

Sequential Information Design

Laura Doval Jeffrey C. Ely

PRELIMINARY - PLEASE DO NOT CIRCULATE WITHOUT AUTHORS' PERMISSION

August 28, 2016*

Abstract

The literature on implementation has focused mainly on cases in which agents have both agency over their actions and over the time at which they commit to their actions. We take as given the set of agents, their set of actions, and their payoffs. We ask what distributions over actions are consistent with the players playing according to some extensive form. The main result of the paper is to show that a distribution over outcomes is implementable as a Perfect Bayesian equilibrium (PBE) of an *admissible* extensive form if, and only if, it is implementable as a PBE of a *canonical* extensive form. The latter is an admissible extensive form, in which there is a randomization device that not only sends (private) recommendations to the agents, but also selects the order in which the agents move; moreover, subsequent recommendations can be made conditional on actions already taken. This result strictly generalizes Aumann's notion of correlated equilibrium, and Bergemann and Morris' [4] notion of Bayes' correlated equilibrium.

*Compiled to update references, and apply cosmetic changes: December 15, 2017

1 INTRODUCTION

When playing a game against other players, an agent faces (potentially) two sources of uncertainty: he may have only partial information about a *state* of the world that affects his payoffs, and he may be uncertain about the actions taken by other players and whether or not these actions have already been taken. Information design in a multi-agent situation has to, therefore, take into account both the release of information about the state of the world, but also about the actions chosen by others and the timing of these actions.

In a setting in which agents are playing a game under uncertainty, that is, where we know the actions available to each agent, their payoffs from each action profile, and the initial information agents have as embodied in a common prior over the state of the world, the way to formalize the problem of the release of information about the state of the world is by looking at the Bayes' correlated equilibria (henceforth, BCE) of the game (Bergemann and Morris [4]). There are two ways of thinking about BCE in this multi-agent framework. The first is along the lines of *persuasion*. There is a principal, whose payoff may depend on the outcome of the game and the state of the world, and designs (and can commit to) an information structure. She observes the state of the world, and makes (private) recommendations to the agents; the players observe their recommendations and *simultaneously* choose their actions. The second way to think about BCE is along the lines of *robustness*: an outside observer knows the ingredients of the game, but does not know what additional information the agents observe about the state of world before they *simultaneously* choose their actions. The set of BCE is the set of Bayes Nash equilibria of the game for some information structure.

In this paper, we study multi-agent information design where the information pertains both the underlying state of the world, and other players' actions and, in particular, their timing. From the persuasion point of view, we ask: what is the set of distributions over outcomes that an interested principal can bring about when she can structure not just the information, but also the timing of actions, and the timing of the revelation of information about the state, and about

others' actions? That is, we study the case in which the principal, when designing the game the agents play, can now control the extensive form. From the robustness point of view, we ask what is the set of distributions over outcomes that is consistent with the players facing some information structure *and* playing according to some extensive form?

The main result of the paper shows that a distribution over outcomes is implementable as a Perfect Bayesian equilibrium (PBE) of an *admissible* extensive form (see example below, and Conditions 1, 2, and 3 in Section 2.1) if, and only if, it is implementable as a PBE of a *canonical* extensive form. The latter is an extensive form where there is a correlation device that, as a function of the state of the world, randomly orders the agents, and (privately) contacts them sequentially. The agents receive recommendations as a function of the state of the world, and the actions taken by others who have already been contacted. Each agent only observes his recommendation, and not who has moved before him or what actions they have taken. Thus, a principal who designs the agents' information structure and can choose any of the admissible extensive forms can, without loss of generality, restrict attention to the canonical extensive form. Similarly, if we don't know what extensive form the agents are playing, we can summarize the set of possible distributions over outcomes by characterizing the PBE of the canonical extensive form, as long as we know that the extensive form they are playing is admissible.

To illustrate the main ideas of the paper, consider the following example. There are two players, 1 and 2. They can choose to contribute or not to the production of a public good. If at least one player contributes, the public good is produced. Contributing costs $c > 0$. The value of the public good to the players depends on the state of the world, $\theta \in \{L, H\}$, where p denotes the probability of the state being H . If the state is H (= high), the good is worth v_H , and if the state is L (= low), the good is worth v_L . We assume that $0 < v_L < c < v_H$. Table 1 summarizes the payoffs of the game.

		P_2	
		C	NC
P_1	C	$v_\theta - c, v_\theta - c$	$v_\theta - c, v_\theta$
	NC	$v_\theta, v_\theta - c$	$0, 0$

Table 1: Payoffs when state is $\theta \in \{L, H\}$.

Players choose their actions without knowing the realization of θ . Assuming that p is such that $pv_H + (1 - p)v_L > c$, there are three Nash equilibria of the game: two in which only one player contributes, and one in which each player contributes with probability $\alpha \equiv \frac{\mathbb{E}_\theta v_\theta - c}{\mathbb{E}_\theta v_\theta}$.

Consider now a principal, whose objective is to maximize the number of contributions to the public good. Her most preferred Nash equilibrium is the mixed strategy one if $\alpha \geq \frac{1}{2}$, which leads to a level of contribution of 2α , and (either of) the pure strategy ones if $\alpha < \frac{1}{2}$. Suppose that she can observe the state, and design the agents' information structure. That is, we can think of the problem from the persuasion point of view, and, in what follows, we look for the BCE that maximizes the number of contributions to the good.

Figure 1 below provides a graphical depiction of the correlation device for the principal (the width and the height are of length one):



Figure 1: State space

The upper region represents the state in which the good is valuable, and its height, p , the probability with which $\theta = H$. We use the width to represent the probability with which a player is recommended an action conditional on the

state being θ . Figure 2 shows the mixed strategy for player 1 (red player), and for player 2 (blue player), where the colored areas represent the probability that each contributes in each state:

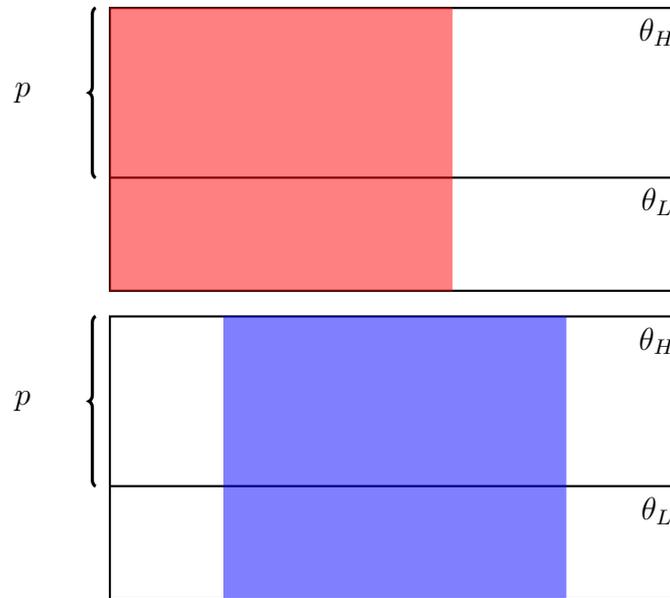


Figure 2: A mixed strategy for the red player (top), and the blue player (bottom)

Figure 3 shows the mixed strategy profile. In Figure 3, the (purple) overlap represents the event in which both players contribute (the independence assumption of Nash equilibrium gets reflected in that the area of the overlap is α^2 .) In the Nash equilibrium, the red player contributes with total probability α , and, with some probability, he is the only one contributing, that is he is pivotal. The probability of being pivotal is the one which determines his incentives to contribute.

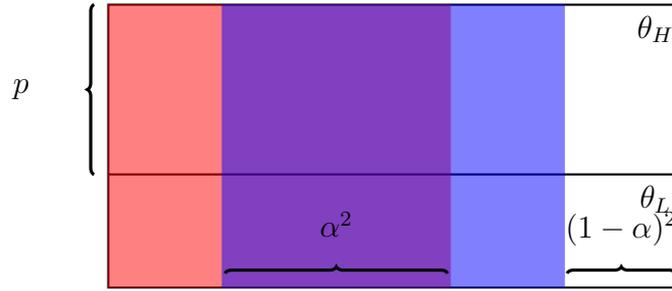


Figure 3: The mixed strategy Nash equilibrium

Note that the principal can easily improve upon the (mixed strategy) Nash equilibrium. First, the principal can keep the total contribution level the same by just shifting the blue region to the right, as depicted in Figure 4. By doing so, he relaxes the incentive constraint for the red player: even though he contributes with the same total probability, the conditional probability that he is pivotal has just gone up. Hence, the principal can use the slack to increase the probability with which each player contributes, as depicted at the bottom of Figure 4

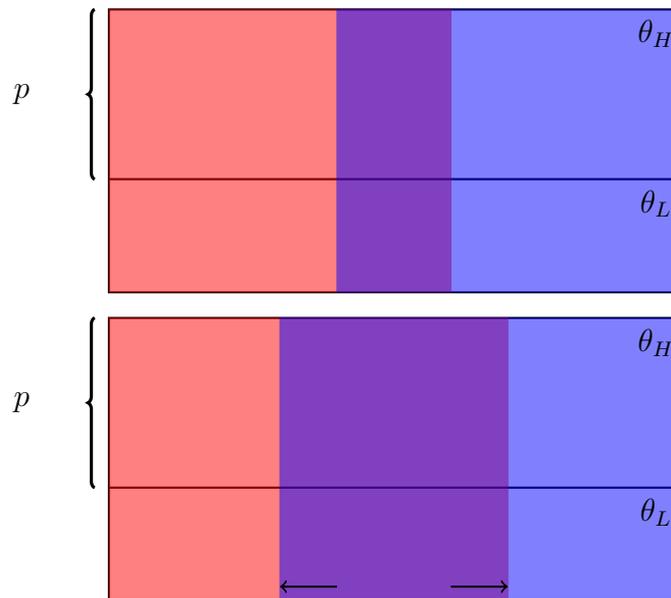


Figure 4: Equally good (top) and improved (bottom) correlated equilibrium

However, the principal is still not using the full machinery of Bayes' correlated

equilibrium since, as the above figures show, her recommendations are independent of the state of the world. Hence, while keeping the total probability with which the blue player contributes the same as in Figure 4, consider increasing the probability with which he contributes in the low state, and decreasing the probability with which he contributes in the high state, as depicted in Figure 5:

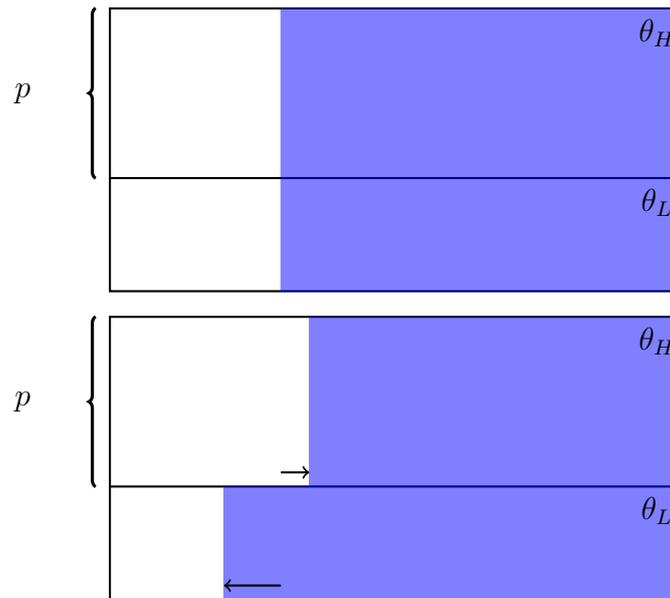


Figure 5: Bayes' correlated equilibrium - blue player

Notice that, if we use the same correlation device for the red player as before, he is pivotal with the same total probability as before, but with higher conditional probability in the high state (see Figure 6 below). This strengthens the incentives for the red player to contribute, which the principal can exploit by increasing the probability with which he is asked to contribute.

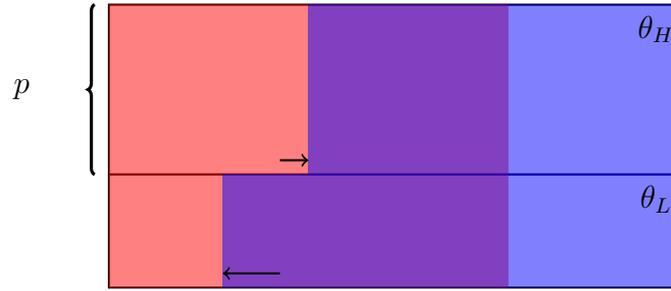


Figure 6: Bayes' correlated equilibrium: for red, conditional on being pivotal, higher chance of state being θ_H

Under some parametric conditions¹, the best Bayes' correlated equilibrium for the principal has both agents contribute with probability 1 in the low state, and with probability greater than a half in the high state:

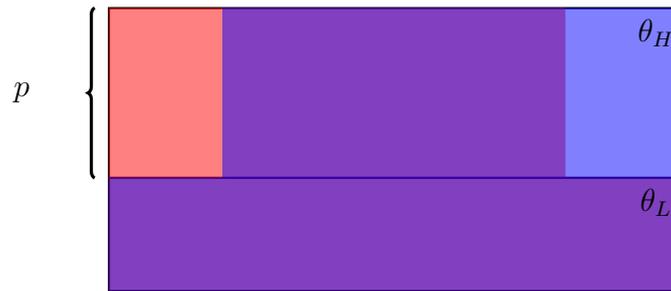


Figure 7: Best Bayes' correlated equilibrium

Consider the following (sequential) implementation of the above BCE. The principal observes the state of the world, and uses the above correlation device to draw a recommendation for the red player. She makes a recommendation to the red player (which is obeyed), and she observes the action taken by the red player (while the blue player does not). Then, she approaches the blue player, and makes a recommendation. Notice that incentives are unchanged from the BCE for the blue player: since he does not know the action taken by the red player, and, since the principal's recommendation is independent of the action taken by the red player, he faces the same considerations as in the BCE. The only

¹All conditions and proofs of results stated in the introduction are in Appendix A

difference is that, when the blue player is approached, there is a chance that the red player has already contributed -which, if known by the blue player, would do away with his incentive to contribute.

We now modify the above mechanism, without altering the incentives for the blue player (see Figure 8 below). Consider reducing the probability with which the red player contributes in the high state. Clearly, the blue player now strictly prefers to contribute. Hence, we can ask the blue player to contribute with higher probability, until the ratio of the area of the purple set to the blue set (the conditional probability of being pivotal and the state being high) is the same as before. By doing this, we increase the probability that both agents contribute in the high state: every time we move the red area to the left, we can move the blue one to the left, and increase the purple one (the (purple) overlap represents the additional contribution beyond one that the principal gets).

However, taking as given the behavior of the blue player in this new device, the red player, now, does not want to contribute when asked to do so (this is why this is not a BCE). But the red player moves first, so that the behavior of the blue player is not given. The principal can, actually, approach the red player first, and, conditional on whether the red player follows the recommendation or not, give different information to the blue player. Thus, consider the following mechanism. The principal approaches the red player first, and gives him a recommendation. If the red player complies, then the blue player is given recommendations according to the bottom of Figure 8. If he chooses not to contribute, when told to do so, then the blue player is told that the red player has already contributed. Note that this has positive probability, and leads the blue player to find it optimal not to contribute. This gets the red player to comply, knowing that if he does not contribute, the good will not be funded at all. Under some parametric restrictions, the optimal two-stage mechanism has the blue player contribute with probability 1 in both states, and the probability with which red contributes is chosen so that blue is indifferent between contributing and not contributing.² This is shown in Figure 9:

²To see that the blue player obeys the recommendation on, and off-path, notice that since the blue player is indifferent between contributing, and not contributing, we could consider an al-

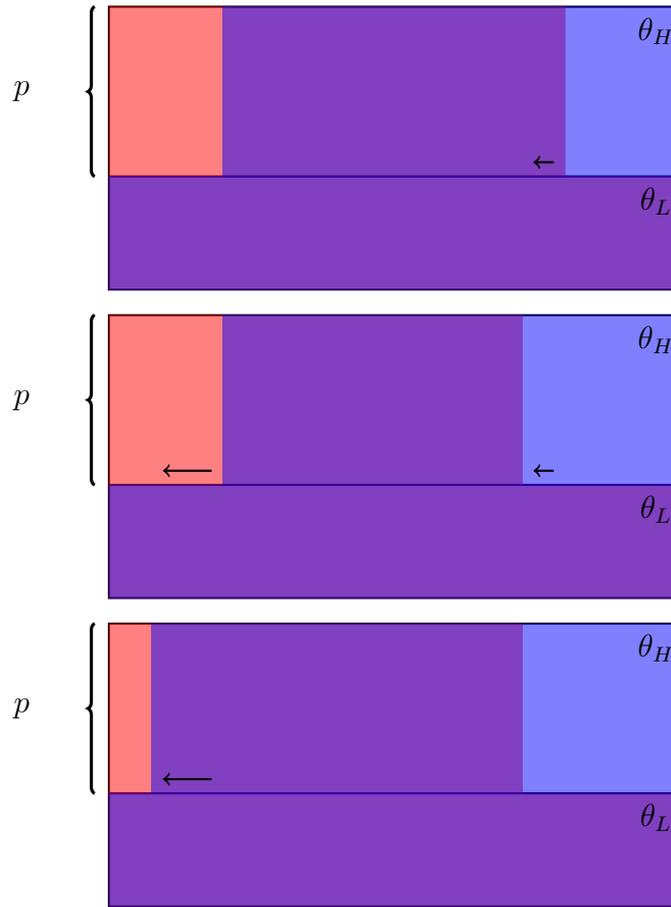


Figure 8: Reduce probability red contributes in high state (top and middle) to increase probability blue contributes in high state (middle and bottom) to increase the overlap.



Figure 9: Optimal two-stage mechanism

ternative mechanism in which, when the red player obeys, the blue player is asked to contribute with probability $1 - \epsilon$, and not to contribute with probability ϵ (think of a red sliver of red in Figure 9). This does not affect red's incentives, and is incentive compatible for blue. Now, when red disobeys, blue can't tell this apart from the on-path recommendation not to contribute. Then, think of making ϵ approach 0, and you obtain the mechanism described above.

Note how, by approaching the agents sequentially, and by conditioning the recommendations to the blue player on red's actions, we uncouple both players' incentive constraints, and increase the probability that both contribute in state H , as compared to the best BCE.

The above is not the best the principal can do; it is the best he can do when restricted to protocols in which player 1 moves first, and player 2 moves second. We could take a three-stage mechanism. The principal chooses whether to approach player 1 first, or not. If she does so, player 1 is the red player, and player 2 is the blue player. If she doesn't, player 2 will be the red player, and player 1 is contacted in the third stage, and is the blue player. Player 2 does not know whether player 1 was contacted in the first stage or not, that is, in the second stage, player 2 does not know if he is the blue player or the red player. Recall that the red player's incentives were slack in the optimal two-stage mechanism: if he disobeyed a recommendation to contribute, the good remained unfunded. By not telling player 2 whether he is the red or the blue player, we pool the slack incentive constraint for the red player with the tight incentive constraint of the blue player. Hence, we generate a slack in the overall incentive constraint of player 2, and we can now increase the probability with which he contributes, thus, improving upon the number of contributions in the two-stage mechanism.

In Appendix A, we show that the principal can continue to improve the overall contribution level by adding stages to this mechanism, and we characterize the value, as the number of stages goes to infinity, that the principal can achieve. This raises the question as to whether or not the principal can improve upon this value.

It turns out that the answer is no. However, in order to answer that question, we need to say what the admissible extensive forms for the principal are. An admissible extensive form is a finite extensive form with perfect recall where the terminal nodes are labeled with a state of the world, and an action profile, e.g. (θ, C, NC) . The interpretation is that the agents are playing this game, and that, at the end of the game, an action profile is determined. Moreover, to be admissible, the extensive form has to satisfy two properties. First, each player

controls his own action, that is, each player, at each point in time, has a way to make sure that his action is going to be the one he wants. We are ruling out extensive forms in which player 2 can force player 1 to contribute, or in which the players' only (trivial) choice is to contribute. Second, each player knows what action he is choosing. That is, there can be information sets in the extensive form, but it cannot be the case that there is a move in an information set which, at one node in the information set, makes him contribute, but in a different node in that information set that same move makes him not contribute.

If we had to solve for the optimal mechanism by optimizing over the communication possibilities in *all* extensive forms, the problem would be intractable. However, building on the ideas of Myerson [17], and Gershkov and Szentes [12] (see also the related and independent work by Salcedo [21]), we show that there exists a *canonical* extensive form. That is, an extensive form that satisfies the aforementioned properties, and that implements any distribution over outcomes that can be implemented by an extensive form in this class. In this canonical extensive form, each player moves exactly once, and when he does he can choose amongst any of his actions. The principal contacts players sequentially, and privately recommends them what action to take. These recommendations might depend on the state of the world, the actions recommended to others, *and* the actions other players have already chosen. The principal does not reveal to the players the order in which they are being contacted, nor the actions players moving before them have selected. However, the players know the mechanism being used by the principal, and calculate their beliefs using it.

With this result in hand, we can reformulate the problem of the principal as that of choosing, as a function of the state of the world, an order in which to contact the players, a distribution of action recommendations for the first (red) player, and a distribution of action recommendations for the second (blue) player, conditional, also, on the recommendation received and the action taken by the first player.

It turns out that in this example, since everything is symmetric, the best the principal can do is contact each player first with probability $\frac{1}{2}$, independently of

the state of the world. With probability $\pi^\theta(C)$, that player is asked to contribute when the state is θ . The player who is contacted second is recommended to play C with probability $\pi^\theta(C|a^{1R}, a^{1C})$, where (a^{1R}, a^{1C}) denote the recommendation received, and the action taken by the first player. Hence, if player i is asked to contribute, he finds it optimal to do so if, and only if,

$$\begin{aligned} & \sum_{\theta \in \{L, H\}} p_\theta \left(\frac{1}{2} \pi^\theta(C) + \frac{1}{2} (\pi^\theta(C) \pi^\theta(C|C, C) + (1 - \pi^\theta(C)) \pi^\theta(C|NC, NC)) \right) (v_\theta - c) \\ & \geq \sum_{\theta \in \{L, H\}} p_\theta \left(\frac{1}{2} \pi^\theta(C) \pi^\theta(C|C, NC) + \frac{1}{2} \pi^\theta(C) \pi^\theta(C|C, C) \right) v_\theta \end{aligned}$$

As before, we can choose $\pi^\theta(C|C, NC) = 0$ to maximize the punishment to the player that moves first. Consider a mechanism that has the first mover contribute with probability 1 if the state is L , and has the second mover contribute with probability 1, regardless of the state of the world. For this to be incentive compatible we need:

$$1 - \pi^H(C) \geq \frac{2c - \mathbb{E}[v]}{cp}$$

If the right hand side is negative, then we can set $\pi^H(C) = 1$, and if the RHS is positive, we set $\pi^H(C)$ to solve that expression for equality. Since this is the highest contribution level that the principal can obtain as a PBE of the canonical extensive form, our result implies that this is the highest contribution level that she can obtain when using any of the admissible extensive forms. In Appendix [A](#), we show that this is the number of contributions implemented by the multiple stage mechanism, as the number of stages goes to infinity.

1.1 Related Literature

The two papers closest to our work are Gershkov and Szentes [\[12\]](#) and Salcedo [\[21\]](#). To the best of our knowledge, Gershkov and Szentes [\[12\]](#) are the first to consider a communication protocol that in their application leads them to de-

rive the canonical extensive form.³ They analyze a binary voting model where (ex-ante symmetric) agents can acquire information about the underlying state of the world, and they have to be incentivized to do so. They show that it is without loss of generality to consider protocols in which agents are not told when they are being called to acquire information, and they are called to move only if they have to acquire information. Our results make transparent what is the class of communication protocols that can be accommodated by the canonical extensive form. Salcedo [21] tackles a related problem from the point of view of complete information. That is, he starts from a game without payoff uncertainty, and uses correlated equilibrium as a benchmark. He proposes an equilibrium notion, interdependent choice equilibrium (ICE). He shows that a distribution over actions is an ICE if, and only if, it is a Nash equilibrium for some extensive form in a (strict) subset of our admissible set. Since we are interested in information design, we start from a game with payoff uncertainty, and use BCE as our benchmark. Moreover, our focus is on understanding what is the largest class of mechanisms a principal can offer to the agents so that, without loss of generality, she can restrict her choice to the canonical extensive forms. In that regard, our results can be seen as a version of the revelation principle for settings in which the principal can structure the release of information and the timing at which agents take their actions. A consequence of the results in both papers is that, under complete information, if a distribution over actions is implementable as a PBE of the canonical extensive form, then it is an ICE.⁴

Another related paper is Peters [19]. In a competing mechanisms game, Peters characterizes the distributions over outcomes an outside observer can expect when he does not know the extensive form the agents are playing, but he knows the extensive form is *regular*. The difference between *regular* and *admissible* extensive forms is that, in the former, any given player, say, i , can guarantee that, for any strategy of his opponents, he has a strategy that (i) allows him to

³Bognar et al. [5] also highlight the benefit of having agents move sequentially, as opposed to simultaneously, in a voting game.

⁴More recently, Gallice and Monzon [11] consider a the benefits of using a sequential move protocol in a public good game. Contrary to Gershkov and Szentes, Salcedo, and our work, their protocol allows later movers to observe the contribution choices of earlier movers.

reach an information set from which he can commit to any of his actions in the base game, and (ii) conditional on reaching that information set, the distribution over others' actions player i faces does not depend of his choice of action at that information set. Admissible extensive forms don't require that (ii) holds. When players don't have private information, Peters' result implies that the set of distributions consistent with play of a regular extensive form is the set of correlated equilibria, while the set of distributions consistent with the play of an admissible extensive form is larger than the set of correlated equilibria.

More generally, our paper relates to the literatures of communication equilibria, persuasion, asynchronous and revision games, endogenous commitment in games, and normal form representation of extensive form games. Within the literature on *communication equilibria*, Myerson [17] characterizes the set of correlated equilibria (Aumann [1, 2]) as the set of distributions over actions that are generated by a communication protocol in which the players and the mediator exchange cheap talk messages first, and then, when all communication is done, simultaneously commit to their actions. Our contribution is to characterize the equilibrium distributions that arise from more general communication protocols in which agents can commit to their actions at different points during the process.⁵ Since correlated equilibria can always be implemented by our canonical extensive form, our result, when there is a unique underlying state of the world, strictly generalizes correlated equilibrium. Similarly, when there is uncertainty about the underlying state of the world, our paper can be seen as a contribution to the literature on Bayes' correlated equilibrium (Bergemann and Morris [4]). Our contribution to the literature on *persuasion* (Kamenica and Gentzkow [14], Ely, Kamenica, and Frankel [9], Ely [10]) is to identify an additional tool to incentivize agents, which is useful even in the absence of uncertainty about an underlying state of the world: communication with players who still have not committed to their actions can be used as an incentive for players who are now choosing what action to commit to. As in the literature on *asynchronicity in games*

⁵In light of our results, another way of interpreting correlated equilibria is as the set of distribution over action profiles that arise in a communication protocol where agents do not necessarily commit to their actions simultaneously, but the mediator cannot communicate with agents conditional on previous commitments.

(Lagunoff and Matsui [15], Kamada and Kandori [13], Calcagno et al [6]), we exploit the timing at which agents commit to their actions to change the set of equilibrium distribution over actions in a normal form game. However, contrary to that literature (with the exception of Kamada and Kandori [13]), we exploit the asynchronicity in the players' commitments to enlarge, rather than to shrink, the set of action distributions that can be implemented. There is a literature that studies the role of commitment in determining the outcomes of a (complete information) base game, by looking at the equilibrium play of a fixed extensive form, the play of which determines an outcome in the base game. For instance, Renou [20] and Bade et al. [3] consider two stage extensive form games where, in the first stage, the agents commit to a subset of their action space, and in the second stage they play the game induced by their reduced action sets (Dutta and Ishii [8] consider an extensive form that is similar in spirit, though much richer in the set of commitments it allows the players). Caruana and Einav [7] consider a finite, multi-stage game, where at each stage agents simultaneously announce the action they will play in the base game, and pay a switching cost whenever the announcement differs from the one in the previous stage. There are two main differences with this literature. First, even though our construction exploits that players commit to their actions sequentially, rather than simultaneously, our focus is not on the equilibrium set of a fixed extensive form, but on characterizing the set of distribution over outcomes that can be induced when one can select from extensive forms that allow for such commitments. Second, the extensive forms we allow for are different. In the extensive forms of the first three papers, players are allowed to first commit to play a subset of their actions in the base game, and then choose which action to play. In an admissible extensive form, however, as long as a player has not committed to any of her actions in the base game, all her commitments are still available to her. As we show in Section 4.3, without this assumption, there are distribution over outcomes that cannot be implemented by the canonical extensive form. We do allow for extensive forms as in Caruana and Einav; however, in our model, the agents' payoffs only depend on the outcome of the base game, and not on the sequence of announcements that led to that outcome. Finally, the paper is related to the literature on normal form representations of extensive form games. In particular, the canonical

extensive form is *non-playable*, as defined in Mailath, Samuelson, and Swinkels [16].

2 MODEL

We consider a base game $G = \langle N, \{A_i\}_{i \in N}, \Theta, p, \{u_i\}_{i \in N} \rangle$, where $N = \{1, \dots, n\}$ is the set of players, A_i is player i 's (finite) set of actions, Θ is a finite set of states of the world, $p \in \Delta_{++}(\Theta)$ is the prior distribution over Θ , and $u_i : \prod_{i \in N} A_i \times \Theta \mapsto \mathbb{R}$ is player i 's payoff as a function of the action profiles, and the realized state of the world.

We deal with finite extensive forms with outcomes in G , and we denote them by Γ . To define Γ formally, we first need to introduce some notation.⁶ Given a set X , and $m \in \mathbb{N}_0$, let $X^m = \prod_{i=0}^m X$, with $X^0 = \{\emptyset\}$. We say $y = (x_1, \dots, x_m) \in X^m$ is a sequence of length m , and \emptyset is the empty sequence from X . If $y \in X^m$, then $|y| = m$ is the length of y . For $m \in \mathbb{N}_0$, $y \in X^m$, and $n \leq m$, $y^n = (x_1, \dots, x_n)$, and $y_n = x_n$. Say y' precedes y if $y = (y', y_{|y'|+1}, \dots, y_{|y|})$, or $y' = \emptyset$. If y' precedes y , and $|y'| < |y|$, we denote this by $y' < y$; if y' precedes y , we denote this by $y' \leq y$.

Definition 2.1. Fix a set X . $V \subset \bigcup_{l \in \mathbb{N}_0} X^l$ is a tree if the following hold:

1. $\emptyset \in V$, and
2. if $v \in V$ and $v' < v$, then $v' \in V$.

We are now ready to define an extensive form. In the definition below, c refers to chance. Chance can be nature which, for example, determines the state of the world, but also the mediator.

Definition 2.2. An extensive form is given by $\Gamma = \langle N \cup \{c\}, M, V, \{H_i\}_{i \in N \cup \{c\}}, \sigma_c, \Theta \times A, \gamma, \iota \rangle$, where:

1. $\Theta \subset M$ is a set of moves,
2. $V \subset \bigcup_{l \in \mathbb{N}_0} M^l$ is a finite set of sequences of moves, which satisfies:

⁶This follows the notation and conventions of Battigalli, Friedenberg, and Siniscalchi.

- (a) V is a tree,
 - (b) $\exists \vartheta : \{v \in V : |v| = 1\} \mapsto \Theta$, one-to-one and onto, and
 - (c) $(\forall v \in V : |v| > 1)(\forall l > 1)v_l \notin \Theta$.
3. $\gamma : \{v \in V : (\nexists v' \in V)v < v'\} \mapsto \Theta \times A$ is the outcome function, where $\gamma(v) = (\theta, a)$ if, and only if, $v_1 = \theta$ and γ is onto.
 4. $\iota : V \setminus \{v \in V : (\nexists v' \in V)v < v'\} \mapsto N \cup \{c\}$ is the player function, and satisfies $\iota(\emptyset) = c$,
 5. H_i is player i 's information partition; i.e a partition of $\{v \in V : \iota(v) = i\}$. Each element h_i of the partition satisfies that if $v \in H_i, (v, m) \in V$, then $(v', m) \in V$ for all $v' \in h_i$.
 6. σ_c specifies for each $h \in H_c$ a distribution over $\{m \in M : (\exists v \in h)(v, m) \in V\}$, and $\sigma_c(\emptyset)(\vartheta^{-1}(\theta)) = p_\theta$.

The extensive forms we consider implement outcomes in the base game; that is, an outcome of the extensive form is a pair (θ, a) . Note that, in the extensive forms we consider, chance moves first and determines the state of the world according to the prior distribution $\{p_\theta\}$. From now on, Z is used to denote the set of terminal nodes, that is the set $\{v \in V : (\nexists v' \in V)v < v'\}$.

We consider finite extensive forms of perfect recall:

Definition 2.3. An extensive form Γ satisfies *perfect recall* if, for each player i , and each $h, h' \in H_i$, if $v', w' \in h'$, and $v \in h$ is such that (v, m_i) precedes v' , then there exists $w \in h$ such that (w, m_i) precedes w' .

An extensive form is *finite* if $\max\{|v| : v \in Z\} < \infty$.

2.1 Conditions

We consider extensive forms that satisfy the conditions described below. In what follows, given any extensive form Γ , for any $v \in V$, we let $\gamma(v) = \bigcup_{v' \in Z: v < v'} \{\gamma(v')\}$. Also, for any $v \in V$, for any $i \in N$, we let $\gamma_i(v) = \text{proj}_{A_i} \gamma(v)$. Finally, for any

$j \in N \cup \{c\}$, for any $h \in H_j$, let $M_j(h) = \{m \in M : (\exists v \in h)(v, m) \in V\}$ denote the set of moves available to player j at information set h .

The first condition we ask of the extensive forms is that only player i has moves in the extensive form that determine his action in the base game:

Condition 1 (No delegation). $(\forall i \in N)(\forall j \in (N \cup \{c\}) \setminus \{i\})(\forall h \in H_j)(\forall m \in M_j(h))(\forall v \in h)\gamma_i(v) = \gamma_i(v, m)$

The next condition is a measurability constraint: even if player i does not know in which node he is at within an information set, if one of his moves allows him to commit to playing action a_i following some node in the information set, then that move commits him to playing action a_i at all nodes of the information set.

Condition 2 (Know your action). $(\forall i \in N)(\forall h \in H_i) : (\exists v \in h)(\exists m \in M_i(h)) : \gamma_i(v, m) = a_i$, then $(\forall v' \in h) \gamma_i(v', m) = a_i$.

Condition 3 below says that as long as player i hasn't chosen a move in the extensive form that commits him to play an action, say a_i , in the base game, he can't rule out playing any of his actions in the base game. Formally,

Condition 3 (No partial commitments). $(\forall i \in N)(\forall h \in H_i)$ if $|\gamma_i(h)| > 1$, then $\gamma_i(h) = A_i$.

Definition 2.4. An extensive form Γ is *admissible* if it satisfies Conditions 1, 2, and 3.

2.2 The canonical extensive form

The canonical extensive form is an extensive form that satisfies Conditions 1, 2, 3, and works as follows. Nature moves first, and determines the state of the world. As a function of the state of the world, the mediator selects (possibly at random) a player, and a recommendation to make to that player. The chosen player selects which action to commit to. The mediator, observing the action taken by the first player, determines what player to contact next, and what recommendation to send to him. This is repeated until all players have been contacted. The canonical extensive form has the property that players will follow

their recommendations, on and off-path.

Formally, the mechanism used by the mediator can be seen as selecting a distribution over the order in which players are contacted, and recommendations as a function of the state of the world, previous recommendations, and previously taken actions. Let $O = \{o : \{1, \dots, n\} \mapsto N : o \text{ is bijective}\}$ denote the set of orders on N . For each $\theta \in \Theta$, let $q_\theta \in \Delta(O)$. For each (possibly empty) subset $N' \subset N$, for each $i \notin N'$, let:

$$\begin{aligned} \pi_i^{N'} : \Theta \times A_{N'} \times A_{N'} \times A_i &\mapsto [0, 1] \\ (\forall(\theta, a_{N'}, \hat{a}_{N'})) \sum_{a_i \in A_i} \pi_i^{N'}(a_i | \theta, a_{N'}, \hat{a}_{N'}) &= 1 \end{aligned}$$

That is $\pi_i^{N'}(a_i | \theta, a_{N'}, \hat{a}_{N'})$ is the probability that player i is recommended to play action a_i when the state of the world is θ , players in N' have already been contacted, recommended to play $a_{N'}$, and actually played $\hat{a}_{N'}$. Thus, the mechanism used by the mediator can be seen as a choice of $\{q_\theta\}, \{\pi_i^{N'}\}_{i \notin N' \subset N}$.

It is useful to understand players' incentives in the canonical extensive form. Towards that end, fix such a $\{q_\theta\}, \{\pi_i^{N'}\}_{i \notin N' \subset N}$. For each $\theta, (a_i, a_{-i})$, and $\hat{a}_i \neq a_i$ define:

$$\alpha(a_i, a_{-i} | \theta) = \sum_{o \in O} q_\theta(o) \prod_{j=1}^n \pi_{o(j)}^{o(j)} \left(a_{o(j)} | \theta, (a)_{o(1)}^{o(j-1)}, (a)_{o(1)}^{o(j-1)} \right),$$

to be the probability that, when the state is θ , action profile (a_i, a_{-i}) is recommended -when players obey their recommendations-. Also, define

$$\begin{aligned} \alpha(a_i, a_{-i} | \theta, \hat{a}_i) = \sum_{o \in O} q_\theta(o) &\prod_{j=1}^{o^{-1}(i)} \pi_{o(j)}^{o(j)} \left(a_{o(j)} | \theta, (a)_{o(1)}^{o(j-1)}, (a)_{o(1)}^{o(j-1)} \right) \\ &\times \prod_{j=o^{-1}(i)+1}^n \pi_{o(j)}^{o(j)} \left(a_{o(j)} | \theta, (a)_{o(1)}^{o(j-1)}, (a_{o(1)}, \dots, \hat{a}_i, \dots, a_{o(j-1)}) \right) \end{aligned}$$

to be the probability that action profile (a_i, a_{-i}) is recommended when all players but i obey the recommendation, and i plays \hat{a}_i .

Pick i and a_i such that $\sum_{\theta} \sum_{a_{-i}} \alpha(a_i, a_{-i} | \theta) > 0$. When player i is contacted by the mediator, he obeys the recommendation to play a_i if, for all $\hat{a}_i \neq a_i$, the following holds:

$$\sum_{\theta \in \Theta} p_{\theta} \sum_{a_{-i} \in A_{-i}} \alpha(a_i, a_{-i} | \theta) u_i(a_i, a_{-i}, \theta) \geq \sum_{\theta \in \Theta} p_{\theta} \sum_{a_{-i} \in A_{-i}} \alpha(a_i, a_{-i} | \theta, \hat{a}_i) u_i(\hat{a}_i, a_{-i}, \theta)$$

Note that the only difference between the above incentive constraints and the ones in Bergemann and Morris [4] is that, when a player deviates from his recommendation, the mediator can have the players who have not yet moved respond to the deviation by player i . This is why Bayes' Correlated Equilibria are always a subset of the distribution over actions implemented by the canonical extensive form.

2.3 Equilibrium assessments

Our main result shows that a distribution over outcomes in the base game induced by a Perfect Bayesian equilibrium assessment in an extensive form that satisfies conditions 1, 2, and 3 can be implemented by a Perfect Bayesian equilibrium in the canonical extensive form. We define in this section what we mean by an assessment and by Perfect Bayesian equilibrium. To do this, the following piece of notation is useful. Given two information sets h, h' , say h precedes h' if for every $v \in h$, there exists $v' \in h'$ such that $v < v'$. In a slight abuse of notation, we denote this by $h < h'$.

Definition 2.5. A behavioral strategy for player i is a map $\sigma_i : H_i \mapsto \bigcup_{h \in H_i} \Delta(M_i(h))$ such that $\sigma_i(h) \in \Delta(M_i(h))$.

Definition 2.6. A belief system is a map $\mu : \bigcup_{i \in N \cup \{c\}} H_i \mapsto \bigcup_{h \in \bigcup_{i \in N \cup \{c\}} H_i} \Delta(\{v \in V : v \in h\})$ such that $\mu(h) \in \Delta(\{v : v \in h\})$.

Given a belief system μ^* , and a strategy profile σ_{-i}^* , the payoff for agent i , at

information set $h \in H_i$, from following σ_i is given by:

$$U_i(\sigma_i, \sigma_{-i}^* | h) = \sum_{\{v \in h\}} \mu^*(v|h) \sum_{\{v' \in Z: v < v'\}} Pr(v' | \sigma_i, \sigma_{-i}^*, v) u_i(\gamma(v')),$$

where

$$Pr(v' | \sigma_i, \sigma_{-i}, v) = \prod_{k=|v|}^{|v'|-1} \sigma_{\iota(v^k)}(v'_{k+1})$$

σ_i is sequentially rational at h for beliefs $\mu^*(\cdot|h)$ if:

$$(\forall \sigma'_i) U_i(\sigma_i, \sigma_{-i}^* | h) \geq U_i(\sigma'_i, \sigma_{-i}^* | h)$$

Definition 2.7. An assessment $\langle \sigma^*, \mu^* \rangle$ is a *Perfect Bayesian equilibrium* if the following hold:

1. For all $i \in N$, σ_i^* is sequentially rational for player i at all information sets $h \in H_i$, and

2. μ^* is determined by Bayes' rule from σ^* , whenever possible, that is:⁷

(a) For all $h \in \cup_{i \in N} H_i$ such that $\sum_{v \in h} Pr(v | \sigma^*) > 0$, then $\mu^*(v|h) = Pr(v | \sigma^*) / (\sum_{v \in h} Pr(v | \sigma^*))$.

(b) For all $i \in N$, for all $h, h' \in H_i$ such that h' is an immediate successor of h , if $\sum_{v \in h} \sum_{v' \in h'} \mu^*(v|h) Pr(v' | \sigma^*, v) > 0$, then $\mu^*(\cdot|h')$ is obtained from $\mu^*(\cdot|h)$ and σ^* applying Bayes' rule.

3 RESULTS

We now state our main theorem:

Theorem 3.1. Fix an admissible extensive form Γ , and a Perfect Bayesian equilibrium $\langle \sigma^*, \mu^* \rangle$ of Γ . Then, there exists a canonical extensive form, and a Perfect Bayesian equilibrium $\langle \sigma^{*'}, \mu^{*'} \rangle$ of it that implements the same distribution over outcomes in the base game as $\langle \sigma^*, \mu^* \rangle$.

⁷Watson [22] calls this minimal consistency.

The proof of the result is in Appendix B. We provide a sketch in what follows. Say that an extensive form has all commitments available if, whenever an agent reaches an information set in which he can determine his action in the base game, he can pick to commit to any of his actions. The proof proceeds in two steps. The first step shows that an admissible Γ can always be transformed into an extensive form where all commitments are available, and in a way in which $\langle \sigma^*, \mu^* \rangle$ (properly adapted to the new extensive form) remains a Perfect Bayesian equilibrium. The second step shows that the Perfect Bayesian equilibrium of an extensive form that has all commitments available can be mapped into the canonical extensive form.

Theorem 3.1 says that when characterizing the distribution over outcomes in the base game that can be sustained by a Perfect Bayesian equilibrium in any of the admissible extensive forms, it suffices to characterize the equilibria of the canonical extensive form. In that sense, Theorem 3.1 provides an analogue of the revelation principle for mechanisms that satisfy our conditions.

4 CONDITIONS REVISITED

We illustrate that Conditions 1, 2, 3 are necessary for equilibrium distribution over outcomes of extensive forms to be replicated by the canonical extensive form. In what follows, since $|\Theta| = 1$, we omit the move by chance at the beginning of the tree.

4.1 An extensive form that fails ‘No delegation’

Example 4.1. Consider the following G , where $\mathcal{N} = \{1, 2\}$, $A_i = \{C, D\}$, and payoffs are as follows:

	C	D
C	5,5	-10,10
D	10,-10	0,0

Table 2: Prisoner’s dilemma

Consider the following extensive form Γ . Edges labeled as actions represent moves that fix that player's action in the terminal history to be the one in the label.

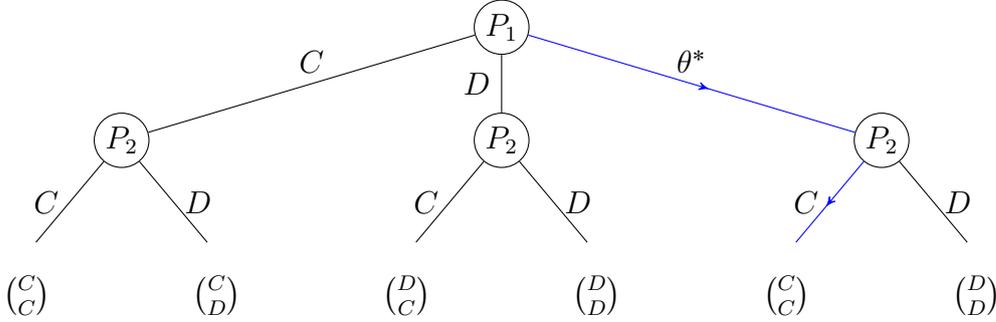


Figure 10: An extensive form that satisfies Conditions 2 and 3

The unique subgame perfect Nash equilibrium of this game is illustrated in blue in Figure 10, and it leads to $\{C, C\}$ being played with probability 1. This only depends on $u_i(C, C) > u_i(D, D)$ for $i \in \{1, 2\}$. Letting p denote the probability that player 1 moves first, the canonical extensive form can implement $\{C, C\}$ with probability 1 only if $p(u_1(C, C) - u_1(D, D)) \geq (1 - p)(u_1(D, C) - u_1(C, C))$, and $(1 - p)(u_2(C, C) - u_2(D, D)) \geq p(u_2(D, C) - u_2(C, C))$. These conditions are clearly stronger than $u_i(C, C) > u_i(D, D)$ for $i \in \{1, 2\}$.

4.2 An extensive form that fails 'Know your own action'

Example 4.2 (Know your action). Consider the following 3-player base game:

	<i>D</i>	<i>E</i>		<i>D</i>	<i>E</i>
<i>A</i>	10,1,-1	2,-2,0	<i>A</i>	1,-1,1	1,-3,-1
<i>B</i>	0,-1,1	0,-3,-1	<i>B</i>	0.5,1,-1	0,-2,0
	<i>F</i>			<i>G</i>	

Table 3: Normal - form game

Consider the following extensive form:

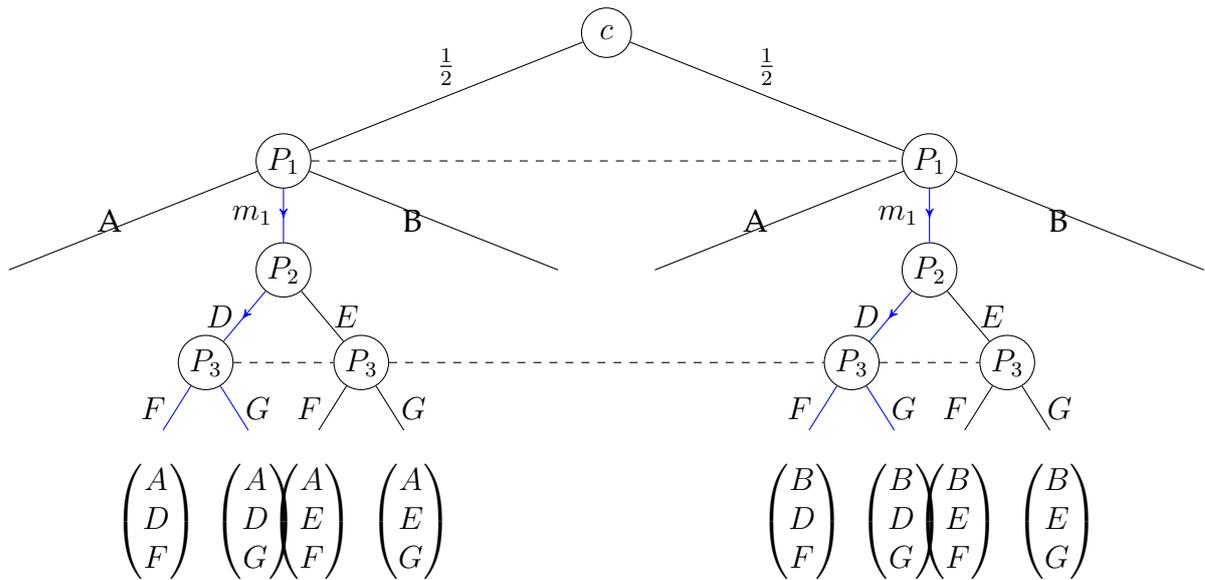


Figure 11: Extensive form that satisfies Conditions 1 and 3

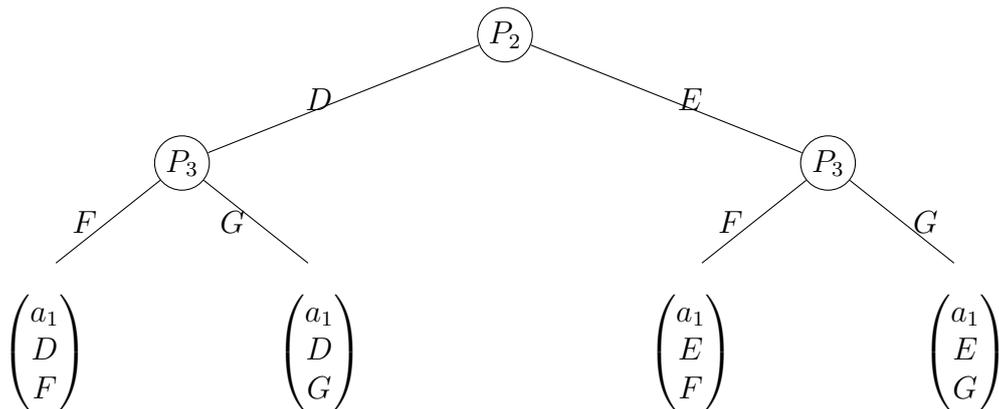


Figure 12: Subgame starting after moves $a_1 \in \{A, B\}$

There is a SPNE, depicted in blue in Figure 11, in which player 1 chooses m_1 , player 2 chooses D , and player 3 mixes with probability 0.5 between F and G . Off path, when player 1 chooses A , player 2 chooses D , and player 3 chooses G . However, player 1 will never follow a recommendation to play B : by using the move m_1 , player 1 can commit to a randomization -induced by the moves of player 3- that he himself is not willing to carry out.

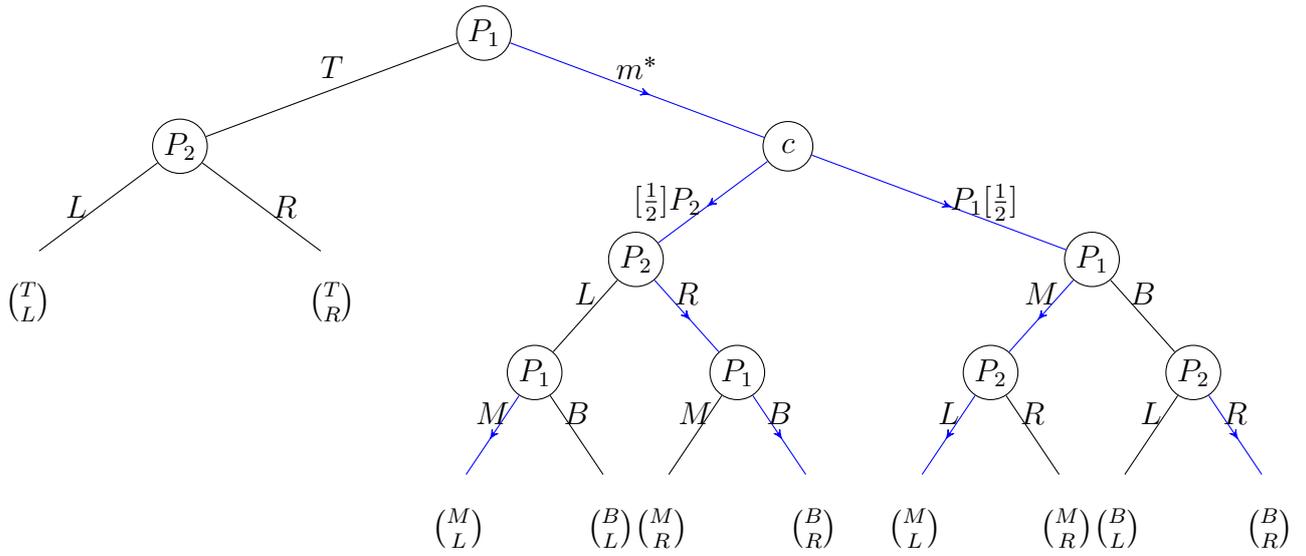


Figure 13: Extensive form that satisfies no delegation and know your own action

4.3 An extensive form that fails 'No partial commitments'

Example 4.3. The following example is from Myerson [18]. Consider the following base game where $N = \{1, 2\}$, $A_1 = \{T, M, B\}$, and $A_2 = \{L, R\}$. Payoffs are as follows:

	L	R
T	2, 2	2, 2
M	5, 1	0, 0
B	0, 0	1, 5

Table 4: Base game

Consider the following extensive form that satisfies Conditions 1 and 2:

In blue, we depict the unique subgame perfect Nash equilibrium of the game. However, the distribution over action profiles that it implements can't be implemented in a canonical extensive form. When Player 1 is recommended to play B , if he deviates, and plays T instead, he can guarantee a payoff of 2, which is higher than the maximum payoff he can obtain by playing B (1).

5 CONCLUSIONS

We study multi-agent information design, where the primitive of the model is a base game which is defined by the state space, the set of players, the set of actions available to them, their payoffs as a function of the state and the action profile, and their initial information, as embodied by a common prior over the state space. In this setting, Bayes' correlated equilibria has been used to formalize information design. On the one hand, it captures the set of distributions an interested principal, who observes the state of the world, could bring about when designing the release of information about the state of nature. On the other hand, it captures the set of distributions over actions an outsider could observe when he knows the previous ingredients, but does not know what additional information the agents observe before choosing their actions. Under both interpretations, the assumption is that players first observe the information, and then *simultaneously* choose their actions.

However, in this setting, for any given player, what others are playing, and whether or not they have already committed to their actions is another source of uncertainty. Thus, a principal could potentially structure both the release of information about the state of the world, and the timing of play when thinking of achieving a certain objective.

In the paper, we study the multi-agent information design problem faced by a principal who can design, not only the information structure, but also the extensive form faced by the agents. We show that, as long as the principal can choose among *admissible* extensive forms, then any distribution over action profiles that she can implement in a Perfect Bayesian equilibrium of one of these extensive forms, she can also implement in a Perfect Bayesian equilibrium of a *canonical*

extensive form. Thus, in the spirit of the revelation principle, we simplify the problem of finding the adequate (for the principal's objective) admissible extensive form to the one of finding the adequate PBE in the canonical extensive form.

A EXAMPLE

A.1 Optimal Bayes Correlated Equilibrium

The principal chooses, for each θ , $\pi^\theta \in \Delta(A_1 \times A_2)$, to maximize:

$$2c(p\pi^H(C, C) + (1-p)\pi^L(C, C)) + c(p(\pi^H(C, NC) + \pi^H(NC, C)) + (1-p)(\pi^L(C, NC) + \pi^L(NC, C)))$$

Since the environment is symmetric, any incentive compatible (asymmetric) recommendation system $\{\pi^\theta(a_1, a_2)\}_{\theta \in \{H, L\}, a_i \in \{C, NC\}}$ can be made into a symmetric, and incentive compatible one by setting $\pi^{\theta'}(C, NC) = \pi^{\theta'}(NC, C) = \frac{\pi^\theta(C, NC) + \pi^\theta(NC, C)}{2}$. Hence, from now on we focus on symmetric recommendation devices, and we denote:

$$q_\theta = \pi^\theta(C, NC) = \pi^\theta(NC, C),$$

$$\pi^L = \pi^L(NC, NC).$$

Note that it has to be the case that $\pi^H(NC, NC) = 0$, and hence $\pi^H(C, C) = 1 - 2q_H$. Note that $\pi^L(C, C) = 1 - q_L - \pi^L$. With this notation, we can write the incentive compatibility constraint for contributing as:

$$p(v_H + c)\left[q_H - \frac{c}{c + v_H}\right] \geq (1-p)(v_L + c)\left(\frac{(1 - \pi_L)c}{v_L + c} - q_L\right)$$

That is, incomplete information implies that, conditional on the state being H , the probability that either individual contributes on their own is greater than the same probability under complete information, while the opposite holds when the state is L .

Set the following Lagrangian:

$$\begin{aligned}\mathcal{L} &= 2c(p(1 - 2q_H) + (1 - p)(1 - \pi_L - 2q_L)) + c(p2q_H + (1 - p)2q_L) \\ &+ \lambda[p(v_H + c)[q_H - \frac{c}{c + v_H}] - (1 - p)(v_L + c)(\frac{(1 - \pi_L)c}{v_L + c} - q_L)] \\ &+ (1 - p)\bar{\phi}_L(1 - \pi_L - 2q_L) + p\bar{\phi}_H(1 - 2q_H) + p\underline{\phi}_H q_H + \underline{\phi}_L^\pi \pi_L + \underline{\phi}_L^q q_L\end{aligned}$$

The first order conditions are:

$$\begin{aligned}(q_H) &: -2c + \lambda(v_H + c) - 2\bar{\phi}_H + \underline{\phi}_H = 0 \\ (q_L) &: -2c + \lambda(v_L + c) - 2\bar{\phi}_L + \underline{\phi}_L^q = 0 \\ (\pi_L) &: -2c + \lambda c - \bar{\phi}_L + \underline{\phi}_L^\pi = 0\end{aligned}$$

Notice that in all solutions $\lambda > 0$. Otherwise, if $\lambda = 0$, we have that $q_H = q_L = \pi_L = 0$. It cannot be that IC holds, a contradiction. Hence, from now on $\lambda > 0$.

Also, it has to be that $\underline{\phi}_H = 0$. Otherwise, $\underline{\phi}_H > 0 \Rightarrow q_H = 0$, and $\bar{\phi}_H = 0$. Then,

$$\begin{aligned}\lambda(v_H + c) + \underline{\phi}_H &= 2c \\ \lambda(v_L + c) + \underline{\phi}_L^q &= 2c + 2\bar{\phi}_L\end{aligned}$$

Hence, $\lambda(v_H - v_L) + \underline{\phi}_H = \underline{\phi}_L^q - 2\bar{\phi}_L$, and since the LHS is strictly positive, then $\underline{\phi}_L^q > 0$. Since the IC constraint can't hold if $q_L = q_H = 0$, this is a contradiction. Hence, from now on $\underline{\phi}_H = 0$. We focus on solutions in which $\pi^L = q_L = 0$, and $q_H < \frac{1}{2}$. In this case, $\bar{\phi}_L = 0$, and we have:

$$\begin{aligned}(q_H) &: -2c + \lambda(v_H + c) - 2\bar{\phi}_H = 0 \\ (q_L) &: -2c + \lambda(v_L + c) + \underline{\phi}_L^q = 0 \\ (\pi_L) &: -2c + \lambda c + \underline{\phi}_L^\pi = 0.\end{aligned}$$

This implies that $\lambda v_L + \underline{\phi}_L^q = \underline{\phi}_L^\pi$. Since $q_L = 0$, from the incentive compatibility constraint we obtain:

$$q_H = \frac{c}{p(v_H + c)},$$

which is feasible if, and only if, $pv_H - c \geq c(1 - p)$. In that case, take $\bar{\phi}^H = 0$, $\underline{\phi}_L^1 = \lambda(v_H - v_L)$, $\lambda = \frac{2c}{v_H + c}$.

A.2 Optimal 2-stage mechanism.

The two stage mechanism is determined by the following variables. In the first stage, the principal contacts player 1, and asks him to contribute in state θ with probability $\pi^\theta(C)$. In the second stage, in state θ , the principal contacts player 2, and asks him to contribute with probability $\pi^\theta(C|a_1^R, a_1^C)$, when player 1 was recommended action a_1^R , and a_1^C was his realized action.

For player 1, it is optimal to follow the recommendations if:

$$\begin{aligned} p\pi^H(c)(v_H - c) + (1 - p)\pi^L(C)(v_L - c) &\geq p\pi^H(C|C, NC)v_H + (1 - p)\pi^L(C|C, NC)v_L, \\ p\pi^H(NC)\pi^H(C|NC, NC)v_H + (1 - p)\pi^L(NC)\pi^L(C|NC, NC)v_L &\geq \\ p\pi^H(NC)(v_H - c) + (1 - p)\pi^L(NC)(v_L - c). \end{aligned}$$

For player 2, if he believes that player 1 follows his recommendation, the mechanism is incentive compatible if:

$$\begin{aligned} \sum_{\theta} p_{\theta} \sum_{a_1} \pi^{\theta}(C|a_1, a_1)(v_{\theta} - c) &\geq p\pi^H(C|C, C)v_H + (1 - p)\pi^L(C|C, C)v_L, \\ \sum_{\theta} p_{\theta} \pi^{\theta}(NC|C, C)v_{\theta} &\geq \sum_{\theta} p_{\theta} \pi^{\theta} \sum_{a_1} \pi^{\theta}(NC|a_1, a_1)(v_{\theta} - c) \end{aligned}$$

Notice that it should always be the case that: $\pi^H(C|NC, NC) = 1$, which implies that the incentive constraint for not contributing holds vacuously for both

players. Moreover, notice that setting $\pi^\theta(C|C, NC) = 0$ relaxes the incentive constraint for player 1, without affecting player 2's incentives, and without reducing total contributions. Finally, notice that $\pi^L(C) = 1 = \pi^\theta(C|C, C)$, and $\pi^H(C) = \frac{pv_H - c}{pv_H}$, satisfies both incentive constraints as long as:

$$\frac{pv_H - c}{pv_H} \geq \frac{(1-p)(c - v_L)}{p(v_H - c)} \quad (1)$$

We now show that as long as equation (1) holds, the above is, in fact, the optimal two-stage mechanism. To do so, we set up the following Lagrangian:

$$\begin{aligned} \mathcal{L} = & c \sum_{\theta} p_{\theta} \pi^{\theta} \pi^{\theta}(C|C) + c(1-p)(\pi^L + (1-\pi^L)\pi^L(C|NC)) \\ & + \lambda_1 \sum_{\theta} p_{\theta} \pi^{\theta}(v_{\theta} - c) \\ & + \lambda_2 \left[p(1-\pi^H)(v_H - c) + (1-p)(1-\pi^L)\pi^L(C|NC)(v_L - c) \right. \\ & \quad \left. - c(p\pi^H\pi^H(C|C) + (1-p)\pi^L\pi^L(C|C)) \right] \\ & + \sum_{\theta} p_{\theta} (\underline{\phi}_1^{\theta} \pi^{\theta} + \bar{\phi}_1^{\theta} (1 - \pi^{\theta})) \\ & + \sum_{\theta} p_{\theta} \pi^{\theta} (\underline{\phi}_2^{\theta, C} \pi^{\theta}(C|C) + \bar{\phi}_2^{\theta, C} (1 - \pi^{\theta}(C|C))) \\ & + (1 - \pi^L)(1 - p) [\underline{\phi}_2^{L, NC} \pi^L(C|NC) + \bar{\phi}_2^{L, NC} (1 - \pi^L(C|NC))], \end{aligned}$$

where we use the shorthand notation $\pi^{\theta} \equiv \pi^{\theta}(C)$, $\pi^{\theta}(C|C) \equiv \pi^{\theta}(C|C, C)$, and $\pi^L(C|NC) \equiv \pi^L(C|NC, NC)$.

The first order conditions are:

$$(\pi^H) : \pi^H(C|C)(c - \lambda_2 c + \underline{\phi}_2^{H, C} - \bar{\phi}_2^{H, C}) + (\lambda_1 - \lambda_2)(v_H - c) + \underline{\phi}_1^H - \bar{\phi}_1^H + \bar{\phi}_2^{H, C} = 0$$

$$\begin{aligned} (\pi^L) : & \pi^L(C|C)(c - \lambda_2 c + \underline{\phi}_2^{L, C} - \bar{\phi}_2^{L, C}) + \bar{\phi}_2^{L, C} + c + \lambda_1(v_L - c) + \underline{\phi}_1^L - \bar{\phi}_1^L \\ & - \pi^L(C|NC)(c - \lambda_2(v_L - c) + \underline{\phi}_2^{L, C} - \bar{\phi}_2^{L, C}) - \bar{\phi}_2^{L, NC} = 0 \end{aligned}$$

$$(\pi^{\theta}(C|C)) : \pi^{\theta}(c - \lambda_2 c + \underline{\phi}_2^{\theta, C} - \bar{\phi}_2^{\theta, C}) = 0, \theta \in \{L, H\}$$

$$(\pi^L(C|NC)) : (1 - \pi^L)(c - \lambda_2(v_L - c) + \underline{\phi}_2^{L, C} - \bar{\phi}_2^{L, C}) = 0.$$

We now show that when $\frac{pv_H - c}{pv_H} \geq \frac{(1-p)(c - v_L)}{p(v_H - c)}$, the above mechanism is a solution to the Lagrangian. In order to do so, let $\lambda_1 = \underline{\phi}_1^H = \overline{\phi}_1^H = 0$, since in the suggested solution the first player's IC slacks, and $\pi^H \in (0, 1)$. Note that $\pi^H > 0$ implies that $c - \lambda_2 c + \underline{\phi}_2^{H,C} - \overline{\phi}_2^{H,C} = 0$. Then, from (π^H) , we get:

$$\lambda_2(v_H - c) = \overline{\phi}_2^{H,C}.$$

Conjecturing that $\lambda_2 > 0$, we get that $\overline{\phi}_2^{H,C} = 0$, and hence $\pi^H(C|C) = 1$. Replacing in $(\pi^H(C|C))$, we obtain that:

$$\lambda_2 = \frac{c}{v_H} > 0,$$

which corroborates the conjecture. From the condition for $\pi^L(C|C)$, we get:

$$\pi^L \left(c \frac{v_H - c}{v_H} + \underline{\phi}_2^{L,C} - \overline{\phi}_2^{L,C} \right) = 0.$$

In our solution, $\pi^L > 0$, and hence $\overline{\phi}_2^{L,C} > 0$, which implies $\pi^L(C|C) = 1$. If $\pi^L < 1$, from the FOC for $\pi^L(C|NC)$, we get:

$$0 = \frac{c}{v_H}(v_H - c + v_L) + \underline{\phi}_2^{L,C} - \overline{\phi}_2^{L,C},$$

which implies $\pi^L(C|NC) = 1$. Replacing in π^L , we get:

$$c \left(1 - \frac{v_L}{v_H} \right) + \underline{\phi}_1^L - \overline{\phi}_1^L = 0,$$

and, hence, $\pi^L = 1$. The solution is feasible if, indeed, the constraint for player 1 is satisfied at $\pi^H = \frac{pv_H - c}{pv_H}$. This is if, and only if, equation (1) holds.

A.3 Derivation of the T -stage mechanism.

Consider the following extension of the sequential procedure discussed in Section 1. It lasts for T stages. The mediator approaches the players sequentially: in odd stages, she can contact player 1, while in even stages, she can contact player 2. Each player has to be contacted exactly once. Once a player is contacted, he chooses whether to contribute or not.

We make two observations. It is without loss of generality to consider mechanisms where (a) if a player is approached before the last stage at which he can move, he is asked to contribute, and (b) if player i is approached in stage t (i.e. player i becomes the red player), then player j is approached in $t + 1$ (i.e., he becomes the blue player, if he hasn't been already approached). The reason why (a) is true is that all recommendations to not contribute can be delayed to the last stage without affecting incentives.

Consider the following recommendation procedure. If $1 \leq t \leq T - 2$, then, if at $t - 1$ a player was contacted, the player on the move at t is asked to contribute with probability 1 regardless of the state of the world; otherwise, the player on the move at t is asked to contribute with probability q_{T-1-t} , unconditional on the state of the world. For $t = T - 1$, if the player on the move at $T - 1$ has not yet been contacted, then he is asked to contribute with probability 1 if the state is L , with probability π if the state is H , in case he is the first to move, and he contributes with probability 1 if he is the second to move. For $t = T$, if the player that moves at this stage has not been contacted, then he is asked to contribute with probability 1. Whenever a player disobeys, if the other player has yet to move, he is contacted next, and told not to contribute.

Note that conditional on being contacted at stage $t \leq T - 2$, the player on the move at that stage only cares about $q_{T-1-(t-1)}$, and q_{T-1-t} , that is he only cares about whether or not someone has already contributed, and what is the probability with which he is being asked to contribute first. To see this, note that, conditional on being contacted in stage t , the player contributes if, and

only if, the following hold:

$$(1 - q_{T-1-(t-1)})q_{T-1-t}(p(v_H - c) + (1 - p)(v_L - c)) \geq cq_{T-1-(t-1)}$$

If someone has already contributed, which conditional on being contacted at t can only happen in stage $t-1$, the player loses c with probability 1 -he is asked to contribute with probability 1, but the good has already been produced. If until t no one has contributed, then he obtains the expected value of the good minus the contribution costs with probability 1.

As in the two-stage mechanism π is chosen so that the stage T incentive constraint binds, and, recursively, $q_{T-1-(t-1)}$ so that the stage t incentive constraint binds. Player 1's incentive constraint always slacks in stage 1 (when he is the red player).

Conditional on being contacted at stage T , the player on the move at T faces the same incentive constraint than the blue player in the 2-stage mechanism: he knows that if the state is L , then the other player has already contributed, and if the state is H the other player has already contributed with probability π . He contributes if, and only if,

$$p(1 - \pi)(v_H - c) + (1 - p)(v_L - c) \geq c(p\pi + (1 - p)).$$

Then, as before, setting $\pi = \frac{pv_H - c}{pv_H}$ makes this player indifferent between contributing, and not contributing. Having π , we can solve backwards for the highest values of $\{q_{T-1-t}\}$ consistent with incentive compatibility. We get:

$$q_1 = \frac{p\pi(v_H - c) + (1 - p)(v_L - c)}{c + p\pi(v_H - c) + (1 - p)(v_L - c)}$$

$$(t \in \{2, \dots, T - 2\})q_t = \frac{q_{t-1}(\mathbb{E}[v] - c)}{c + q_{t-1}(\mathbb{E}[v] - c)}$$

Under this mechanism, the good is always produced in both states, and both players contribute when the state is L . Hence, it suffices to show that the probability of exactly one player contributing in state H goes down as T increases to show that adding stages (strictly) benefits the principal.

Note first that, if $T = 3$, the probability that a single player contributes in H is $(1 - q_1)(1 - \pi) < (1 - \pi)$, where the latter is the probability that only one player contributes in the 2-stage mechanism. Hence, the 3-stage mechanism improves upon the optimal 2-stage mechanism. More generally, the probability that only one player contributes when the state is H in the T -stage mechanism is: $\prod_{t=1}^{T-2} (1 - q_t)(1 - \pi) < \prod_{t=1}^{T-3} (1 - q_t)(1 - \pi)$, where the latter is the probability that only one player contributes, when the state is H , in the $T - 1$ -stage mechanism. Hence, the number of contributions collected increases with T .

We now want to calculate, as a function of T , the probability that exactly one player contributes when the state is H . Manipulation of the above equations shows that:

$$\prod_{t=1}^{T-2} (1 - q_t)(1 - \pi) = (1 - \pi) \frac{(1 - q_1)q_{T-2}}{q_1} \frac{c^{T-3}}{(\mathbb{E}[v] - c)^{T-3}}.$$

Recall that:

$$q_2 = \frac{q_1(\mathbb{E}[v] - c)}{c + q_1(\mathbb{E}[v] - c)}$$

$$q_3 = \frac{q_2(\mathbb{E}[v] - c)}{c + q_2(\mathbb{E}[v] - c)} = \frac{q_1(\mathbb{E}[v] - c)^2}{c(q_1(\mathbb{E}[v] - c) + q_1(\mathbb{E}[v] - c)^2)}.$$

Let $d_1 = 1$, and define recursively for $t \in \{2, \dots, T - 2\}$:

$$d_t = q_1(\mathbb{E}[v] - c)^{t-1} + cd_{t-1}.$$

Applying the above definition recursively, we can write:

$$q_{T-2} = \frac{q_1(\mathbb{E}[v] - c)^{T-3}}{d_{T-2}},$$

where

$$d_{T-2} = q_1 \sum_{t=0}^{T-4} (\mathbb{E}[v] - c)^{T-3-t} c^t + c^{T-3}.$$

Consider first, the case in which $\mathbb{E}[v] - c < c$. Then, we can write the probability that exactly one agent contributes as:

$$\begin{aligned} \frac{(1 - q_1)(1 - \pi)c^{T-3}}{q_1 \sum_{t=0}^{T-4} (\mathbb{E}[v] - c)^{T-3-t} c^t + c^{T-3}} &= \frac{(1 - q_1)(1 - \pi)}{q_1 \sum_{t=0}^{T-4} \left(\frac{\mathbb{E}[v]-c}{c}\right)^{T-3-t} + 1} \\ &= \frac{(1 - q_1)(1 - \pi)}{q_1 \frac{\mathbb{E}[v]-c}{c} \sum_{t=0}^{T-4} \left(\frac{\mathbb{E}[v]-c}{c}\right)^t + 1}. \end{aligned}$$

As $T \rightarrow \infty$, the last term goes to:

$$\frac{(1 - q_1)(1 - \pi)}{q_1 \frac{\frac{\mathbb{E}[v]-c}{c}}{1 - \frac{\mathbb{E}[v]-c}{c}} + 1}.$$

Expanding the expressions for π and q_1 , one gets that the above equals $1 - \pi^H(C)$ in the canonical extensive form.

Suppose now that $\mathbb{E}[v] - c > c$. Then, the probability that exactly one player contributes when the state is H can be written as:

$$\frac{(1 - q_1)(1 - \pi)c^{T-3}}{q_1 \sum_{t=0}^{T-4} (\mathbb{E}[v] - c)^{T-3-t} c^t + c^{T-3}} = \left(\frac{c}{\mathbb{E}[v] - c}\right)^{T-3} \frac{(1 - q_1)(1 - \pi)}{q_1 \sum_{t=0}^{T-4} \left(\frac{c}{\mathbb{E}[v]-c}\right)^t + \left(\frac{c}{\mathbb{E}[v]-c}\right)^{T-3}}.$$

As $T \rightarrow \infty$, the second term is bounded, while the first term goes to 0. Hence, in this case, the probability that both agents contribute when the state is H goes to 1.

B PROOF OF THEOREM 3.1

Before proving the main theorem, we introduce some notation and definitions that are useful in what follows.

An information set h precedes h' if for all $v \in h$, there exists $v' \in h'$ such that $v < v'$. We denote this by $h < h'$. Perfect recall implies that for all $h, h' \in H_i$ either $h < h'$, $h' < h$, or neither hold.

For $v \in V$, we denote by $h(v)$ the information set to which v belongs. As in the main text, we use Z to denote the terminal nodes, that is the set $\{v \in V : (\nexists v' \in V)v < v'\}$. Also, for $v \in Z$, let $\gamma_i(v) = \text{proj}_{A_i}\gamma(v)$, that is $\gamma_i(v)$ is player i 's action in the action profile $\gamma(v)$.

Fix an admissible extensive form, Γ , that is, one that satisfies Conditions 1, 2, and 3, and an equilibrium assessment (σ^*, μ^*) of Γ . For each $i \in N$, partition the set of i 's information sets as follows:

$$H_i = H_i^1 \cup H_i^2 \cup H_i^3,$$

where

$$\begin{aligned} H_i^1 &= \{h \in H_i : (\forall v \in h)(\forall m \in M_i(h))\gamma_i(v) = \gamma_i(v, m)\} \\ H_i^2 &= \left\{ h \in H_i : \begin{array}{l} \gamma_i(h) = A_i \wedge [(\exists a_i \in A_i)(\exists m \in h)\gamma_i(v, m) = a_i, v \in h] \\ \wedge [(\exists a'_i \in A_i)(\forall m \in h)a'_i \in \gamma(v, m) \Rightarrow \gamma_i(v, m) = A_i] \end{array} \right\} \\ H_i^3 &= \{h \in H_i : \gamma_i(h) = A_i \wedge [(\forall a_i \in A_i)(\exists m \in M_i(h))\gamma_i(v, m) = a_i, v \in h]\} \end{aligned}$$

Information sets in H_i^1 are those in which player i 's moves are cheap talk: they don't commit him to any of his actions. Information sets in H_i^2 are those in which player i can commit to some of his actions in the base game, like a_i , but he cannot commit to all of them, like a'_i . Those information sets necessarily have a cheap talk move, that is, a move that allows player i not to commit to any of his actions. Information sets in H_i^3 are those in which player i has the ability to commit to any of his actions. Note that these information sets may have a cheap talk move

- you can think of it as a move that allows player i to pass on the opportunity of committing to any of his actions-, but since the extensive form is finite, these information sets always precede another information set in which player i can commit to any of his actions, and where no cheap talk moves are available.

For $h \in H_i^2 \cup H_i^3$, if player i has a move that allows him to commit to $a_i \in A_i$ at h , then this move is denoted by $m_{a_i}(h)$.

The proof proceeds in two steps. **Step 1** shows that, for any such pair $\langle \Gamma, (\sigma^*, \mu^*) \rangle$, if there exists $i \in N$ such that $H_i^2 \neq \emptyset$, then we can find $\langle \Gamma', (\sigma'^*, \mu'^*) \rangle$ such that for Γ' , $H_i^2 = \emptyset$, and (σ'^*, μ'^*) generates the same distribution over action profiles in G that (σ^*, μ^*) did.⁸ **Step 2** shows that any pair $\langle \Gamma, (\sigma^*, \mu^*) \rangle$ such that $(\forall i \in N) H_i^2 = \emptyset$ can be embedded in a canonical extensive form.

B.1 Step 1

We begin by making two preliminary observations, which are useful in what follows:

Observation 1 If $h \in H_i^2$ allows to commit player i to a_i , but not to $a'_i \in A_i \setminus \{a_i\}$, then there exists $h' \in H_i$, such that $h < h'$, and $(\forall v' \in h') \cdot \gamma_i(v') = A_i$

To see this, note that $h \in H_i^2$ implies that $\gamma_i(h) = \bigcup_{v \in h} \gamma_i(v) = A_i$. Perfect recall, and know your action imply that for all $v \in h$, $\gamma_i(v) = A_i$. Since $a'_i \in \gamma_i(v)$ for all $v \in h$, then for each $v \in h$ there exists $m_v \in M_i(h)$ such that $a'_i \in \gamma_i(v, m_v)$. Noticing that we can always take $m \in M_i(h)$ not to depend on v , we have that there exists $m \in M_i(h)$ such that $(\forall v \in h) \gamma_i(v, m) = A_i$. No delegation, and perfect recall imply that there exists $h < h'$ such that $\gamma_i(h') = A_i$.

Observation 2 If $h \in H_i^2$, there exists $\bar{h} \in H_i^3$ such that $h < \bar{h}$.

This follows from Observation 1, and the finiteness of the extensive form.

Fix an admissible Γ , and an equilibrium assessment (σ^*, μ^*) . Take $i \in N$ such

⁸Step 1 can, of course, be skipped if Γ satisfies that $(\forall i \in N) H_i^2 = \emptyset$

that $H_i^2 \neq \emptyset$. Define

$$H_i^{2,<} = \{h' \in H_i^2 : (\forall h \in H_i^3) : h' < h, \text{ if } \tilde{h} \in H_i \text{ is such that } h' < \tilde{h} < h, \text{ then } \tilde{h} \in H_i^1\},$$

to be the set of information sets of player i in H_i^2 such that, if he chooses not to commit to one of his actions, his next opportunity to do so allows him to commit to any of his actions. Fix $h' \in H_i^{2,<}$, and let a'_i denote the action player i cannot commit to at h' . Let $m_{h'}^*$ be defined by:

$$m_{h'}^* \in \arg \max_{\{m \in M(h') : \gamma_i(v) = \gamma_i(v, m)\}} U_i(\sigma_i^m, \sigma_{-i}^* | h'),$$

where

$$\sigma_i^m(h) = \begin{cases} \sigma_i^*(h) & \text{if } h \neq h' \\ \delta_m & \text{otherwise} \end{cases},$$

and

$$U_i(\sigma_i^m, \sigma_{-i}^* | h') = \sum_{v' \in h'} \mu^*(v' | h') \sum_{\{v \in Z : v' < v\}} Pr(v | \sigma_i^m, \sigma_{-i}^*, v') u_i(\gamma(v)).$$

Define:

$$H_i^3(h') = \{h \in H_i^3 : (\forall v' \in h') (\exists v \in h) : (v', m_{h'}^*) \leq v \wedge (\nexists \tilde{h} \in H_i^3) : h' < \tilde{h} < h\},$$

to be the set of information sets in H_i^3 that are immediate successors - amongst other information sets in H_i^3 - of h' . Also, let

$$V^<(H_i^3(h')) = \{v \in V : (\exists v' \in h') (v', m_{h'}^*) \leq v \wedge (\exists h \in H_3(h')) (\exists \bar{v} \in h) v \leq \bar{v}\},$$

denote the set of histories that follow $(v', m_{h'}^*)$, for some $v' \in h'$, and precede $\bar{v} \in h \in H_i^3(h')$. Finally, let

$$V^>(H_i^3(h')) = \{v \in V : (\exists h \in H_i^3(h')) (\exists \bar{v} \in h) (\bar{v}, m_{a'_i}(h)) \leq v\},$$

denote the histories that follow from $h \in H_i^3(h')$, when player i chooses to commit to a'_i at h .

For each $v \in V^<(H_i^3(h'))$, let

$$\lambda_1^>(v) = \begin{cases} \emptyset & \text{if } v \in \bigcup_{v' \in h'} \{(v', m_{h'}^*)\} \\ v \setminus (v', m_{h'}^*) & \text{if } (v', m_{h'}^*) < v \end{cases},$$

and construct

$$\lambda_1(v) = (v', m_{a'_i}(h'), \lambda_1^>(v)),$$

for each one. Similarly, for each $v \in V^>(H_i^3(h'))$, letting $\bar{v}(v)$ denote the node at $h \in H_i^2(h')$ such that $(\bar{v}(v), m_{a'_i}(h)) \leq v$, construct:

$$\lambda_2(v) = (\lambda_1(\bar{v}(v)), m_{a'_i}^d, v \setminus (\bar{v}(v), m_{a'_i}(h))),$$

where the superscript d in $m_{a'_i}^d$ signifies that this will be a dummy move, not a commitment move to play a'_i .

We are now ready to define the transformation Γ' of Γ . The tree is now given by:

$$V' = V \cup \bigcup_{v \in V^<(H_i^3(H'))} \{\lambda_1(v)\} \cup \bigcup_{v \in V^>(H_i^3(H'))} \{\lambda_2(v)\}.$$

For each $v \in Z \cap V^>(H_i^3(h'))$, let $\gamma(\lambda_2(v)) = \gamma(v)$.⁹

Let $M'_i(h') = M_i(h) \cup \{m_{a'_i}(h')\}$. For each $v \in V^<(H_i^3(h'))$,

1. If $\iota(v) = j \neq i$, then $\iota(\lambda_1(v)) = j$, and $\lambda_1(v) = h_j(v)$. Extend $\mu^*(\cdot|h_j(v))$ so that $\mu^*(\lambda_1(v)|h_j(v)) = 0$.
2. If $\iota(v) = i$, and $v \notin h \in H_i^3(h')$, set $\iota(\lambda_1(v)) = c$, $M_c(\lambda_1(v)) = M_i(v)$. If $\iota(v) = i$, and $v \in H_i^3(h')$, set $\iota(\lambda_1(v)) = c$, $M_c(\lambda_1(v)) = \{m_{a'_i}^d\}$.
3. If $\iota(v) = c$, then $\iota(\lambda_1(v)) = c$, and $\lambda_1(v) \in h_c(v)$.

We now add the following information sets to H_c . For each $v \in V^<(H_i^3(h'))$ such that $\iota(v) = i$, let

$$h_c(\lambda_1(v)) = \{\lambda_1(\tilde{v}) : \tilde{v} \in h_i(v)\},$$

and extend σ_c so that $\sigma_c(\cdot|h_c(\lambda_1(v))) = \sigma_i^*(\cdot|h_i(v))$.

For each $v \in V^>(H_i^3(h'))$,

1. If $\iota(v) = j \neq i$, then $\iota(\lambda_2(v)) = j$, $\lambda_2(v) \in h_j(v)$, and extend $\mu^*(\cdot|h_j(v))$ so that $\mu^*(\lambda_2(v)|h_j(v)) = 0$.
2. If $\iota(v) = i$, then $\iota(\lambda_2(v)) = c$, and $M_c(\lambda_2(v)) = M_i(v)$.
3. If $\iota(v) = c$, then $\iota(\lambda_2(v)) = c$, and $\lambda_2(v) \in h_c(v)$.

As before, we add the following information sets to H_c : for each $v \in V^>(H_i^3(h'))$ such that $\iota(v) = i$, let

$$h_c(\lambda_2(v)) = \{\lambda_2(\tilde{v}) : \tilde{v} \in h_i(v)\},$$

and extend σ_c so that $\sigma_c(\cdot|h_c(\lambda_2(v))) = \sigma_i^*(\cdot|h_i(v))$.

We now argue that (σ^*, μ^*) continues to be an equilibrium assessment of Γ' , where in a slight abuse of notation, we keep the same labels for the assessment, even though the domain of σ_i^* , and of μ^* have changed.

Note that, for players other than i , their strategy continues to be a best response to σ_i^* since their beliefs assign probability 0 to the new nodes in their information sets. Therefore, we only need to show that player i 's incentives at h' have not changed. Indeed, in what follows, we show that the payoff agent i obtains from playing $m_{a_i}(h')$ coincides with the expected payoff, starting from h' , of playing $m_{a_i}(h)$ for $h \in H_i^3(h')$ (expectations taken with respect to the probability of reaching h , conditional on being at h' .) Since σ_i^* satisfies sequential rationality, for each $h \in H_i^3(h')$, the payoff of playing $m_{a_i}(h)$ at h is bounded above by the payoff of playing $\sigma_i^*(h)$. Thus, the payoff of playing $m_{a_i}(h')$ at h' is bounded above by the payoff of playing $m_{h'}^*$. This implies that player i 's incentives have not changed, and hence σ_i^* continues to be a best response to σ_{-i}^* in Γ' .

We start by showing that:

$$U_i(\sigma_i^{m_{h'}^*}, \sigma_{-i}^*|h') = \sum_{h \in H_i^3(h')} Pr(h|h', \sigma_i^{m_{h'}^*}, \sigma_{-i}^*) U_i(\sigma_i^*, \sigma_{-i}^*|h),$$

where

$$\begin{aligned} Pr(h|h', \sigma_i^{m_{h'}^*}, \sigma_{-i}^*) &= \sum_{v \in h} Pr(v|h', \sigma_i^{m_{h'}^*}, \sigma_{-i}^*) \\ &= \sum_{v \in h} \mu^*(v'(v)|h') \prod_{k=|v'(v)|}^{|v|-1} \sigma_{\iota(v^k)}(v_{k+1}), \end{aligned}$$

where $v'(v)$ is the element of h' that precedes v . To see this note that:

$$\begin{aligned}
U_i(\sigma_i^{m_{h'}^*}, \sigma_{-i}^* | h') &= \sum_{\{v \in Z : (\exists v' \in h') (v', m_{h'}^*) < v\}} Pr(v | \sigma_i^{m_{h'}^*}, \sigma_{-i}^*) u_i(\gamma(v)) \\
&= \sum_{h \in H_i^3(h')} \sum_{\tilde{v} \in h} \sum_{\{v \in Z : \tilde{v} < v\}} Pr(v | \sigma_i^{m_{h'}^*}, \sigma_{-i}^*, h') u_i(\gamma(v)) \\
&= \sum_{h \in H_i^3(h')} \sum_{\tilde{v} \in h} \sum_{\{v \in Z : \tilde{v} < v\}} \mu^*(v'(\tilde{v}) | h') \prod_{k=|v'(\tilde{v})|}^{|\tilde{v}|-1} \sigma_{\iota(\tilde{v}^k)}(\tilde{v}_{k+1}) \prod_{k=|\tilde{v}|}^{|v|-1} \sigma_{\iota(v^k)}(v_{k+1}) u_i(\gamma(v)) \\
&= \sum_{h \in H_i^3(h')} \sum_{\tilde{v} \in h} \mu^*(v'(\tilde{v}) | h') \prod_{k=|v'(\tilde{v})|}^{|\tilde{v}|-1} \sigma_{\iota(\tilde{v}^k)}(\tilde{v}_{k+1}) \sum_{\{v \in Z : \tilde{v} < v\}} \prod_{k=|\tilde{v}|}^{|v|-1} \sigma_{\iota(v^k)}(v_{k+1}) u_i(\gamma(v)) \\
&= \sum_{h \in H_i^3(h')} Pr(h | h', \sigma_i^{m_{h'}^*}, \sigma_{-i}^*) U_i(\sigma_i^*, \sigma_{-i}^* | h),
\end{aligned}$$

where the second equality follows from noting that $\{v \in Z : (\exists v' \in h') : (v', m_{h'}^*) < v\} = \bigcup_{h \in H_i^3(h')} \bigcup_{\tilde{v} \in h} \{v \in Z : \tilde{v} < v\}$, and the last equality follows from noting that: (i)

$\sum_{h \in H_i^3(h')} Pr(h | h', \sigma_i^{m_{h'}^*}, \sigma_{-i}^*) = 1$, (ii) if $Pr(h | h', \sigma_i^{m_{h'}^*}, \sigma_{-i}^*) = 0$, then $Pr(h | h', \sigma_i^{m_{h'}^*}, \sigma_{-i}^*) U_i(\sigma_i^*, \sigma_{-i}^* | h) = 0$, and (iii) if $Pr(h | h', \sigma_i^{m_{h'}^*}, \sigma_{-i}^*) \neq 0$ then multiplying and dividing through the

term corresponding to h by $Pr(h | h', \sigma_i^{m_{h'}^*}, \sigma_{-i}^*)$, we get :

$$\sum_{\tilde{v} \in h} \frac{\mu^*(v'(\tilde{v}) | h')}{Pr(h | h', \sigma_i^{m_{h'}^*}, \sigma_{-i}^*)} \prod_{k=|v'(\tilde{v})|}^{|\tilde{v}|-1} \sigma_{\iota(\tilde{v}^k)}(\tilde{v}_{k+1}) \sum_{\{v \in Z : \tilde{v} < v\}} \prod_{k=|\tilde{v}|}^{|v|-1} \sigma_{\iota(v^k)}(v_{k+1}) u_i(\gamma(v)) = U_i(\sigma_i^*, \sigma_{-i}^* | h).$$

Consider now the strategy for player i , $\hat{\sigma}_i$, defined as follows:

$$\hat{\sigma}_i(h) = \begin{cases} \sigma_i^*(h) & \text{if } h \notin \{h'\} \cup H_i^3(h') \\ \delta_{m_{h'}^*} & \text{if } h = h' \\ \delta_{m_{\alpha_i}^*}(h) & \text{if } h \in H_i^3(h') \end{cases}.$$

Sequential rationality of σ_i^* implies that for all $h \in H_i^3(h')$,

$$U_i(\sigma_i^*, \sigma_{-i}^* | h) \geq U_i(\hat{\sigma}_i, \sigma_{-i}^* | h).$$

Moreover, since $\hat{\sigma}_i$ only changes $\sigma_i^{m_{h'}^*}$ after h' , we have that:

$$\begin{aligned} U_i(\sigma_i^{m_{h'}^*}, \sigma_{-i}^* | h') &= \sum_{h \in H_i^3(h')} Pr(h | h', \sigma_i^{m_{h'}^*}, \sigma_{-i}^*) U_i(\sigma_i^*, \sigma_{-i}^* | h) \\ &\geq \sum_{h \in H_i^3(h')} Pr(h | h', \sigma_i^{m_{h'}^*}, \sigma_{-i}^*) U_i(\hat{\sigma}_i, \sigma_{-i}^* | h) = U_i(\hat{\sigma}_i, \sigma_{-i}^* | h'). \end{aligned}$$

Finally, we show that $U_i(\hat{\sigma}_i, \sigma_i^* | h') = U_i(\sigma_i^{m_{a_i'}(h')}, \sigma_{-i}^* | h')$. To see this, note that:

$$\begin{aligned} U_i(\hat{\sigma}_i, \sigma_{-i}^* | h') &= \sum_{\{v \in Z: (\exists v' \in h') v' < v\}} Pr(v | \hat{\sigma}_i, \sigma_{-i}^*, h') u_i(\gamma(v)) \\ &= \sum_{h \in H_i^3(h')} \sum_{\tilde{v} \in h} \sum_{\{v \in Z: \tilde{v} < v\}} Pr(v | \hat{\sigma}_i, \sigma_{-i}^*, h') u_i(\gamma(v)) \\ &= \sum_{h \in H_i^3(h')} \sum_{\tilde{v} \in h} \sum_{\{v \in Z: \tilde{v} < v\}} \left(\begin{aligned} &\mu^*(v'(\tilde{v}) | h') \delta_{m_{h'}^*} \prod_{k=|v'(\tilde{v})|+1}^{|\tilde{v}|-1} \sigma_{\iota(\tilde{v}^k)}(\tilde{v}_{k+1}) \delta_{m_{a_i'}(h)} \\ &\times \prod_{k=|\tilde{v}|}^{|v|-1} \sigma_{\iota(v^k)}(v_{k+1}) u_i(\gamma(v)) \end{aligned} \right) \\ &= \sum_{h \in H_i^3(h')} \sum_{\tilde{v} \in h} \sum_{\{v \in Z: \tilde{v} < v\}} \left(\begin{aligned} &\mu^*(v'(\tilde{v}) | h') \delta_{m_{a_i'}(h')} \prod_{k=|v'(\tilde{v})|+1}^{|\tilde{v}|-1} \sigma_{\iota(\lambda_1(\tilde{v}^k))}(\lambda_1(\tilde{v}_{k+1})) \\ &\times \delta_{m_{a_i'}^d} \prod_{k=|\tilde{v}|}^{|v|-1} \sigma_{\iota(\lambda_2(v^k))}(\lambda_2(v_{k+1})) u_i(\gamma(\lambda_2(v))) \end{aligned} \right) \end{aligned}$$

$$\begin{aligned}
&= \sum_{h \in H_i^3(h')} \sum_{\tilde{v} \in h} \sum_{\{v \in Z: \tilde{v} < v\}} Pr(\lambda_2(v) | \sigma_i^{m_{a'_i}(h')}, \sigma_{-i}^*, h') u_i(\gamma(\lambda_2(v))) \\
&= \sum_{\{v \in Z: (\exists v' \in h') (v', m_{a'_i}(h')) < v\}} Pr(v | \sigma_i^{m_{a'_i}(h')}, \sigma_{-i}^*, h') u_i(\gamma(v)) \\
&= U_i(\sigma_i^{m_{a'_i}(h')}, \sigma_{-i}^* | h'),
\end{aligned}$$

where the fourth, and fifth equalities hold because if $v \in Z$ has positive probability under $\hat{\sigma}_i, \sigma_{-i}^*$ starting from h' , then $v \in V^>(H_i^3(h'))$, and by construction $Pr(v | \hat{\sigma}_i, \sigma_{-i}^*, h') = Pr(\lambda_2(v) | \sigma_i^{m_{a'_i}(h')}, \sigma_{-i}^*, h')$, and $\gamma(v) = \gamma(\lambda_2(v))$.

This completes the proof that (σ^*, μ^*) is an equilibrium assessment of Γ' . Repeat the same step for each $a'_i \in A_i$ that player i can't commit to at h' , and then for each $\tilde{h} \in H_i^{2,<} \setminus \{h'\}$. In each step, player i will not have any incentive to change his move from what was originally prescribed by his strategy at \tilde{h} . Do this for each agent $j \neq i$, such that $H_j^{2,<} \neq \emptyset$. If the final extensive form satisfies that $(\forall i) H_i^{2,<} = \emptyset$, move to Step 2. Otherwise, repeat Step 1. Since the original extensive form Γ is finite, one proceeds to Step 2 in a finite number of rounds.

B.2 Step 2

Fix a pair $\langle \Gamma, (\sigma^*, \mu^*) \rangle$, where Γ is an admissible extensive form that satisfies $(\forall i \in N) H_i^2 = \emptyset$, and (σ^*, μ^*) is an equilibrium assessment of Γ . The proof of Step 2 follows the same ideas as the one in Gershkov and Szentes (2009). We first explain the transformations in words, and then we expand on them more formally.

Modify $\langle \Gamma, (\sigma^*, \mu^*) \rangle$ as follows:

- i. For each $i \in N$, for each $h_i \in H_i^1$, assign that information set to chance, and extend chance's moves so that $\sigma_c(\cdot | h) = \sigma_i^*(\cdot | h)$. This transformation does not change incentives.¹⁰ From now on, we deal with Γ' where $H_i^1 = \emptyset$ for all $i \in N$, that is, for all $i \in N$, $H_i^3 \equiv H_i$. Denote the new pair $\langle \Gamma', (\sigma^{*'}, \mu^{*'}) \rangle$.

¹⁰If we started from a sequential equilibrium, then we endow chance with trembles at these information sets. Here chance represents the mediator, and not nature.

- ii. If at an information set, $h \in H_i$, player i makes a move that does not fix his action in the base game with probability α , introduce a new information set for chance $h_c(h)$. At $h_c(h)$, chance moves, and makes the same (random) move as player i . For every node preceding h , information set h is reached with probability $1 - \alpha$, and information set $h_c(h)$ is reached with probability $1 - \alpha$. Do this for every such $h \in H_i$, and for every player i . Let $\langle \Gamma'', (\sigma^{*''}, \mu^{*''}) \rangle$ denote the new pair, and note that players' incentives have been preserved, i.e. the pair implements the same distribution over action profiles in G as $\langle \Gamma', (\sigma^{*'}, \mu^{*'}) \rangle$.
- iii. $\langle \Gamma'', (\sigma^{*''}, \mu^{*''}) \rangle$ is transformed so that players use pure strategies. Let $\langle \Gamma''', (\sigma^{*'''}, \mu^{*'''}) \rangle$ denote the new pair.
- iv. The pair obtained in [iii](#). satisfies that whenever an agent moves in the extensive form, he commits to an action in the base game G . We now transform $\langle \Gamma''', (\sigma^{*'''}, \mu^{*'''}) \rangle$ by unifying all information sets of player i in which player i commits to the same action.

Formally, for each i , let H_i^{****} be defined as follows:

$$h_i^{a_i} \in H_i^{****} \Leftrightarrow (\exists a_i \in A_i) : h_i^{a_i} = \bigcup_{\{\tilde{h} \in H_i^{***} : \sigma_i^{*'''}(\tilde{h}) = \delta_{m_{a_i}(h)}\}} \tilde{h}$$

$$M_i(h_i) = A_i$$

For each $i \in N$, let $\sigma_i^{*'''}(\cdot | h_i^{a_i}) = \delta_{a_i}$. If $H_i^{a_i} \equiv \{\tilde{h} \in H_i^{***} : \sigma_i^{*'''}(\tilde{h}) = \delta_{m_{a_i}(h)}\}$ has positive probability under $\sigma^{*''}$, define for $v \in h_i^{a_i}$:

$$\mu^{*'''}(v | h_i^{a_i}) = \frac{Pr(v | \sigma^{*'''})}{Pr(H_i^{a_i} | \sigma^{*'''})},$$

whereas if $H_i^{a_i} \neq \emptyset$ has zero probability under $\sigma^{*''}$, define, for $v \in h_i^{a_i}$:

$$\mu^{*'''}(v | h_i^{a_i}) = \frac{\mu^{*'''}(v | h(v))}{|H_i^{a_i}|},$$

where $h(v) \in H_i^{a_i}$ is the information set in which $v \in h_i^{a_i}$ belongs in Γ''' .¹¹

The extensive form we obtain after the successive transformations satisfies that along each path of play, each player moves exactly once, and when he does, he can choose to commit to any of his actions in the base, and he does so. Hence, the distribution over action profiles in the base game G implemented by Γ can also be implemented in the canonical extensive form.

We now go through items **i.-iii.** formally. With regards to item **i.**, consider $\langle \Gamma', (\sigma^{*'}, \mu^{*'}) \rangle$ defined as follows:

$$\begin{aligned}
(\forall i \in N) H_i' &= H_i^3 \\
(\forall i \in N)(\forall h \in H_i^1) \iota'(h) &= c \wedge M_{\iota(h)}(h) = M_i(h) \\
H_c' &= H_c \cup \bigcup_{i \in N} H_i^1 \\
\sigma_c'(h) &= \begin{cases} \sigma_c(h) & \text{if } h \in H_c \\ \sigma_i^*(h) & \text{if } h \in H_i^1 \end{cases} \\
(\forall i \in N)(\forall h \in H_i^3) \sigma_i^{*'}(h) &= \sigma_i^*(h) \\
(\forall i \in N)(\forall h \in H_i^3) \mu^{*'}(|h) &= \mu^*(|h)
\end{aligned}$$

Remark B.1. If (σ^*, μ^*) is a sequential equilibrium, we need to endow the mediator with the trembles of the players at $h \in \bigcup_{i \in N} H_i^1$, so as to be able to preserve agents' beliefs off-path.

We now proceed to item **ii.** For each $i \in N$, write $H_i' = \{h \in H_i' : (\exists m \in M_i(h))(\forall v \in h)\gamma_i(v) = \gamma_i(v, m)\} \cup \{h \in H_i' : (\forall m \in M_i(h))(\forall v \in h)\gamma_i(v) \neq \gamma_i(v, m)\} \equiv H_i^{3,P} \cup H_i^{3,C}$. Recall that in Γ' , all information sets of player i are such that he can commit to any of his actions. Information sets in $H_i^{3,P}$ have cheap talk moves, that is, moves that allow him to postpone committing to one of his actions, whereas the only thing he can do in an information set in $H_i^{3,C}$ is

to commit to one of his actions.

Suppose there exists $i \in N$, such that $H_i^{3,P} \neq \emptyset$. Fix $h_i \in H_i^{3,P}$. Recall the convention that if $m \in M_i(h_i)$ is such that for $v \in h_i$, $\gamma_i(v) \neq \gamma_i(v, m) = a_i$, then we label $m_{a_i}(h_i) \equiv m$. Denote by $M^*(h_i) = M_i(h_i) \setminus \bigcup_{a_i \in A_i} \{m_{a_i}(h_i)\}$. For each $m \in M_i(h_i)$, define:

$$V_m^>(h_i) = \{v' \in V' : (\exists v \in h_i)(v, m) \leq v'\}.$$

We transform V' as follows. For each $v' \in V_m^>(h_i)$, let $v(v') \in h_i$ be such that $v(v') < v'$. Define:

$$\lambda(v') = \begin{cases} (v(v'), m_i, v' \setminus v(v')) & \text{if } v' \in \bigcup_{a_i \in A_i} V_{m_{a_i}(h_i)}^>(h_i) \\ (v(v'), m_c, v' \setminus v(v')) & \text{otherwise} \end{cases}.$$

Let the new tree be given by:

$$V'' = (V' \setminus \bigcup_{m \in M_i(h_i)} V_m^>(h_i)) \cup \bigcup_{v \in h_i} \{(v, m_i), (v, m_c)\} \cup \bigcup_{\substack{v' \in \\ m \in M_i(h_i)}} \bigcup_{V_m^>(h_i)} \{\lambda(v')\}.$$

For each $v \in h_i$, let $\iota''(v) = c$, and $M_c''(h_i) = \{m_i, m_c\}$. For all $v \in h_i$, let $\iota''(v, m_i) = i$, $h'_i = \bigcup_{v \in h_i} \{(v, m_i)\}$, $M_i(h'_i) = M_i(h_i) \setminus M^*(h_i)$, and $\mu^{*''}((v, m_i) | h'_i) = \mu^{*'}(v | h_i)$. Moreover, for all $v \in h_i$, let $\iota''(v, m_c) = c$, $h'_c(v, m_c) = \{(v, m_c)\}$, $M_c(h'_c(v, m_c)) = M^*(h_i)$.

For each $v' \in \bigcup_{m \in M_i(h_i)} V_m^>(h_i)$, let:

$$\begin{aligned} \iota''(\lambda(v')) &= \iota(v'), \\ M_{\iota''(\lambda(v'))}(\lambda(v')) &= M_{\iota(v')}(v'), \\ h_{\iota''(\lambda(v'))}(\lambda(v')) &= \{\lambda(\tilde{v}) : \tilde{v} \in h_{\iota(v')}(v')\}, \\ \sigma_{\iota''(h(\lambda(v')))}^{*''}(|h(\lambda(v')))) &= \sigma_{\iota(h(v'))}^{*'}(|h(v')), \\ \mu^{*''}(|h(\lambda(v')))) &= \mu^{*'}(|h(v')). \end{aligned}$$

Finally, define:

$$\sigma_c''(m|h_i) = \begin{cases} \sum_{a_i \in A_i} \sigma_i^{*'}(m_{a_i}(h_i)|h_i) & \text{if } m = m_i \\ 1 - \sum_{a_i \in A_i} \sigma_i^{*'}(m_{a_i}(h_i)|h_i) & \text{otherwise} \end{cases}.$$

We have to consider three cases:

1. $\sigma_c^{*''}(m_i|h_i) \in (0, 1)$. Then, let $\sigma_i^{*''}(m_{a_i}(h_i)|h'_i) = \frac{\sigma_i^*(m_{a_i}(h_i)|h_i)}{\sigma_c^{*''}(m_i|h_i)}$, and for all $v \in h_i$, for all $m \in M^*(h_i)$, let $\sigma_c^{*''}(m|(v, m_c)) = \frac{\sigma_i^{*'}(m|h_i)}{1 - \sigma_c^{*''}(m_i|h_i)}$.
2. $\sigma_c^{*''}(m_i|h_i) = 1$. Then, $\sigma_i^{*''}(m_{a_i}(h_i)|h'_i) = \sigma_i^*(m_{a_i}(h_i)|h_i)$, and $(\forall v \in h_i) \sigma_c^{*''}(|(v, m_c)) \in \Delta(M^*(h_i))$.

Remark B.2. Note that if we had a sequential equilibrium, then we could pick (up to a subsequence):

$$\sigma_c^{*''}(m|(v, m_c)) = \lim_{k \rightarrow \infty} \frac{\sigma_i^{*',k}(m|h_i)}{\sum_{m' \in M^*(h_i)} \sigma_i^{*',k}(m'|h_i)}$$

where $\sigma_i^{*',k}(m|h_i)$ are the trembles associated with the sequential equilibrium.

3. $\sigma_c^{*''}(m_i|h_i) = 0$. Then, let $\sigma_c^{*''}(m|(v, m_c)) = \sigma_i^{*'}(m|h_i)$ for all $v \in h_i, m \in M^*(h_i)$, and pick $\sigma_i^{*''}(|h'_i)$ so that it is a best response at h'_i to beliefs $\mu^{*''}(|h'_i)$.

Remark B.3. In the case we were given a sequential equilibrium, we may want to give the mediator $|A_i|$ moves that lead to player i moving, so that the trembles associated with player i playing $\{m_{a_i}(h_i)\}_{a_i \in A_i}$ can be replicated.

Note that we have not altered incentives for any agent. Hence, $(\sigma^{*''}, \mu^{*''})$ constitute an equilibrium assessment of Γ'' .

Repeat this for every $h \in H_i^{3,P}$, and for every $i \in N$. In a slight abuse of notation, let $\langle \Gamma'', (\sigma^{*''}, \mu^{*''}) \rangle$ be the resulting extensive form.

Remark B.4. Note that in Γ'' , for all $i \in N$, $H_i^{3,P} = \emptyset$, and, hence, all $h_i, h'_i \in H_i''$ satisfy that neither $h_i < h'_i$, nor $h'_i < h_i$.

We now proceed to Step [iii.](#), that is, we turn all (possibly mixed) behavioral strategies of players in $\sigma^{*''}$ into pure behavioral strategies.

For each $i \in N$, let $H_i^{'',M} = \{h_i \in H_i'' : |\text{supp}\sigma^{*''}(\cdot|h_i)| > 1\}$. Fix $h_i \in H_i^{'',M}$, and as before, let:

$$V_{a_i}^>(h_i) = \{v'' \in V'' : (\exists v \in h)(v, m_{a_i}(h)) \leq v''\}.$$

Transform V'' as follows. For each $v'' \in V_{a_i}^>(h_i)$, for each $\tilde{a}_i \in A_i$ such that $\sigma_i^{*''}(m_{\tilde{a}_i}(h_i)|h_i) > 0$, let:

$$\lambda(v'') = (v(v''), \tilde{a}_i, v'' \setminus v(v'')),$$

where $v(v'') \in h_i$ is such that $v(v'') < v''$.

Define the new tree to be:

$$V''' = (V'' \setminus \bigcup_{a_i \in A_i} V_{a_i}^>(h_i)) \cup \bigcup_{v \in h_i, m_{a_i} \in \text{supp}\sigma_i^{*''}(|h_i)} \{(v, a_i)\} \cup \bigcup_{v'' \in \bigcup_{a_i} V_{a_i}^>(h_i)} \bigcup_{m_{\tilde{a}_i} \in \text{supp}\sigma_i^{*''}(|h_i)} \{\lambda_{\tilde{a}_i}(v'')\}.$$

For each $v \in h_i$,

$$\begin{aligned} t'''(v) &= c, \\ M_c(h_i) &= \{a_i : \sigma_i^{*''}(m_{a_i}(h_i)|h_i) > 0\} \\ \sigma_c'''(a_i|h_i) &= \sigma_i^{*''}(m_{a_i}(h_i)|h_i), \\ t'''(v, a_i) &= i, \\ M_i(v, a_i) &= \bigcup_{a_i \in A_i} \{m_{a_i}(h_i)\}. \end{aligned}$$

Define:

$$\begin{aligned} h_{i,\tilde{a}_i}''' &= \cup_{v \in h_i} \{(v, \tilde{a}_i)\}, \\ \sigma_i^{*'''}(h_{i,\tilde{a}_i}''') &= \delta_{m_{\tilde{a}_i}(h_i)}, \\ \mu^{*'''}((v, \tilde{a}_i)|h_{i,\tilde{a}_i}''') &= \mu^{*''}(v|h_i). \end{aligned}$$

For each $v'' \in \bigcup_{a_i \in A_i} V_{a_i}^>(h_i)$, let:

$$\begin{aligned} l'''(\lambda_{\tilde{a}_i}(v'')) &= l''(v''), \\ h_{l'''(\lambda_{\tilde{a}_i}(v''))}(\lambda_{\tilde{a}_i}(v'')) &= \{\lambda(\tilde{v}'') : \tilde{v}'' \in h(v'')\} \\ M_{l'''(\lambda(v''))}(\lambda(v'')) &= M_{l''(v'')}(v'') \\ \sigma_{l'''(h(\lambda(v'')))}^{*'''}(h_{l'''(\lambda(v''))}(\lambda(v''))) &= \sigma_{l''(h(v''))}^{*''}(\cdot|h_{l''(v'')}(v'')), \\ \mu^{*'''}(\lambda_{\tilde{a}_i}(v'')|h_{l'''(\lambda(v''))}) &= \begin{cases} \mu^{*''}(v''|h_{l''(v'')}(v'')) & \text{if } v'' \in V_{\tilde{a}_i}^>(h_i) \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

Note that player i , after moving at h'_i , moves no longer in the game. The beliefs of other players are specified so that they assign zero probability to the mediator recommending \tilde{a}_i , and player i choosing $a_i \neq \tilde{a}_i$. Thus, for players moving after i , the modification leaves their incentives unchanged, whereas player i was indifferent between all of the moves in $\text{supp} \sigma_i^{*''}(\cdot|h_i)$, and hence, the new strategy is still a best response.

REFERENCES

- [1] R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of mathematical Economics*, 1(1):67–96, 1974. pages 15
- [2] R. J. Aumann. Correlated equilibrium as an expression of bayesian rationality. *Econometrica: Journal of the Econometric Society*, pages 1–18, 1987. pages 15
- [3] S. Bade, G. Haeringer, and L. Renou. Bilateral commitment. *Journal of Economic Theory*, 144(4):1817–1831, 2009. pages 16

- [4] D. Bergemann and S. Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016. pages 1, 2, 15, 21
- [5] K. Bognar, T. Börgers, and M. Meyer-ter Vehn. An optimal voting procedure when voting is costly. *Journal of Economic Theory*, 159:1056–1073, 2015. pages 14
- [6] R. Calcagno, Y. Kamada, S. Lovo, and T. Sugaya. Asynchronicity and coordination in common and opposing interest games. *Theoretical Economics*, 9(2):409–434, 2014. pages 16
- [7] G. Caruana and L. Einav. A theory of endogenous commitment. *The Review of Economic Studies*, 75(1):99–116, 2008. pages 16
- [8] R. Dutta and R. Ishii. Dynamic commitment games, efficiency and coordination. *Journal of Economic Theory*, 163:699–727, 2016. pages 16
- [9] J. Ely, A. Frankel, and E. Kamenica. Suspense and surprise. *Journal of Political Economy*, 123(1):215–260, 2015. pages 15
- [10] J. C. Ely. Beeps. *The American Economic Review*, 107(1):31–53, 2017. pages 15
- [11] A. Gallice and I. Monzon. Cooperation in social dilemmas through position uncertainty. 2017. pages 14
- [12] A. Gershkov and B. Szentes. Optimal voting schemes with costly information acquisition. *Journal of Economic Theory*, 144(1):36–68, 2009. pages 12, 13
- [13] Y. Kamada and M. Kandori. Asynchronous revision games. Technical report, mimeo, 2011. pages 16
- [14] E. Kamenica and M. Gentzkow. Bayesian persuasion. *American Economic Review*, 101:2590–2615, 2011. pages 15
- [15] R. Lagunoff and A. Matsui. Asynchronous choice in repeated coordination games. *Econometrica*, 65(6):1467–1477, 1997. pages 16

- [16] G. J. Mailath, L. Samuelson, and J. M. Swinkels. Normal form structures in extensive form games. *Journal of Economic Theory*, 64(2):325–371, 1994. pages 17
- [17] R. B. Myerson. Optimal coordination mechanisms in generalized principal–agent problems. *Journal of mathematical economics*, 10(1):67–81, 1982. pages 12, 15
- [18] R. B. Myerson. Multistage games with communication. *Econometrica: Journal of the Econometric Society*, pages 323–358, 1986. pages 26
- [19] M. Peters. Reciprocal contracting. *Journal of Economic Theory*, 158:102–126, 2015. pages 14
- [20] L. Renou. Commitment games. *Games and Economic Behavior*, 66(1):488–505, 2009. pages 16
- [21] B. Salcedo. Implementation without commitment in moral hazard environments. *arXiv preprint arXiv:1303.0916*, 2013. pages 12, 13, 14
- [22] J. Watson. A general, practicable definition of perfect bayesian equilibrium. 2017. pages 22