Privately Designed Correlated Equilibrium*

Guilherme Carmona† Krittanai Laohakunakorn‡

University of Surrey University of Surrey

May 16, 2025

Abstract

We consider a setting where each player of a simultaneous-move game privately designs an information structure before playing the game. One of these designs is chosen at random to determine the distribution of the private messages that players receive. These messages allow players to correlate their actions; however, private information design implies a push from correlated to Nash equilibria. Indeed, the sequential equilibrium payoffs of the private information design extensive-form game are correlated equilibrium payoffs of the underlying simultaneous-move game, but not all correlated equilibrium payoffs are sequential equilibrium payoffs. In generic 2-player games, the latter are specific convex combinations of two Nash equilibrium payoffs.

*We wish to thank Gabriel Carroll, Hervé Moulin, Constantine Sorokin and seminar participants at the University of Glasgow, the 2023 SAET conference (Paris), the 2023 EWET (Naples) and the 2023 Portuguese Economic Journal conference (Braga) for helpful comments. Any remaining errors are, of course, ours.

†Address: University of Surrey, School of Economics, Guildford, GU2 7XH, UK; email: g.carmona@surrey.ac.uk.

‡Address: University of Surrey, School of Economics, Guildford, GU2 7XH, UK; email: k.laohakunakorn@surrey.ac.uk.

1 Introduction

It is well-known since Aumann (1974) that all players in a normal-form game can obtain a payoff higher than in any of its Nash equilibria by correlating their play, i.e. in a correlated equilibrium.[1] Achieving correlated equilibrium payoffs requires lotteries over a set of messages that are privately observed by the players and which can be thought of as being chosen by an outside mediator. Since the assumption of an impartial mediator may not always be appropriate, there is an interest in the payoffs that can be achieved through unmediated interaction between the players. Bárány (1992), Ben-Porath (1998), Urbano and Vila (2002) and Gerardi (2004) among others have shown that (nearly) all correlated equilibrium payoffs can be obtained through unmediated interaction.[2]

Aumann and Hart's (2003) results already imply that for two player games, pre-play cheap talk can achieve the entire convex hull of Nash payoffs (but no more). Thus, in the above papers, either the number of players is assumed to be greater than two or players have access to richer communication technology than cheap talk (e.g. balls and urns, public verification). In this paper, we consider this question from a different perspective. We focus on 2-player games and allow players access to fully mediated communication as long as they can agree on the mediation. However, although the *technology* of mediated communication is available, we assume that players can *manipulate* this technology in a general way. This addresses a difficulty with some of the above results, which is that certain deviations are ruled out by assumption. For example, in Ben-Porath's (1998) result for 2-player games, player 2 lets player 1 choose a ball from an urn. But player 1 cannot deviate by secretly manipulating the content of the urn before choosing from it. Do conceivable manipulations such as this one matter for the correlated equilibrium payoffs that can be obtained through unmediated interaction?

[1]Note, however, that Neyman (1997) defines a class of games having a smooth concave potential such that any correlated equilibrium is a convex combination of pure strategy Nash equilibria.

[2]Other related papers include Ben-Porath (2003), Krishna (2007), Wagner (2011), Rivera (2018) and Blume (2024).

To answer the above question, we introduce a model of unmediated (or, more accurately, self-mediated) interaction that builds on the information design literature popularized by a large number of papers since the influential work of Kamenica and Gentzkow (2011). As in this literature, the payoff of each player depends on the lottery over privately-observed message profiles, on the information structure for short, hence it is likely that he will try to design it in an optimal way. Our model is one of private information design in the sense that each player chooses an information design. Furthermore, it is guided by the observation that there are many actions that players can take to influence the information structure: for example, one player may anticipate that another will tamper with an agreed upon randomization device and respond by including additional safeguards. The other may anticipate this and secretly hide backdoors in the device. It is difficult to model explicitly each possible manipulation and its effect on the resulting information structure. On the other hand, we do not wish to rule out any kind of manipulation by assumption.

Thus, our aim is to provide a reduced form model that captures the idea that players are able to try to manipulate the information structure in any way they desire. We achieve this by letting each player choose the information structure directly. Our model is also a reduced form model of conflict as it specifies what information structure actually determines message profiles when different players choose different information structures. Our specification is that each player's chosen information structure is the one that actually determines message profiles with a strictly positive probability, i.e. each player i's information design is chosen with probability $\beta_i > 0$ (with $\sum_i \beta_i = 1$) to determine the message profile that players receive. This specification is a tractable way of obtaining that (i) if all players choose the same information structure, then message profiles are drawn from such common information structure, and (ii) each player is, with strictly positive probability, successful in attempting to manipulate the information structure however he wishes; this strictly positive probability can be thought of as the relative power that each player has in determining the information structure that actually determines message profiles.

We focus on 2-player simultaneous-move games and analyze the extensive-form game where players first choose an information design and then play the simultaneous-move game. We show that the set of (Nash or sequential) equilibrium payoffs of the extensive-form is a specific subset of the convex hull of the Nash equilibrium payoffs of the simultaneous-move game. For generic 2-player simultaneous games, the only achievable payoffs are specific convex combinations of two Nash equilibrium payoffs. These results are in contrast with, e.g. Aumann and Hart (2003) or Ben-Porath (1998), and show that the details of what is allowed for players to choose in unmediated interaction matter for the payoffs that can be achieved in equilibrium.[3] In particular, this paper shows that when information is designed optimally by the individuals involved in a strategic situation, very few correlated equilibrium payoffs can be achieved and there is a push from correlated to Nash equilibria.

The paper is organized as follows. Section 2 introduces our model of private information design and characterizes the equilibrium outcomes of the information design extensive-form game. Section 3 contains a motivating example suggesting an interpretation for our model and results. Related literature is discussed in Section 4, along with extensions and concluding remarks. Proofs of our main results can be found in the Appendix. Some details of the extensions we consider in Section 4 are left to the supplementary material.[4]

2 Privately designed correlated equilibrium

Consider a 2-player simultaneous-move game $G = (A_i, u_i)_{i \in N}$ where $N = \{1, 2\}$ is the set of players and, for each $i \in N$, A_i is a finite set of player i's actions and $u_i : A \to \mathbb{R}$ is player i's payoff function, where $A = \prod_{i \in N} A_i$. Let $N(G)$ denote the set of Nash

[3] See Section 4 for a detailed comparison with these papers. Our view is that there are many plausible models of unmediated interaction, but we find the one we put forward in this paper interesting because (i) it allows certain deviations that are ruled out in other models and, consequently, (ii) it leads to the extreme result that only a few correlated equilibrium payoffs can be achieved.

[4] Available at https://klaohakunakorn.com/idsm.pdf

equilibria of G and $C(G)$ the set of correlated equilibria of G.[5]

Before the game G is played, each player chooses an information design, which sends messages to both players. We thus consider a setting with *private information design* which is formalized by the following extensive-form game G_{id}. At the beginning of the game, each player $i \in N$ chooses simultaneously an information design which is a probability distribution over messages. The set of messages each player $i \in N$ can potentially receive is $M_i = \mathbb{N}$. An *information design* is a finitely supported probability measure on $M = M_1 \times M_2 = \mathbb{N}^2$. Let S be the set of information designs. Thus, each player i chooses an information design $\phi_i \in S$. After all players have chosen their information designs, a profile of messages $m \in M$ is realized according to $\phi \in \Delta(M)$ defined by setting, for each $m \in M$,

$$\phi[m] = \sum_{i \in N} \beta_i \phi_i[m],$$

where $\beta_i > 0$ for each $i \in N$ and $\sum_{i \in N} \beta_i = 1$; the probabilities β_1 and β_2 are exogenous and fixed throughout the paper, and one interpretation for them is that the information design of each $i \in N$ is chosen by nature with probability β_i. Each player $i \in N$ observes his coordinate $m_i \in M_i$ of the realized message profile m and his choice $\phi_i \in S$ but not the other player's coordinate $m_j \in M_j$ of the realized message profile m or choice $\phi_j \in S$, where $j \neq i$. Then each player i chooses an action $a_i \in A_i$ conditional on the observed (m_i, ϕ_i). Player i's payoff is then $u_i(a_1, a_2)$.

The information design is private in the sense that (i) it is done by the players, (ii) each player's choice of information design is his own private information and (iii) no player observes the aggregate information design. Assuming that information designs have finite support implies that each player always has the choice of knowing whether his information design is the one that was chosen by nature; indeed, the set of messages he can receive if his opponent's design is chosen is the finite subset $\text{supp}(\phi_{j,M_i})$ of $\mathbb{N}$ and, hence, he can choose ϕ_i such that $\text{supp}(\phi_{i,M_i})$ belongs to the

[5]Given a metric space X, $\Delta(X)$ denotes the set of Borel probability measures on X. For each $\mu \in \Delta(X)$, $\text{supp}(\mu)$ denotes the support of μ. When $X = \prod_{j \in J} X_j$ for some finite set J, μ_{X_j} denotes the marginal of μ on X_j for each $j \in J$. In particular, $N(G) \subseteq \Delta(A_1) \times \Delta(A_2)$ and $C(G) \subseteq \Delta(A)$.

complement of $\text{supp}(\phi_{j,M_i})$.[6]

A (behavioral) strategy for player $i \in N$ is $\pi_i = (\pi_i^1, \pi_i^2)$ such that $\pi_i^1 \in \Delta(S)$ and $\pi_i^2 : M_i \times S \to \Delta(A_i)$ is measurable.[7] A strategy is $\pi = (\pi_1, \pi_2)$ and let Π^* be the set of strategies. We focus mostly on strategies where players do not mix over the choice of information structures.[8] Let Π be the set of strategies π such that $\pi_i^1 \in S$ (i.e. π_i^1 is pure) for each $i \in N$.

For each strategy $\pi \in \Pi$ and for each $i \in N$, $m_i \in M_i$ and $\phi_i \in S$, we often write $\phi_i^* = \pi_i^1$, $\pi_i(m_i, \phi_i) = \pi_i^2(m_i, \phi_i)$ and $\pi_i(m_i) = \pi_i^2(m_i, \phi_i^*)$. For each $m \in M$ and $\phi \in S^2$, we let $\pi^2(m, \phi)$ be defined by $\pi_i^2(m, \phi) = \pi_i^2(m_i, \phi_i)$ for each $i \in N$; we also write $\pi(m, \phi)$ for $\pi^2(m, \phi)$ and use $\pi_{-i}^2(m_{-i}, \phi_{-i})$ and $\pi_{-i}(m_{-i}, \phi_{-i})$ for the vector of mixed actions $\pi(m, \phi)$ without the ith coordinate.

For each $\pi \in \Pi$, we also write $u_i(\pi) = \sum_{m \in M} \phi^*[m] u_i(\pi(m))$ for each $i \in N$, where $\phi^*[m] = \sum_{i \in N} \beta_i \phi_i^*[m]$, $\pi(m) \in \Delta(A)$ is defined by $\pi(m)[a] = \prod_{i \in N} \pi_i(m_i)[a_i]$ for each $a \in A$ and, for each $\sigma \in \Delta(A)$, $u_i(\sigma) = \sum_{a \in A} \sigma[a] u_i(a)$. We sometimes abuse notation and also let $\pi(m) = (\pi_1(m_1), \pi_2(m_2))$.

We use Nash equilibrium and sequential equilibrium as solution concepts. Sequential equilibrium is defined analogously to Myerson and Reny (2020) (MR henceforth): a strategy $\pi \in \Pi$ is a *sequential equilibrium* if it is a perfect conditional ε-equilibrium for each $\varepsilon > 0$.[9]

In contrast to our setting, if an impartial mediator sends messages according to some exogenously given information design $\phi \in S$, the set of equilibrium action distributions that result from varying ϕ is exactly the set of correlated equilibria of G, as shown by Aumann (1987).[10] With private information design there will, in general, be a reduction in the set of equilibrium outcomes. The reason is that the messages $m \in \text{supp}(\phi_i^*)$ that each player i sends must be optimal for player i. This is

[6]The assumption that the support of information designs is finite is also technically convenient since then expected values are finite sums.

[7]The set S is endowed with the topology of the weak convergence of probability measures.

[8]See Section 4.6 for an extension of our results to the case where players can mix over the information design.

[9]See A.1 in the Appendix for the definition of perfect conditional ε-equilibrium in our setting.

[10]This result is also implied by Myerson (1982, Proposition 2).

established in Theorem 1 which fully characterizes the set of sequential equilibrium outcomes of G_{id}.

The following notation is used in the statement of Theorem 1. The *outcome* of a strategy $\pi \in \Pi$ is $\left(\phi_i^*, (\pi_i(m_i))_{m_i \in \text{supp}(\phi^*_{M_i})}\right)_{i \in N}$; it consists of the information design for each player and, for each message that he may receive with strictly positive probability, the action he will choose in response. Let $M^* = \prod_{i \in N} \text{supp}(\phi^*_{M_i})$ be the product of the set of messages that each player may receive with strictly positive probability. For each $i \in N$, $j \neq i$ and $\delta \in \Delta(A_j)$, let $v_i(\delta) = \max_{\alpha \in \Delta(A_i)} u_i(\alpha, \delta)$ and $BR_i(\delta) = \{\alpha \in \Delta(A_i) : u_i(\alpha, \delta) = v_i(\delta)\}$ be, respectively, player i's value function and best-reply correspondence.

Theorem 1. *For each 2-player game G, the following conditions are equivalent:*

1. $\left(\phi_i^*, (\pi_i(m_i))_{m_i \in \text{supp}(\phi^*_{M_i})}\right)_{i \in N}$ *is the outcome of a Nash equilibrium of G_{id}.*
2. $\left(\phi_i^*, (\pi_i(m_i))_{m_i \in \text{supp}(\phi^*_{M_i})}\right)_{i \in N}$ *is the outcome of a sequential equilibrium of G_{id}.*
3. $\left(\phi_i^*, (\pi_i(m_i))_{m_i \in \text{supp}(\phi^*_{M_i})}\right)_{i \in N}$ *is such that, for each $i, j \in N$ and $j \neq i$,*

$$v_i(\pi_j(m_j)) = \max_{m'_j \in M^*_j} v_i(\pi_j(m'_j)) \text{ and } \pi_i(m_i) \in BR_i(\pi_j(m_j)) \tag{1}$$

for each $m \in \text{supp}(\phi_i^*)$, and

$$\pi_i(m_i) \text{ solves } \max_{\alpha_i \in \Delta(A_i)} \sum_{m_j} \frac{\phi_j^*[m_i, m_{-i}]}{\phi^*_{j,M_i}[m_i]} u_i(\alpha_i, \pi_j(m_j)) \tag{2}$$

for each $m_i \in \text{supp}(\phi^*_{j,M_i})$.

Theorem 1 shows that Nash and sequential equilibrium outcomes of the private information design game coincide. These are characterized by the optimality of the messages each player sends and of the actions he chooses. Each message profile sent by a player is optimal in the sense that the payoff of the action profile it induces is the highest amongst the action profiles belonging to the outcome. The optimality of the actions chosen by each player $i \in N$ consists of $\pi_i(m_i)$ maximizing his expected payoff conditional on his information design not being chosen when m_i is a message that he receives with strictly positive probability from the information design of the other player. The two optimality conditions imply that, for each player i, $\pi_i(m_i)$ maximizes his expected payoff conditional on his information design being chosen when $m_i \in \text{supp}(\phi^*_{i,M_i}) \setminus \text{supp}(\phi^*_{j,M_i})$ and maximizes his expected payoff conditional on his information design not being chosen when $m_i \in \text{supp}(\phi^*_{j,M_i}) \setminus \text{supp}(\phi^*_{i,M_i})$; indeed, in the former case, player i can only have received message m_i if ϕ^*_i has been chosen and, in the latter case, only if ϕ^*_i has not been chosen. Furthermore, in the remaining case where $m_i \in \text{supp}(\phi^*_{i,M_i}) \cap \text{supp}(\phi^*_{j,M_i})$, it turns out that $\pi_i(m_i)$ must satisfy the two criteria. Briefly, this happens because player i can always make sure that the message he sends to himself is different from the ones he may receive from the other player (we will elaborate on conditions (1) and (2) in the proof of the necessity part of Theorem 1 in Section A.2.1).

The following corollary of Theorem 1 characterizes the action distributions of sequential (and Nash) equilibria of G_{id}. For each strategy $\pi \in \Pi$, the *action distribution* of π is $\sigma_\pi \in \Delta(A)$ such that, for each $a \in A$,

$$\sigma_\pi[a] = \sum_{m \in M^*} \phi^*[m]\pi(m)[a].$$

Let

$$A(G) = \{\sigma_\pi : \pi \in \Pi \text{ is a sequential equilibrium of } G_{id}\}$$

be the set of action distributions of the sequential equilibria of G_{id}.

Corollary 1 characterizes each equilibrium action distribution as the expected value of a specific distribution over Nash equilibria of G. Let $\Delta_f(\Delta(A_1) \times \Delta(A_2))$ be the set of finitely supported distributions over $\Delta(A_1) \times \Delta(A_2)$ and, for each $\psi \in \Delta_f(\Delta(A_1) \times \Delta(A_2))$, let $\psi^A \in \Delta(A)$ be defined by setting, for each $a \in A$,

$$\psi^A[a] = \sum_{(\alpha_1,\alpha_2) \in \text{supp}(\psi)} \psi[\alpha_1, \alpha_2]\alpha_1[a_1]\alpha_2[a_2];$$

i.e. ψ^A is the reduced lottery of ψ or, alternatively, the expected value of ψ. Corollary

1 then shows that $A(G)$ equals the following set:

$$\mathcal{A} = \Big\{ (\beta_1\psi_1 + \beta_2\psi_2)^A : \forall i \in N, \text{ there exists } L_i, (\eta^{i,l})_{l=1}^{L_i}, (\sigma^{i,l})_{l=1}^{L_i} \text{ such that}$$
$$\psi_i = \sum_{l=1}^{L_i} \eta^{i,l} 1_{\sigma^{i,l}}, \eta^i \geq 0, \sum_{l=1}^{L_i} \eta^{i,l} = 1,$$
$$\sigma^{i,l} \in N(G) \text{ and } u_i(\sigma^{i,k}) = u_i(\sigma^{i,l}) \geq u_i(\sigma^{j,r})$$
$$\forall k, l \in \{1, \ldots, L_i\}, j \in N \text{ and } r \in \{1, \ldots, L_j\} \Big\}.$$

Corollary 1. *For each 2-player game G, $A(G) = \mathcal{A}$.*

Corollary 1 characterizes the equilibrium action distributions of G_{id} for 2-player games. It shows that when player i's information design is chosen, there is a resulting distribution ψ_i over Nash equilibria of G, all of which give the same payoff to player i. Furthermore, this common payoff is no less than the payoff player i obtains in each of the Nash equilibria of G in the support of ψ_j, $j \neq i$. In other words, player i weakly prefers any Nash equilibria of G in the support of ψ_i to any of them in the support of ψ_j.

The characterization of equilibrium action distributions in Corollary 1 implies an analogous characterization of the set of equilibrium payoffs of G_{id}. Let

$$U(G) = \{u(\pi) : \pi \in \Pi \text{ is a sequential equilibrium of } G_{id}\}$$

be the set of sequential equilibrium payoffs of G_{id}. Corollary 2 shows that $U(G)$ equals the following set:

$$\mathcal{U} = \Big\{ \beta_1 u^1 + \beta_2 u^2 : \forall i \in N, \text{ there exists } L_i, (\eta^{i,l})_{l=1}^{L_i}, (\sigma^{i,l})_{l=1}^{L_i} \text{ such that}$$
$$u^i = \sum_{l=1}^{L_i} \eta^{i,l} u(\sigma^{i,l}), \eta^i \geq 0, \sum_{l=1}^{L_i} \eta^{i,l} = 1,$$
$$\sigma^{i,l} \in N(G) \text{ and } u_i(\sigma^{i,k}) = u_i(\sigma^{i,l}) \geq u_i(\sigma^{j,r})$$
$$\forall k, l \in \{1, \ldots, L_i\}, j \in N \text{ and } r \in \{1, \ldots, L_j\} \Big\}.$$

Corollary 2. *For each 2-player game G, $U(G) = \mathcal{U}$.*

Thus, in general, not all correlated equilibrium payoffs of G can be achieved in G_{id} when information is designed privately. Indeed, equilibrium payoffs of G_{id} form a particular subset of the convex hull of the Nash equilibrium payoffs of G. For the battle of the sexes,

$1\backslash 2$	A	B
A	$2,1$	$0,0$
B	$0,0$	$1,2$

Corollary 2 implies that $U(G) = u(N(G)) \cup \{\beta_1(2,1) + \beta_2(1,2)\}$.

The characterization of $U(G)$ is simpler in generic games, such as the battle of the sexes, since then the payoff resulting after each information design is chosen is that of a Nash equilibrium. Let $\mathcal{G}$ be the set of games such that, for each Nash equilibria σ and σ' of G, if $u_i(\sigma) = u_i(\sigma')$ for some $i \in N$, then $u_j(\sigma) = u_j(\sigma')$ for $j \neq i$ (equivalently, if $u_i(\sigma) \neq u_i(\sigma')$ for some $i \in N$ then $u_j(\sigma) \neq u_j(\sigma')$ for $j \neq i$). We regard $\mathcal{G}$ as a subset of $\mathbb{R}^{2|A|}$. A subset B of an Euclidean space is *generic* if the closure of its complement has Lebesgue measure zero.

Corollary 3. *The set $\mathcal{G}$ is generic and, for each 2-player game $G \in \mathcal{G}$,*

$$U(G) = \{\beta_1 u(\sigma) + \beta_2 u(\sigma') : \sigma, \sigma' \in N(G), u_1(\sigma) \geq u_1(\sigma'), u_2(\sigma') \geq u_2(\sigma)\}.$$

The proof of Corollary 3 actually shows that the set of games such that $u_i(\sigma) \neq u_i(\sigma')$ for each $i \in N$ and $\sigma, \sigma' \in N(G)$ such that $\sigma \neq \sigma'$ is generic. This set is contained in $\mathcal{G}$ and contains all games with a unique equilibrium as well as the battle of the sexes. It is clear from Corollary 3 that $U(G) = u(N(G))$ for each 2-player game G with a unique Nash equilibrium.

3 An example

We motivate our model of private information design in the context of the following game of "chicken":

1\2	A	B
A	$6,6$	$1,7$
B	$7,1$	$0,0$

We interpret this game as representing a scenario where two competing firms have the option to enter a new market (e.g. a new drug). This new market has the potential to be very profitable if and only if both firms enter; this can happen, for example, if the demand in the new market is high if and only if there is a large investment in R&D which is beyond the capability of a single firm. Thus, A stands for entry in the new market and B for no entry. If no firm enters the new market, each gets its profit in the old market. If only one firm enters, then that firm gets a small increase of $1 million in profits whereas the other firm obtains a large increase of $7 million in profits, for instance, because the latter firm becomes dominant in the old market. If both enter, then each obtains an increase of $6 million in profits.

An alternative interpretation of the chicken game is as a tariff war between two countries. In this interpretation, choosing B means increasing tariffs whereas A means keeping them at the current level. If only one country imposes high tariffs, that country gains while the other loses but trade between the two still occurs; in contrast, if both countries impose high tariffs, then trade between them collapses.

The Nash equilibria are (A, B), (B, A) and $(\frac{1}{2}1_A + \frac{1}{2}1_B, \frac{1}{2}1_A + \frac{1}{2}1_B)$ giving payoffs $(7, 1)$, $(1, 7)$ and $(\frac{7}{2}, \frac{7}{2})$ respectively. It is well-known that there are correlated equilibria with payoffs outside the convex hull of the Nash equilibrium payoffs. In the scenario of two competing firms who seek to collaborate, it is often the case that the effort to improve joint payoffs is assisted by a consulting firm, which may play the role of mediator by providing information and recommendations to both parties.[11] In

[11] For example, as described in Marshall and Marx (2012), the consulting firm AC-Treuhand was found guilty of participating in two cartels between 1993 and 2000. According to a European Court ruling, AC-Treuhand "acted as a moderator in case of tensions between members of the agreement and encouraged the parties to find compromises." In Carmona and Laohakunakorn (2023), we show that a consulting firm may help support cooperative payoffs in a repeated game by aggregating information. In contrast, our results in the current paper suggest that a consulting firm may be less effective in facilitating correlated equilibrium payoffs in a one-shot interaction.

the scenario of a tariff war, it is often the case that the effort to improve joint payoffs is carried out by negotiators that each country sends to put forward its position.

Suppose then that players 1 and 2 are competing firms and hire a consulting firm to act as a mediator to try to achieve a payoff of $4\frac{2}{3}$ each. For example, if the action profiles (A, A), (A, B) and (B, A) are recommended, each with probability $\frac{1}{3}$, then each player will find it optimal to follow the recommendation and the outcome is a correlated equilibrium with payoffs $(4\frac{2}{3}, 4\frac{2}{3})$. However, the consulting firm may not be an impartial mediator but may be susceptible to influence by the players themselves. In this case, can the players still achieve higher payoffs with its help?

Consider a player who has received recommendation A. He finds it optimal to play A because he believes that his opponent has received recommendations A and B with equal probability. But if he could acquire some additional information about the message of his opponent, then he will find it optimal to play B after receiving any information that increases his belief that his opponent has received a recommendation to play A. Thus, such player would have an incentive to find out from the consulting firm what recommendation it made to his opponent. Moreover, each player may try to influence the consulting firm into sending the recommendation for himself to play B more often. In other words, players may have incentives (i) to acquire additional information from the mediator and (ii) to influence the recommendations of the mediator. We assume that the set of possible messages for each player is $\mathbb{N}$ in order to allow the players to receive sufficiently rich information from the consulting firm if they desire.

As a reduced form representation of the interaction between the players and the consulting firm, we assume that each player chooses an information structure $\phi_i \in \Delta(\mathbb{N} \times \mathbb{N})$ and the consulting firm releases information according to $\beta_1\phi_1 + \beta_2\phi_2$. For example, each player may try to persuade the consulting firm to provide information according to his wishes; according to its unmodelled preferences, the consulting firm favors player i with some fixed probability β_i.

Another interesting interpretation of our joint information design problem, illustrated in the tariff war example, is that each player sends a negotiator to discuss the possible increase of their tariffs. The negotiators are tasked with conveying a particular message to the other player. With probability β_i, the negotiator for player i "wins" and each reports the message as suggested by the winning negotiator back to the players.

After having received his private message from the consulting firm (resp. its negotiator), each player then chooses whether or not to enter the new market (resp. to increase its tariffs). Note that although each player i does not observe ϕ_{-i}, m_{-i} or the realized information structure, each player has the option to choose an information structure such that its support is disjoint from the support of the equilibrium information structure chosen by the other player. We view this as a simple way of giving each player the option to learn when the consulting firm implements the information structure he requests (resp. when its negotiator succeeds). This implies that the players will either learn which information structure was implemented or their actions must be optimal conditional on both information structures.

Our results imply that the set of (Nash or sequential) equilibrium payoffs of the extended game is:

$$\left\{(7,1),(1,7),(\tfrac{7}{2},\tfrac{7}{2}),\beta_1(7,1)+\beta_2(1,7),\beta_1(7,1)+\beta_2(\tfrac{7}{2},\tfrac{7}{2}),\beta_1(\tfrac{7}{2},\tfrac{7}{2})+\beta_2(1,7)\right\}.$$

In particular, $(4\frac{2}{3},4\frac{2}{3})$ is not a sequential equilibrium payoff and the action distribution $\frac{1}{3}1_{(A,A)}+\frac{1}{3}1_{(A,B)}+\frac{1}{3}1_{(B,A)}$ is not the action distribution of a sequential equilibrium of the information design extensive-form game. This payoff profile and action distribution could be obtained with $\phi_1=\phi_2=\frac{1}{3}1_{(1,1)}+\frac{1}{3}1_{(1,2)}+\frac{1}{3}1_{(2,1)}$ and $\pi_i(1,\phi_i)=A$ and $\pi_i(2,\phi_i)=B$ for each i. But then player 1 would gain by deviating to $\phi_1'=1_{(2,1)}$ thereby increasing the probability that his preferred action profile, (B,A), is played.

Thus, the consulting firm cannot act as an effective mediator. The possibility that the consulting firm can be influenced or manipulated by the players implies that the achievable outcomes are only specific convex combinations of Nash outcomes. In particular, each player will optimally choose to bring about his favorite outcome of the ones that may result from the mediation. The consulting firm cannot improve the players' payoffs relative to Nash equilibrium, and, in particular, there is no sequential equilibrium of the extended game where the good outcome (A, A) is achieved without the bad outcome (B, B) also being inadvertently played. Likewise, the negotiators that countries employ cannot act as effective mediators since there is no sequential equilibrium of the extended game where the good outcome (A, A) of low tariffs is achieved without also triggering the bad outcome (B, B) of a trade war with positive probability.[12]

Correlated equilibrium is justified in Aumann (1987) as the result of Bayesian rationality — each player is maximizing his utility given his information. But where does this information come from? If the information is chosen optimally by players who have the ability to privately manipulate the information structure, then it is possible that only a very specific subset of the convex hull of Nash payoffs can be achieved in the chicken game.

4 Related literature and discussion

Many papers have considered whether correlated equilibrium payoffs can be sustained as the outcome of an extended game where players can take "cheap" pre-play actions. For 2-player games, we find that only a very restricted set of outcomes is achievable when each player has the ability to influence and manipulate the information structure in a general way. The distinguishing feature of our model is that we allow each

[12]Experimental evidence on correlated equilibria with an impartial mediator, e.g. Cason and Sharma (2007) and Duffy and Feltovich (2010), suggests that individuals often do not follow the mediator's recommendation and, thus, do not achieve the correlated equilibrium the mediator aims at obtaining. We expect the same problem to occur in experimental designs of our setting. However, the main point of our setting is how individuals chose information designs assuming that players do follow recommendations (more precisely, that each follows some equilibrium strategy upon observing his message). This could be tested by adapting Cason and Sharma's (2007) approach of having robots who always follow recommendations to choose the action profile. Specifically, in such an experiment, (i) some strategy profile describing how each player chooses his action for each possible message would be fixed, (ii) players in the experiment would to choose only an information design and (iii) the action profile would then be determined by the fixed strategy and the realized message profile.

player to choose any information structure he desires, and with some probability the information structure he chooses is the one that actually determines the joint distribution of the messages of all players. This section provides a discussion of these features and how they relate with alternative formalizations in the literature.

Our specification captures certain reasonable features of unmediated interaction that are missing from other models. We emphasise that our point is not that our specification is right or that others are wrong, but simply that these modelling details matter greatly for the question of which payoffs can be supported in equilibrium. Our specification is of interest because it leads to the stark result that very few payoffs can be supported.

4.1 Cheap talk

Several papers have studied the question of whether players can achieve correlated equilibrium outcomes by directly communicating with each other via cheap talk before playing the game.

For 2-player games, Aumann and Hart's (2003) results imply that any payoff in the convex hull of the Nash equilibrium payoffs can be achieved as the outcome of an extended game where players talk for as long as they like before playing the game. In Aumann and Hart (2003), messages are common knowledge so there is no possibility of getting payoffs outside of the convex hull, but cheap talk is enough for players to reach any outcome achievable using publicly observed lotteries. On the other hand, in our model, there are privately observed lotteries but nevertheless players can only get payoffs in $\mathrm{co}(u(N(G)))$ and not even all of those (even if we were to vary β). We discuss this connection further in the next subsection.

Other papers (e.g. Bárány (1992), Ben-Porath (1998) or Gerardi (2004)) attempt to achieve the entire set of correlated equilibrium payoffs via cheap talk, which requires more than two players.[13] Although our focus is primarily on 2-player games, we discuss the extension of our results to more than two players in Section 4.7, which

[13] See Forges (2020), Section 4, for a survey of these papers.

implies that, in general, payoffs from private information design are a strict subset of the set of correlated equilibrium payoffs.

Beyond cheap talk, Ben-Porath (1998) attempts to achieve correlated equilibrium payoffs by relaxing the cheap talk assumption, i.e. by introducing additional communication protocols. This will be discussed in Section 4.2.

4.1.1 Public cheap talk

Consider a (one-shot) cheap talk extension of a 2-player game G where each player $i \in N = \{1, 2\}$ chooses a message $m^i \in \mathbb{N}$ (possibly at random). Each player $i \in N$ then observes $m = (m^1, m^2)$ and chooses a mixed action $\pi_i(m) \in \Delta(A_i)$. This corresponds to a special case of Aumann and Hart's (2003) model with complete information and communication restricted to one period only. Nevertheless, in this model, every rational convex combination of Nash payoffs can be achieved using jointly controlled lotteries, in contrast to our model where only a few specific convex combination of Nash payoffs can be achieved. In Section 4.3, we show that this special case can be captured in our framework using an alternative aggregation function, i.e. by specifying the appropriate mapping $(\phi_1, \phi_2) \mapsto \phi$ that determines the true information structure from the choices of the players.

The key difference between our specification and the setting of Aumann and Hart (2003) is that each player in the latter is sure that his opponent receives the message he sends, he knows what this message is, and his opponent cannot do anything to influence this message. On the other hand, according to our specification, there is always a possibility that each player gets to determine the messages of both players. For example, if player 2 benefits from player 1 sending some message m^1, then player 2 may want to take certain (unmodelled) actions that increase the likelihood that player 1 will send message m^1.

4.1.2 Cheap talk with private messages

There is a variation of the cheap talk framework of Section 4.1.1 that is related to our model. This happens when each player i is chosen to be a sender with probability

β_i. When i is the sender, he chooses a message for the other player (possibly at random) who then privately observes the message, i.e. player i chooses $\phi \in \Delta(\mathbb{N})$; player j then privately observes the realization of ϕ, while player i observes nothing (for convenience, we represent this as player i receiving a dummy message $\emptyset$).

We can model this one-shot cheap talk extension of a 2-player game G with *private messages* in our framework as follows. Each player $i \in N = \{1, 2\}$ chooses a probability distribution over the empty message for himself and a message in $\mathbb{N}$ for the other player, i.e. player 1 chooses $\phi_1 \in \Delta(\{\emptyset\} \times \mathbb{N})$ and player 2 chooses $\phi_2 \in \Delta(\mathbb{N} \times \{\emptyset\})$. Then nature chooses a message profile $m \in (\mathbb{N} \cup \{\emptyset\})^2$ with probability $\phi[m] = \sum_{i=1}^{2} \beta_i \phi_i[m]$.[14] Each player $i \in N$ observes ϕ_i and m_i, and then chooses a mixed action $\pi_i(m_i, \phi_i) \in \Delta(A_i)$. Given (ϕ_1, ϕ_2, m), each player $i \in N$ receives a payoff of $u_i(\pi_1(\phi_1, m_1), \pi_2(\phi_2, m_2))$.

If player i is the sender and $j \neq i$ is the receiver, then, for some $m_j \in \operatorname{supp}(\phi_{i,\mathbb{N}})$, $(\pi_i(\emptyset), \pi_j(m_j))$ is played. Thus, for each $i \in N$ and $m_j \in \operatorname{supp}(\phi_{i,\mathbb{N}})$, $(\pi_i(\emptyset), \pi_j(m_j))$ is a Nash equilibrium. In particular, $(\pi_i(\emptyset), \sum_{m_j} \phi_{i,\mathbb{N}}[m_j]\pi_j(m_j)))$ is a Nash equilibrium, and the set of (Nash or sequential) equilibrium payoffs of the simultaneous, one-shot cheap talk extension of the 2-player game G with private messages is $\{\beta_1 u(\sigma) + \beta_2 u(\sigma') : \sigma, \sigma' \in N(G)\}$.

4.1.3 Cheap talk with random sender

Another variation is when each player i chooses a message $m^i \in \mathbb{N}$ (possibly at random) and, with probability β_i, the message m^i is publicly observed by both players, i.e. player i becomes the sender with probability β_i. This variation is exactly as in Section 4.1.1 except the publicly observed message profile is (m^i, m^i) with probability β_i instead of (m^1, m^2) with probability 1. It is also closely related to the variation in 4.1.2, the difference being that the message is publicly rather than privately observed.

This setting is, in fact, the same as our model with the additional restriction that players must choose information designs supported on $\{(m_1, m_2) \in \mathbb{N}^2 : m_1 = m_2\}$.

[14] Note that the support of ϕ is contained in $(\{\emptyset\} \times \mathbb{N}) \cup (\mathbb{N} \times \{\emptyset\})$ and, for each $k \in \mathbb{N}$, $\phi[\emptyset, k] = \beta_1 \phi_1[\emptyset, k]$ and $\phi[k, \emptyset] = \beta_2 \phi_2[k, \emptyset]$.

Nevertheless, the set of (Nash or sequential) equilibrium payoffs of this setting is the same set $\mathcal{U}$ that features in Corollary 2.

We emphasize that the equivalence in payoffs between this setting and ours is a result that follows from Corollary 2, which implies that publicly observed lotteries are sufficient to achieve all the payoffs in our model (in the two player case). In our model, the players have access to privately observed lotteries but cannot be incentivized to choose them because of the possibility of manipulation.

4.2 Communication protocols and manipulability

Beyond Aumann and Hart (2003), the literature has focused on whether players can communicate in a more sophisticated manner to achieve correlated equilibrium payoffs. For instance, Ben-Porath (1998) shows that each correlated equilibrium can be approximated by the action distribution of a sequential equilibrium in a specific information design extensive-form game that includes the possibility of credibly revealing messages and (in the case of two players) ball and urns.[15] However, the specification of such extensive form games rules out the possibility of certain manipulations by assumption.

For example, consider once again the chicken game from Section 3. In Ben-Porath (1998), the correlated equilibrium $\phi = \frac{1}{3}1_{(A,A)} + \frac{1}{3}1_{(A,B)} + \frac{1}{3}1_{(B,A)}$ is close to the action distribution of a sequential equilibrium of his information design extensive-form, which works as follows. Player 2 lets player 1 choose a ball from an urn U_1 and the ball which player 1 draws from U_1 determines an action a_1; the induced distribution of player 1's action is $\phi_{A_1} = \frac{2}{3}1_A + \frac{1}{3}1_B$. Player 1 then gives player 2 an urn $U_2(a_1)$ inducing on A_2 the distribution $\frac{1}{2}1_A + \frac{1}{2}1_B$ if $a_1 = A$ and 1_A if $a_1 = B$. After this has occurred, there is a sufficiently high probability that the contents of the urns U_1 and $U_2(a_1)$ are revealed as well as the ball that was chosen by player 1 from U_1. Our point is that there is a (unmodelled) possibility of manipulation by one player in this extensive-form. Specifically, player 1 can send urn $U_2(B)$ to player 2

[15] Gerardi (2004) obtains a stronger result for games with at least five players. See also Urbano and Vila (2002) for 2-player games where players are boundedly rational.

when the ball he draws from U_1 indicates that he should play A; in addition, he could take all the balls from U_1 and put them all inside again, except one ball indicating that he should play B. In this way, he obtains a payoff of 7 instead of $4\frac{2}{3}$.

We share Ben-Porath's (1998) motivation that a reliable mediator who is immune to manipulation by the players is not always available. However, players may wish to manipulate the information structure regardless of whether it is the result of some procedure designed by the players themselves or if it comes from a mediator. As the previous example shows, if players are able to manipulate the communication protocol, then they will do so as well, i.e. Ben-Porath's (1998) results require that certain manipulations are ruled out by assumption.

4.3 Aggregation of information designs

In our model, the true information structure is a convex combination of the ones chosen by the players. A more general way of combining the two information designs is to postulate an abstract aggregation function $\alpha : S^2 \to S$ such that if player 1 chooses information structure $\phi_1 \in S$ and player 2 chooses information structure $\phi_2 \in S$, then the realized information structure is $\alpha(\phi_1, \phi_2) \in S$. Alternative formulations of information design in a setting without an explicit designer can then be obtained by specifying alternative aggregation functions α.

One such alternative is for each player $i \in \{1, 2\}$ to choose $\phi_i \in S$ and then assume that each i receives two messages m_i^1 and m_i^2, where $m^1 = (m_1^1, m_2^1)$ and $m^2 = (m_1^2, m_2^2)$ are independently drawn from ϕ_1 and ϕ_2 respectively. We note that this formulation can be embedded in our framework under an alternative aggregation function α. Indeed, let $\psi : \mathbb{N}^2 \times \mathbb{N}^2 \to \mathbb{N}^2$ be a bijection and let $\alpha(\phi_1, \phi_2) = (\phi_1 \times \phi_2) \circ \psi^{-1}$. If we additionally impose the restriction that players are only allowed to choose ϕ_i such that $\phi_i[m_1^i, m_2^i] = 0$ whenever $m_1^i \neq m_2^i$, then this formulation captures exactly the model of Aumann and Hart (2003) when cheap talk is restricted to take place over a single period.

In particular, for each $i \in \{1, 2\}$, let $m_i = (m_i^1, m_i^2)$ and restrict each player's choice of ϕ_i to distributions over $m^i = (m_1^i, m_2^i)$ such that $\phi_i[m_1^i, m_2^i] = 0$ whenever

$m_1^i \neq m_2^i$. Recall that m^1 and m^2 are independently drawn from ϕ_1 and ϕ_2 respectively. Thus, the message m_i^i just reveals to i the message he sends to j, and m^i can be identified with the message sent by player i in Aumann and Hart (2003). According to this formulation, note that player i has no ability to influence the distribution of m^{-i}. On the other hand, our specification of α does not restrict players from learning about or influencing any aspect of the information structure.

4.4 Privacy

An alternative to our assumption that the information design choices are made privately is to assume that information design is public so that, for example, each player observes the information structure $\beta_1\phi_1 + \beta_2\phi_2$ chosen by nature. To model this, we can let, for each player i, i's action be a function of the message that he receives, his own information design and the information structure $\beta_1\phi_1 + \beta_2\phi_2$, i.e. $\pi_i : M_i \times S \times S \to \Delta(\{A, B\})$. Under this assumption, the payoff $(4\frac{2}{3}, 4\frac{2}{3})$ can be achieved by specifying that $\phi_1 = \phi_2 = \phi^* = \frac{1}{3}1_{(A,B)} + \frac{1}{3}1_{(B,A)} + \frac{1}{3}1_{(A,A)}$, $\pi_i(m_i, \phi_i, \phi^*) = m_i$ for each $\phi_i \in S$ and $m_i \in \text{supp}(\phi^*_{M_i})$; and $\pi_i(m_i, \phi_i, \hat{\phi}) = \frac{1}{2}1_A + \frac{1}{2}1_B$ for each $\phi_i \in S$, $\hat{\phi} \neq \phi^*$ and $m_i \in \text{supp}(\hat{\phi}_{M_i})$. Intuitively, deviations from ϕ^* can be deterred by the threat of reverting to the mixed strategy Nash equilibrium whenever some alternative information structure is realized. The reason we assume that information design is private is because we are interested in how the ability to manipulate the information structure affects the outcomes of the game. When the choice of information is observed, certain information structures can be sustained by the threat of punishment. Our aim is instead to ask which outcomes can arise abstracting away from the possibility of such threats.

4.5 Limits of perfect conditional ε-equilibria

Theorem 1 implies, in particular, that there exists a strategy of G_{id} which is a perfect conditional ε-equilibrium for each $\varepsilon > 0$ for any 2-player game G, i.e. there exists a sequential equilibrium as we have defined it. In finite games, π is a sequential equilibrium if and only if it is a perfect conditional ε-equilibrium for each $\varepsilon > 0$; thus, our definition is the natural extension to infinite games. However, as argued by Myerson and Reny (2020), a drawback of this definition is that a sequential equilibrium may not exist in general. To circumvent this non-existence issue, they focus instead on distributions of outcomes and payoffs that arise as limits of perfect conditional ε-equilibria.

In our setting, and focusing on payoffs, the set of payoffs achievable with limits of perfect conditional ε-equilibria is

$$\begin{aligned} U^{\text{limit}}(G) = \{u \in \mathbb{R}^2 : u = \lim_L u(\pi^L) \text{ for some } \{\pi^L\}_{L=1}^{\infty} \text{ such that, for each } \varepsilon > 0, \\ \text{there exists } \bar{L} \in \mathbb{N} \text{ such that } \pi^L \in \Pi \text{ is a perfect conditional} \\ \varepsilon\text{-equilibrium of } G_{id} \text{ for each } L \geq \bar{L}\}. \end{aligned}$$

We show in the supplementary material to this paper that, for each 2-player game G, $U^{\text{limit}}(G) = \mathcal{U}$. Thus, no change results to the set of equilibrium payoffs from weakening the equilibrium concept from sequential equilibrium as we have defined it (i.e. perfect conditional ε-equilibrium for each $\varepsilon > 0$) to limits of perfect conditional ε-equilibrium as $\varepsilon \to 0$.[16]

4.6 Mixed information designs

We have focused so far in the case where players are not allowed to mix in their choice of an information design. As we argue in this section, allowing for mixed information designs does not significantly change our results.[17]

We focus on Corollary 3 and let

$$U^*(G) = \{u(\pi) : \pi \in \Pi^* \text{ is a sequential equilibrium of } G_{id}\},$$

where, recall, Π^* is the set of mixed strategies of G_{id}. We then have that, for each

[16] The limit notion is weaker because limits of perfect conditional ε-equilibrium strategies may not be a strategy; thus, there may be no sequential equilibrium (as we have defined it) that achieves the limit payoff.

[17] See the supplementary material to this paper for the details for this section.

2-player game $G \in \mathcal{G}$,

$$U(G) \subseteq U^*(G) \subseteq \{\beta_1 u(\sigma) + \beta_2 u(\sigma') : \sigma, \sigma' \in N(G)\}.$$

Thus, in generic 2-player games, sequential equilibrium payoffs of G_{id} when mixed information designs are allowed continue to be specific convex combinations of two Nash equilibrium payoffs.

In games with a unique equilibrium, it then follows that $U(G) = U^*(G) = u(N(G))$. If G is a 2-player game that has more than one Nash equilibrium, then mixed information designs can expand the set of equilibrium payoffs. We illustrate this claim in the battle of the sexes by showing in the supplementary material to this paper that $\beta_1 u(B, B) + \beta_2 u(A, A) \in U^*(G) \setminus U(G)$. The reason why this payoff profile does not belong to $U(G)$ is that e.g. player 2 is obtaining a lower payoff if his information design is chosen than if player's 1 design is chosen. Thus, player 2 could deviate in her information design by sending a message to player 1 from player 1's design which triggers player 1 to choose B. With mixed designs, player 1 can avoid this profitable deviation by player 2 by obfuscating her message, for instance by uniformly randomizing over L information designs, $\phi_1^1, \ldots, \phi_1^L$, with ϕ_1^l sending message l to herself and choosing B if and only if she receives message l and had chosen ϕ_1^l. In this way, if player 2 sends message $l \in \{1, \ldots, L\}$ to player 1, this will trigger B only with probability $1/L$.

Thus, the possibility of mixing in the first period allows additional payoffs to be sustained, but since, for $G \in \mathcal{G}$,

$$U^*(G) \subseteq \{\beta_1 u(\sigma) + \beta_2 u(\sigma') : \sigma, \sigma' \in N(G)\},$$

our main conclusion that only a small subset of correlated equilibrium payoffs can be achieved with private information design continues to holds. It is natural to ask whether also

$$\{\beta_1 u(\sigma) + \beta_2 u(\sigma') : \sigma, \sigma' \in N(G)\} \subseteq U^*(G),$$

i.e. whether the two sets are in fact equal. This is false in general, e.g. when σ involves a weakly dominated action for player 1. Nevertheless, the argument in the previous paragraph can be used to show that each player can only send a message to his opponent that triggers some favorable action for himself with a probability that is essentially equal to zero in the limit as $L \to \infty$. Thus, considering instead the weaker notion of limits of perfect conditional ε-equilibrium payoffs (as discussed in Section 4.5), i.e.

$$\begin{aligned} U^{\text{limit}*}(G) = \{ & u \in \mathbb{R}^2 : u = \lim_L u(\pi^L) \text{ for some } \{\pi^L\}_{L=1}^{\infty} \text{ such that, for each } \varepsilon > 0, \\ & \text{there exists } \bar{L} \in \mathbb{N} \text{ such that } \pi^L \in \Pi^* \text{ is a perfect conditional} \\ & \varepsilon\text{-equilibrium of } G_{id} \text{ for each } L \geq \bar{L}\} \end{aligned}$$

we obtain that, for each 2-player game G,

$$\{\beta_1 u(\sigma) + \beta_2 u(\sigma') : \sigma, \sigma' \in N(G)\} \subseteq U^{\text{limit}*}(G).$$

4.7 More than two players

The extension of our setting to the case of more than two players is straightforward.[18] Theorem 1 extends, with condition (2) in part 3 stating that, for each $i \in N$ and $m_i \in \cup_{j \in N \setminus \{i\}} \text{supp}(\phi^*_{j,M_i})$, $\pi_i(m_i)$ solves

$$\max_{\alpha_i \in \Delta(A_i)} \sum_{m_{-i}} \frac{\sum_{j \in N \setminus \{i\}} \beta_j \phi^*_j[m_i, m_{-i}]}{\sum_{j \in N \setminus \{i\}} \beta_j \phi^*_{j,M_i}[m_i]} u_i(\alpha_i, \pi_{-i}(m_{-i})).$$

Corollaries 1, 2 and 3 do not extend. To see this, consider the following game, Example 2.5 in Aumann (1974), where player 1 chooses the row, player 2 chooses the column, and player 3 chooses the matrix ($A_3 = \{L, M, R\}$):

$1\backslash 2$	A	B
A	$0,0,3$	$0,0,0$
B	$1,0,0$	$0,0,0$

$1\backslash 2$	A	B
A	$2,2,2$	$0,0,0$
B	$0,0,0$	$2,2,2$

$1\backslash 2$	A	B
A	$0,0,0$	$0,0,0$
B	$0,1,0$	$0,0,3$

If $\min\{2\beta_1, 2\beta_2\} \geq \beta_3$, then $(1 - \beta_3)(2,2,2) + \beta_3(0,0,3)$ is a sequential equilibrium payoff. This payoff can be obtained by setting $\phi^*_1 = \phi^*_2 = \frac{1}{2}1_{(m'_1, m'_2, \hat{m}_3)} + \frac{1}{2}1_{(m''_1, m''_2, \hat{m}_3)}$,

[18] See the supplementary material to this paper for the details for this section.

$\phi_3^* = \frac{1}{2}1_{(m_1',m_2',\hat{m}_3')} + \frac{1}{2}1_{(m_1'',m_2'',\hat{m}_3'')}$ and

$$\pi_1(m_1') = A, \pi_1(m_1'') = B,$$
$$\pi_2(m_2') = A, \pi_2(m_2'') = B,$$
$$\pi_3(\hat{m}_3) = M, \pi_3(\hat{m}_3') = L \text{ and } \pi_3(\hat{m}_3'') = R.$$[19]

Thus, even though $u_i \leq 1$ for each $u \in u(N(G))$. Thus, correlation of players' actions through private information design can still significantly improve the payoff to everybody relative to Nash equilibrium payoffs.

Nevertheless, it is also clear that, with private information design, not all correlated equilibrium payoffs can be achieved. For example, consider $(2,2,2) \in u(C(G))$; if $(2,2,2) \in U(G)$, then, for some sequential equilibrium $\pi \in \Pi$,

$$(2,2,2) = \sum_{m \in \text{supp}(\phi^*)} \phi^*[m]u(\pi(m))$$

and, thus, $\pi(m) = (A,A,M)$ or $\pi(m) = (B,B,M)$ for each $m \in \text{supp}(\phi^*)$. But then, for each $m \in \text{supp}(\phi_3^*)$, $\pi_3(m_3)$ is not a best-reply against $\pi_{-3}(m_{-3})$, contradicting (the extension of) Theorem 1.[20]

References

Aliprantis, C., and K. Border (2006): *Infinite Dimensional Analysis*. Springer, Berlin, 3rd edn.

Aumann, R. (1974): "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, 1, 67–96.

——— (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1–18.

[19] Note that $\pi(m_1', m_2', \hat{m}_3) = (A,A,M)$, $\pi(m_1'', m_2'', \hat{m}_3) = (B,B,M)$, $\pi(m_1', m_2', \hat{m}_3') = (A,A,L)$ and $\pi(m_1'', m_2'', \hat{m}_3'') = (B,B,R)$.

[20] Note that $(2,2,2) \in u(C(G))$ cannot be approximated by $u \in U(G)$. Indeed, to get close to $(2,2,2)$, ϕ^* must put small probability on m such that $\pi(m) \notin \{(A,A,M),(B,B,M)\}$. Thus, ϕ_3^* must also put small probability on such m. But then there exists $m' \in \text{supp}(\phi_3^*)$ such that $\pi(m') \in \{(A,A,M),(B,B,M)\}$, which contradicts (the extension of) Theorem 1.

AUMANN, R., AND S. HART (2003): "Long Cheap Talk," *Econometrica*, 71, 1619–1660.

BÁRÁNY, I. (1992): "Fair Distribution Protocols or How the Players Replace Fortune," *Mathematics of Operations Research*, 17, 327–340.

BEN-PORATH, E. (1998): "Correlation without Mediation: Expanding the Set of Equilibrium Outcomes by "Cheap" Pre-play Procedures," *Journal of Economic Theory*, 80, 108–122.

——— (2003): "Cheap Talk in Games with Incomplete Information," *Journal of Economic Theory*, 108, 45–71.

BLUME, A. (2024): "Language Games: Correlation through Non-Understanding, Dialogue, Inarticulateness, and Misunderstanding," University of Arizona.

CARMONA, G., AND K. LAOHAKUNAKORN (2023): "The Folk Theorem for the Prisoner's Dilemma with Endogenous Private Monitoring," *Journal of Economic Theory*, 213, 105731.

CASON, T., AND T. SHARMA (2007): "Recommended Play and Correlated Equilibria: An Experimental Study," *Economic Theory*, 22, 11–27.

DUFFY, J., AND N. FELTOVICH (2010): "Correlated Equilibria, Good and Bad: An Experimental Study," *International Economic Review*, 51, 701–721.

FORGES, F. (2020): "Correlated Equilibria and Communication in Games," in *Complex Social and Behavioral Systems*, ed. by M. Sotomayor, D. Pérez-Castrillo, and F. Castiglione. Springer.

GERARDI, D. (2004): "Unmediated Communication in Games with Complete and Incomplete Information," *Journal of Economic Theory*, 114, 104–131.

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

Krishna, V. (2007): "Communication in Games of Incomplete Information: Two Players," *Journal of Economic Theory*, 132, 584–592.

Marshall, R. C., and L. M. Marx (2012): *The Economics of Collusion: Cartels and Bidding Rings*. The MIT Press.

Myerson, R. (1982): "Optimal coordination mechanisms in generalized principal–agent problems," *Journal of Mathematical Economics*, 10, 67–81.

Myerson, R., and P. Reny (2020): "Perfect Conditional ε-Equilibria of Multi-Stage Games with Infinite Sets of Signals and Actions," *Econometrica*, 88, 495–531.

Neyman, A. (1997): "Correlated Equilibrium and Potential Games," *International Journal of Game Theory*, 26, 223–227.

Rivera, T. (2018): "Incentives and the Structure of Communication," *Journal of Economic Theory*, 175, 201–247.

Urbano, A., and J. Vila (2002): "Computational Complexity and Communication: Coordination in Two-Player Games," *Econometrica*, 70, 1893–1927.

van Damme, E. (1991): *Perfection and Stability of Nash equilibrium*. Springer Verlag, Berlin.

Wagner, A. (2011): "Unmediated Communication with Partially Verifiable Types," *Journal of Mathematical Economics*, 47, 99–107.

Wheeden, R., and A. Zygmund (1977): *Measure and Integral*. Marcel Dekker, New York.

A Appendix

A.1 Definition of perfect conditional ε-equilibrium

A sequential equilibrium $\pi \in \Pi^*$ is, by definition, a perfect conditional ε-equilibrium for each $\varepsilon > 0$. For each $\varepsilon > 0$, $\pi \in \Pi^*$ is a *perfect conditional ε-equilibrium* if there exits a net $\{\pi^\alpha, p^\alpha\}_\alpha$ such that the following properties hold. The first five require that $\{\pi^\alpha\}_\alpha$ is a net of strategies converging to π that assign strictly positive probability to each action and information design beyond a certain order, and that $\{p^\alpha\}_\alpha$ is a net of nature's choices regarding the probability distribution of message profiles for each profile of information designs (ϕ_1, ϕ_2) that converges to $\beta_1\phi_1 + \beta_2\phi_2$ and assigns strictly positive probability to each message profile beyond a certain order:

(i) For each α, π^α is a strategy and $p^\alpha : S^2 \to \Delta(M)$ is measurable,

(ii) For each $i \in N$, $\sup_{B\in\mathcal{B}(S)} |\pi_i^{1,\alpha}[B] - \pi_i^1[B]| \to 0$ and

$$\sup_{(m_i,\phi_i)\in M_i\times S, a_i\in A_i} |\pi_i^{2,\alpha}(m_i,\phi_i)[a_i] - \pi_i(m_i,\phi_i)[a_i]| \to 0,$$ [21]

(iii) For each $i \in N$, $m_i \in M_i$, $\phi_i \in S$ and $a_i \in A_i$, there is $\bar{\alpha}$ such that $\pi_i^{1,\alpha}[\phi_i] > 0$ and $\pi_i^{2,\alpha}(m_i,\phi_i)[a_i] > 0$ for each $\alpha \geq \bar{\alpha}$,

(iv) $\sup_{\phi\in S^2, B\subseteq M} |p^\alpha(\phi)[B] - \sum_{i\in N}\beta_i\phi_i[B]| \to 0$, and

(v) For each $\phi \in S^2$ and $m \in M$, there is $\bar{\alpha}$ such that $p^\alpha(\phi)[m] > 0$ for each $\alpha \geq \bar{\alpha}$.

A final condition requires that, for each α, π^α is such that the payoff that each player obtains by following it at each information set which is reached with strictly positive probability is within ε of his maximum payoff conditional on that information set:

(vi) for each α and $i, j \in N$, with $j \neq i$,

[21] We let $\mathcal{B}(S)$ denote the class of Borel measurable subsets of S and, for each $\phi \in S$, 1_ϕ denote the probability measure on S degenerate at ϕ.

(a) For each $\phi_i' \in S$,

$$\sum_{\phi \in \text{supp}(\pi^{1,\alpha})} \pi^{1,\alpha}[\phi] \left(\sum_m p^\alpha(\phi)[m] u_i(\pi^{2,\alpha}(m,\phi)) \right) \geq$$

$$\sum_{\phi \in \text{supp}(1_{\phi_i'} \times \pi_j^{1,\alpha})} (1_{\phi_i'} \times \pi_j^{1,\alpha})[\phi] \left(\sum_m p^\alpha(\phi)[m] u_i(\pi^{2,\alpha}(m,\phi)) \right) - \varepsilon,$$

where $\pi^{1,\alpha} = \pi_1^{1,\alpha} \times \pi_2^{1,\alpha}$, and

(b) For each $i \in N$, $(m_i, \phi_i) \in M_i \times S$ such that

$$\pi_i^{1,\alpha}[\phi_i] \sum_{\phi_j \in \text{supp}(\pi_j^{1,\alpha})} \pi_j^{1,\alpha}[\phi_j] p_{M_i}^\alpha(\phi_i, \phi_j)[m_i] > 0$$

and $a_i \in A_i$,

$$\frac{\sum_{\phi_j \in \text{supp}(\pi_j^{1,\alpha})} \pi_j^{1,\alpha}[\phi_j] \left(\sum_{m_j} p^\alpha(\phi_i,\phi_j)[m_i,m_j] u_i(\pi^{2,\alpha}(m,\phi)) \right)}{\sum_{\phi_j \in \text{supp}(\pi_j^{1,\alpha})} \pi_j^{1,\alpha}[\phi_j] p_{M_i}^\alpha(\phi_i,\phi_j)[m_i]} \geq$$

$$\frac{\sum_{\phi_j \in \text{supp}(\pi_j^{1,\alpha})} \pi_j^{1,\alpha}[\phi_j] \left(\sum_{m_j} p^\alpha(\phi_i,\phi_j)[m_i,m_j] u_i(a_i, \pi_j^{2,\alpha}(m_j,\phi_j)) \right)}{\sum_{\phi_j \in \text{supp}(\pi_j^{1,\alpha})} \pi_j^{1,\alpha}[\phi_j] p_{M_i}^\alpha(\phi_i,\phi_j)[m_i]} - \varepsilon.$$

A.2 Proof of Theorem 1

Every sequential equilibrium is a Nash equilibrium, hence condition 2 implies condition 1. Thus, it suffices to show that condition 1 implies condition 3 and that condition 3 implies condition 2.

A.2.1 Proof that condition 1 implies condition 3

Let $\pi \in \Pi$ be a Nash equilibrium of G_{id}. Then

$$\sum_m \phi^*[m] u_i(\pi(m)) \geq \sum_m (\phi_i', \phi_j^*)[m] u_i(\pi_i'(m_i, \phi_i'), \pi_j(m_j)), \tag{3}$$

for each $i, j \in N$, $j \neq i$, $\phi_i' \in S$ and $\pi_i' : M_i \times S \to \Delta(A_i)$, where $(\phi_i', \phi_j^*) = \beta_i \phi_i' + \beta_j \phi_j^*$. It follows from (3) that

$$\sum_{m_j} \frac{\phi^*[m]}{\phi_{M_i}^*[m_i]} u_i(\pi(m)) \geq \sum_{m_j} \frac{\phi^*[m]}{\phi_{M_i}^*[m_i]} u_i(a_i, \pi_j(m_j)) \tag{4}$$

for each $i, j \in N$, $j \neq i$, $m_i \in \text{supp}(\phi^*_{M_i})$ and $a_i \in A_i$.

In each Nash equilibrium of G_{id}, any player $i \in N$ must send optimal messages m in the sense that they induce an action profile $\pi(m)$ that maximizes i's payoff function. This is stated in Lemma 1 which is a preliminary result for condition (1).

Lemma 1. *For each $i \in N$, $\text{supp}(\phi^*_i) \subseteq \{m \in M : u_i(\pi(m)) = \sup_{m' \in M} u_i(\pi(m'))\}$.*

Proof. Suppose not; then there is $i \in N$, $m' \in \text{supp}(\phi^*_i)$ and $m^* \in M$ such that $u_i(\pi(m^*)) > u_i(\pi(m'))$. Define ϕ'_i by setting, for each $m \in \text{supp}(\phi^*_i)$,

$$\phi'_i[m] = \begin{cases} 0 & \text{if } m = m', \\ \phi^*_i[m^*] + \phi^*_i[m'] & \text{if } m = m^*, \\ \phi^*_i[m] & \text{otherwise,} \end{cases}$$

and let $\pi'_i : M_i \times S \to \Delta(A_i)$ be such that $\pi'_i(m_i, \phi'_i) = \pi_i(m_i, \phi^*_i)$ for each $m_i \in M_i$. Then

$$\begin{aligned} &\sum_m (\phi'_i, \phi^*_j)[m] u_i(\pi'_i(m_i, \phi'_i), \pi_j(m_j)) - \sum_m \phi^*[m] u_i(\pi(m)) \\ &= \sum_m (\phi'_i, \phi^*_j)[m] u_i(\pi(m)) - \sum_m \phi^*[m] u_i(\pi(m)) \\ &= \sum_m \beta_i \left(\phi'_i[m] - \phi^*_i[m]\right) u_i(\pi(m)) \\ &= \beta_i \phi^*_i[m'] \Big(u_i(\pi(m^*)) - u_i(\pi(m')) \Big) > 0. \end{aligned}$$

But this contradicts (3). ■

The conclusion of Lemma 1 can be strengthened: for a message m to be optimal, $u_i(\pi(m))$ must achieve $\max_{m'_j} v_i(\pi_j(m'_j))$ and, thus, $\pi_i(m_i)$ be a best-reply to $\pi_j(m_j)$.

Lemma 2. *For each $i, j \in N$ with $i \neq j$,*

$$\text{supp}(\phi^*_i) \subseteq \{m \in M : \; v_i(\pi_j(m_j)) = \sup_{m'_j \in M_j} v_i(\pi_j(m'_j)) \textit{ and } \pi_i(m_i) \in BR_i(\pi_j(m_j))\}.$$

Proof. Suppose not; then there is $i \in N$, $j \neq i$, $m' \in \text{supp}(\phi^*_i)$ and $m^* \in M$ such that (i) $v_i(\pi_j(m^*_j)) > v_i(\pi_j(m'_j))$ or (ii) $v_i(\pi_j(m'_j)) = \sup_{\hat{m}_j \in M_j} v_i(\pi_j(\hat{m}_j))$ and $\pi_i(m'_i) \notin BR_i(\pi_j(m'_j))$; in case (ii), let $m^* = m'$. Let $a^*_i \in BR_i(\pi_j(m^*_j))$, $\bar{m}_i \notin$

$\text{supp}(\phi^*_{M_i})$, $\phi'_i = 1_{(\bar{m}_i, m^*_j)}$ and $\pi'_i : M_i \times S \to \Delta(A_i)$ be such that $\pi'_i(\bar{m}_i, \phi'_i) = a^*_i$ and $\pi'_i(m_i, \phi'_i) = \pi_i(m_i, \phi^*_i)$ for each $m_i \neq \bar{m}_i$. Then

$$\begin{aligned}
&\sum_m (\phi'_i, \phi^*_j)[m] u_i(\pi'_i(m_i, \phi'_i), \pi_j(m_j)) - \sum_m \phi^*[m] u_i(\pi(m)) \\
&= \sum_m \beta_i \phi'_i[m] u_i(\pi'_i(m_i, \phi'_i), \pi_j(m_j)) - \sum_m \beta_i \phi^*_i[m] u_i(\pi(m)) \\
&= \beta_i \Big(u_i(a^*_i, \pi_j(m^*_j)) - \sum_{m \in \text{supp}(\phi^*_i)} \phi^*_i[m] u_i(\pi(m)) \Big) \\
&= \beta_i \Big(v_i(\pi_j(m^*_j)) - u_i(\pi(m')) \Big)
\end{aligned}$$

because $u_i(\pi(m)) = u_i(\pi(m'))$ for each $m \in \text{supp}(\phi^*_i)$ by Lemma 1 as $m' \in \text{supp}(\phi^*_i)$. Thus, if $v_i(\pi_j(m^*_j)) > v_i(\pi_j(m'_j))$, then

$$v_i(\pi_j(m^*_j)) - u_i(\pi(m')) \geq v_i(\pi_j(m^*_j)) - v_i(\pi_j(m'_j)) > 0;$$

if $v_i(\pi_j(m^*_j)) = v_i(\pi_j(m'_j))$, then $\pi_i(m'_i) \notin BR_i(\pi_j(m'_j))$ and

$$v_i(\pi_j(m^*_j)) - u_i(\pi(m')) > v_i(\pi_j(m^*_j)) - v_i(\pi_j(m'_j)) \geq 0.$$

It then follows that $\sum_m (\phi'_i, \phi^*_j)[m] u_i(\pi'_i(m_i, \phi'_i), \pi_j(m_j)) - \sum_m \phi^*[m] u_i(\pi(m)) > 0$ in either case. But this contradicts (3). ∎

Lemma 2 implies that $\pi_i(m_i)$ is a best-reply against $\pi_j(m_j)$ whenever $m \in \text{supp}(\phi^*_i)$ and $i, j \in N$ with $i \neq j$. We will now show that if, in addition, $m_i \in \text{supp}(\phi^*_{j,M_i})$, then $\pi_i(m_i)$ solves

$$\max_{\alpha_i \in \Delta(A_i)} \sum_{m_j} \frac{\phi^*_j[m_i, m_j]}{\phi^*_{j,M_i}[m_i]} u_i(\alpha_i, \pi_j(m_j)).$$

Thus, whenever $m_i \in \text{supp}(\phi^*_{i,M_i}) \cap \text{supp}(\phi^*_{j,M_i})$, $\pi_i(m_i)$ solves player i's expected payoff conditional on his information design ϕ^*_i being chosen and also conditional on it not being chosen. The reason for this is that player i can always differentiate the messages he receives from himself from those that he receives from the other players: if $m \in \text{supp}(\phi^*_i)$ is such that $\pi_i(m_i)$ does not maximize i's expected payoff conditional on his information design ϕ^*_i not being chosen, then player i would gain by deviating from ϕ^*_i by simply sending a message $(\bar{m}_i, m_j)$ with probability one for some $\bar{m}_i \notin \text{supp}(\phi^*_{M_i})$. If he receives message m_i, then he can be sure that his information design has not been chosen and can choose a solution to that problem in response to m_i; if he receives message $\bar{m}_i$, then the can be sure that his information design has been chosen and choose $\pi_i(m_i)$, which is a best-reply against m_j, in response to $\bar{m}_i$.

Lemma 3. *For each $i, j \in N$ with $i \neq j$,*

$$\mathrm{supp}(\phi_i^*) \subseteq \Big\{ m \in M : m_i \notin \mathrm{supp}(\phi_{j,M_i}^*) \text{ or } \pi_i(m_i) \text{ solves} \\ \max_{\alpha_i \in \Delta(A_i)} \sum_{m_j} \frac{\phi_j^*[m_i, m_j]}{\phi_{j,M_i}^*[m_i]} u_i(\alpha_i, \pi_j(m_j)) \Big\}.$$

Proof. Suppose not; then there is $i \in N$ and $m' \in \mathrm{supp}(\phi_i^*)$ such that $m_i' \in \mathrm{supp}(\phi_{j,M_i}^*)$, $j \neq i$, and $\pi_i(m_i')$ does not solve

$$\max_{\alpha_i \in \Delta(A_i)} \sum_{m_j} \frac{\phi_j^*[m_i', m_j]}{\phi_{j,M_i}^*[m_i']} u_i(\alpha_i, \pi_j(m_j)). \tag{5}$$

Let a_i^* be a solution to problem (5), $\bar{m}_i \notin \mathrm{supp}(\phi_{M_i}^*)$, $\phi_i' = 1_{(\bar{m}_i, m_j')}$ and $\pi_i' : M_i \times S \to \Delta(A_i)$ be such that

$$\pi_i'(m_i, \phi_i') = \begin{cases} a_i^* & \text{if } m_i = m_i', \\ \pi_i(m_i') & \text{if } m_i = \bar{m}_i, \\ \pi_i(m_i) & \text{otherwise.} \end{cases}$$

Then

$$\begin{aligned} &\sum_m (\phi_i', \phi_j^*)[m] u_i(\pi_i'(m_i, \phi_i'), \pi_j(m_j)) - \sum_m \phi^*[m] u_i(\pi(m)) \\ &= \beta_i \Big(u_i(\pi(m')) - \sum_{m \in \mathrm{supp}(\phi_i^*)} \phi_i^*[m] u_i(\pi(m)) \Big) \\ &+ \beta_j \sum_{m_j} \phi_j^*[m_i', m_j] \Big(u_i(a_i^*, \pi_j(m_j)) - u_i(\pi_i(m_i'), \pi_j(m_j)) \Big) \\ &= \beta_j \sum_{m_j} \phi_j^*[m_i', m_j] \Big(u_i(a_i^*, \pi_j(m_j)) - u_i(\pi_i(m_i'), \pi_j(m_j)) \Big) \end{aligned}$$

where the last equality follows by Lemma 1 since $m' \in \mathrm{supp}(\phi_i^*)$. Since $\pi_i(m_i')$ does not solve problem (5) but a_i^* does, it follows that

$$\sum_{m_j} \frac{\phi_j^*[m_i', m_j]}{\phi_{j,M_i}^*[m_i']} \Big(u_i(a_i^*, \pi_j(m_j)) - u_i(\pi_i(m_i'), \pi_j(m_j)) \Big) > 0$$

and, since $m_i' \in \text{supp}(\phi^*_{j,M_i})$,

$$\sum_{m_j} \phi^*_j[m_i', m_j]\Big(u_i(a^*_i, \pi_j(m_j)) - u_i(\pi_i(m_i'), \pi_j(m_j))\Big) > 0.$$

Hence, $\sum_m (\phi_i', \phi^*_j)[m] u_i(\pi_i'(m_i, \phi_i'), \pi_j(m_j)) - \sum_m \phi^*[m] u_i(\pi(m)) > 0$. But this contradicts (3). ■

It follows by Lemmas 2 and 3 that, for each Nash equilibrium outcome, $i, j \in N$, $i \neq j$, and $m \in \text{supp}(\phi^*_i)$, condition (1) in Theorem 1 holds and $\pi_i(m_i)$ solves

$$\max_{\alpha_i \in \Delta(A_i)} \sum_{m_j} \frac{\phi^*_j[m_i, m_j]}{\phi^*_{j,M_i}[m_i]} u_i(\alpha_i, \pi_j(m_j))$$

whenever $m_i \in \text{supp}(\phi^*_{j,M_i})$ and, hence, $m_i \in \text{supp}(\phi^*_{i,M_i}) \cap \text{supp}(\phi^*_{j,M_i})$. In fact, regarding (1), note that if $i \in N$ and $m \in \text{supp}(\phi^*_i)$, then $m_k \in \text{supp}(\phi^*_{M_k})$ for each $k \in N$ and, thus, $m \in M^*$. Hence,

$$v_i(\pi_j(m_j)) \leq \max_{m_j' \in M^*_j} v_i(\pi_j(m_j')) \leq \sup_{m_j' \in M_j} v_i(\pi_j(m_j')) = v_i(\pi_j(m_j)).$$

Condition (4) implies that, for each $i \in N$, $\pi_i(m_i)$ solves

$$\max_{\alpha_i \in \Delta(A_i)} \sum_{m_j} \frac{\phi^*_j[m_i, m_j]}{\phi^*_{j,M_i}[m_i]} u_i(\alpha_i, \pi_j(m_j))$$

whenever $m_i \in \text{supp}(\phi^*_{j,M_i}) \setminus \text{supp}(\phi^*_{i,M_i})$. This, together with what has been shown in the previous paragraph, shows that condition (2) in Theorem 1 holds.

A.2.2 Proof that condition 3 implies condition 2

Let $\Big((\phi^*_i)_{i \in \text{supp}(\beta)}, \Big((\pi_i(m_i))_{m_i \in \text{supp}(\phi^*_{M_i})}\Big)_{i \in N}\Big)$ be such that conditions (1) and (2) in Theorem 1 hold; we will show that it is the outcome of a sequential equilibrium.

We will construct a sequential equilibrium π with the desired outcome. Let $i \in N$ and $j \neq i$. Set $\pi^1_i = \phi^*_i$ and $\pi^2_i(m_i, \phi^*_i) = \pi_i(m_i)$ for each $m_i \in \text{supp}(\phi^*_{M_i})$ since the goal is to define a strategy with outcome $\Big((\phi^*_i)_{i \in \text{supp}(\beta)}, \Big((\pi_i(m_i))_{m_i \in \text{supp}(\phi^*_{M_i})}\Big)_{i \in N}\Big)$.

We will specify the remaining values of π^2_i as follows. Let

$$\bar{m}_i \in \text{supp}(\phi^*_{i,M_i}).$$

Informally, we will define $\{\pi^\alpha, p^\alpha\}_\alpha$ such that player i, after choosing ϕ_i and receiving m_i, believes that $\phi_j = \phi_j^*$ and that m_j occurs with probability $\frac{\phi_i^*[\bar{m}_i, m_j]}{\phi_{i,M_i}^*[\bar{m}_i]}$. In this case, we set player i's action to be $\pi_i(\bar{m}_i)$, which is a best-reply against the action $\pi_j(m_j)$ of player j for each m_j such that $\phi_i^*[\bar{m}_i, m_j] > 0$.

The above belief is only possible when $\beta_i\phi_{i,M_i}[m_i]+\beta_i\phi_{i,M_j}[m_j] = 0$ since otherwise, player i has to assign probability $\frac{\beta_i\phi_i[m_i,m_j]+\beta_j\phi_j^*[m_i,m_j]}{\beta_i\phi_{i,M_i}[m_i]+\beta_j\phi_{j,M_i}^*[m_i]}$ to (m_j, ϕ_j^*). In this case, we specify player i's action to be a best-reply against the expected action of player j.

The formal details are as follows. For each $m_i \in M_i$ and $\phi_i \in S$ such that $\beta_i\phi_{i,M_i}[m_i] + \beta_j\phi_{j,M_i}^*[m_i] = 0$, let $\pi_i^2(m_i, \phi_i) = \pi_i(\bar{m}_i)$. In particular, $\pi_i^2(m_i, \phi_i^*) = \pi_i(\bar{m}_i)$ if $m_i \notin \text{supp}(\phi_{M_i}^*)$ since then $\beta_i\phi_{i,M_i}^*[m_i] + \beta_j\phi_{j,M_i}^*[m_i] = 0$. Thus, $\pi_i(m_i) = \pi_i^2(m_i, \phi_i^*)$ is defined for each $i \in N$ and $m_i \in M_i$.

For each $m_i \in M_i$ and $\phi_i \neq \phi_i^*$ such that $\beta_i\phi_{i,M_i}[m_i]+\beta_j\phi_{j,M_i}^*[m_i] > 0$, let $\pi_i^2(m_i, \phi_i)$ be a best-reply against

$$\sum_{m_j} \frac{\beta_i\phi_i[m_i, m_j] + \beta_j\phi_j^*[m_i, m_j]}{\beta_i\phi_{i,M_i}[m_i] + \beta_j\phi_{j,M_i}^*[m_i]}\pi_j(m_j).$$

We may assume that $\pi_i : M_i \times S \to \Delta(A_i)$ is measurable. Note first that $M_i \times S = \cup_{r=1}^3 B_r$ with

$$B_1 = \{(m_i, \phi_i) : \phi_i = \phi_i^*\},$$
$$B_2 = \{(m_i, \phi_i) : \phi_i \neq \phi_i^* \text{ and } \beta_i\phi_{i,M_i}[m_i] + \beta_j\phi_{j,M_i}^*[m_i] = 0\} \text{ and}$$
$$B_3 = \{(m_i, \phi_i) : \phi_i \neq \phi_i^* \text{ and } \beta_i\phi_{i,M_i}[m_i] + \beta_j\phi_{j,M_i}^*[m_i] > 0\}.$$

For each $r \in \{1, 2, 3\}$, B_r is measurable. Indeed, B_1 is closed, B_2 is the intersection of an open set, $\{(m_i, \phi_i) : \phi_i \neq \phi_i^*\}$, with a closed set, $\{(m_i, \phi_i) : \beta_i\phi_{i,M_i}[m_i] + \beta_j\phi_{j,M_i}^*[m_i] = 0\}$, and B_3 is open. Then, for each measurable $B \subseteq \Delta(A_i)$, $\pi_i^{-1}(B) \cap B_1$ is measurable since $\pi_i^{-1}(B) \cap B_1$ is countable. Regarding $\pi_i^{-1}(B) \cap B_3$: Let $f : M_i \times S \to \Delta(A_j)$ be defined by setting, for each $(m_i, \phi_i) \in B_3$, $f(m_i, \phi_i) = \sum_{m_j} \frac{\beta_i\phi_i[m_i,m_j]+\beta_j\phi_j^*[m_i,m_j]}{\beta_i\phi_{i,M_i}[m_i]+\beta_j\phi_{j,M_i}^*[m_i]}\pi_j(m_j)$. Letting $BR_i : \Delta(A_j) \rightrightarrows \Delta(A_i)$ be player i's best-reply correspondence in G, define $\Psi : M_i \times S \rightrightarrows \Delta(A_i)$ by setting, for each $(m_i, \phi_i) \in B_3$, $\Psi(m_i, \phi_i) = BR_i(f(m_i, \phi_i))$. Since $\Delta(A_i)$ is compact, f is continuous and BR_i

is upper hemicontinuous, it follows that Ψ is upper hemicontinuous and, hence, measurable (and, thus, weakly measurable). Hence, Ψ has a measurable selection by the Kuratowski-Ryll-Nardzewski Selection Theorem (e.g. Aliprantis and Border (2006, Theorem 18.13, p. 600)). Finally, for each measurable $B \subseteq \Delta(A_i)$, $\pi_i^{-1}(B) = B_2$ if $\pi_i(\bar{m}_i) \in B$ and $\pi_i^{-1}(B) = \emptyset$ otherwise; thus $\pi_i^{-1}(B) \cap B_2$ is measurable.

We define $\{\pi^\alpha, p^\alpha\}_\alpha$ as follows. The index set consists of $\alpha = (k, F, \hat{F})$ such that $k \in \mathbb{N}$, F is a finite subset of $\mathbb{N}$ and $\hat{F}$ is a finite subset of S; this set is partially ordered by defining $(k', F', \hat{F}') \geq (k, F, \hat{F})$ if $k' \geq k$, $F \subseteq F'$ and $\hat{F} \subseteq \hat{F}'$. For each $m_j \in M_j$, let

$$\bar{q}_i[m_j] = \frac{\phi_i^*[\bar{m}_i, m_j]}{\phi_{i,M_i}^*[\bar{m}_i]}$$

be the ϕ_i^*-probability of m_j conditional on $\bar{m}_i$. For each $\alpha = (k, F, \hat{F})$, we define $p^\alpha(\phi)$ such that the probability distribution of message profiles is $\beta_1\phi_1 + \beta_2\phi_2$ with probability $1 - k^{-1}$; with probability k^{-3}, it equals a probability distribution τ^α that assigns strictly positive probability to each message in F and in the support of information designs in $\hat{F}$; and, with the remaining probability of $k^{-1}(1 - k^{-2})$, it equals a probability distribution q^α such that the probability of m_j conditional on m_i equals $\bar{q}_i[m_j]$. Let

$$\begin{aligned}
\tau_i^\alpha &= \frac{\sum_{l \in F \cup \left(\cup_{\phi \in \hat{F}} \mathrm{supp}(\phi_{M_i})\right)} 2^{-l} 1_l}{\sum_{l \in F \cup \left(\cup_{\phi \in \hat{F}} \mathrm{supp}(\phi_{M_i})\right)} 2^{-l}}, \\
q_i^\alpha &= \tau_i^\alpha \times \bar{q}_i, \\
\tau^\alpha &= \tau_1^\alpha \times \tau_2^\alpha, \\
q^\alpha &= (q_1^\alpha + q_2^\alpha)/2, \\
\mu^\alpha &= (1 - k^{-2}) q^\alpha + k^{-2} \tau^\alpha, \text{ and} \\
p^\alpha(\phi) &= (1 - k^{-1})(\beta_1\phi_1 + \beta_2\phi_2) + k^{-1}\mu^\alpha.
\end{aligned}$$

Furthermore, let $\upsilon_X \in \Delta(X)$ be uniform on X whenever X is a finite set and let

$$\pi_i^{1,\alpha} = (1 - k^{-3}) 1_{\phi_i^*} + k^{-3}\upsilon_{\hat{F}} \text{ and } \pi_i^{2,\alpha}(m_i, \phi_i) = (1 - k^{-1})\pi_i(m_i, \phi_i) + k^{-1}\upsilon_{A_i}$$

for each $(m_i, \phi_i) \in M_i \times S$. Thus, all information designs in $\hat{F}$ and all actions are played with strictly positive probability. Furthermore, the probability of an information design different from ϕ_i^* is much smaller than the probability of q^α (i.e. their ratio is $k^{-3}/k^{-1}(1-k^{-2})$ and goes to zero), which implies that if player i receives a message that is neither in the support of the design he choose nor in the support of ϕ_j^*, then player i believes that this happened because the message was drawn from q^α and not because player j chose a design different from ϕ_j^*.

Let $\varepsilon > 0$. We have that the conditions (i)–(v) in the definition of perfect conditional ε-equilibrium hold by construction. We will show that condition (vi) holds for some subnet of $\{\pi^\alpha, p^\alpha\}_\alpha$. Some technical details of this argument are simplified by our construction of $\{\pi^\alpha, p^\alpha\}_\alpha$ which is such that $\text{supp}(\pi^{1,\alpha})$ and $\text{supp}(p^\alpha)$ are finite for each α. We define

$$S_i(F,\hat{F}) = \Big(\Big(F \cup (\cup_{\phi\in\hat{F}}\text{supp}(\phi_{M_i})) \cup (\text{supp}(\phi^*_{j,M_i}))\Big) \times \hat{F}\Big)$$
$$\cup \Big(\Big(F \cup (\cup_{\phi\in\hat{F}}\text{supp}(\phi_{M_i})) \cup (\text{supp}(\phi^*_{M_i}))\Big) \times \{\phi_i^*\}\Big)$$

which is the set of pairs (m_i, ϕ_i) that can occur with strictly positive probability. Indeed, if $(m,\phi) \in \mathbb{N}^2 \times S^2$ is such that $\pi^{1,\alpha}[\phi] > 0$ and $\sum_{\phi'\in\text{supp}(\pi^{1,\alpha})} p^\alpha(\phi')[m] > 0$, then $(m_i, \phi_i) \in S_i(F,\hat{F})$ for each $i \in N$.

Recall that $\alpha = (k, F, \hat{F})$. In what follows, we will often fix F and $\hat{F}$ and take limits as $k \to \infty$. Regarding condition (vi) (a), let $i, j \in N$, $j \neq i$ and $\phi_i' \in S$. We have that, for each finite subsets F and $\hat{F}$ of $\mathbb{N}$ and S, respectively,

$$\lim_k \sum_{\phi\in\text{supp}(\pi^{1,\alpha})} \pi^{1,\alpha}[\phi]\left(\sum_m p^\alpha(\phi)[m]u_i(\pi^{2,\alpha}(m,\phi))\right) = \sum_m \phi^*[m]u_i(\pi(m))$$

and that

$$\lim_k \sum_{\phi\in\text{supp}(1_{\phi_i'}\times\pi_j^{1,\alpha})} (1_{\phi_i'} \times \pi_j^{1,\alpha})[\phi]\left(\sum_m p^\alpha(\phi)[m]u_i(\pi^{2,\alpha}(m,\phi))\right) =$$
$$\sum_m (\phi_i', \phi_j^*)[m]u_i(\pi_i(m_i,\phi_i'), \pi_j(m_j)).$$

Hence, by considering α such that $k \geq k_0$ for some $k_0 \in \mathbb{N}$, it is enough to show that

$$\sum_m \phi^*[m]u_i(\pi(m)) \geq \sum_m (\phi_i', \phi_j^*)[m]u_i(\pi_i(m_i,\phi_i'), \pi_j(m_j)),$$

which is equivalent to

$$\sum_m \phi_i^*[m]u_i(\pi(m)) \geq \sum_m \phi_i'[m]u_i(\pi_i(m_i,\phi_i'), \pi_j(m_j)). \tag{6}$$

For each $m_j \in M_j$, $\pi_j(m_j) \in \{\pi_j(m'_j) : m'_j \in \text{supp}(\phi^*_{M_j})\}$ since $\pi_j(m_j) = \pi_j(\bar{m}_j)$ whenever $m_j \notin \text{supp}(\phi^*_{M_j})$. Thus, by (1),

$$\sum_m \phi'_i[m]u_i(\pi_i(m_i,\phi'_i),\pi_j(m_j)) \leq \sum_m \phi'_i[m]v_i(\pi_j(m_j))$$
$$\leq \max_{m_j\in M^*_j} v_i(\pi_j(m_j)) = \sum_m \phi^*_i[m]u_i(\pi(m))$$

and, hence, (6) holds. It then follows that condition (vi) (a) also holds.

Consider next condition (vi) (b). For each $i,j \in N$, $i \neq j$, finite subset F of $\mathbb{N}$, finite subset $\hat{F}$ of S, $(m_i,\phi_i) \in S_i(F,\hat{F})$ and $\gamma_i \in \Delta(A_i)$, we have that

$$\lim_k \frac{\sum_{\phi_j\in\text{supp}(\pi_j^{1,\alpha})} \pi_j^{1,\alpha}[\phi_j]\left(\sum_{m_j} p^\alpha(\phi_i,\phi_j)[m_i,m_j]u_i(\gamma_i,\pi_j^{2,\alpha}(m_j,\phi_j))\right)}{\sum_{\phi_j\in\text{supp}(\pi_j^{1,\alpha})} \pi_j^{1,\alpha}[\phi_j]p^\alpha_{M_i}(\phi_i,\phi_j)[m_i]}$$
$$= \sum_{m_j} \frac{\phi^*_i[\bar{m}_i,m_j]}{\phi^*_{i,M_i}[\bar{m}_i]}u_i(\gamma_i,\pi_j(m_j))$$

if $\beta_i\phi_{i,M_i}[m_i] + \beta_j\phi^*_{j,M_i}[m_i] = 0$, and

$$\lim_k \frac{\sum_{\phi_j\in\text{supp}(\pi_j^{1,\alpha})} \pi_j^{1,\alpha}[\phi_j]\left(\sum_{m_j} p^\alpha(\phi_i,\phi_j)[m_i,m_j]u_i(\gamma_i,\pi_j^{2,\alpha}(m_j,\phi_j))\right)}{\sum_{\phi_j\in\text{supp}(\pi_j^{1,\alpha})} \pi_j^{1,\alpha}[\phi_j]p^\alpha_{M_i}(\phi_i,\phi_j)[m_i]} =$$
$$\sum_{m_j} \frac{\beta_i\phi_i[m_i,m_j] + \beta_j\phi^*_j[m_i,m_j]}{\beta_i\phi_{i,M_i}[m_i] + \beta_j\phi^*_{j,M_i}[m_i]}u_i(\gamma_i,\pi_j(m_j))$$

if $\beta_i\phi_{i,M_i}[m_i] + \beta_j\phi^*_{j,M_i}[m_i] > 0$. The latter case is clear since all terms in the denominator of the fraction converge to zero except the one that converges to $\beta_i\phi_{i,M_i}[m_i] + \beta_j\phi^*_{j,M_i}[m_i]$ and similarly regarding the numerator.

In the former case, both the numerator and the denominator converge to zero since $\beta_i\phi_{i,M_i}[m_i] + \beta_j\phi^*_{j,M_i}[m_i] = 0$. Multiplying each by k, it follows that all terms converge to zero except the ones corresponding to the case where $\pi_j^{1,\alpha} = \phi^*_j$ and $p^\alpha(\phi_i,\phi^*_j) = q^\alpha$. Furthermore, for each $m_j \in M_j$,

$$q^\alpha[m_i,m_j] = 2^{-1}(q^\alpha_i[m_i,m_j] + q^\alpha_j[m_i,m_j]),$$
$$q^\alpha_i[m_i,m_j] = \tau^\alpha_i[m_i]\bar{q}_i[m_j] \text{ and}$$
$$q^\alpha_j[m_i,m_j] = 0,$$

the latter since $m_i \notin \text{supp}(\phi^*_{j,M_i})$. Hence, $q^\alpha[m_i, m_j] = 2^{-1}\tau_i^\alpha[m_i]\bar{q}_i[m_j]$ and $q^\alpha_{M_i}[m_i] = 2^{-1}\tau_i^\alpha[m_i]$. Thus,

$$\frac{q^\alpha[m_i, m_j]}{q^\alpha_{M_i}[m_i]} = \bar{q}_i[m_j] = \frac{\phi_i^*[\bar{m}_i, m_j]}{\phi^*_{i,M_i}[\bar{m}_i]}.$$

We will next show that $\pi_i(m_i, \phi_i)$ solves

$$\max_{\gamma_i \in \Delta(A_i)} \lim_k \frac{\sum_{\phi_{-i} \in \text{supp}(\pi^{1,\alpha}_{-i})} \pi^{1,\alpha}_{-i}[\phi_{-i}] \left(\sum_{m_{-i}} p^\alpha(\phi_i, \phi_{-i})[m_i, m_{-i}] u_i(\gamma_i, \pi^{2,\alpha}_{-i}(m_{-i}, \phi_{-i})) \right)}{\sum_{\phi_{-i} \in \text{supp}(\pi^{1,\alpha}_{-i})} \pi^{1,\alpha}_{-i}[\phi_{-i}] p^\alpha_{M_i}(\phi_i, \phi_{-i})[m_i]} \tag{7}$$

for each $i \in N$ and $(m_i, \phi_i) \in S_i(F, \hat{F})$.

If $\beta_i \phi_{i,M_i}[m_i] + \beta_j \phi^*_{j,M_i}[m_i] = 0$, then

$$\begin{aligned}
&\lim_k \frac{\sum_{\phi_j \in \text{supp}(\pi^{1,\alpha}_j)} \pi^{1,\alpha}_j[\phi_j] \left(\sum_{m_j} p^\alpha(\phi_i, \phi_j)[m_i, m_j] u_i(\gamma_i, \pi^{2,\alpha}_j(m_j, \phi_j)) \right)}{\sum_{\phi_j \in \text{supp}(\pi^{1,\alpha}_j)} \pi^{1,\alpha}_j[\phi_j] p^\alpha_{M_i}(\phi_i, \phi_j)[m_i]} \\
&= \sum_{m_j} \frac{\phi_i^*[\bar{m}_i, m_j]}{\phi^*_{i,M_i}[\bar{m}_i]} u_i(\gamma_i, \pi_j(m_j)).
\end{aligned}$$

Since $\pi_i(m_i, \phi_i) = \pi_i(\bar{m}_i)$ and $\pi_i(\bar{m}_i) \in BR_i(\pi_j(m_j))$ for each $m_j \in M_j$ such that $(\bar{m}_i, m_j) \in \text{supp}(\phi_i^*)$ by (1), it follows that (7) holds in this case.

If $\beta_i \phi_{i,M_i}[m_i] + \beta_j \phi^*_{j,M_i}[m_i] > 0$ and $\phi_i \neq \phi_i^*$, then

$$\begin{aligned}
&\lim_k \frac{\sum_{\phi_j \in \text{supp}(\pi^{1,\alpha}_j)} \pi^{1,\alpha}_j[\phi_j] \left(\sum_{m_j} p^\alpha(\phi_i, \phi_j)[m_i, m_j] u_i(\gamma_i, \pi^{2,\alpha}_j(m_j, \phi_j)) \right)}{\sum_{\phi_j \in \text{supp}(\pi^{1,\alpha}_j)} \pi^{1,\alpha}_j[\phi_j] p^\alpha_{M_i}(\phi_i, \phi_j)[m_i]} \\
&= \sum_{m_j} \frac{\beta_i \phi_i[m_i, m_j] + \beta_j \phi_j^*[m_i, m_j]}{\beta_i \phi_{i,M_i}[m_i] + \beta_j \phi^*_{j,M_i}[m_i]} u_i(\gamma_i, \pi_j(m_j)) \\
&= u_i \left(\gamma_i, \sum_{m_j} \frac{\beta_i \phi_i[m_i, m_j] + \beta_j \phi_j^*[m_i, m_j]}{\beta_i \phi_{i,M_i}[m_i] + \beta_j \phi^*_{j,M_i}[m_i]} \pi_j(m_j) \right).
\end{aligned}$$

Since $\pi_i(m_i, \phi_i)$ is optimal against $\sum_{m_j} \frac{\beta_i \phi_i[m_i, m_j] + \beta_j \phi_j^*[m_i, m_j]}{\beta_i \phi_{i,M_i}[m_i] + \beta_j \phi^*_{j,M_i}[m_i]} \pi_j(m_j)$, it follows that (7) holds in this case.

Finally, consider the case where $\phi_i = \phi_i^*$ and $\beta_i \phi_{i,M_i}[m_i] + \beta_j \phi^*_{j,M_i}[m_i] > 0$. Note that it is enough to show that

$$\sum_{m_j} \phi^*[m] \big(u_i(\pi(m)) - u_i(a_i, \pi_j(m_j)) \big) \geq 0 \tag{8}$$

for each $a_i \in A_i$ and that

$$\sum_{m_j} \phi^*[m]\big(u_i(\pi(m)) - u_i(a_i, \pi_j(m_j))\big) =$$
$$\sum_{m_j} \beta_i \phi_i^*[m]\big(u_i(\pi(m)) - u_i(a_i, \pi_j(m_j))\big) + \sum_{m_j} \beta_j \phi_j^*[m]\big(u_i(\pi(m)) - u_i(a_i, \pi_j(m_j))\big).$$

We have that $u_i(\pi(m)) \geq u_i(a_i, \pi_j(m_j))$ for each m_j such that $\phi_i^*[m] > 0$ by (1); moreover, for each m_j such that $\phi_j^*[m] > 0$, then $m_i \in \text{supp}(\phi^*_{j,M_i})$ and, hence, $\sum_{m_j} \beta_j \phi_j^*[m]\big(u_i(\pi(m)) - u_i(a_i, \pi_j(m_j))\big) \geq 0$ by (2). Thus, (8) holds and so does (7).

The above arguments show that, for each finite subsets F of $\mathbb{N}$ and $\hat{F}$ of S, condition (vi) holds whenever k is sufficiently high. Specifically, condition (vi) (a) holds for each $i \in N$ whenever $k \geq k_0$. For each $i \in N$ and $(m_i, \phi_i) \in S_i(F, \hat{F})$, there is $k(m_i, \phi_i)$ such that condition (vi) (b) holds whenever $k \geq k(m_i, \phi_i)$. Thus, let

$$k(F, \hat{F}) = \max\left\{k_0, \max_{i \in N} \max_{(m_i, \phi_i) \in S_i(F, \hat{F})} k(m_i, \phi_i)\right\}.$$

Since condition (vi) (b) is trivially satisfied when

$$\pi_i^{1,\alpha}[\phi_i] \sum_{\phi_j \in \text{supp}(\pi_j^{1,\alpha})} \pi_j^{1,\alpha}[\phi_j] p^\alpha_{M_i}(\phi_i, \phi_j)[m_i] = 0,$$

i.e. when $i \in N$ and $(m_i, \phi_i) \notin S_i(F, \hat{F})$, it follows that condition (vi) holds whenever $k \geq k(F, \hat{F})$. This allows us to define the following subnet $\{\pi^{\varphi(\eta)}, p^{\varphi(\eta)}\}_\eta$ of $\{\pi^\alpha, p^\alpha\}_\alpha$ such that condition (vi) holds.

The index set of the subnet $\{\pi^{\varphi(\eta)}, p^{\varphi(\eta)}\}_\eta$ is the same as the one in the net $\{\pi^\alpha, p^\alpha\}_\alpha$. The function $\varphi : \eta \mapsto \alpha$ is defined by setting, for each $\eta = (k, F, \hat{F})$,

$$\varphi(\eta) = \left(\max\left\{k, k(F, \hat{F})\right\}, F, \hat{F}\right).$$

It is then clear that condition (vi) holds and that, as required by the definition of a subnet, for each α_0, there exists η_0, e.g. $\eta_0 = \alpha_0$, such that $\varphi(\eta) \geq \alpha_0$ for each $\eta \geq \eta_0$.

A.3 Proof of Corollary 1

We first show that $A(G) \subseteq \mathcal{A}$. Let $\pi \in \Pi$ be a sequential equilibrium of G_{id} and $\sigma = \sigma_\pi$ be its action distribution. Then $\left(\phi_i^*, (\pi_i(m_i))_{m_i \in \text{supp}(\phi^*_{M_i})}\right)_{i \in N}$ satisfies the conditions of Theorem 1. Define $\phi_i \in \Delta_f(\Delta(A_1) \times \Delta(A_2))$ such that $\phi_i[\alpha] = \sum_{m:\pi(m)=\alpha} \phi_i^*[m]$. Then

$$\sigma[a] = \sum_{i,m} \beta_i \phi_i^*[m]\pi(m)[a] = \sum_{i,\alpha} \beta_i \phi_i[\alpha]\alpha[a].$$

Thus, $\sigma = (\sum_i \beta_i \phi_i)^A$.

Claim 1. *For each $i \in N$ and $\alpha_j \in \text{supp}(\phi_{i,\Delta(A_j)})$,*

$$\left(\sum_{\alpha_i \in \text{supp}(\phi_{i,\Delta(A_i)})} \frac{\phi_i[\alpha_i, \alpha_j]}{\phi_{i,\Delta(A_j)}[\alpha_j]} \alpha_i, \alpha_j \right)$$

is a Nash equilibrium of G and

$$u_i(\alpha_i, \alpha_j) = \max_{\alpha' \in \cup_l \text{supp}(\phi_l)} u_i(\alpha')$$

for each $\alpha_i \in \Delta(A_i)$ such that $\phi_i[\alpha_i, \alpha_j] > 0$.

Proof. Since $\alpha_j \in \text{supp}(\phi_{i,\Delta(A_j)})$, there exists $m_j \in \text{supp}(\phi_{i,M_j}^*)$ such that $\alpha_j = \pi_j(m_j)$. Let $\hat{M}_j$ be the set of such m_j and note that, for each $a_i \in A_i$,

$$\sum_{m_i} \sum_{m_j \in \hat{M}_j} \phi_i^*[m_i, m_j]\pi_i(m_i)[a_i] = \sum_{\alpha_i} \phi_i[\alpha_i, \alpha_j]\alpha_i[a_i].$$

Thus, it follows that $\sum_{m_i} \frac{\phi_i^*[\{m_i\} \times \hat{M}_j]}{\phi_{i,M_j}^*[\hat{M}_j]} \pi_i(m_i)[a_i] = \sum_{\alpha_i} \frac{\phi_i[\alpha_i,\alpha_j]}{\phi_{i,\Delta(A_j)}[\alpha_j]} \alpha_i[a_i]$. Since, for each $m_j \in \hat{M}_j$ and $\hat{\alpha}_j \in \Delta(A_j)$,

$$\sum_{m_i} \phi_i^*[m_i, m_j]u_j(\alpha_j, \pi_i(m_i)) \geq \sum_{m_i} \phi_i^*[m_i, m_j]u_j(\hat{\alpha}_j, \pi_i(m_i))$$

by (2), it follows that

$$\sum_{m_i} \sum_{m_j \in \hat{M}_j} \phi_i^*[m_i, m_j]u_j(\alpha_j, \pi_i(m_i)) \geq \sum_{m_i} \sum_{m_j \in \hat{M}_j} \phi_i^*[m_i, m_j]u_j(\hat{\alpha}_j, \pi_i(m_i))$$

and that α_j maximises $u_j(\cdot, \sum_{\alpha_i} \frac{\phi_i[\alpha_i,\alpha_j]}{\phi_{i,\Delta(A_j)}[\alpha_j]} \alpha_i)$.

For each $\alpha_i \in \Delta(A_i)$ such that $\phi_i[\alpha_i, \alpha_j] > 0$, there exists $m \in \text{supp}(\phi_i^*)$ such that $\pi(m) = (\alpha_i, \alpha_j)$. Furthermore, for each $\alpha' \in \cup_l \text{supp}(\phi_l)$, there exists $m' \in M^*$ such that $\pi(m') = \alpha'$. It then follows by (1) that $\alpha_i = \pi_i(m_i) \in BR_i(\pi_j(m_j)) = BR_i(\alpha_j)$ and

$$u_i(\alpha_i, \alpha_j) = v_i(\alpha_j) = v_i(\pi_j(m_j)) \geq v_i(\pi_j(m'_j)) \geq u_i(\pi(m')) = u_i(\alpha').$$

Thus, $\sum_{\alpha_i \in \text{supp}(\phi_{i,\Delta(A_i)})} \frac{\phi_i[\alpha_i,\alpha_j]}{\phi_{i,\Delta(A_j)}[\alpha_j]}\alpha_i \in BR_i(\alpha_j)$. □

Note that ϕ_i can be written as:

$$\phi_i = \sum_{\alpha_j} \phi_{i,\Delta(A_j)}[\alpha_j] \left(\left(\sum_{\alpha_i} \frac{\phi_i[\alpha_i, \alpha_j]}{\phi_{i,\Delta(A_j)}[\alpha_j]} 1_{\alpha_i} \right) \otimes 1_{\alpha_j} \right)$$

Hence, define ψ_i as:

$$\psi_i = \sum_{\alpha_j} \phi_{i,\Delta(A_j)}[\alpha_j] 1_{\sigma^{i,\alpha_j}},$$

where, for each $\alpha_j \in \text{supp}(\phi_{i,\Delta(A_j)})$,

$$\sigma^{i,\alpha_j} = \left(\sum_{\alpha_i} \frac{\phi_i[\alpha_i, \alpha_j]}{\phi_{i,\Delta(A_j)}[\alpha_j]} \alpha_i, \alpha_j \right)$$

is a Nash equilibrium of G by Claim 1. It is clear that $\psi_i^A = \phi_i^A$ and, hence,

$$\left(\sum_i \beta_i \psi_i \right)^A = \sigma.$$

Then let $L_i = |\text{supp}(\phi_{i,\Delta(A_j)})|$ and, writing $\text{supp}(\phi_{i,\Delta(A_j)}) = \{\alpha_j^1, \ldots, \alpha_j^{L_i}\}$, let $\eta^{i,l} = \phi_{i,\Delta(A_j)}[\alpha_j^l]$ and $\sigma^{i,l} = \sigma^{i,\alpha_j^l}$ for each $l \in \{1, \ldots, L_i\}$.

For each $\alpha_j \in \text{supp}(\phi_{i,\Delta(A_j)})$, it follows by Claim 1 that

$$u_i(\sigma^{i,\alpha_j}) = \sum_{\alpha_i} \frac{\phi_i[\alpha_i, \alpha_j]}{\phi_{i,\Delta(A_j)}[\alpha_j]} u_i(\alpha_i, \alpha_j) = \max_{\hat{\alpha} \in \cup_l \text{supp}(\phi_l)} u_i(\hat{\alpha}).$$

Thus, $u_i(\sigma^{i,\alpha_j}) = u_i(\sigma^{i,\alpha'_j})$ for each $\alpha'_j \in \text{supp}(\phi_{i,\Delta(A_j)})$. Furthermore, for each $\alpha_i \in \text{supp}(\phi_{j,\Delta(A_i)})$,

$$u_i(\sigma^{j,\alpha_i}) = \sum_{\alpha_j} \frac{\phi_j[\alpha_i, \alpha_j]}{\phi_{j,\Delta(A_i)}[\alpha_i]} u_i(\alpha_i, \alpha_j) \leq \max_{\hat{\alpha} \in \cup_l \text{supp}(\phi_l)} u_i(\hat{\alpha});$$

thus, $u_i(\sigma^{i,\alpha_j}) \geq u_i(\sigma^{j,\alpha_i})$. This completes the proof that $A(G) \subseteq \mathcal{A}$.

We now show that $\mathcal{A} \subseteq A(G)$. Let $\sigma = (\beta_1\phi_1 + \beta_2\phi_2)^A \in \mathcal{A}$, i.e. $\phi_i = \sum_{l=1}^{L_i} \eta^{i,l} 1_{\sigma^{i,l}}$, $\sum_{l=1}^{L_i} \eta^{i,l} = 1$, $\eta^{i,l} \geq 0$, $\sigma^{i,l} \in N(G)$, $u_i(\sigma^{i,l}) = u_i(\sigma^{i,k}) \geq u_j(\sigma^{j,r})$ for each $i, j \in N$ with $i \neq j$, $l, k \in \{1, \ldots, L_i\}$ and $r \in \{1, \ldots, L_j\}$.

For each $i \in N$, $l \in \{1, \dots, L_1\}$ and $k \in \{1, \dots, L_2\}$, pick $m_i^{1,l}$ and $m_i^{2,k}$ in M_i such that $m_i^{1,l} \neq m_i^{1,r}$, $m_i^{2,k} \neq m_i^{2,s}$ and $m_i^{1,l} \neq m_i^{2,k}$ for each $l, r \in \{1, \dots, L_1\}$ and $k, s \in \{1, \dots, L_2\}$. Set $\phi_1^* = \sum_{l=1}^{L_1} \eta^{1,l} 1_{m^{1,l}}$, $\phi_2^* = \sum_{l=1}^{L_2} \eta^{2,l} 1_{m^{2,l}}$ and, for each $i \in N$, $j \in N$ and $l \in \{1, \dots, L_j\}$, $\pi_i(m_i^{j,l}) = \sigma_i^{j,l}$. It is then clear that σ is the action distribution of the outcome $\left(\phi_i^*, (\pi_i(m_i))_{m_i \in \text{supp}(\phi_{M_i}^*)}\right)_{i \in N}$, i.e. $\sigma[a] = \sum_{m \in M^*} \phi^*[m]\pi(m)[a]$ for each $a \in A$.

Fix $i \in N$ and let $j \neq i$. For each $m_j \in M_j^*$, $\pi_j(m_j) = \sigma_j^{k,l}$ for some $k \in \{i, j\}$ and $l \in \{1, \dots, L_k\}$. Since $v_i(\sigma_j^{i,l}) = u_i(\sigma^{i,l}) \geq u_i(\sigma^{j,r}) = v_i(\sigma_j^{j,r})$ for each $l \in \{1, \dots, L_i\}$ and $r \in \{1, \dots, L_j\}$, it follows that $\max_{m_j \in M_j^*} v_i(\pi_j(m_j)) = u_i(\sigma^{i,l})$ for each $l \in \{1, \dots, L_i\}$. Since, for each $m \in \text{supp}(\phi_i^*)$, there exists $l \in \{1, \dots, L_i\}$ such that $\pi_i(m_i) = \sigma_i^{i,l}$ and $\pi_j(m_j) = \sigma_j^{i,l}$, it follows that $\text{supp}(\phi_i^*) \subseteq \{m \in M : v_i(\pi_{-i}(m_{-i})) = \max_{m'_{-i} \in M_{-i}} v_i(\pi_{-i}(m'_{-i}))$ and $\pi_i(m_i) \in BR_i(\pi_{-i}(m_{-i}))\}$. Thus, (1) holds.

Moreover, for each $m_i \in \text{supp}(\phi_{j,M_i}^*)$, $m_i = m_i^{j,l}$ for some $l \in \{1, \dots, L_j\}$ and hence $\pi_i(m_i) = \sigma_i^{j,l}$ solves

$$\max_{\alpha_i \in \Delta(A_i)} \sum_{m_j} \frac{\phi_j^*[m_i, m_j]}{\phi_{j,M_i}^*[m_i]} u_i(\alpha_i, \pi_j(m_j)) = \max_{\alpha_i \in \Delta(A_i)} u_i(\alpha_i, \sigma_j^{j,l}).$$

Thus, (2) holds. It then follows by Theorem 1 that $(\phi_i^*, (\pi_i(m_i))_{m_i \in \text{supp}(\phi_{M_i}^*)})_{i \in N}$ is the outcome of a sequential equilibrium of G_{id} and, thus, that $\sigma \in A(G)$.

A.4 Proof of Corollary 3

The characterization of $U(G)$ follows from the definition of $\mathcal{G}$ and Corollary 2.

Standard results (e.g. Theorems 2.5.5 and 2.6.1 in van Damme (1991) and their proofs) imply that there is an open set O of $\mathbb{R}^{2|A|}$ such that its complement has Lebesgue measure zero and, for each $u \in O$, there is an open neighborhood V_u of u and $|N(u)|$ continuous functions, $f_k : V_u \to \Delta(A_1) \times \Delta(A_2)$ with $k \in \{1, \dots, |N(u)|\}$ such that, for each $u' \in V_u$, $N(u') = \left\{f_k(u') : k \in \{1, \dots, |N(u)|\}\right\}$ and $f_k(u) \neq f_l(u)$ for each $k, l \in \{1, \dots, |N(u)|\}$ with $k \neq l$.[22] Shrinking V_u if needed, we may assume that, for each $a \in A$, $k, l \in \{1, \dots, |N(u)|\}$ and $u' \in V_u$, $f_k(u')[a] \neq f_l(u')[a]$ if $f_k(u)[a] \neq f_l(u)[a]$.

[22]The set $N(u)$ denotes the set of Nash equilibria of the game with payoff function u.

We have that $\mathbb{R}^{2|A|}$ is separable, hence, there is a countable collection $\{V_{u_j}\}_{j=1}^{\infty}$ such that $O = \cup_{j=1}^{\infty} V_{u_j}$. Define, for each $j \in \mathbb{N}$, $I_j = \{1, \ldots, |N(u_j)|\}$ and

$$O_j = \cap_{(i,k,l) \in N \times I_j^2 : k \neq l} \{u \in V_{u_j} : u_i(f_k(u)) \neq u_i(f_l(u))\}.$$

Then O_j is open and $\cup_{j=1}^{\infty} O_j \subseteq \mathcal{G}$. It thus suffices to show that $C_{j,i,k,l} = \{u \in V_{u_j} : u_i(f_k(u)) = u_i(f_l(u))\}$ has Lebesgue measure zero for each $j \in \mathbb{N}$ and $(i,k,l) \in N \times I_j^2$ such that $k \neq l$.

Let $j \in \mathbb{N}$ and $(i,k,l) \in N \times I_j^2$ be such that $k \neq l$. Since $f_k(u_j) \neq f_l(u_j)$, let $a \in A$ be such that $f_k(u)[a] \neq f_l(u)[a]$ for each $u \in V_{u_j}$. Then

$$C_{j,i,k,l} \subseteq \left\{ u \in V_{u_j} : u_i(a) = \frac{\sum_{a' \neq a} u_i(a') \big(f_l(u)[a'] - f_k(u)[a'] \big)}{f_k(u)[a] - f_l(u)[a]} \right\}.$$

It then follows by Tonelli's Theorem (e.g. Wheeden and Zygmund (1977, Theorem 6.10, p. 92)) that $C_{j,i,k,l}$ has Lesbegue measure zero.