h_i) > 0$ and by the definition of Bayes' rule where possible we conclude $a_i = \sigma(y) = \sigma^*(y)$.

We conclude the proof by showing that the measure $\pi \in \Delta(\Theta \times Q)$ defined by

$$\pi(\theta, q) = \mu(\phi_k^{-1}(\theta, q) \mid \Sigma \setminus \overline{F}^k).$$

is a deviant coordinated equilibrium relative to $\overline{E}^k$.

Sequential rationality implies that for any information set h_i, and alternative strategy σ_i'

$$U_i(\sigma_i^*, \mu \mid h_i) \geq U_i(\sigma_i', \mu \mid h_i).$$

Pick $a_i \notin \overline{E}^k$, consider any alternative action b_i and let σ_i' be the strategy which is identical to σ_i^* except that at every information set that is decisive for a_i, it follows the continuation (provided by Lemma 4) which determines b_i. Then

$$\sum_{h_i \in (i, a_i)} \mu(h_i \mid \Sigma \setminus \overline{F}^k) \left[U_i(\sigma_i^*, \mu \mid h_i) - U_i(\sigma_i', \mu \mid h_i)\right] \geq 0$$

$$\sum_{h_i \in (i, a_i)} \mu(h_i \mid \Sigma \setminus \overline{F}^k) \sum_{\sigma_{-i} \in h_i} \mu(\sigma_{-i} \mid h_i) \left[U_i(\sigma_i^*, \sigma_{-i} \mid h_i) - U_i(\sigma_i', \sigma_{-i} \mid h_i)\right] \geq 0$$

$$\sum_{h_i \in (i, a_i)} \mu(h_i \mid \Sigma \setminus \overline{F}^k) \sum_{\sigma_{-i} \in h_i} \sum_{\sigma_i \in \Sigma_i} \mu(\sigma \mid h_i) \left[U_i(\sigma_i^*, \sigma_{-i} \mid h_i) - U_i(\sigma_i', \sigma_{-i} \mid h_i)\right] \geq 0$$

The next step uses the fact that $\mu(\sigma \mid h_i) > 0$ only if $\sigma \in y$ for some $y \in h_i$.

$$\sum_{h_i \in (i,a_i)} \mu(h_i \mid \Sigma \setminus \overline{F}^k) \sum_{y \in h_i} \sum_{\sigma \in y} \mu(\sigma \mid h_i) \left[U_i(\sigma_i^*, \sigma_{-i} \mid h_i) - U_i(\sigma_i', \sigma_{-i} \mid h_i)\right] \geq 0$$

$$\sum_{h_i \in (i,a_i)} \sum_{y \in h_i} \sum_{\sigma \in y} \mu(\sigma \mid \Sigma \setminus \overline{F}^k) \left[U_i(\sigma_i^*, \sigma_{-i} \mid h_i) - U_i(\sigma_i', \sigma_{-i} \mid h_i)\right] \geq 0$$

$$\sum_{\sigma \in \langle i,a_i \rangle} \mu(\sigma \mid \Sigma \setminus \overline{F}^k) \left[U_i(\sigma_i^*, \sigma_{-i} \mid h_i) - U_i(\sigma_i', \sigma_{-i} \mid h_i)\right] \geq 0$$

$$\sum_{(\theta,q) \in \Theta \times \langle a_i \rangle} \mu(\phi_k^{-1}(\theta, q) \mid \Sigma \setminus \overline{F}^k) \left[U_i(\sigma_i^*, \sigma_{-i} \mid h_i) - U_i(\sigma_i', \sigma_{-i} \mid h_i)\right] \geq 0$$

The next step uses Equation 8.

$$\sum_{(\theta,q) \in \Theta \times \langle a_i \rangle} \mu(\phi_k^{-1}(\theta, q) \mid \Sigma \setminus \overline{F}^k) \left[u_i(a_i, q, \theta) - u_i(b_i, q, \theta)\right] \geq 0$$

$$\sum_{\theta \in \Theta, q \in \langle a_i \rangle} \pi(\theta, q) \left[u_i(a_i, q, \theta) - u_i(b_i, q, \theta)\right] \geq 0.$$

$\square$

D Proof of Lemma 4

Since $|\gamma_i(y)| > 1$, no partial commitments implies that $\gamma_i(y) = A_i$. Letting (y, m) denote the node that follows after taking move m at y, we obtain that $a_i' \in \gamma_i(y) = \cup_{m \in M(y)} \gamma_i(y, m)$, where $M(y)$ denotes the set of moves available at node y. Thus, there exists $m^* \in M(y)$ such that $a_i' \in \gamma_i(y, m^*)$.[17] We claim that $\gamma_i(y, m^*) = A_i$. Clearly, $|\gamma_i(y, m^*)| > 1$. Let $\overline{y}$ denote the longest length node y' that succeeds (y, m^*) and such that $|\gamma_i(y')| > 1$. No delegation implies that $\overline{y} \in \overline{h} \in H_i$. No partial commitments implies that $\gamma_i(\overline{y}) = A_i$. Then, $A_i = \gamma_i(\overline{y}) \subseteq \gamma_i(y, m^*)$.

[17]Note that m^* does not depend on y. First, no delegation, know your action, and no partial commitments implies that since $\gamma_i(y) = A_i$, then for all nodes $y' \in h(y)$, $\gamma_i(y') = A_i$, where $h(y)$ denotes the information set to which y belongs. Since $a_i \in \gamma_i(y') = \cup_{m \in M(y)} \gamma_i(y', m)$, then there exists $m_{y'}$ such that $a_i \in \gamma_i(y', m_{y'})$. Suppose that $a_i \notin \gamma_i(y, m_{y'})$. Hence, $|\gamma_i(y, m_{y'})| = 1$ by no partial commitments. But then, know your action implies that $\gamma_i(y, m_{y'}) = \gamma_i(y', m_{y'})$, which implies that $a_i \notin \gamma_i(y', m_{y'})$; a contradiction.

Let $H_i^{\prec}$ denote the information sets of player i that succeed h after taking move m^*. Note that $H_i^{\prec}$ can be partitioned into the following sets:

1. Information sets that are decisive for a_i',
2. Information sets h' such that player i has still not chosen his action in the base game (i.e., $\gamma_i(h') = A_i$), but they are not decisive for a_i', and
3. Information sets such that player i has already chosen his action in the base game, that is, there exists $\tilde{a}_i$ such that $\{\tilde{a}_i\} = \gamma_i(h')$.

Denote these sets of information sets by $H_i^{\prec,D}(a_i'), H_i^{\prec,ND}$, and $H_i^{\prec}(\tilde{a}_i)$, respectively. We make the following observations:

1. For all $h'' \in H_i^{\prec}(\tilde{a}_i)$, there exists $h' \in H_i^{\prec,ND} \cup H_i^{\prec,D}(a_i')$ such that $h \prec h' \prec h''$.
2. $H_i^{\prec,D}(a_i') \neq \emptyset$. This follows from noting that the set $\{y' : (y, m^*) \prec y'$ and $\gamma_i(y') = \{a_i'\}\}$ is non-empty and no delegation.
3. If $h' \in H_i^{\prec,ND}$, then there exists $h'' \in H_i^{\prec,D}(a_i')$ such that $h' \prec h''$. Perfect recall implies there exists a move, $m_{h',h''}$, in $M(h')$ that satisfies that for all $y'' \in h''$, there exists $y' \in h'$ such that $(y', m_{h',h''}) \preceq y''$. Since $h'' \in H_i^{\prec,ND}$, then $\gamma_i(y', m_{h',h''}) = A_i$ for all $y'' \in h''$.

Let σ_i be such that $\sigma_i(h_i) = m^*$, $\sigma_i(h') = m_{h',h''}, h' \in H_i^{\prec,ND}, h'' \in H_i^{\prec,D}(a_i')$ and $\sigma_i(h') = m_i(a_i'), h' \in H_i^{\prec,D}(a_i')$. Note that any such σ_i precludes reaching information sets in $H_i^{\prec}(\tilde{a}_i)$ for $\tilde{a}_i \neq a_i'$. At $h' \in H_i^{\prec}(\tilde{a}_i)$, σ_i can be specified arbitrarily.

Let $\sigma_{-i} \in \Sigma_{-i}(y)$ and consider $(\sigma_i, \sigma_{-i})|_{\Sigma(y)}$. Let z denote the terminal history reached by the strategy starting from y and suppose that $\gamma_i(z) \neq \{a_i'\}$. Since $\gamma_i(y) = A_i$, the set $\{y' : y' \prec z \wedge |\gamma_i(y')| \neq 1\} \neq \emptyset$. Let $\overline{y}$ be the node of longest length in that set. Then, no delegation implies that $\overline{y} \in \overline{h} \in H_i$. We claim that $h \in H_i^{\prec,D}(a_i')$. Clearly, $h \in H_i^{\prec,D}(a_i')$. If $\overline{h} \in H_i^{\prec,ND}$, then it cannot be that $|\gamma_i(\overline{y}, \sigma_i(\overline{y}))| = 1$. Then, $\overline{h} \in H_i^{\prec,D}(a_i')$ and by definition $\gamma_i(\overline{y}, \sigma_i(\overline{y})) = \{a_i'\}$.