Mediated Subgame Perfect Equilibrium*

Christian Ewerhart[†] Haoyuan Zeng[‡]

First version: December 15, 2025

—To download the most recent version, click here.—

Abstract. In multi-stage games with observed actions, correlated action profiles can be implemented using an autonomous device that sends private messages to the players at the beginning of each stage. In this paper, we assume that such messages, while initially confidential, are publicly revealed at the end of each stage. We call the resulting equilibrium concept mediated subgame perfect equilibrium (MSPE). We establish a revelation principle and derive necessary and sufficient conditions for the implementability of a given payoff profile as an MSPE. These conditions coincide with the standard perfect folk-theorem conditions for two-player games and for games that satisfy the NEU condition, but are more flexible otherwise.

Keywords. Multi-stage games; revelation principle; perfect folk theorem; correlated equilibrium; mediation

JEL classification. C72 - Noncooperative Games; C73 - Stochastic and Dynamic Games, Evolutionary Games, Repeated Games

*) For useful discussions, we are grateful to Wojciech Olszewski and Marek Pycia.

†) (corresponding) Department of Economics, University of Zurich; Schönberggasse 1, 8001 Zurich, Switzerland; e-mail: christian.ewerhart@econ.uzh.ch.

‡) Zurich Graduate School of Economics, University of Zurich; Blümlisalpstrasse 10, 8006 Zurich; e-mail: haoyuan.zeng@econ.uzh.ch.

1 Preliminaries

1.1 Introduction

Folk theorems shed light on the role of social norms in long-term relationships and the game-theoretic mechanisms that render such norms sustainable. In their seminal work, Fudenberg and Maskin (1986) not only delineated conditions under which the perfect folk theorem applies to infinitely repeated games with perfect monitoring, but also pointed out inherent limitations of those conditions. Specifically, they noted that sanctions against an individual player may be hard to implement if the set of feasible and strictly individually rational payoff profiles is not full-dimensional. Indeed, as they show with an example, the conclusion of the folk theorem can fail in that case. Subsequent work by Abreu et al. (1994) further clarified the role of the dimensionality assumption for the scope of the perfect folk theorem, highlighting that effective sanctions presuppose that players' payoff functions satisfy the *NEU condition*, i.e., no two players possess equivalent utility functions. Building on this observation, Wen (1994) introduced the notion of an *effective minimax value*. This led to a version of the folk theorem that applies to all finite stage games, i.e., even if they do not satisfy NEU.[1]

While these contributions focus on enforcement of cooperation through unsupervised equilibrium play, a complementary line of research examines how the introduction of informational coordination can facilitate the implementation of socially desirable outcomes. Thus, there has been an interest in understanding the implications of assuming that a central entity provides players with information or action recommendations. By a *mediator*, we mean an (autonomous) device that commu-

[1] The ability to penalize individual members of groups or alliances with closely aligned interests can be socially beneficial in a wide range of settings. Examples include criminal organizations and clans, oligarchic and kleptocratic elites, cartels, racketeering enterprises, and terrorist organizations.

nicates confidential messages to the players at each stage, just before decisions are made (Aumann, 1974, 1987; Forges, 1986; Myerson, 1986). It is not hard to see that mediation may have an impact on the set of implementable payoffs in repeated games. For instance, with three or more players, mediation enables players enforcing punishments to coordinate their actions, which may strictly lower the utility level to which an individual player can be held (Hart, 1979; Renault and Tomala, 1998). In the theory of infinitely repeated games, this corresponds to an expansion of the set of penalty payoffs, down to the *correlated minimax value*. Along these lines, the Nash folk theorem (where sequential rationality is not imposed at information sets off the equilibrium path) extends to settings with mediated play in an essentially straightforward way (Sorin, 1992, p. 86). However, to our knowledge, no corresponding extension of the perfect folk theorem exists. Indeed, as pointed out by Sugaya and Wolitzky (2021), combining mediation and sequential rationality in dynamic games can lead to unexpected pitfalls. For instance, the revelation principle no longer holds in general for the solution concept of sequential equilibrium.

The present paper takes this observation as motivation to revisit the idea of mediation in a framework that remains as close as possible to the standard model. Specifically, in addition to perfect monitoring, we assume that messages are *ex post observable*. Thus, all private messages sent by the mediator are publicly revealed at the end of each stage. Moreover, we assume that the mediator is perfectly informed about all prior (pure) action profiles, i.e., no input messages are required. The resulting solution concept for repeated games, *mediated subgame perfect equilibrium (MSPE)*, requires that, after any history, players' behavioral strategies induce conditionally optimal choices for every message profile received from the mediator at that stage. This approach turns out to be well-behaved in several ways. We prove a revelation principle for our solution concept and characterize the set of payoffs implementable as an MSPE for sufficiently patient players. We also show that mediation can strictly expand the set of implementable payoff vectors.

While our key assumption, the ex post observability of messages, departs from the main thrust of the literature on communication in dynamic games, it has two notable benefits. First, assuming that mediation is ex post transparent allows us to rely on subgame perfect equilibrium (Selten, 1965) as a solution concept and to keep explicit modeling of beliefs at a minimum. Second, ex post observability also has applied appeal: for instance, it is common in public administration, law enforcement, procurement, and hierarchical relationships, where it has the potential to improve accountability, reduce corruption, and foster trust. ex post observability may also arise from inadvertent or opportunistic leaks of confidential information, which have become harder to prevent in the digital era.[2]

1.2 Preview of results

The main results of this paper fall into four groups. First, we derive a *revelation principle* for our solution concept. The principle says that, without loss of generality, messages sent by the mediator can be assumed to convey recommendations for each player to choose a specific pure action, and players are obedient in the sense that they follow those recommendations (except possibly after recommendations that occur with zero probability given the history). Forges (1986) observed that the classic revelation principle extends with little change to multi-stage games. Indeed, a player who is directly informed about the action she is supposed to choose at any on-path information set is merely deprived of knowledge that is irrelevant for her decision. Our proof draws from that very same intuition, but extends the analysis from Nash equilibrium to subgame perfection. This extension requires two novel arguments.

[2]For an analysis of mechanism design with information leakage, see Häfner et al. (2025).

First and foremost, because the canonical device may be in an informational state that is limited compared to that of the original device, the canonical mechanism will be assumed to compute conditional probabilities to recover the correct probability distribution for randomized recommendations.[3] Second, to deal with subgames that are reached by (counterfactual) deviations of the canonical device, we revert to a correlated equilibrium. Along these lines, we expand the scope of the revelation principle to reflect subgame perfect rationality in a large class of repeated games with perfect monitoring.

Then, turning to conditions necessary for the implementability of a payoff profile as an MSPE, we introduce the notion of an *effective correlated minimax* value w_i^{cor} of player i. To define the concept, one starts from a correlated action profile α in the stage game, induced w.l.o.g. by a direct correlation device, and determines the highest payoff that player i could possibly realize from an alternative outcome of the stage game in which either player i herself or some other player $j \neq i$ with equivalent utility has the discretion to deviate from the recommended action.[4] Then, minimizing over all α yields w_i^{cor}. This concept relates to existing concepts in a natural way. In games with just two players, a player's effective correlated minimax value coincides with her reservation value. In games that satisfy NEU, the effective correlated minimax value coincides with the correlated minimax value (Wen, 1994). We show that a condition necessary for the implementability of a payoff profile as an MSPE is that each player obtains at least her effective correlated minimax value. The proof of that result crucially exploits the revelation principle.

[3]In Forges (1988, p. 197), the canonical mediator remains *connected* to the non-canonical device and, as we explain in Section 6.2, the same assumption is made in Forges (1986, Prop. 1). Similarly, in Sugaya and Wolitzky (2021), the mediator can send *confidential messages* to her future selves. Sticking with our transparency assumption, however, we exclude such possibilities in the present analysis.

[4]As usual, two players i and j have *equivalent utilities* if player j's utility is a positive affine transformation of player i's utility (Abreu et al., 1994).

We go on and derive sufficient conditions for implementability. An easy catch is the adaptation of the perfect folk theorem of Friedman (1971) to allow for a correlated equilibrium, rather than a Nash equilibrium, as a threat point. More interestingly, we establish a perfect folk theorem with ex post observable mediation, where the individual rationality condition takes the form that the expected discounted payoff for each player strictly exceeds her effective correlated minimax value. From this general result, alternative versions akin to the theorems of Fudenberg and Maskin (1986, Thms. 1 & 2) and Abreu et al. (1994) are obtained as corollaries. The proof of our main sufficient conditions follows established lines, with one important exception. Specifically, given mediation, the reference to "ultimately dispensable" assumptions, such as the observability of mixed actions and the access to a public randomization device (Fudenberg and Maskin, 1986, p. 536), is not needed. Given that a rigorous development of these topics remains technically demanding (Sorin, 1986; Fudenberg and Maskin, 1991; Fudenberg et al., 2007), the ability to avoid them is, in our view, an additional advantage of our approach. Overall, relative to the classic treatment, the analysis of sufficient conditions achieves a substantial simplification.

Finally, we offer a discussion, covering various examples and extensions. We show that Fudenberg and Maskin (1986, Ex. 3) is robust to the introduction of mediation and, hence, to the introduction of a public randomization device. We present an example that illustrates the need for internal confidential messages in the setup of Forges (1986). We develop an analogue to Friedman's (1971) folk theorem involving correlated threats. And we review an example due to Forges et al. (1986), showing that its implications matter also in our setup.

1.3 Related literature

Following the Nash reversion result by Friedman (1971), the seminal work on infinitely repeated games and the *folk theorem* includes Aumann and Shapley (1976) and Rubinstein (1979). The present paper contributes to the "discounted payoffs approach" (Fudenberg and Maskin, 1986; Abreu et al., 1994; Wen, 1994) that has been reviewed above.[5]

The idea of *mediation* in game theory has its origin in the study of correlated equilibrium (Aumann, 1974, 1987) and communication equilibria in one-shot games (Myerson, 1982). These fundamental concepts have been extended to extensive-form games by Forges (1986) and Myerson (1986).[6] None of those concepts, however, features ex post observable messages. Closest in spirit to the present analysis is probably Prokopovych and Smith (2004), who defined a subgame perfect correlated equilibrium. In contrast to our assumptions, however, mediator and players in their model cannot condition their choices on messages exchanged in prior periods, but only on the publicly observable history of action profiles.[7] Starting with Lehrer (1991) and Matsushima (1991), the role of communication has been studied predominantly in repeated games with private monitoring (Ben-Porath and Kahneman, 1996; Compte, 1998; Kandori and Matsushima, 1998).[8] More recently, Sugaya and Wolitzky (2017, 2018) stressed the role of mediation under perfect monitoring for the determination of the equilibrium set under imperfect private monitoring. In contrast to the present

[5]Gossner and Tomala (2020) offer a succinct survey of the literature. For more comprehensive accounts, see Sorin (1992), Mailath and Samuelson (2006), and Mertens et al. (2015).

[6]These and other extensions of correlated equilibrium to dynamic games have been surveyed by von Stengel and Forges (2008, Sec. 2.4). See also Forges (2020) and references given there.

[7]Interestingly, in their conclusion, Prokopovych and Smith (2004) mention the possibility of adding confidential messages to the public history, yet only as a means to implement public randomization effects into their model, and without further elaborating on the implications of making that assumption.

[8]See also Obara (2009), Cherry and Smith (2010), and Awaya and Krishna (2016).

study, Sugaya and Wolitzky (2017) assumed that messages are *not* ex post observable. Therefore, in the minimax phase, only the mediator and the deviator know that someone is minimaxed, while all the other players do not know that they are minimaxing someone. As a result of their assumptions that differ from ours, the usual full-dimensionality condition for the perfect monitoring folk theorem is not needed in their analysis.[9]

The *revelation principle* for Nash equilibrium in multi-stage games appeared first in Forges (1986). In contrast, Myerson (1986) assumed sequential rationality relative to conditional probability systems, but did not include a formal statement of the revelation principle. Townsend (1988) derived a revelation principle in a two-stage insurance market. His model is one of pure adverse selection with ex post unobservable messages. The mechanism sends internal messages to itself, and the second-stage report in the canonical device concerns signals obtained in both stages (i.e., there is the possibility to "confess"). His solution concept reflects optimizing behavior in both stages conditional on beliefs, with posteriors determined by Bayes' rule whenever possible. A failure of the revelation principle for the ε-perfect equilibrium was noted by Dhillon and Mertens (1996). Prokopovych and Smith (2004) obtained a revelation principle for subgame perfect correlated equilibrium. However, owing to their simpler informational set-up, in which messages exchanged in prior periods are erased from the public history, their proof is more straightforward than ours. The most comprehensive analysis of the revelation principle in multi-stage games, allowing for both adverse selection and moral hazard, is Sugaya and Wolitzky (2021). They showed that the

[9]The literature is divided regarding the *plausibility* of mediated play. Lehrer (1992, p. 175) lauded correlated equilibrium as a solution concept "more attractive than Nash equilibrium." In line with this positive assessment, the role of third parties, such as trade associations or specialized consultants, for collusion in oligopolistic markets and bid rigging in auctions has been acknowledged by a number of contributions (Aoyagi, 2005; Rahman, 2014; Ortner et al., 2024). On the other hand, Sugaya and Wolitzky (2017, p. 692) decided to "not take a position on the realism of allowing a mediator," and we take a similar stance in this paper.

communication revelation principle may fail for sequential equilibrium, whereas it holds for conditional probability perfect Bayesian equilibrium.[10] The main difference from the present study is in the information structure. Specifically, messages are ex post unobservable in their setting, whereas messages are ex post observable in our setting.[11]

The *full dimensionality* condition introduced by Fudenberg and Maskin (1986, Ex. 3) is not only relevant in the class of games considered in the present paper, but also in other classes of games such as finitely repeated games (Benoit and Krishna, 1985), OLG models (Kandori, 1992; Smith, 1995), and infinitely repeated games with random matching (Deb et al., 2020). The NEU condition is strictly less stringent (Abreu et al., 1994). More recently, Sekiguchi (2022) pointed out that the conclusion of the perfect folk theorem can be obtained in games in which all players have equivalent utilities when monitoring is both endogenous and unobservable.

The *correlated minimax value* appeared in Renault and Tomala (1998, 2011), Tomala (1999, 2009, 2013), Gossner and Hörner (2010), Laclau (2014), and Bavly and Peretz (2019). Renault and Tomala (2011, Ex. 2.7) illustrated the fact that the correlated minimax may be strictly lower than the minimax. However, their example does not feature common interests. More generally, communication allows players to depress a deviator's payoff below the minimax value (Gossner and Tomala, 2007).

1.4 Overview

The rest of the paper is structured as follows. Section 2 introduces infinitely repeated games with mediation. The revelation principle is established in Section 3. Section

[10] See also the discussion in Makris and Renou (2023).

[11] Notably, making messages ex post observable is feasible in Sugaya and Wolitzky (2021) as well. For that, the mediator would send carbon copies of the messages to herself and publish those in the subsequent period. While interesting, that observation unfortunately does not simplify our proofs.

4 derives necessary conditions, while Section 5 deals with sufficient conditions. A discussion can be found in Section 6. Section 7 concludes. Technical proofs are relegated to an Appendix.

2 Infinitely repeated games with mediation

This section prepares the main analysis. We first introduce our equilibrium concept (Section 2.1), and then derive its basic properties (Section 2.2).

2.1 Mediated subgame perfect equilibrium (MSPE)

For a finite set of players $N = \{1, 2, \ldots, n\}$, let $G = \{A_i, u_i\}_{i \in N}$ be a *stage game*, where for each player $i \in N$, A_i is the finite set of player i's actions, and $u_i : A \equiv \bigtimes_{i \in N} A_i \to \mathbb{R}$ is player i's payoff function. At any stage $t \in \{0, 1, 2, \ldots\}$, each player i chooses an action $a_i^t \in A_i$. Let M_i denote the finite, nonempty set of messages for player i. By a *history* (of length t), we mean a finite sequence

$$h \equiv h^t = (m^0, a^0; m^1, a^1; \ldots; m^{t-1}, a^{t-1}), \tag{1}$$

where, for any $\tau \in \{0, \ldots, t-1\}$, $m^\tau \in M \equiv \bigtimes_{i \in N} M_i$ is a profile of messages, and $a^\tau \in A$ is a profile of actions. The set of histories will be denoted by H, where $h^0 = \varnothing$ is the initial history. A *device* is a commitment $\mu : H \to \Delta(M)$,[12] and we denote by $\mu(\cdot \,|\, h)$ and $\mu_i(\cdot \,|\, h)$, respectively, the probability distributions over $m \equiv m^t$ and $m_i \equiv m_i^t$ at history h. A pair (h, m_i) with a history $h \in H$ and a message $m_i \in M_i$ will be called an *information set* for player i. Let I_i denote the set of player i's information sets. While these definitions already capture our main assumption, we repeat it here for clarity:

Assumption (ex post transparency of mediation). *At each stage* $t \in \{0, 1, \ldots\}$,

[12] For any finite set X, we denote by $\Delta(X)$ the set of probability distributions on X.

(i) *the device cannot condition its recommendation profile on any information other than* h^t, *and*

(ii) *each player* i*'s information consists of* h^t *and* m_i^t.

Thus, the device cannot send confidential messages to its future selves (or, in any case, not make use of them), and any private message communicated in some stage is made public by the beginning of the next stage.

A *system of beliefs* β specifies a mapping $\beta_i : I_i \to \Delta(M_{-i})$ for each $i \in N$, where $\beta_i(\cdot \mid h, m_i) \in \Delta(M_{-i})$ denotes the belief over m_{-i} at the information set (h, m_i).[13] A *behavior strategy* for player i is a mapping $\sigma_i : I_i \to \Delta(A_i)$, with $\sigma_i(\cdot \mid h, m_i)$ denoting the probability distribution over player i's actions $a_i \in A_i$ at (h, m_i). Let Σ_i be the set of behavior strategies for player i, and $\Sigma = \bigtimes_{i\in N} \Sigma_i$. Given an information set $(h, m_i) \in I_i$, the conditional distribution $\beta_i(\cdot \mid h, m_i) \in \Delta(M_{-i})$, a device $\mu : H \to \Delta(M)$, and a profile of behavior strategies $\sigma \in \Sigma$, the mediated interaction of the players induces a distribution over *conditional outcomes* $\{(\bar{m}^\tau, \bar{a}^\tau)\}_{\tau=0}^{\infty}$.[14] Let $\delta \in (0, 1)$ denote the discount factor. Then, after normalization, player i's *continuation payoff* at the information set (h, m_i) is given as

$$U_i(\sigma \mid h, m_i) = \mathbb{E}\left[(1-\delta)\sum_{\tau=0}^{\infty} \delta^\tau u_i(\bar{a}^{t+\tau})\right],$$

where we suppress the dependence on μ and β. Similarly, we let

$$U_i(\sigma) = \sum_{m_i^0 \in M_i} \mu_i(m_i^0 \mid h^0)\, U_i(\sigma \mid h^0, m_i^0).$$

denote player i's *expected discounted payoff.* This defines the *infinitely repeated game with mediation.*

[13] As usual, the index $-i$ refers to the tuple of elements with index $j \neq i$.

[14] Formally, in line with the history h, let $(\bar{m}^\tau, \bar{a}^\tau) = (m^\tau, a^\tau)$, for any $\tau \in \{0, \ldots, t-1\}$. Then, let $\bar{m}_i^t = m_i$, draw $\bar{m}_{-i}^t$ according to $\beta_i(\cdot \mid h, m_i)$, and draw $\bar{a}_j^t$ according to $\sigma_j(\cdot \mid h, \bar{m}_j^t)$, for any $j \in N$. Finally, for any stage $\tau \in \{t+1, t+2, \ldots\}$, we construct the history $\bar{h}^\tau = (\bar{m}^0, \bar{a}^0; \ldots; \bar{m}^{\tau-1}, \bar{a}^{\tau-1})$, draw $\bar{m}^\tau$ according to $\mu(\cdot \mid \bar{h}^\tau)$, and draw $\bar{a}_j^\tau$ according to $\sigma_j(\cdot \mid \bar{h}^\tau, \bar{m}_j^\tau)$, for any $j \in N$.

Definition 1. *A* mediated subgame perfect equilibrium (MSPE) *is a triple* (μ, σ, β) *such that:*

(i) $U_i(\sigma \mid h, m_i) \geq U_i(\sigma_i', \sigma_{-i} \mid h, m_i)$, *for any player* $i \in N$, *any information set* $(h, m_i) \in I_i$ *and any behavior strategy* $\sigma_i' \in \Sigma_i$*;*

(ii) $\beta_i(\cdot \mid h, m_i)$ *is derived from* $\mu(\cdot \mid h) \in \Delta(M)$ *via Bayes' rule,*[15] *for any player* $i \in N$ *and any information set* $(h, m_i) \in I_i$ *such that* $\mu_i(m_i \mid h) > 0$.

We emphasize that the set of messages M_i for each player is part of the MSPE rather than exogenously given. The following example illustrates the MSPE concept.[16]

	L	**R**
T	$6, 6$	$2, 7$
B	$7, 2$	$0, 0$

Figure 1: Using a correlated threat

Example 1. *Consider the game* G^1 *shown in Figure 1. As long as no deviation occurs, the mediator recommends* $(\mathbf{T}, \mathbf{L})$ *and* $(\mathbf{T}, \mathbf{R})$ *with equal probability* $\frac{1}{2}$, *leading to the payoff profile* $(4, 6.5)$. *If a deviation occurs, the mediator recommends* $(\mathbf{T}, \mathbf{R})$ *with probability* $\frac{4}{5}$, *and* $(\mathbf{B}, \mathbf{L})$ *with probability* $\frac{1}{5}$, *leading to the Pareto inferior payoff profile* $(3, 6)$. *Hence, for* δ *sufficiently close to one, we have an MSPE.*

Generalizing from the example above, one can derive a perfect folk theorem that uses correlated equilibria as threat points. See Section 6.

[15] I.e., $\beta_i(m_{-i} \mid h, m_i) = \mu(m_i, m_{-i} \mid h)/\mu_i(m_i \mid h)$ for all $m_{-i} \in M_{-i}$.
[16] For a richer illustration that uses private messages, see Example 3 below.

2.2 Basic properties of MSPE

Recall that, by definition, a *public correlation device* makes an independent draw from the unit interval at the start of each stage $t \in \{0, 1, \ldots\}$ (Hart, 1979). The following lemma collects a few simple consequences of our equilibrium concept.

Lemma 1. *The MSPE solution concept has the following properties:*

(i) *The one-stage deviation principle applies.*

(ii) *Any subgame-perfect public randomization equilibrium is an MSPE (but not vice versa).*

(iii) *The unconditional repetition of any correlated equilibrium of the stage game is an MSPE.*

(iv) *An MSPE exists.*

Proof. See Appendix A.1. □

By part (i), the sequential rationality requirement in Definition 1 can be conveniently checked by considering changes of a behavior strategy at single information sets only. Part (ii) is immediate for any SPE *without* public randomization, by assuming a trivial message structure (i.e., all message spaces are singletons). The fact that also any SPE *with* a public randomization device can be understood as an MSPE requires a proof, however. Indeed, while the device can send only finitely many messages, the public randomization device admits an uncountably infinite number of messages. For games in strategic form, this additional flexibility can be seen to be of little value by a reference to Carathéodory's Theorem (e.g., Rockafellar, 1970). Indeed, any distribution over payoff profiles induced by a continuum signal can be replaced, without affecting the expectation, by a distribution with finite support, hence by finitely many messages. In the Appendix, we offer a refined argument valid for infinitely repeated games. Part (iii) is again essentially immediate. However, to ensure sequential rationality at an information set that corresponds to an action a_i unused at a history h, we assume that the respective player i selects another action a_i' used with positive probability at h, and holds the same belief and chooses the same action as if she had been recommended to play a_i' rather than a_i. Finally, recalling that any stage game admits a correlated equilibrium (Hart and Schmeidler, 1989), one obtains part (iv).

3 A revelation principle

In this section, we derive a revelation principle for MSPE. As mentioned before, the revelation principle enters the derivation of necessary conditions for the folk theorem. However, given the recent findings by Sugaya and Wolitzky (2021), it might also be of independent interest.

Definition 2. *An MSPE (μ, σ, β) is called* canonical *if*

(i) $M_i = A_i$, for every $i \in N$, and

(ii) for any history $h \in H$ and any message profile $m \in M$ such that $\mu(m \,|\, h) > 0$, the mixed action $\sigma_i(\cdot \mid h, m_i) \in \Delta(A_i)$ assigns probability one to $a_i = m_i$, for any $i \in N$.

The first condition characterizes the device as *direct*, while the second condition requires all players to be *obedient*. It should be noted that Definition 2 does not require player i's obedience in response to a message m_i that arises with probability zero given a history h.[17]

[17]This avoids, in particular, that players choose strictly dominated actions after a malfunction of the device.

In a canonical MSPE, any history of length t is of the form

$$h = (\hat{a}^0, a^0; \ldots; \hat{a}^{t-1}, a^{t-1}),$$

where $\hat{a}_i^\tau \in A_i$ is the action *recommended* for player i at stage τ, and $a_i^\tau \in A_i$ is the action actually *chosen* by player i in stage τ, for any $\tau \in \{0, \ldots, t-1\}$. In a canonical MSPE, any deviation from the recommended action profile is detectable at the end of the stage where both the recommendation and the chosen profile become public information. As usual, an *outcome* is a probability distribution over paths of play in the infinitely repeated game.

Theorem 1 (Revelation Principle). *For any MSPE, there exists an outcome-equivalent MSPE that is canonical.*

Proof. See Appendix A.2. □

Thus, it is w.l.o.g. to assume that, for each player $i \in N$, the set of messages M_i corresponds in a one-to-one fashion to the set of actions A_i, and player i's behavior strategy σ_i reflects obedience in the sense explained above.

The proof of Theorem 1 proceeds conceptually in three steps, each of which constructs an MSPE (or some generalization of it) that is outcome-equivalent to the original MSPE:

Step 1 An augmented device complements the private message m_i sent to player i at stage t by including a recommended pure action $\hat{a}_i$. If $\hat{a}_i$ is consistent both with m_i and with equilibrium play in the original MSPE (conditional on reaching the history), then player i inherits the belief from the original MSPE and is obedient. If, however, $\hat{a}_i$ is not consistent in this way, then the recommendation (but not the message) is interpreted as a tremble and player i's belief and choice are as if she had received some recommendation consistent with m_i.[18] Both the device and the players ignore recommendations that have been made in earlier stages. In the augmented MSPE, all players play pure strategies.

Step 2 A further refinement of the device reports only the recommended action $\hat{a}_i$ to player i, for all $i \in N$. Thus, the message m_i is no longer communicated to player i. However, to allow keeping the law governing the distribution of messages in later periods, the profile of messages m is still communicated confidentially to all the future selves of the device. As such "internal" messages are not allowed under our definition of MSPE, this requires a strict generalization of the equilibrium concept. If the recommended action is on-path given the history, players are obedient. Indeed, this is sequentially rational because (i) players are merely deprived of information not relevant for them, i.e., they *know less* just as in Forges (1986), and (ii) players assign probability zero to trembles unless they have evidence to the opposite. If the recommended action is off-path given the history, the recommended action is interpreted as a tremble and player i's belief and choice is as if she had received some recommendation in the support of the on-path distribution.

Step 3 A canonical device does away with the confidential messages from earlier selves by speculating, in a Bayesian fashion, over the message profiles in the original MSPE. If there is no evidence that one of its earlier selves of the device trembled, i.e., if no earlier self of the device made a recommendation of probability zero under its own rule, this allows the canonical device to replicate the randomized recommendation profile without making use of the confidential messages. If, however, the device realizes that one of its earlier selves trembled, it recommends

[18] This trick has been used in the proof of Lemma 1 already.

an arbitrary correlated equilibrium $\alpha^* \in \Delta(A)$.[19] Player i is obedient unless she has evidence that the *current* device trembles, in which case she interprets the unexpected recommendation as a tremble, dealing with the situation as above by assuming belief and choice from some on-path recommendation.

These steps establish the revelation principle for MSPE. Step 1 is a simple but important "purification" argument. Step 2 is familiar from the proof of the revelation principle for extensive-form correlated equilibrium (Forges, 1986, Prop. 1). Step 3 is needed in our framework because the MSPE does not allow the mediator to keep information undisclosed across stages. The reason why Step 3 is dispensable in the analysis of extensive-form correlated equilibrium is illustrated with an example in Appendix 6.2.

Sugaya and Wolitzky (2021, Online Appendix, pp. 81-87) established a communication revelation principle for finitely repeated multi-stage games with pure moral hazard. In contrast to the present setting, however, they allow for undisclosed forms of mediation, i.e., they assume that the mediator's messages in later stages can be based on private information from earlier stages that is not accessible to the players. Specifically, in their setting, the mediator is able to condition its messages in some stage on any "fictitious message" created by herself through private randomization in some earlier stage. This is feasible under their assumptions because the mediator can send private messages to itself, i.e., can keep a hidden record of information that may serve as a basis for later recommendations. In our setting, however, the device cannot condition on private information it held in earlier stages. As a result, we cannot make direct use of Sugaya and Wolitzky (2021, Prop. 4) in our proof of Theorem 1.

[19] Alternatively, the device could be programmed to "reset" itself in such a case, by erasing its inconsistent history and starting from scratch.

4 Necessary conditions

This section derives conditions necessary for a payoff profile to be implementable as an MSPE. We first recall the definition of various minimax values (Section 4.1), and then discuss effective minimax values (Section 4.2).

4.1 Minimax values

Let $G = \{A_i, u_i\}_{i \in N}$ be an arbitrary finite game, and $i \in N$ be a player. Then, player i's *independent minimax value* is defined as

$$v_i^{\mathtt{ind}} = \min_{\alpha_{-i} \in \times_{j \neq i} \Delta(A_j)} \max_{a_i \in A_i} \mathbb{E}_{\alpha_{-i}}[u_i(a_i, a_{-i})],$$

where $\mathbb{E}_{\alpha_{-i}}[\cdot]$ denotes the expectation with respect to the profile of mixed actions α_{-i}. Similarly, player i's *correlated minimax value* is defined as

$$v_i^{\mathtt{cor}} = \min_{\alpha_{-i} \in \Delta(\times_{j \neq i} A_j)} \max_{a_i \in A_i} \mathbb{E}_{\alpha_{-i}}[u_i(a_i, a_{-i})],$$

It follows from the definitions that $v_i^{\mathtt{cor}} \leq v_i^{\mathtt{ind}}$, with equality if $n = 2$. For $n \geq 3$ players, however, the ability to correlate sanctions may allow the opponents of player i to hold her expected payoff strictly below the independent minimax payoff, so that $v_i^{\mathtt{cor}} < v_i^{\mathtt{ind}}$ becomes a possibility (Hart, 1979; Renault and Tomala, 1998).

Using these definitions, we can easily state the following necessary conditions for implementability.

Lemma 2. *Let $U_i(\sigma)$ be the expected discounted payoff resulting for player $i \in N$ from an MSPE (SPE). Then, $U_i(\sigma) \geq v_i^{\mathtt{cor}}$ ($U_i(\sigma) \geq v_i^{\mathtt{ind}}$).*

Proof. See Appendix A.3. □

While Lemma 2 offers a lower bound on the payoff implementable by an MSPE, this bound can be further improved in certain games. This will be discussed next.

4.2 Effective minimax values

Let $G = \{A_i, u_i\}_{i\in N}$ be an arbitrary finite game, as before. Recall that two players $i, j \in N$ have *equivalent utilities*, formally $i \sim j$, if there exist scalars c, d such that $d > 0$ and $u_i(a) = c + du_j(a)$ for all $a \in A$ (Abreu et al., 1994). The game G satisfies *NEU* if no pair of distinct players has equivalent utilities. Following Wen (1994), let player i's *effective independent minimax value* be defined as

$$w_i^{\texttt{ind}} = \min_{\alpha\in\times_{k\in N}\Delta(A_k)} \max_{\substack{j\in N\\ \text{s.t. } j\sim i}} \max_{\hat{a}_j\in A_j} \mathbb{E}_{\alpha_{-j}}[u_i(\hat{a}_j, a_{-j})],$$

where $\alpha_{-i} = (\alpha_j)_{j\neq i}$. It can be readily verified that this value coincides with the independent minimax value v_i^{ind} if G satisfies the NEU condition. Indeed, the maximization over the equivalence class of player i is trivial in this case.[20] If G does not satisfy the NEU condition, however, then there may be some player j equivalent to i that is able to raise her and therefore i's utility even more against any joint minimax action profile α, so that we have, in general, only $v_i^{\texttt{ind}} \leq w_i^{\texttt{ind}}$. In fact, a well-known example by Fudenberg and Maskin (1986) shows precisely that this inequality can be strict if NEU does not hold, i.e., players cannot, in general, be held down to their independent minimax values in such games.[21]

In analogy to the development above, we introduce the following variant of the correlated minimax value. Let player i's *effective correlated minimax* value be defined as

$$w_i^{\texttt{cor}} = \min_{\alpha\in\Delta(A)} \max_{\substack{j\in N\\ \text{s.t. } j\sim i}} \mathbb{E}_{\alpha_j}\left[\max_{\hat{a}_j\in A_j} \mathbb{E}_{\alpha}[u_i(\hat{a}_j, a_{-j})|\, a_j]\right],$$

[20] Further, as pointed out by Smith (1995), if $n = 2$, one may assume without loss of generality that $u_1 = u_2$, in which case $w_1^{\texttt{ind}} = w_2^{\texttt{ind}} = \max\{v_1^{\texttt{ind}}, v_2^{\texttt{ind}}\}$.

[21] Fudenberg and Maskin (1986) concluded from their example that the dimensionality assumption cannot be dispensed with in their statement of the perfect folk theorem. Strictly speaking, however, this conclusion would require to show that, *even using a public randomization device*, no player i can be held down to v_i^{ind}. In Section 6, we show that this is indeed the case for their example (but not in general).

where α_j denotes the marginal distribution of α on A_j, and the inner expectation is conditional on a_j, i.e., on the j's component of the realization of α.[22] The following lemma collects some basic properties of effective minimax values.

Lemma 3. *Let $G = \{A_i, u_i\}_{i \in N}$ be an arbitrary finite game, and $i \in N$. Then, the following holds true:*

(i) The minimax values and effective minimax values introduced above satisfy

$$\begin{array}{ccc} v_i^{\texttt{ind}} & \leq & w_i^{\texttt{ind}} \\ \rotatebox{90}{$\geq$} & & \rotatebox{90}{$\geq$} \\ v_i^{\texttt{cor}} & \leq & w_i^{\texttt{cor}}. \end{array}$$

(ii) If G satisfies NEU, then $v_i^{\texttt{ind}} = w_i^{\texttt{ind}}$ and $v_i^{\texttt{cor}} = w_i^{\texttt{cor}}$.

(iii) If $n = 2$, then $v_i^{\texttt{ind}} = v_i^{\texttt{cor}}$ and $w_i^{\texttt{ind}} = w_i^{\texttt{cor}}$.

Proof. See Appendix A.4. □

The horizontal inequalities in part (i) say that any player's effective minimax payoff is weakly larger than her minimax payoff, not only in the independent case discussed above, but also in the correlated case. The vertical inequalities say that correlation may allow to depress the minimax payoffs. All four inequalities may be strict, as will be explained later in the paper.[23] By part (ii), the horizontal inequalities are equalities if G satisfies NEU. Similarly, by part (iii), the vertical inequalities are equalities if $n = 2$.

[22] Here and below, we use the following convention: If the marginal distribution α_j assigns probability zero to some action $a_j \in A_j$, then the conditional expectation $\mathbb{E}_\alpha[u_i(\hat{a}_j, a_{-j}) | \, a_j]$ is replaced by the unconditional expectation $\mathbb{E}_{\alpha_{-j}}[u_i(\hat{a}_j, a_{-j})]$. In fact, any alternative convention delivers the same value because the outer expectation gives weight zero to such a_j.

[23] In particular, an example with $v_i^{\texttt{cor}} < w_i^{\texttt{cor}} < w_i^{\texttt{ind}}$ is presented as Lemma 4(ii).

Using the concepts introduced above, we can strengthen the conclusions of Lemma 2 as follows.

Theorem 2 (Necessary Conditions). *If U_i^* is the expected discounted payoff resulting for player i from an MSPE (SPE), then $U_i^* \geq w_i^{\texttt{cor}}$ ($U_i^* \geq w_i^{\texttt{ind}}$).*

Proof. See Appendix A.5. □

The bracketed claim, due to Wen (1994), is based on the intuition that player i's payoff cannot be depressed strictly below w_i^{ind} if there is another player j with equivalent utility who can avert this outcome. The proof of the unbracketed claim of Theorem 2 requires an additional step. Specifically, one notes that, by the revelation principle, the mediator gives, w.l.o.g., recommendations in the form of pure actions. Therefore, the stage payoff of the "most fortunate deviator" boils down to the effective correlated minimax value.

5 Sufficient conditions

Given the concepts developed in the previous sections, the derivation of a perfect folk theorem in the tradition of Wen (1994), yet with mediation, becomes straightforward. As usual, the set of *feasible* payoff vectors V is the convex hull of $\{(u_1(a), \ldots, u_n(a)) : a \in A\}$.

Theorem 3. *For any $v \in V$ such that $v_i > w_i^{\texttt{cor}}$ ($v_i > w_i^{\texttt{ind}}$) for all $i \in N$, there exists $\underline{\delta} \in (0, 1)$ such that for all $\delta \in (\underline{\delta}, 1)$, there exists an MSPE (SPE with observability of mixed actions[24]) in the infinitely repeated game in which player i's expected discounted payoff is v_i for all $i \in N$.*

[24]The MSPE does *not* assume that mixed actions are observable. However, as pointed out by Fudenberg et al. (2007, Fn. 10), the sufficient conditions in Wen (1994) crucially depend on that assumption. We therefore make this dependence explicit in the bracketed case.

Proof. See Appendix A.6. □

The proof of Theorem 3 follows the steps of the corresponding result, Wen (1994, Thm. 2), for SPE reported above in brackets. The main change in the statement of the result is that the effective independent minimax value is replaced by the effective correlated minimax.

There is an important distinction in the proof, however. To derive sufficient conditions for the folk theorem under perfect monitoring, Wen (1994) imposed simplifying assumptions that (i) players have access to a public randomization device, and that (ii) mixed actions are ex post observable. While a follow-up paper (Wen, 2002) seems to explain why these assumptions are ultimately dispensable for sufficiently patient players, the corresponding proofs are somewhat deep and spread out over multiple papers. Indeed, to replace public and private randomizations by deterministic sequences of pure actions, three techniques are employed. First, a time-averaging argument is used to represent arbitrary feasible payoff profiles as discounted averages of pure-strategy outcomes (Sorin, 1986). Second, cycling over action profiles is used to keep incentives from one-shot deviations small (Fudenberg and Maskin, 1991). Third and finally, a form of long-term accounting is used to make players truly indifferent between pure actions during the minimax phase (Fudenberg and Maskin, 1986, Sect. 6). While, to our understanding, all of these arguments are important and reflect considerations that arise similarly in real-world applications, the literature has tended to treat them as a black box, thereby obfuscating the precise relationship between necessary and sufficient conditions. Notably, allowing for mediation makes those arguments obsolete. Indeed, the reference to ultimately dispensable assumptions is absent from the proof of the non-bracketed part of Theorem 3.[25] Given Lemma 2,

[25]Mediation also circumvents pitfalls that have been identified in the use of public randomization in asymmetric settings (Olszewski, 1998) and in limits (Yamamoto, 2010).

the NEU assumption obviates the need for working with the effective minimax value. The following version of the folk theorem with mediation is an analogue to Abreu et al. (1994, Thm. 1).

Corollary 1. *Suppose that NEU holds. Then, any* $v \in V$ *such that* $v_i > v_i^{\mathtt{cor}}$ *(*$v_i > v_i^{\mathtt{ind}}$*) for all* $i \in N$ *is an MSPE (SPE) payoff profile in the infinitely repeated game when players are sufficiently patient.*

Proof. Immediate from Theorem 3 and Lemma 3(ii). □

Finally, we turn to two-player games. Mediation does not affect the set of SPE payoffs in two-player games.

Corollary 2. *Suppose that* $n = 2$*. Then, any* $v \in V$ *such that* $v_i > v_i^{\mathtt{ind}}$ *for all* $i \in N$ *is an MSPE (or SPE) payoff profile in the infinitely repeated game when players are sufficiently patient.*

Proof. See Appendix A.7. □

Forges et al. (1986) presented an example of a game showing that the strict inequality $v_i > v_i^{\mathtt{ind}}$ cannot be simply replaced by a weak inequality in the statement of the standard perfect folk theorem without losing the conclusion. The same example shows that, also in Theorem 3 and Corollary 1, the strict inequality is needed, i.e., $v_i > v_i^{\mathtt{cor}}$. See Section 6.4.

6 Discussion

This section presents additional findings related to the MSPE concept.

6.1 The example of Fudenberg and Maskin (1986)

Many of the results of this paper can be illustrated with the help of the game $G^2(r)$ shown in Figure 2. For $r = 1$, this game corresponds to Fudenberg and Maskin (1986, Ex. 3). Regardless of $r > 0$, we have $v_i^{\texttt{ind}} = v_i^{\texttt{cor}} = 0$, for each player $i \in N$. Indeed, if two players decide to choose different pure actions, then the payoff of the third player is held down to zero.

Lemma 4. *Consider the family of games $G^2(r)$.*

(i) If $r = 1$, we have $w_i^{\texttt{cor}} = w_i^{\texttt{ind}} = \frac{1}{4}$, for each $i \in N$.

(ii) For $r > 0$, $r \neq 1$, we have $w_i^{\texttt{cor}} < w_i^{\texttt{ind}}$, for each $i \in N$.

Proof. See Appendices A.8 and A.9. □

The fact that $w_i^{\texttt{ind}} = \frac{1}{4}$ for $r = 1$ is due to Fudenberg and Maskin (1986).[26] The statement that $w_i^{\texttt{cor}} = \frac{1}{4}$ for $r = 1$ is considerably stronger. It says that, even in the presence of mediation, it is not feasible to lower any player's payoff to less than $\frac{1}{4}$. This fact has a corollary, which we state separately for clarity.

F

	F	**S**
F	r, r, r	$0, 0, 0$
S	$0, 0, 0$	$0, 0, 0$

S

	F	**S**
F	$0, 0, 0$	$0, 0, 0$
S	$0, 0, 0$	$1, 1, 1$

Figure 2: The game $G^2(r)$

Corollary 3. *Fudenberg and Maskin (1986, Ex. 3) is robust with respect to the introduction of a public randomization device.*

[26] We briefly recall the proof. Let α_i denote the probability that player i chooses **F**. Then, w.l.o.g., $\alpha_1 \geq \alpha_2 \geq \alpha_3$. If $\alpha_2 \geq \frac{1}{2}$, then player 3 has a payoff of at least $\frac{1}{4}$ from choosing **F**. Otherwise, i.e., if $\alpha_2 < \frac{1}{2}$, then player 1 has a payoff of at least $\frac{1}{4}$ from choosing **S**. Conversely, if $\alpha_1 = \alpha_2 = \alpha_3 = \frac{1}{2}$, then all players obtain a payoff of $\frac{1}{4}$.

Proof. Given that mediation is more flexible than access to public randomization (cf. Lemma 1(ii)), the claim follows from Lemma 4(i). □

Thus, even if public randomization is allowed, the conclusion of the perfect folk theorem for games with $n \geq 3$ players may become invalid if the set of feasible payoff vectors is of a dimension strictly lower than n. This may be of interest because the perfect folk theorem has often been presented under the simplifying assumption that players have access to a public randomization device. In this sense, Corollary 3 closes a potentially important gap in the literature.

The proof of part (i) exploits the double symmetry of $G^2(1)$. Indeed, players have identical payoffs, while payoffs are invariant under a simultaneous renaming of the actions for all players. Another observation used in the proof is the fact that the maximum of linear functions is convex. Combining these two observations, one may compute the effective correlated minimax by restricting attention to doubly symmetric probability distributions over pure action profiles. The thereby restricted problem, however, is tractable, recovering the uniform distribution as its unique solution.[27]

Part (ii) of Lemma 4 shows that the set of implementable payoff vectors may be strictly larger if players' actions are orchestrated by a mediator. The example shows that mediation can facilitate equilibrium play and can be sustained with less patience from players because of the harsher punishments feasible via coordinated sanctions. Additional payoff profiles implementable as an MSPE but not as a SPE may require the prospect of lower payoffs for players within an equivalence class, while benefiting players outside that equivalence class. Therefore, the feasibility of $w_i^{\texttt{cor}} < w_i^{\texttt{ind}}$ for some player i established above implies that ex post observable mediation can indeed be an effective instrument in the implementation of social norms.

[27]Lemma 4 may also help to see why a dimensionality assumption is imposed in Tomala (2009, Thm. 4.4).

6.2 On the Revelation Principle for Extensive Form Correlated Equilibrium

This section elaborates on the revelation principle for extensive-form correlated equilibrium. In Forges (1986), the use of a canonical device can be assumed without loss of generality because players interacting with it receive less information, and therefore have fewer ways to deviate. The proof provided in the paper does not explicitly discuss the point, however, that the device also receives less information if replaced by a canonical version.[28] In the example below, the device proposes one of two mixed Nash equilibria in the first stage. Since the two equilibria have identical support, the information regarding the equilibrium recommendation gets partly lost when the players receive direct messages, information only about their pure choices. A canonical device *without internal memory* might then lack sufficiently precise information that would allow it to make recommendations in the same as in the non-canonical setup. The example therefore illustrates the point in Forges et al. (1986) that the multiple selves of the device need to be "connected," i.e., earlier selves can send confidential messages to their future selves. The example also illustrates the fact that the ability to communicate confidential messages is stronger than the assumption that future selves have perfect recall of messages sent in prior periods.

Example 2. *There are three players, $n = 1, 2, 3$, and two stages $t = 1, 2$. Players 1 and 2 make their choices at stage $t = 1$, player 3 makes her choice at stage $t = 2$. There are no moves of nature. Neither player 3 nor the device can observe the choices made at stage $t = 1$ by players 1 and 2. Payoffs in the resulting two-stage G^3 are*

[28] Neither does it discuss explicitly the "purification" argument, which corresponds to Step 1 in the proof of Theorem 1.

given as follows:

E_1

	L	R
T	$1,2,0$	$0,0,0$
B	$0,0,12$	$2,1,0$

E_2

	L	R
T	$2,1,0$	$0,0,12$
B	$0,0,0$	$1,2,0$

E_3

	L	R
T	$-1,-1,0$	$-1,-1,9$
B	$-1,-1,9$	$-1,-1,0$

Player 1 chooses rows, player 2 chooses columns, and player 3 chooses boxes.

Consider the following *extensive-form correlated equilibrium*:

- At stage $t = 1$, the device sends messages e_1 and e_2 with equal probability. These messages are observed only by players 1 and 2.
- Upon observing e_1, players 1 and 2 play $(\frac{1}{3}T + \frac{2}{3}B,\ \frac{2}{3}L + \frac{1}{3}R)$, which is a Nash equilibrium between the two players if player 3 chooses E_1.
- Upon observing e_2, players 1 and 2 play $(\frac{2}{3}T + \frac{1}{3}B,\ \frac{1}{3}L + \frac{2}{3}R)$, which is a Nash equilibrium between the two players if player 3 chooses E_2.
- At stage $t = 2$, the device recalls its earlier message (either e_1 or e_2) and informs player 3 correspondingly. Player 3 chooses E_1 if the message is e_1, and E_2 if the message is e_2. This is optimal for her if players 1 and 2 adhere to their strategies, because $12 \cdot \frac{4}{9} > 9 \cdot \frac{5}{9}$.

This equilibrium is not canonical. Now suppose that players 1 and 2 are merely told at stage $t = 1$ which pure strategies they are supposed to play. Then the device recalls only which of the four pairs (T, L), (T, R), (B, L), and (B, R) it recommended in stage $t = 1$. The respective probabilities are

Conditional on e_1

	L	R
T	$\frac{2}{9}$	$\frac{1}{9}$
B	$\frac{4}{9}$	$\frac{2}{9}$

Conditional on e_2

	L	R
T	$\frac{2}{9}$	$\frac{4}{9}$
B	$\frac{1}{9}$	$\frac{2}{9}$

Aggregate

	L	R
T	$\frac{2}{9}$	$\frac{5}{18}$
B	$\frac{5}{18}$	$\frac{2}{9}$

If now the device would not inform player 3 what do do, she would find it optimal to respond with E_3, breaking the equilibrium. By the definition of a canonical device, it is not allowed to give a recommendation, at stage $t = 1$ to player 3, because she moves at stage $t = 2$ only. So, also in the case of the extensive form correlated equilibrium, the device needs to apply Bayes' rule to come up with the correct conditional probabilities before recommending actions to player 3.

6.3 Correlated threats

The following result is an analogue to Friedman's (1971) perfect folk theorem for MSPE, obtained in straightforward generalization of Example 1.

Theorem 4. *Let $\alpha^* \in \Delta(A)$ be a correlated equilibrium (Nash equilibrium) in the stage game G, with payoff profile $u^* \in \mathbb{R}^n$. Then, any feasible payoff profile that strictly dominates u^* in the Pareto sense results from an MSPE (SPE) for δ sufficiently close to one.*

Proof. The bracketed claim is due to Friedman (1971). The non-bracketed claim follows from Lemma 1. □

The following example illustrates Theorem 4.

	L	**C**	**R**
T	$-1, 6$	$7, 4$	$-2, 0$
M	$1, 1$	$1, 3$	$5, 5$
B	$6, 6$	$3, 2$	$4, 9$

Figure 3: Using a correlated threat

Example 3. *Consider the game G^4 in Figure 3. Suppose that, on the equilibrium path, the mediator recommends* $(\mathbf{B}, \mathbf{R})$ *in every stage, which yields the payoff profile* $U^* = (4, 9)$. *If player 1 deviates, however, the mediator reverts to the correlated equilibrium $\alpha \in \Delta(A)$, defined by*

$$\alpha(\mathbf{T}, \mathbf{C}) = \tfrac{3}{14}, \quad \alpha(\mathbf{T}, \mathbf{R}) = \tfrac{2}{14}, \quad \alpha(\mathbf{M}, \mathbf{C}) = \tfrac{3}{14}, \quad \alpha(\mathbf{M}, \mathbf{R}) = \tfrac{6}{14},$$

which yields a continuation payoff of $\frac{25}{7} < 4$ for player 1. Hence, for $\delta \geq \frac{7}{10}$, we have an MSPE.[29]

6.4 The example of Forges et al. (1986)

Finally, we revisit an example that has been used to explain why sufficient conditions for folk theorems commonly require *strict* individual rationality for all players.

Example 4. *Consider the stage game G^5 shown in Figure 4. As before, player 1 chooses rows, player 2 columns, and player 3 matrices. The payoff profile* $v = (1, 0, 0)$ *is feasible by alternating between* $(\mathbf{F}, \mathbf{F}, \mathbf{F})$ *and* $(\mathbf{S}, \mathbf{S}, \mathbf{S})$. *Moreover,* $v_i^{\mathtt{ind}} = 0$ *for* $i \in \{1, 2, 3\}$.[30] *However, v is not implementable as an SPE (even if public randomization is available).*[31]

Turning to the possibility of mediation, the very same arguments show that $v_i^{\mathtt{cor}} = 0$ for $i \in \{1, 2, 3\}$ and that $(1, 0, 0)$ cannot be implemented as an MSPE either. The consideration of effective minimax values is obsolete as a consequence of Lemma 3(ii)

[29] U^* cannot be enforced by a Nash threat. Indeed, G^1 has a unique Nash equilibrium, namely $(\mathbf{M}, \mathbf{R})$, yielding the payoff profile $(5, 5)$. Further, U^* is not a correlated-equilibrium payoff either.

[30] Indeed, players 2 and 3 may hold player 1's payoff down to zero by choosing $a_2 = \mathbf{S}$ and $a_3 = \mathbf{F}$, while player 1 always has a nonnegative payoff. Similarly, players 1 and 3 may hold player 2's payoff down to zero by choosing $a_1 = \mathbf{F}$ and $a_3 = \mathbf{S}$, while player 2 may choose her weakly dominant action $a_2 = \mathbf{F}$ to ensure a nonnegative payoff. Finally, players 1 and 2 may hold player 3's payoff down to zero by choosing $a_1 = \mathbf{F}$ and $a_2 = \mathbf{S}$, while player 3 may choose her weakly dominant action $a_3 = \mathbf{S}$ to ensure a nonnegative payoff.

[31] Indeed, if $a_1 = a_3 = \mathbf{F}$ ($a_1 = a_2 = \mathbf{S}$) is anticipated with positive probability at some stage, then player 2 (player 3) could realize a strictly positive payoff by choosing $a_2 = \mathbf{F}$ ($a_3 = \mathbf{S}$) in the first such stage, and choosing her weakly dominant action in all subsequent stages.

F

	F	**S**
F	1, 1, −1	0, 0, 0
S	0, 0, 0	0, 0, 0

S

	F	**S**
F	0, 0, 0	0, 0, 0
S	0, 0, 0	1, −1, 1

Figure 4: The game G^5

because G^5 satisfies NEU. Thus, mediation does not obviate the need for the strictness assumption.

7 Concluding remarks

Enforcing sanctions against an individual player is more challenging when some or all of the involved players share equivalent utilities. We show that ex post observable mediation cannot, in general, resolve that challenge. However, we also identify specific scenarios in which mediation strengthens the impact of sanctions. In those cases, ex post observable mediation makes punishments more effective and better coordinated, with potential benefit to society at large.

Another contribution of this paper may be found in the simplification of standard theory. Indeed, the use of the MSPE solution concept lets us dispense with auxiliary assumptions such as the observability of mixed actions.

Finally, although we introduced the MSPE for repeated games with perfect monitoring, the concept appears to be more flexible. Specifically, we expect that the approach extends in a largely straightforward way to finitely repeated games, settings with imperfect public monitoring, the theory of coarse correlated equilibrium, and general extensive-form games. Given restrictions in space, we leave these extensions for future work.

A Appendix

This appendix contains material omitted from the body of the paper.

A.1 Proof of Lemma 1

(i) Take any triple (μ, σ, β), and suppose that, for any player $i \in N$, and any of her information sets $(h, m_i) \in I_i$, replacing σ_i by the modified behavior strategy σ_i' that is obtained from σ_i by replacing the mixed action $\sigma_i(\cdot \mid h, m_i) \in \Delta(A_i)$ by the pure action $a_i \in A_i$ does not raise player i's continuation payoff, i.e., $U_i(\sigma_i, \sigma_{-i} \mid h, m_i) \geq U_i(\sigma_i', \sigma_{-i} \mid h, m_i)$. Thus, the deviating strategy σ_i' differs from σ_i only at the information set (h, m_i). If σ_{-i} is considered exogenous for the moment, then each information set $(h, m_i) \in I_i$ starts a new subgame in the one-player game. By the one-stage deviation principle for infinite multi-stage games with discounting (Fudenberg and Tirole, 1991), σ_i induces a best response in any subgame. Therefore, σ_i induces a best response at any stage t. As the choice of the player i was arbitrary, this proves the claim.

(ii) Let $\theta^t \in [0,1]$ denote the public signal drawn in stage $t \in \{0, 1, \ldots\}$. Focus on stage $t = 0$. For any signal $\theta^0 \in [0,1]$, the SPE in the repeated game induces a payoff vector $u(\theta^0) \in \mathbb{R}^n$. Let $\mathcal{U}^0 \subseteq \mathbb{R}^n$ denote the resulting set of payoff vectors. Clearly, the SPE payoff profile satisfies $u^* = \mathbb{E}[u(\theta^0)]$. Hence, by Carathéodory's Theorem (Rockafellar, 1970, Thm. 17.1), there are signals $\theta_1^0, \ldots, \theta_{n+1}^0 \in [0,1]$ and weights $\lambda_1, \ldots, \lambda_{n+1} \geq 0$ with $\sum_{m=1}^{n+1} \lambda_m = 1$, such that

$$u_i^* = \sum_{m=1}^{n+1} \lambda_m u_i(\theta_m^0)$$

for all $i \in N$. Hence, we may replace the public randomization device at stage $t = 0$, without affecting expected payoffs or the SPE property, by an autonomous device sending public messages (i.e., identical private messages) from the finite set $M = \{1, \ldots, n+1\}$, where the probability of message m is λ_m. Applying the same reduction iteratively for any subgame starting at stage $t = 1, 2, \ldots$, we can construct, for any horizon $T \geq 1$, an infinitely repeated game $\mathcal{G}^T$ where the public randomization device has been replaced by a finite device. Consider now the limit game $\mathcal{G}^\infty$, in which this replacement has been done at *all* stages, and the corresponding limit strategy profile. We need to check that no player has an incentive to deviate. By the one-stage-deviation principle (see part (i) above), it suffices to check deviations at a single stage. But the respective continuation payoffs resulting from equilibrium play and deviation in $\mathcal{G}^\infty$ are identical to the corresponding continuation payoffs in $\mathcal{G}^T$, for T sufficiently large. Therefore, a deviation is not profitable, proving the claim.

(iii) Let $\alpha^* \in \Delta(A)$ be a correlated equilibrium in G. Then, there exists a collection of message spaces $M_1, \ldots, M_n$, a probability distribution μ^* on $M = M_1 \times \ldots \times M_n$, and a message-dependent mixed action $\alpha_i^*(m_i)$ for each player $i \in N$ such that player i's expected payoff conditional on m_i is maximized by adhering to α_i^*. Define an MSPE (μ, σ, β) as follows. Let $\mu(\cdot \mid h) = \mu^*(\cdot)$ for any $h \in H$. For any $i \in N$ and $h \in H$, define beliefs and actions as follows. For messages $m_i \in M_i$ such that $\mu_i(m_i \mid h) > 0$, define $\beta_i(\cdot \mid h, m_i)$ by Bayes rule, and let $\sigma_i(\cdot \mid h, m_i) = \alpha_i^*(\cdot \mid m_i)$. For messages m_i such that $\mu_i(m_i \mid h) = 0$, select a message m_i' with $\mu_i(m_i' \mid h) > 0$ and define $\beta_i(\cdot \mid h, m_i) = \beta_i(\cdot \mid h, m_i')$ and $\sigma_i(\cdot \mid h, m_i) = \alpha_i^*(\cdot \mid m_i')$. Then, sequential rationality for player i at any information set (h^t, m_i^t) follows directly from the optimality condition for player i in the correlated equilibrium α^*.

(iv) See the paragraph following the statement of the lemma. □

A.2 Proof of Theorem 1

Take any MSPE (μ, σ, β). As outlined in the body of the paper, we construct an outcome-equivalent canonical MSPE $(\hat{\mu}, \hat{\sigma}, \hat{\beta})$ in three steps.

Step 1. Delegating randomization to the device. Let $\tilde{M}_i = M_i \times A_i$, with typical element $\tilde{m}_i = (m_i, \hat{a}_i)$, so that $\hat{a}_i$ is the action *recommended* to player i, to be distinguished from a_i, the action actually *chosen* by player i. We denote the corresponding set of histories by $\tilde{H}$, with typical element

$$\tilde{h} \equiv \tilde{h}^t = (\tilde{m}^0, a^0; \ldots; \tilde{m}^{t-1}, a^{t-1}),$$

where $t \in \{0, 1, \ldots\}$. In the sequel, we will make tacit use of the natural projection from $\tilde{H}$ onto H, which yields a history $h = (m^0, a^0; \ldots; m^{t-1}, a^{t-1})$ in the original MSPE for any $\tilde{h}$. Define the device $\tilde{\mu} : \tilde{H} \to \Delta(\tilde{M})$ by

$$\tilde{\mu}(\tilde{m} \,|\, \tilde{h}) = \mu(m \,|\, h) \cdot \prod_{i \in N} \sigma_i(\hat{a}_i \,|\, h, m_i).$$

Intuitively, $\tilde{\mu}$ mimics the players' randomization in the original MSPE. Conditional on $\tilde{h}$, the recommendation $\hat{a}_i$ is made by the augmented device with the probability that this action is chosen in the original MSPE. Moreover, the augmented device does not condition on the profiles of actions recommended by $\tilde{m}$, which means that it "overlooks" disobedient pure-strategy choices in earlier periods that are not categorized as deviations in the original MSPE (because they are consistent with the randomized actions foreseen in equilibrium play). The corresponding set of player i's information sets is $\tilde{I}_i = \tilde{H} \times \tilde{M}_i$. We define player i's behavior strategy $\tilde{\sigma}_i : \tilde{I}_i \to \Delta(A_i)$ as follows. There are two cases:

(Case A) Suppose first that $(\tilde{h}, \tilde{m}_i) \in \tilde{I}_i$ satisfies $\sigma_i(\hat{a}_i \,|\, h,\ m_i) > 0$, i.e., the action $\hat{a}_i$ is used with positive probability at the information set (h, m_i) in the original MSPE. Then, we require that $\tilde{\sigma}_i(\tilde{h}, \tilde{m}_i) \in \Delta(A_i)$ puts all weight on $\hat{a}_i$.

(Case B) Suppose next that $\sigma_i(\hat{a}_i \,|\, h, m_i) = 0$, i.e., the action $\hat{a}_i$ is never used at the information set (h, m_i) in the original MSPE. Then, we choose some action $a_i^{\#} \in \text{supp}(\sigma_i(\cdot \mid h, m_i))$ and require that $\tilde{\sigma}_i(\tilde{h}, \tilde{m}_i)$ puts all weight on $a_i^{\#}$. Thus, $\tilde{\sigma}_i$ reflects obedience at the information set $(\tilde{h}, \tilde{m}_i)$ if the action $\hat{a}_i$ recommended by $\tilde{\mu}$ (potentially by a tremble) is consistent with equilibrium play at (h, m_i), and otherwise specifies a pure action $a_i^{\#}$ to which σ_i assigns positive probability at (h, m_i). To define beliefs, fix some information set $(\tilde{h}, \tilde{m}_i)$. Then, the belief $\widetilde{\beta}_i(\tilde{h}, \tilde{m}_i) \in \Delta(M_{-i} \times A_{-i})$ is derived from $\tilde{\mu}(\cdot \mid \tilde{h})$ via Bayes' rule if $\tilde{\mu}_i(\tilde{m}_i \,|\, \tilde{h}) > 0$,[32] and arbitrarily otherwise. By construction, at each stage t, the pair (m^t, a^t) is drawn according to the same joint distribution as under (μ, σ), hence the infinite sequence of action profiles $(a^t)_{t \geq 0}$ has the same law in $(\tilde{\mu}, \tilde{\sigma}, \tilde{\beta})$ as in the original MSPE and $U_i(\tilde{\sigma}) = U_i(\sigma)$ for all i.

It remains to check sequential rationality. Fix i and an information set $(\tilde{h}, \tilde{m}_i)$, where $\tilde{m}_i = (m_i, \hat{a}_i)$, and $\hat{a}_i$ is the action recommended for player i at history $\tilde{h}$. By sequential rationality in the original MSPE and obedience of $\tilde{\sigma}_i$, player i attains the maximal continuation payoff by being obedient if $\hat{a}_i \in \text{supp}(\sigma_i(\cdot \mid h, m_i))$. If, however, $\hat{a}_i \notin \text{supp}(\sigma_i(\cdot \mid h, m_i))$, then the maximal continuation payoff is obtained if player i chooses $a_i^{\#} \in \text{supp}(\sigma_i(\cdot \mid h, m_i))$, as specified above. This establishes sequential rationality of $\tilde{\sigma}$. Hence $(\tilde{\mu}, \tilde{\sigma}, \tilde{\beta})$ is an MSPE, outcome-equivalent to (μ, σ, β). We may therefore assume w.l.o.g. that messages are of the form $\tilde{m}_i = (m_i, \hat{a}_i)$ and that player i chooses $\hat{a}_i$ if $\tilde{\mu}_i(\tilde{m}_i \,|\, \tilde{h}) > 0$.[33]

Step 2 and 3. Dropping the uninformative parts from players' messages, and making

[32] Thus, writing $\tilde{m}_{-i} = (m_{-i}, \hat{a}_{-i})$, we have

$$\tilde{\beta}_i\big(\tilde{m}_{-i} \,|\, \tilde{h}, \tilde{m}_i\big) = \beta_i(m_{-i} \,|\, h, m_i) \prod_{j \neq i} \sigma_j(\hat{a}_j \,|\, h, m_j).$$

[33] By construction, players are obedient even if $\tilde{\mu}_i(\tilde{m}_i \,|\, \tilde{h}) = 0$ provided that $\hat{a}_i \in \text{supp}(\sigma_i(\cdot \mid h, m_i))$, but we will not make use of this fact in the sequel.

internal messages obsolete.[34] For the outcome-equivalent MSPE $(\tilde{\mu}, \tilde{\sigma}, \tilde{\beta})$ constructed in the previous stage, we have $\tilde{M}_i = M_i \times A_i$, where $\tilde{\sigma}_i(\tilde{h}, \tilde{m}_i)$ puts all weight on $\hat{a}_i$ if the augmented message $\tilde{m}_i = (m_i, \hat{a}_i)$ is sent by $\tilde{\mu}$ with positive probability at $\tilde{h}$. To define an outcome-equivalent *canonical* MSPE $(\hat{\mu}, \hat{\sigma}, \hat{\beta})$, let $\hat{M}_i = A_i$ for all $i \in N$, and let $\hat{H}$ denote the corresponding set of histories, with typical element

$$\hat{h} \equiv \hat{h}^t = (\hat{a}^0, a^0; \ldots; \hat{a}^{t-1}, a^{t-1}),$$

where $\hat{a}^\tau, a^\tau \in A$, for $\tau = 0, 1, \ldots t-1$. Take some $\hat{h}$, and put yourself into the mediator's shoes at stage t who needs to determine $\hat{\mu}(\hat{h}) \in \Delta(A)$. The mediator speculates over the sequence $(m^0, \ldots, m^{t-1})$ consistent with the observed recommendation history $\hat{h}$, or equivalently, over the history $\tilde{h} = (\tilde{m}^0, a^0; \ldots; \tilde{m}^{t-1}, a^{t-1}) \in \tilde{H}$, with $\tilde{m}^\tau = (m^\tau, \hat{a}^\tau)$, in the MSPE $(\tilde{\mu}, \tilde{\sigma}, \tilde{\beta})$ that canonically reduces to $\hat{h}$.[35] The (unnormalized) likelihood of $\tilde{h}$ conditional on $\hat{h}$ is defined recursively as

$$q(\tilde{h} \,|\, \hat{h}) = \tilde{\mu}(\tilde{m}^0 \,|\, \varnothing) \cdot \tilde{\mu}(\tilde{m}^1 \,|\, \tilde{h}^1) \cdots \tilde{\mu}(\tilde{m}^{t-1} \,|\, \tilde{h}^{t-1}),$$

where $\tilde{h}^\tau = (\tilde{m}^0, a^0; \ldots; \tilde{m}^{\tau-1}, a^{\tau-1}) \in \tilde{H}$ are the histories in the MSPE $(\tilde{\mu}, \tilde{\sigma}, \tilde{\beta})$ that arise from truncating $\tilde{h}$ at stage $\tau \in \{1, 2, \ldots, t-1\}$.[36] For any $\tilde{h} \in \tilde{H}$ that does not canonically reduce to $\hat{h}$, let $q(\tilde{h} \,|\, \hat{h}) = 0$. The mediator can now be in one of two situations:

(Case A$'$) Suppose first that $q(\tilde{h} \,|\, \hat{h})$ vanishes for all $\tilde{h} \in \tilde{H}$ that canonically reduce to $\hat{h}$. Intuitively, this means that the mediator realizes at stage t that at least one of her former selves has trembled. We will assume that the mediator then implements some correlated equilibrium $\alpha^* \in \Delta(A)$ of the stage game. Formally, $\hat{\mu}(\hat{h}) = \alpha^*$ if $\sum_{\tilde{h} \in \tilde{H}} q(\tilde{h} \,|\, \hat{h}) = 0$.

[34]The two steps are merged because the Step 2 produces a construct that is not a MSPE (because the device conditions its recommendations on confidential messages from former selves).

[35]Thus, there is another natural projection, this time from $\tilde{H}$ onto $\hat{H}$.

[36]For $t = 0$, let $q(\tilde{h}^0 \mid \hat{h}^0) = 1$.

(Case B$'$) Suppose next that $q(\tilde{h} \,|\, \hat{h}) > 0$ for at least one $\tilde{h} \in \tilde{H}$ that canonically reduces to $\hat{h}$. Then, to determine the probability of a specific recommendation $\hat{a} \in A$ at stage t, the mediator randomly draws one of such $\tilde{h} \in \tilde{H}$ according to the conditional probability distribution $q(\cdot \,|\, \hat{h}) / \sum_{\tilde{h} \in \tilde{H}} q(\tilde{h} \,|\, \hat{h})$. Thus, the probability that $\hat{a}$ is chosen in this case is

$$\hat{\mu}(\hat{a} \,|\, \hat{h}) = \frac{\sum_{\tilde{h} \in \tilde{H}} q(\tilde{h} \,|\, \hat{h}) \sum_{m \in M} \tilde{\mu}\left((m, \hat{a}) \,|\, \tilde{h}\right)}{\sum_{\tilde{h} \in \tilde{H}} q(\tilde{h} \,|\, \hat{h})},$$

where the denominator is positive under the present conditions.

Players' information sets are $\hat{I}_i = \hat{H} \times A_i$, i.e., they *know less than before* (Forges, 1986). Next, we define a system of beliefs $\hat{\beta}$ and a profile of behavior strategies $\hat{\sigma}$. Fix some player $i \in N$, a history $\hat{h} \in \hat{H}$, and a recommendation $\hat{a}_i \in A_i$. Again, we distinguish two cases:

(Case A$''$) Suppose first that, for player i, the recommendation $\hat{a}_i$ is on-path conditional on the history $\hat{h}$. Formally, this is the case if $\hat{\mu}_i(\hat{a}_i \,|\, \hat{h}) > 0$. Then, player i's belief $\hat{\beta}_i(\hat{h}, \hat{a}_i) \in \Delta(A_{-i})$ conditional on $\hat{a}_i$ may be derived via Bayes's rule from $\hat{\mu}(\cdot \,|\, \hat{h})$. Moreover, we require obedience, i.e., $\hat{\sigma}_i(\hat{h}, \hat{a}_i) \in \Delta(A_i)$ assigns probability one to $\hat{a}_i$. To see that $\hat{\sigma}_i(\hat{h}, \hat{a}_i)$ is sequentially rational, distinguish two subcases. If $\sum_{\tilde{h} \in \tilde{H}} q(\tilde{h} \mid \hat{h}) > 0$, then at the information set $(\hat{h}, \hat{a}_i)$ the continuation problem is a mixture of continuation problems $(\tilde{h}, (m_i, \hat{a}_i))$ in the MSPE $(\tilde{\mu}, \tilde{\sigma}, \tilde{\beta})$, and obedience is optimal in each such continuation problem. Hence obedience remains optimal after averaging. If instead $\sum_{\tilde{h} \in \tilde{H}} q(\tilde{h} \mid \hat{h}) = 0$, then obedience is optimal because α^* is a correlated equilibrium of the stage game and the mediator continues to implement α^* thereafter.

(Case B$''$) Suppose next that for player i, the recommendation $\hat{a}_i$ is off-path conditional on the history $\hat{h}$, i.e., $\hat{\mu}_i(\hat{a}_i \,|\, \hat{h}) = 0$. In this case, player i realizes that the device has trembled at the current stage t.[37] In this case, we interpret the unexpected recommendation $\hat{a}_i$ as a tremble of the device at stage t. Choose any action $a_i^{\#} \in A_i$ such that $\alpha_i^*(a_i^{\#}) > 0$, and define player i's belief at $(\hat{h}, \hat{a}_i)$ by $\hat{\beta}_i(a_{-i} \mid \hat{h}, \hat{a}_i) = \alpha^*(a_i^{\#}, a_{-i})/\alpha_i^*(a_i^{\#})$, for all $a_{-i} \in A_{-i}$. Finally, require that $\hat{\sigma}_i(\hat{h}, \hat{a}_i) \in \Delta(A_i)$ assigns probability one to $a_i^{\#}$. Since α^* is a correlated equilibrium of the stage game, $a_i^{\#}$ is a best reply to the conditional distribution $\alpha^*(\cdot \mid a_i^{\#})$, and hence $\hat{\sigma}_i(\hat{h}, \hat{a}_i)$ is sequentially rational at $(\hat{h}, \hat{a}_i)$. Moreover, since $\hat{\mu}_i(\hat{a}_i \mid \hat{h}) = 0$, any extension history $\hat{h}^{t+1}$ that records the off-path recommendation $\hat{a}_i$ cannot be canonically reduced from any $\tilde{h}$ with positive weight, i.e.,

$$\sum_{\tilde{h} \in \tilde{H}} q(\tilde{h} \mid \hat{h}^{t+1}) = 0.$$

Hence, by Case A′, the device sets $\hat{\mu}(\hat{h}^{t+1}) = \alpha^*$, and the same holds at all subsequent histories (since any further extension still has zero total q-weight). Moreover, from stage $t+1$ onward the mediator implements α^*, and this continuation is independent of player i's stage-t action once the off-path recommendation has occurred. Therefore, sequential rationality at $(\hat{h}, \hat{a}_i)$ reduces to a one-shot best-reply condition, which holds by the correlated-equilibrium property of α^*.

Summing up, $(\hat{\mu}, \hat{\sigma}, \hat{\beta})$ is a canonical MSPE, outcome-equivalent to (μ, σ, β) . □

A.3 Proof of Lemma 2

Only the non-bracketed claim requires a proof. At any stage t, player i, after observing her message m_i, can correctly anticipate the correlated action profile of her opponents, allowing her to ensure a stage payoff of at least $v_i^{\texttt{cor}}$. Therefore, the expected discounted payoff for player i in any MSPE is at least $v_i^{\texttt{cor}}$. □

[37] Note that the original MSPE does not necessarily suggest a belief for this situation (either because there are no unused messages conditional on h, or there are several such messages and the corresponding beliefs $\beta_i(h, \hat{a}_i)$ differ).

A.4 Proof of Lemma 3

(i) As noted before, the inequality $v_i^{\texttt{ind}} \leq w_i^{\texttt{ind}}$ is obvious. To prove $v_i^{\texttt{cor}} \leq w_i^{\texttt{cor}}$, we first remark that, since $\sim$ is reflexive,

$$\mathbb{E}_{\alpha_i}\left[\max_{\hat{a}_i \in A_i} \mathbb{E}_\alpha[u_i(\hat{a}_i, a_{-i}) | \, a_i]\right] \leq \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \mathbb{E}_{\alpha_j}\left[\max_{\hat{a}_j \in A_j} \mathbb{E}_\alpha[u_i(\hat{a}_j, a_{-j}) | \, a_j]\right],$$

for any $\alpha \in \Delta(A)$. Taking the minimum over all $\alpha \in \Delta(A)$ on both sides yields

$$\min_{\alpha \in \Delta(A)} \mathbb{E}_{\alpha_i}\left[\max_{\hat{a}_i \in A_i} \mathbb{E}_\alpha[u_i(\hat{a}_i, a_{-i}) | \, a_i]\right] \leq w_i^{\texttt{cor}}.$$

Since the maximum depends on the conditional distribution over a_{-i} only,

$$\begin{aligned} w_i^{\texttt{cor}} &\geq \min_{\alpha_i \in \Delta(A_i)} \min_{\alpha_{-i} \in \Delta(A_{-i})} \mathbb{E}_{\alpha_i}\left[\max_{\hat{a}_i \in A_i} \mathbb{E}_{\alpha_{-i}}[u_i(\hat{a}_i, a_{-i})]\right] \\ &= \min_{\alpha_{-i} \in \Delta(\times_{j \neq i} A_j)} \max_{\hat{a}_i \in A_i} \mathbb{E}_{\alpha_{-i}}[u_i(\hat{a}_i, a_{-i})] \\ &= v_i^{\texttt{cor}}. \end{aligned}$$

The claim follows. This proves the horizontal inequalities. As for the vertical inequalities, $v_i^{\texttt{cor}} \leq v_i^{\texttt{ind}}$ is again obvious. To prove that $w_i^{\texttt{cor}} \leq w_i^{\texttt{ind}}$, note that for any product distribution $\alpha \in \times_{k \in N} \Delta(A_k)$,

$$\mathbb{E}_\alpha[u_i(\hat{a}_j, a_{-j}) | \, a_j] = \mathbb{E}_{\alpha_{-j}}[u_i(\hat{a}_j, a_{-j})],$$

because a_{-j} is independent of a_j. Therefore,

$$\begin{aligned} w_i^{\texttt{cor}} &= \min_{\alpha \in \Delta(A)} \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \mathbb{E}_{\alpha_j}\left[\max_{\hat{a}_j \in A_j} \mathbb{E}_\alpha[u_i(\hat{a}_j, a_{-j}) | \, a_j]\right] \\ &\leq \min_{\alpha \in \times_{k \in N} \Delta(A_k)} \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \mathbb{E}_{\alpha_j}\left[\max_{\hat{a}_j \in A_j} \mathbb{E}_{\alpha_{-j}}[u_i(\hat{a}_j, a_{-j})]\right] \\ &= \min_{\alpha \in \times_{k \in N} \Delta(A_k)} \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \max_{a_j \in A_j} \mathbb{E}_{\alpha_{-j}}[u_i(a_j, a_{-j})] \\ &= w_i^{\texttt{ind}}, \end{aligned}$$

which proves the claim.

(ii) If G satisfies NEU, player i's equivalence class is a singleton. Hence,

$$\begin{aligned} w_i^{\texttt{ind}} &= \min_{\alpha \in \times_{k\in N} \Delta(A_k)} \max_{\substack{j\in N \\ \text{s.t. } j\sim i}} \max_{\hat{a}_j \in A_j} \mathbb{E}_{\alpha_{-j}}[u_i(\hat{a}_j, a_{-j})] \\ &= \min_{\alpha \in \times_{k\in N} \Delta(A_k)} \max_{\hat{a}_i \in A_i} \mathbb{E}_{\alpha_{-i}}[u_i(\hat{a}_i, a_{-i})] \\ &= v_i^{\texttt{ind}}. \end{aligned}$$

Hence, $v_i^{\texttt{ind}} = w_i^{\texttt{ind}}$, as claimed. Similarly,

$$\begin{aligned} w_i^{\texttt{cor}} &= \min_{\alpha \in \Delta(A)} \max_{\substack{j\in N \\ \text{s.t. } j\sim i}} \mathbb{E}_{\alpha_j}\left[\max_{\hat{a}_j \in A_j} \mathbb{E}_{\alpha}[u_i(\hat{a}_j, a_{-j}) \mid a_j]\right] \\ &= \min_{\alpha \in \Delta(A)} \mathbb{E}_{\alpha_i}\left[\max_{\hat{a}_i \in A_i} \mathbb{E}_{\alpha}[u_i(\hat{a}_i, a_{-i}) \mid a_i]\right] \\ &= v_i^{\texttt{cor}}, \end{aligned}$$

where the last equality has been shown in part (i). Hence, $v_i^{\texttt{cor}} = w_i^{\texttt{cor}}$, proving this part of the lemma.

(iii) For $n = 2$, $v_i^{\texttt{ind}} = v_i^{\texttt{cor}}$ is obvious. We claim that $w_i^{\texttt{ind}} = w_i^{\texttt{cor}}$. By part (i), it suffices to show $w_i^{\texttt{cor}} \geq w_i^{\texttt{ind}}$. Note that

$$\begin{aligned} \mathbb{E}_{\alpha_j}\left[\max_{\hat{a}_j \in A_j} \mathbb{E}_{\alpha}[u_i(\hat{a}_j, a_{-j}) \mid a_j]\right] &\geq \max_{\hat{a}_j \in A_j} \mathbb{E}_{\alpha_j}[\mathbb{E}_{\alpha}[u_i(\hat{a}_j, a_{-j}) \mid a_j]] \\ &= \max_{\hat{a}_j \in A_j} \mathbb{E}_{\alpha_{-j}}[u_i(\hat{a}_j, a_{-j})], \end{aligned}$$

for any $\alpha \in \Delta(A)$ and any $j \in N$. Taking first the maximum over all $j \in N$ such that $j \sim i$ and, subsequently, the minimum over all $\alpha \in \Delta(A)$ on both sides yields

$$w_i^{\texttt{cor}} \geq \min_{\alpha \in \Delta(A)} \max_{\substack{j\in N \\ \text{s.t. } j\sim i}} \max_{\hat{a}_j \in A_j} \mathbb{E}_{\alpha_{-j}}[u_i(\hat{a}_j, a_{-j})].$$

But, since $n = 2$, α_{-j} is just a mixed action, so that the RHS equals

$$\min_{(\alpha_1,\alpha_2) \in \Delta(A_1)\times\Delta(A_2)} \max_{\substack{j\in N \\ \text{s.t. } j\sim i}} \max_{\hat{a}_j \in A_j} \mathbb{E}_{\alpha_{-j}}[u_i(\hat{a}_j, a_{-j})] = w_i^{\texttt{ind}}.$$

This proves the last part and, hence, the lemma. □

A.5 Proof of Theorem 2

The following lemma prepares the proof of the theorem.

Lemma A.1. *The mapping*

$$\alpha \mapsto \mathbb{E}_{\alpha_i}\left[\max_{\hat{a}_i \in A_i} \mathbb{E}_\alpha[u_i(\hat{a}_i, a_{-i}) \mid a_i]\right]$$

is continuous on $\Delta(A)$*.*

Proof. Let $\mathrm{pr}(a_i \,|\, \alpha)$ denote the marginal probability of a_i given $\alpha \in \Delta(A)$. Similarly, let $\mathrm{pr}(a_i, a_{-i} \,|\, \alpha)$ denote the probability of (a_i, a_{-i}) given $\alpha \in \Delta(A)$. Then,

$$\begin{aligned}
&\mathbb{E}_{\alpha_i}\left[\max_{\hat{a}_i \in A_i} \mathbb{E}_\alpha[u_i(\hat{a}_i, a_{-i}) \mid a_i]\right] \\
&= \sum_{\substack{a_i \in A_i \\ \text{s.t. } \mathrm{pr}(a_i \,|\, \alpha) > 0}} \mathrm{pr}(a_i \,|\, \alpha) \cdot \max_{\hat{a}_i \in A_i} \sum_{a_{-i} \in A_{-i}} \frac{\mathrm{pr}(a_i, a_{-i} \,|\, \alpha)}{\mathrm{pr}(a_i \,|\, \alpha)} u_i(\hat{a}_i, a_{-i}) \\
&= \sum_{a_i \in A_i} \max_{\hat{a}_i \in A_i} \sum_{a_{-i} \in A_{-i}} \mathrm{pr}(a_i, a_{-i} \,|\, \alpha) u_i(\hat{a}_i, a_{-i}),
\end{aligned} \tag{A.1}$$

because $\mathrm{pr}(a_i \,|\, \alpha) = 0$ implies $\mathrm{pr}(a_i, a_{-i} \,|\, \alpha) = 0$ for all a_{-i}. Since the maxima on the RHS of equation (A.1) are continuous as functions of α, this proves the lemma. □

To prove the theorem, let (μ, σ, β) be an MSPE. By Theorem 1, we may assume w.l.o.g. that $M_i = A_i$, for any $i \in N$, and that σ^* corresponds to obedient execution of any recommendation expected with positive probability. Then, in stage $t = 0$, message profiles are drawn from A according to the probability distribution $\mu^0 \equiv \mu(\emptyset) \in \Delta(A)$. For any recommended action $a_i^0 \in A_i$ for player i, let $\mu_{-i}^0 \in \Delta(A_{-i})$ denote the conditional distribution of μ^0 on A_{-i} given $a_i^0 \in A_i$. Then, by player i's optimality condition at her information set corresponding to the recommended action $m_i^0 = a_i^0$ at stage $t = 0$, player i's expected continuation payoff satisfies

$$U_i(\sigma \mid h^0, m_i^0) \geq (1 - \delta) \max_{\hat{a}_i \in A_i} \mathbb{E}_{\mu_{-i}^0}[u_i(\hat{a}_i, a_{-i}) \,|\, a_i^0] + \delta L_i,$$

where L_i denotes the infimum over the set of expected discounted payoffs for player i resulting from *any* MSPE. We observe that L_i is finite because an MSPE exists by Lemma 1. Because the repeated game is stationary, the continuation at any stage t induces an MSPE in the subgame, so that any continuation payoff is at least L_i. Let $\mu_i^0 \in \Delta(A_i)$ denote the marginal distribution of μ^0 on A_i. Taking expectations over a_i^0 according to μ_i^0 shows that i's expected discounted payoff resulting from (μ, σ, β) satisfies

$$\mathbb{E}_{\sigma^*}\left[U_i\right] \geq (1-\delta)\mathbb{E}_{\mu_i^0}\left[\max_{\hat{a}_i \in A_i} \mathbb{E}_{\mu_{-i}^0}\left[u_i(\hat{a}_i, a_{-i}) \,|\, a_i^0\right]\right] + \delta L_i.$$

Consider now a sequence $\{e^\nu\}_{\nu=0}^\infty$ of MSPE with player i's equilibrium payoff converging to L_i. Without loss of generality, we replace $\{e^\nu\}_{\nu=0}^\infty$ by a subsequence such that the corresponding sequence of $\mu^{0,\nu} \in \Delta(A)$ converges as well. Then, we may replace (μ, σ, β) by e^ν in the above derivation. Taking the limit $\nu \to \infty$ and subsequently rearranging yields, in view of Lemma A.1,

$$L_i \geq \mathbb{E}_{\mu_i^0}\left[\max_{\hat{a}_i \in A_i} \mathbb{E}_{\mu_{-i}^0}\left[u_i(\hat{a}_i, a_{-i})|\, a_i^0\right]\right].$$

Let $j \in N$ such that $j \sim i$. Then, $u_j = c_{ji} \cdot u_i + d_{ji}$ for constants $c_{ji} > 0$ and $d_{ji} \in \mathbb{R}$. In particular, as the above inequality holds analogously for L_j,

$$\begin{aligned} L_i &= c_{ij} \cdot L_j + d_{ij} \\ &\geq \mathbb{E}_{\mu_j^0}\left[\max_{\hat{a}_j \in A_j} \mathbb{E}_{\mu_{-j}^0}\left[u_i(\hat{a}_j, a_{-j})|\, a_j^0\right]\right]. \end{aligned}$$

It follows that

$$L_i \geq \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \mathbb{E}_{\mu_j^0}\left[\max_{\hat{a}_j \in A_j} \mathbb{E}_{\mu_{-j}^0}\left[u_i(\hat{a}_j, a_{-j}) \ \,|\, a_j^0\right]\right].$$

Recalling that $\mu^0 \in \Delta(A)$ is just one correlated action profile shows that indeed $L_i \geq w_i^{\mathtt{cor}}$, proving the non-bracketed claim. The bracketed claim is due to Wen (1994, Thm. 1). This completes the proof. □

A.6 Proof of Theorem 3

We prepare the proof with two lemmas.

Lemma A.2. *Let $v \in V$ such that $v_i > w_i^{\texttt{cor}}$ for all $i \in N$. Then, there exist payoff vectors $v^1, \ldots, v^n \in V$ that satisfy:*

(i) $v_i^i > w_i^{\texttt{cor}}$ *for all* $i \in N$*;*

(ii) $v_i^i < v_i$ *for all* $i \in N$*;*

(iii) $v_i^i < v_i^j$ *for all* $i, j \in N$ *such that* $i \not\sim j$*.*

Proof. Select one player i from each equivalence class and denote the set of selected players as $N' \subseteq N$. Consider the modified game G' in which equivalent players collude, i.e., the set of players is N', the action sets of player $i \in N'$ is $A_i' = \bigtimes_{j \sim i} A_j$, and the payoff function of player $i \in N'$ is the same as in G. Note that NEU holds for G'. Moreover, given that $v_i > w_i^{\texttt{cor}}$, no player $i \in N'$ is indifferent over all action profiles in G'. Therefore, by Abreu et al. (1994, Lem. 1 and 2), there exist vectors $\{v_i'\}_{i \in I'}$ that satisfy $v_i'^i < v_i'^j$ for all $i, j \in N'$ such that $j \neq i$. Let now $v_i = v_i'$, where $i' \in N'$ such that $i' \sim i$. Now, following Abreu et al. (1994, p. 942), let $a_i^{\min} \in \arg\min_{a \in A} u_i(a)$ denote an action profile that minimizes player i's payoff in G. Define

$$x^i = \varepsilon(1 - \eta)u(a_i^{\min}) + \eta\varepsilon v^i + (1 - \varepsilon)v,$$

where $\varepsilon > 0$ and $\eta > 0$ are constants independent of i. Recall that $v_i > w_i^{\texttt{cor}}$. Hence, for $\varepsilon > 0$ small enough, condition (i) is satisfied. Next, for $\eta > 0$ small enough, condition (ii) is satisfied. Finally, since both $\varepsilon > 0$ and $\eta > 0$, condition (iii) holds true as well. $\square$

Lemma A.3. *Let $w^i \in \Delta(A)$ be an effective correlated minimax profile against player $i \in N$. Then,*

$$\mathbb{E}_{w_i^i}\left[\max_{\hat{a}_i \in A_i} \mathbb{E}_{w_{-i}^i}[u_i(\hat{a}_i, a_{-i}) \mid a_i]\right] \leq w_i^{\texttt{cor}}.$$

Proof. By the reflexivity of equivalence and the definition of w^i,

$$\begin{aligned}\mathbb{E}_{w_i^i}\left[\max_{\hat{a}_i \in A_i} \mathbb{E}_{w_{-i}^i}[u_i(\hat{a}_i, a_{-i}) \mid a_i]\right] &\leq \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \mathbb{E}_{w_j^i}\left[\max_{\hat{a}_j \in A_j} \mathbb{E}_{w_{-j}^i}[u_i(\hat{a}_j, a_{-j}) \mid a_j]\right] \\ &= w_i^{\texttt{cor}}.\end{aligned}$$

This proves the lemma. □

After these preparations, we adopt the normalization that $w_i^{\texttt{cor}} = 0$ for all $i \in N$. Take a feasible payoff vector $v = (v_1, v_2, \ldots, v_n) \in V$, such that $v_i > 0$. Choose some distribution $\hat{\alpha} \in \Delta(A)$ such that $\mathbb{E}_{\hat{\alpha}}[u_i(a)] = v_i$ for all $i \in N$. Let $w^i \in \Delta(A)$ be an effective correlated minimax distribution against (the equivalence class of) player i. Finally, let $\alpha^i \in \Delta(A)$ be a correlated action profile that implements the v^i identified in Lemma A.2 above. For a given discount factor $\delta \in (0, 1)$ and positive integer T, consider the following candidate MSPE.

Phase A. If, in any prior stage, all players have been obedient, then the mediator recommends playing $\hat{\alpha}$.[38]

Phase B. If, in some prior stage, precisely one player has been disobedient, player i has been the last such deviator, and this happened at most T stages before, then the mediator recommends playing w^i.

Phase C. If, in any prior stage, precisely one player has been disobedient, player i has been the last such deviator, and this happened more than T stage before, then the mediator recommends playing α^i.

[38]As usual, histories reflecting deviations by two or more players can be ignored in the analysis because they cannot be reached by a unilateral deviation.

By Lemma A.3 above, player i cannot receive a stage payoff strictly larger than $w_i^{\texttt{cor}} = 0$ by deviating from w^i. Therefore, T may be chosen sufficiently large so that a one-stage deviation benefit of any kind is strictly smaller than $T \cdot u_i(\alpha^i)$ for all $i \in N$. Then $u_i(\alpha^j) > u_i(\alpha^i) > 0$ for $i \not\sim j$ implies that there exists $\underline{\delta} < 1$ such that the following inequalities hold for all $\delta \in (\underline{\delta}, 1)$,

- $(1-\delta^T)\underline{u} + \delta^T u_i(\alpha^j) > u_i(\alpha^i)$ for all $i \not\sim j$;
- $(1-\delta)\overline{u} + \delta^{T+1} u_i(\alpha^i) < u_i(\alpha^i)$ for all $i \in N$,

where $\overline{u}$ and $\underline{u}$ are the maximum and the minimum payoff to a player in the game G, respectively. Following the steps of the proof in Wen (1994), one may now show that μ is an MSPE for $\delta \in (0,1)$ sufficiently close to one. Indeed, player i will not deviate from $\hat{\alpha}$ nor from α^j for all j. If player i deviates, her payoff is at most $(1-\delta)\overline{u} + \delta^{T+1} u_i(\alpha^i)$, which is less than $u_i(\alpha^i)$, and hence less than or equal to $u_i(\hat{\alpha})$ and $u_i(\alpha^j)$ for all j. Therefore, the proposed strategy profile is an MSPE with average equilibrium payoff v. This proves the non-bracketed claim. The bracketed claim is Wen (1994, Thm. 2). □

A.7 Proof of Corollary 2

If G satisfies NEU, then the claim follows directly from Lemma 3 and Theorem 3. If, however, G does not satisfy NEU, then there exist scalars c, d such that $d > 0$ and $u_1(a) = c + du_2(a)$ for all $a \in A$. The following proof follows Wen (1994). Let $\alpha^i = (\alpha_1^i, \alpha_2^i) \in \Delta(A_1) \times \Delta(A_2)$ be a minimax action profile against player i, i.e., $u_i(\alpha^i) = \mathbb{E}_{\alpha^i}[u_i(a)] = v_i^{\texttt{ind}}$. Let $v = (v_1, v_2)$ be feasible with $v_i > v_i^{\texttt{ind}} = u_i(\alpha^i)$. Note that $v_1 = c + dv_2 > c + du_2(\alpha^2) = u_1(\alpha^2)$. Therefore, $v_i > \max\{u_i(\alpha^1), u_i(\alpha^2)\}$ and

$$\left\{ (v_1, v_2) \middle|\, v_i > v_i^{\texttt{ind}} \right\} = \left\{ (v_1, v_2) \middle|\, v_i > \max\left\{u_i(\alpha^1), u_i(\alpha^2)\right\} \right\}.$$

Consider the mutual minimax action profile $\alpha^* = (\alpha_1^2, \alpha_2^1)$. By the definition of $w_i^{\texttt{ind}}$, we have

$$\begin{aligned} w_i^{\texttt{ind}} &\leq \max_{j\in\{1,2\}} \max_{\alpha_j\in\Delta(A_j)} \mathbb{E}_{\alpha_j,\alpha^*_{-j}}[u_i(a_j, a_{-j})] \\ &= \max\left\{ \max_{\alpha_1\in\Delta(A_1)} \mathbb{E}_{\alpha_1,\alpha_2^1}[u_i(a_1,a_2)], \max_{\alpha_2\in\Delta(A_2)} \mathbb{E}_{\alpha_1^2,\alpha_2}[u_i(a_1,a_2)] \right\} \\ &= \max\left\{u_i(\alpha^1), u_i(\alpha^2)\right\} \end{aligned}$$

WLOG, assume that $u_1(\alpha^1) \geq u_1(\alpha^2)$.[39] From Lemma 3, we have

$$w_1^{\texttt{ind}} \geq v_1^{\texttt{ind}} = \max\left\{u_1(\alpha^1), u_1(\alpha^2)\right\}$$

When there are two players and they have equivalent utilities,

$$w_i^{\texttt{cor}} = w_i^{\texttt{ind}} \geq \max\left\{u_i(\alpha^1), u_i(\alpha^2)\right\}.$$

Then, $w_i^{\texttt{cor}} = \max\{u_i(\alpha^1), u_i(\alpha^2)\}$. Hence,

$$\left\{ (v_1, v_2) \middle| \, v_i > v_i^{\texttt{ind}} \right\} = \left\{ (v_1, v_2) \middle| \, v_i > w_i^{\texttt{cor}} \right\}$$

and the claim follows from Theorem 3.[40] □

A.8 Proof of Lemma 4(i)

Take some distribution $\alpha \in \Delta(A)$ induced by the mediator. Let $p_{a_1a_2a_3} \in [0,1]$ denote the corresponding probability that the triple $(a_1, a_2, a_3) \in A$ is chosen. Then,

$$\sum_{a_1\in A_1}\sum_{a_2\in A_2}\sum_{a_3\in A_3} p_{a_1a_2a_3} = 1. \tag{A.2}$$

We use $\text{pr}\{a_1 = \mathbf{F}\}$ to denote the marginal probability of player 1 receiving the recommendation to choose $\mathbf{F}$, i.e. $\text{pr}\{a_1 = \mathbf{F}\} = \sum_{a_2\in A_2}\sum_{a_3\in A_3} p_{\mathbf{F}a_2a_3}$. Analogous

[39] Otherwise, we have $u_2(\alpha^2) \geq u_2(\alpha^1)$ and $w_2^{\texttt{ind}} = \max\left\{u_2(\alpha^1), u_2(\alpha^2)\right\}$.

[40] Here comes an alternative proof. For $n = 2$, we have $v_i^{\texttt{cor}} = v_i^{\texttt{ind}}$ from Lemma 3. Moreover, by Fudenberg and Maskin (1986, Thm. 1), v is implementable as a subgame perfect equilibrium with public randomization and observable mixed actions for any δ sufficiently close to one. As any payoff vector resulting from such equilibrium can be implemented by an MSPE, the claim follows.

notation will be used for players 2 and 3, as well as for action **S**. For each player i, when recommended action a_i by the mediator, i forms a conditional expectation about actions taken by others when they follow the recommendations from the mediator. Since players only obtain nonzero payoff when they coordinate on either **F** or **S**, we can focus on the probability of coordination. For example, when recommended to choose $a_1 = \mathbf{F}$, player 1 infers that player 2 and player 3 coordinate on **F** with probability $\frac{p_{\mathbf{FFF}}}{\mathrm{pr}\{a_1=\mathbf{F}\}}$, and on **S** with probability $\frac{p_{\mathbf{FSS}}}{\mathrm{pr}\{a_1=\mathbf{F}\}}$. Player 1 optimizes by taking $\hat{a}_1 \in A_1 = \{\mathbf{F}, \mathbf{S}\}$, i.e.

$$\max_{\hat{a}_1 \in A_1} \ \mathbb{E}_\alpha[u_1(\hat{a}_1, a_{-1})|\, a_1 = \mathbf{F}] = \max\{\frac{p_{\mathbf{FFF}}}{\mathrm{pr}\{a_1 = \mathbf{F}\}}, \frac{p_{\mathbf{FSS}}}{\mathrm{pr}\{a_1 = \mathbf{F}\}}\}.$$

Similarly, when $a_1 = \mathbf{S}$, we have

$$\max_{\hat{a}_1 \in A_1} \ \mathbb{E}_\alpha[u_1(\hat{a}_1, a_{-1})|\, a_1 = \mathbf{S}] = \max\{\frac{p_{\mathbf{SFF}}}{\mathrm{pr}\{a_1 = \mathbf{S}\}}, \frac{p_{\mathbf{SSS}}}{\mathrm{pr}\{a_1 = \mathbf{S}\}}\}.$$

Taking the expectation over a_1 yields

$$\begin{aligned}
&\mathbb{E}_{\alpha_1}\left[\max_{\hat{a}_1 \in A_1} \ \mathbb{E}_\alpha[u_1(\hat{a}_1, a_{-1})|\, a_1]\right] \\
&= \mathrm{pr}\{a_1 = \mathbf{F}\} \max_{\hat{a}_1 \in A_1} \ \mathbb{E}_\alpha[u_1(\hat{a}_1, a_{-1})|\, a_1 = \mathbf{F}] \\
&\ + \mathrm{pr}\{a_1 = \mathbf{S}\} \max_{\hat{a}_1 \in A_1} \ \mathbb{E}_\alpha[u_1(\hat{a}_1, a_{-1})|\, a_1 = \mathbf{S}],
\end{aligned}$$

or equivalently,

$$\mathbb{E}_{\alpha_1}\left[\max_{\hat{a}_1 \in A_1} \ \mathbb{E}_\alpha[u_1(\hat{a}_1, a_{-1})|\, a_1]\right] = \max\{p_{\mathbf{FFF}}, p_{\mathbf{FSS}}\} + \max\{p_{\mathbf{SFF}}, p_{\mathbf{SSS}}\}.$$

This relationship holds even if $\mathrm{pr}\{a_1 = \mathbf{F}\} = 0$ or $\mathrm{pr}\{a_1 = \mathbf{S}\} = 0$, i.e., if conditional probabilities are not well-defined. Moreover, the same formula extends analogously to players 2 and 3. Hence, given that all players are in the same equivalence class,

$$\Phi(\alpha) \equiv \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \ \mathbb{E}_{\alpha_j}\left[\max_{\hat{a}_j \in A_j} \ \mathbb{E}_\alpha[u_i(\hat{a}_j, a_{-j})|\, a_j]\right]$$

$$= \max \left\{ \begin{array}{c} \max\{p_{\mathbf{FFF}}, p_{\mathbf{FSS}}\} + \max\{p_{\mathbf{SFF}}, p_{\mathbf{SSS}}\}, \\ \max\{p_{\mathbf{FFF}}, p_{\mathbf{SFS}}\} + \max\{p_{\mathbf{FSF}}, p_{\mathbf{SSS}}\}, \\ \max\{p_{\mathbf{FFF}}, p_{\mathbf{SSF}}\} + \max\{p_{\mathbf{FFS}}, p_{\mathbf{SSS}}\} \end{array} \right\}. \tag{A.3}$$

Recall that player i's effective correlated minimax value $w_i^{\mathtt{cor}}$ is the minimum of $\Phi(\alpha)$ by choice of $\alpha = \{p_{a_1\,a_2\,a_3}\}_{a_1 \in A_1, a_2 \in A_2, a_3 \in A_3}$ subject to (A.2). Note further that the RHS of (A.3) is symmetric both with respect to arbitrary permutations of the set of players and with respect to a simultaneous swap of **F** and **S** in all three action spaces. Therefore, starting from any solution α, we obtain a potentially new solution $\alpha^{\pi,\chi}$ for any permutation π of the player set $N = \{1, 2, 3\}$ and any permutation χ of the common action set $A_1 = A_2 = A_3$. Note, finally, that the RHS of (A.3) is a convex function of α. Hence, the set of minima is convex. In particular,

$$\overline{\alpha} = \frac{1}{12} \sum_{\pi} \sum_{\chi} \alpha^{\pi,\chi}$$

is a doubly symmetric solution, whose corresponding probabilities $\overline{p}_{a_1 a_2 a_3}$ satisfy

$$\overline{p}_{\mathbf{FFF}} = \overline{p}_{\mathbf{SSS}}$$

$$\overline{p}_{\mathbf{FFS}} = \overline{p}_{\mathbf{FSF}} = \overline{p}_{\mathbf{SFF}} = \overline{p}_{\mathbf{FSS}} = \overline{p}_{\mathbf{SFS}} = \overline{p}_{\mathbf{SSF}}.$$

But then, $\Phi(\overline{\alpha}) = 2 \max\{\overline{p}_{\mathbf{FFF}}, \overline{p}_{\mathbf{FSS}}\}$, which indeed has $\frac{1}{4}$ as its minimum. □

A.9 Proof of Lemma 4(ii)

We start with the following auxiliary result.

Lemma A.4. *Let $r > 0$. Then, the following holds true for $G^2(r)$.*

(i) $w_i^{\mathtt{ind}} = \frac{r}{(1+\sqrt{r})^2}$;

(ii) if $r \in (0, 1)$, then $w_i^{\mathtt{cor}} = \frac{r}{r+3}$;

(iii) if $r > 1$, then $w_i^{\mathtt{cor}} = \frac{r}{3r+1}$.

Proof. (i) Suppose that players 1, 2, and 3, choose action $\mathbf{F}$ with respective probabilities $x = \alpha_1(\mathbf{F}), y = \alpha_2(\mathbf{F})$, and $z = \alpha_3(\mathbf{F})$. Then, player 1's expected payoff, if best responding, is

$$\overline{u}(y, z) = \max\{yzr, (1 - y)(1 - z)\}. \tag{A.4}$$

Analogous relationships hold for players 2 and 3. Thus, the "most fortunate deviator" can ensure an expected payoff of

$$U(x, y, z) = \max\{\overline{u}(y, z), \overline{u}(x, z), \overline{u}(x, y)\}. \tag{A.5}$$

Combining equations (A.4) and (A.5) yields

$$\begin{aligned} &U(x, y, z) \\ &= \max \left\{ \begin{array}{c} \max\{xyr, (1 - x)(1 - y)\}, \\ \max\{yzr, (1 - y)(1 - z)\}, \\ \max\{zxr, (1 - z)(1 - x)\} \end{array} \right\} \\ &= \max\{xyr, yzr, zxr, (1 - x)(1 - y), (1 - y)(1 - z), (1 - z)(1 - x)\}. \end{aligned}$$

Since $U(x, y, z)$ is symmetric with respect to arbitrary permutations of the variables x, y, and z, we may assume without loss of generality that the triple (x, y, z) minimizing $U(x, y, z)$ satisfies $x \geq y \geq z$. But then, $xy \geq \max\{yz, zx\}$ and

$$(1 - y)(1 - z) \geq \max\{(1 - x)(1 - y), (1 - z)(1 - x)\}.$$

Therefore,

$$\begin{aligned} w_i^{\texttt{ind}} &= \min_{\substack{(x,y,z)\in[0,1]^3 \\ x\geq y\geq z}} U(x, y, z) \\ &= \min_{\substack{(x,y,z)\in[0,1]^3 \\ x\geq y\geq z}} \max\{xyr, (1 - y)(1 - z)\}. \end{aligned}$$

Moreover, given that xyr is monotone increasing in x, and $(1 - y)(1 - z)$ does not depend on x, we may assume w.l.o.g. that $x = y$. For similar reasons, we may assume w.l.o.g. that $y = z$. Hence,

$$w_i^{\texttt{ind}} = \min_{y \in [0,1]} \max\{y^2 r, (1-y)^2\}.$$

One notes that $y^2 r$ is strictly increasing in y, while $(1-y)^2$ is strictly declining in y. The minimum is, therefore, obtained at $x = y = z = 1/(\sqrt{r}+1)$. The claim follows.

(ii) Take some distribution $\alpha \in \Delta(A)$. In analogy to the derivation in the proof of Lemma 4, one shows that

$$\begin{aligned}\Phi(\alpha) &\equiv \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \mathbb{E}_{\alpha_j}\left[\max_{\hat{a}_j \in A_j} \mathbb{E}_\alpha[u_i(\hat{a}_j, a_{-j}) | \, a_j]\right] \\ &= \max \left\{ \begin{array}{c} \max\{rp_{\mathbf{FFF}}, p_{\mathbf{FSS}}\} + \max\{rp_{\mathbf{SFF}}, p_{\mathbf{SSS}}\}, \\ \max\{rp_{\mathbf{FFF}}, p_{\mathbf{SFS}}\} + \max\{rp_{\mathbf{FSF}}, p_{\mathbf{SSS}}\}, \\ \max\{rp_{\mathbf{FFF}}, p_{\mathbf{SSF}}\} + \max\{rp_{\mathbf{FFS}}, p_{\mathbf{SSS}}\} \end{array} \right\}. \end{aligned} \tag{A.6}$$

Then, player i's effective correlated minimax value is given as $w_i^{\texttt{cor}} = \min_{\alpha \in \Delta(A)} \Phi(\alpha)$. The RHS of (A.6) is symmetric with respect to arbitrary permutations of the set of players.[41] Therefore, starting from any solution α, we obtain a potentially new solution α^π for any permutation π of the player set $I = \{1, 2, 3\}$. Note, finally, that the RHS of (A.6) is a convex function of α. Hence, the set of minima is convex. In particular,

$$\overline{\alpha} = \frac{1}{6} \sum_\pi \alpha^\pi$$

is a symmetric solution, whose corresponding probabilities $\overline{p}_{a_1 a_2 a_3}$ satisfy conditions

$$\overline{p}_{\mathbf{FFS}} = \overline{p}_{\mathbf{FSF}} = \overline{p}_{\mathbf{SFF}} \equiv q_1,$$

$$\overline{p}_{\mathbf{FSS}} = \overline{p}_{\mathbf{SFS}} = \overline{p}_{\mathbf{SSF}} \equiv q_2.$$

Similarly, let $q_0 = \overline{p}_{\mathbf{FFF}}$ and $q_3 = \overline{p}_{\mathbf{SSS}}$. Then,

$$w_i^{\texttt{cor}} = \min_{\alpha \in \Delta(A)} \Phi(\alpha) = \min_{\substack{q_0 + 3q_1 + 3q_2 + q_3 = 1 \\ q_0, q_1, q_2, q_3 \in [0,1]}} \{\max\{rq_0, q_2\} + \max\{rq_1, q_3\}\}.$$

[41] Since $r \neq 1$, however, there is no symmetry with respect to a swap of **F** and **S**.

At the minimum, we must have $rq_0 = q_2$ and $rq_1 = q_3$. Otherwise, equalizing rq_0 and q_2 or rq_1 and q_3 lowers the objective function. Hence,

$$w_i^{\texttt{cor}} = \min_{\substack{(3r+1)q_0+(3+r)q_3=1 \\ q_0,q_3\in[0,1]}} r(q_0 + q_3).$$

If $r \in (0, 1)$, then we have $w_i^{\texttt{cor}} = \frac{r}{r+3}$ at $q_0 = 0$ and $q_3 = \frac{1}{3+r}$.

(iii) Similarly, if $r > 1$, then $w_i^{\texttt{cor}} = \frac{r}{3r+1}$ at $q_0 = \frac{1}{3r+1}$ and $q_3 = 0$. This proves the lemma. □

After these preparations, we can now prove Lemma 4(ii). Suppose first that $r \in (0, 1)$. From Lemma A.4, we have $w_i^{\texttt{ind}} > w_i^{\texttt{cor}}$ if and only if

$$\frac{r}{(1+\sqrt{r})^2} > \frac{r}{3+r}.$$

Rearranging, this inequality is seen to be equivalent to $r < 1$. The case where $r > 1$ is analogous and, therefore, omitted. Hence, $w_i^{\texttt{ind}} > w_i^{\texttt{cor}}$, unless $r = 1$. The inequality $w_i^{\texttt{cor}} > v_i^{\texttt{cor}}$ is immediate from $v_i^{\texttt{cor}} = 0$ and part (i). This proves the lemma. □

References

Abreu, Dilip, Prajit K. Dutta, and Lones Smith (1994) "The Folk Theorem for Repeated Games: A Neu Condition," *Econometrica*, 62 (4), 939–948, https://www.jstor.org/stable/2951739. 2, 5, 6, 7, 9, 19, 23, 42

Aoyagi, Masaki (2005) "Collusion through Mediated Communication in Repeated Games with Imperfect Private Monitoring," *Economic Theory*, 25 (2), 455–475, https://doi.org/10.1007/s00199-003-0436-6. 8

Aumann, Robert J. (1974) "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, 1 (1), 67–96, https://www.sciencedirect.com/science/article/pii/0304406874900378. 3, 7

——— (1987) "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55 (1), 1–18, https://www.jstor.org/stable/1911154. 3, 7

Aumann, Robert J. and Lloyd S. Shapley (1976) "Long-Term Competition: A Game-Theoretic Analysis," https://doi.org/10.1007/978-1-4612-2648-2_1, Mimeo, The Hebrew University, 1976. Reprinted in *Essays in Game Theory: In Honor of Michael Maschler*, edited by N. Megiddo, Springer, New York, 1994, pp. 1–15. 7

Awaya, Yu and Vijay Krishna (2016) "On Communication and Collusion," *American Economic Review*, 106 (2), 285–315, https://www.aeaweb.org/articles?id=10.1257/aer.20141469. 7

Bavly, Gilad and Ron Peretz (2019) "Limits of Correlation in Repeated Games with Bounded Memory," *Games and Economic Behavior*, 115, 131–145, https://www.sciencedirect.com/science/article/pii/S0899825619300387. 9

Ben-Porath, Elchanan and Michael Kahneman (1996) "Communication in Repeated Games with Private Monitoring," *Journal of Economic Theory*, 70 (2), 281–297, https://www.sciencedirect.com/science/article/pii/S0022053196900903. 7

Benoit, Jean-Pierre and Vijay Krishna (1985) "Finitely Repeated Games," *Econometrica*, 53 (4), 905–922, https://www.jstor.org/stable/1912660. 9

Cherry, Josh and Lones Smith (2010) "Unattainable Payoffs for Repeated Games of Private Monitoring," SSRN Working Paper 1427602, https://papers.ssrn.com/abstract=1427602. 7

Compte, Olivier (1998) "Communication in Repeated Games with Imperfect Private

Monitoring," *Econometrica*, 66 (3), 597–626, https://www.jstor.org/stable/2998576. 7

Deb, Joyee, Takuo Sugaya, and Alexander Wolitzky (2020) "The Folk Theorem in Repeated Games with Anonymous Random Matching," *Econometrica*, 88 (3), 917–964, https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA16680. 9

Dhillon, Amrita and Jean-François Mertens (1996) "Perfect Correlated Equilibria," *Journal of Economic Theory*, 68 (2), 279–302, 10.1006/jeth.1996.0018. 8

Forges, Françoise (1986) "An Approach to Communication Equilibria," *Econometrica*, 54 (6), 1375–1385, https://www.jstor.org/stable/1914304. 3, 4, 5, 6, 7, 8, 16, 17, 26, 36

——— (1988) "Communication Equilibria in Repeated Games with Incomplete Information," *Mathematics of Operations Research*, 13 (2), 191–231, https://www.jstor.org/stable/3689820. 5

——— (2020) "Correlated equilibria and communication in games," *Complex Social and Behavioral Systems: Game Theory and Agent-Based Models*, 107–118, https://link.springer.com/content/pdf/10.1007/978-1-0716-0368-0_103.pdf. 7

Forges, Françoise, Jean-François Mertens, and Abraham Neyman (1986) "A Counterexample to the Folk Theorem with Discounting," *Economics Letters*, 20 (1), 7, https://www.sciencedirect.com/science/article/pii/0165176586900698. 6, 23, 26, 29

Friedman, James W. (1971) "A Non-cooperative Equilibrium for Supergames," *The Review of Economic Studies*, 38 (1), 1–12, https://doi.org/10.2307/2296617. 6, 7, 28

Fudenberg, Drew, David K. Levine, and Satoru Takahashi (2007) "Perfect Public Equilibrium When Players Are Patient," *Games and Economic Behavior*, 61 (1), 27–49, https://www.sciencedirect.com/science/article/pii/S0899825606001825. 6, 21

Fudenberg, Drew and Eric Maskin (1986) "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54 (3), 533–554, https://www.jstor.org/stable/1911307. 2, 6, 7, 9, 19, 22, 24, 45

——— (1991) "On the Dispensability of Public Randomization in Discounted Repeated Games," *Journal of Economic Theory*, 53 (2), 428–438, https://www.sciencedirect.com/science/article/pii/002205319190163X. 6, 22

Fudenberg, Drew and Jean Tirole (1991) *Game Theory*: MIT Press, https://mitpress.mit.edu/9780262061414/game-theory/. 31

Gossner, Olivier and Johannes Hörner (2010) "When Is the Lowest Equilibrium Payoff in a Repeated Game Equal to the Min Max Payoff?" *Journal of Economic Theory*, 145 (1), 63–84, https://www.sciencedirect.com/science/article/pii/S0022053109001045. 9

Gossner, Olivier and Tristan Tomala (2007) "Secret Correlation in Repeated Games with Imperfect Monitoring," *Mathematics of Operations Research*, 32 (2), 413–424, https://www.jstor.org/stable/25151797. 9

——— (2020) "Repeated Games with Complete Information," in Sotomayor, Marilda, David Pérez-Castrillo, and Filippo Castiglione eds. *Complex Social and Behavioral Systems*, 139–156, New York, NY: Springer, http://link.springer.com/10.1007/978-1-0716-0368-0_451. 7

Häfner, Samuel, Marek Pycia, and Haoyuan Zeng (2025) "Mechanism Design with Information Leakage," 10.48550/arXiv.2511.00715. 4

Hart, Sergiu (1979) "Lecture Notes: Special Topics in Game Theory," Mimeo, Institute for Mathematical Studies in the Social Sciences, Department of Economics, Stanford University, http://www.ma.huji.ac.il/~hart/abs/lect-notes.html. 3, 13, 18

Hart, Sergiu and David Schmeidler (1989) "Existence of Correlated Equilibria," *Mathematics of Operations Research*, 14 (1), 18–25, 10.1287/moor.14.1.18. 14

Kandori, Michihiro (1992) "Repeated Games Played by Overlapping Generations of Players," *The Review of Economic Studies*, 59 (1), 81–92, https://doi.org/10.2307/2297926. 9

Kandori, Michihiro and Hitoshi Matsushima (1998) "Private Observation, Communication and Collusion," *Econometrica*, 66 (3), 627–652, https://www.jstor.org/stable/2998577. 7

Laclau, M. (2014) "Communication in Repeated Network Games with Imperfect Monitoring," *Games and Economic Behavior*, 87, 136–160, https://www.sciencedirect.com/science/article/pii/S0899825614000864. 9

Lehrer, Ehud (1991) "Internal Correlation in Repeated Games," *International Journal of Game Theory*, 19 (4), 431–456, http://link.springer.com/10.1007/BF01766432. 7

——— (1992) "Correlated Equilibria in Two-Player Repeated Games with Nonobservable Actions," *Mathematics of Operations Research*, 17 (1), 175–199, 10.1287/moor.17.1.175. 8

Mailath, George J. and Larry Samuelson (2006) *Repeated Games and Reputations: Long-Run Relationships*: Oxford University Press, https://academic.oup.com/book/26071. 7

Makris, Miltiadis and Ludovic Renou (2023) "Information Design in Multistage Games," *Theoretical Economics*, 18 (4), 1475–1509, https://econtheory.org/ojs/index.php/te/article/view/4769. 9

Matsushima, Hitoshi (1991) "On the Theory of Repeated Games with Private Information: Part II: Revelation through Communication," *Economics Letters*, 35 (3), 257–261, https://www.sciencedirect.com/science/article/pii/016517659190140G. 7

Mertens, Jean-François, Sylvain Sorin, and Shmuel Zamir (2015) *Repeated Games*, Econometric Society Monographs, Cambridge: Cambridge University Press, https://www.cambridge.org/core/books/repeated-games/6E933D5A1B0C1116E90EC7333FF72841. 7

Myerson, Roger B. (1982) "Optimal Coordination Mechanisms in Generalized Principal–Agent Problems," *Journal of Mathematical Economics*, 10 (1), 67–81, https://www.sciencedirect.com/science/article/pii/0304406882900064. 7

——— (1986) "Multistage Games with Communication," *Econometrica*, 54 (2), 323, https://www.jstor.org/stable/1913154?origin=crossref. 3, 7, 8

Obara, Ichiro (2009) "Folk Theorem with Communication," *Journal of Economic Theory*, 144 (1), 120–134, https://www.sciencedirect.com/science/article/pii/S0022053108000525. 7

Olszewski, Wojciech (1998) "Perfect Folk Theorems. Does Public Randomization

Matter?" *International Journal of Game Theory*, 27 (1), 147–156, https://doi.org/10.1007/BF01243200. 22

Ortner, Juan, Takuo Sugaya, and Alexander Wolitzky (2024) "Mediated Collusion," *Journal of Political Economy*, 132 (4), 1247–1289, https://www.journals.uchicago.edu/doi/full/10.1086/727710. 8

Prokopovych, Pavlo and Lones Smith (2004) "Subgame Perfect Correlated Equilibria in Repeated Games," Working Paper, Department of Economics, University of Michigan, https://core.ac.uk/download/pdf/7358381.pdf. 7, 8

Rahman, David (2014) "The Power of Communication," *American Economic Review*, 104 (11), 3737–3751, https://www.aeaweb.org/articles?id=10.1257/aer.104.11.3737. 8

Renault, Jérôme and Tristan Tomala (1998) "Repeated Proximity Games," *International Journal of Game Theory*, 27 (4), 539–559, https://doi.org/10.1007/s001820050089. 3, 9, 18

——— (2011) "General Properties of Long-Run Supergames," *Dynamic Games and Applications*, 1 (2), 319–350, https://doi.org/10.1007/s13235-011-0018-3. 9

Rockafellar, R. Tyrrell (1970) *Convex Analysis*, 28 of Princeton Mathematical Series, Princeton, NJ: Princeton University Press. 13, 31

Rubinstein, Ariel (1979) "Equilibrium in Supergames with the Overtaking Criterion," *Journal of Economic Theory*, 21 (1), 1–9, https://www.sciencedirect.com/science/article/pii/0022053179900024. 7

Sekiguchi, Tadashi (2022) "A Folk Theorem for Infinitely Repeated Games with Equivalent Payoffs under Optional Monitoring," Working Paper, Kyoto Institute of Economic Research, Kyoto University, https://www.game.kier.kyoto-u.ac.jp/2022/Sekiguchi.pdf. 9

Selten, Reinhard (1965) “Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit: Teil I: Bestimmung des dynamischen Preisgleichgewichts,” *Zeitschrift für die gesamte Staatswissenschaft / Journal of Institutional and Theoretical Economics*, 121 (2), 301–324, https://www.jstor.org/stable/40748884. 4

Smith, Lones (1995) “Necessary and Sufficient Conditions for the Perfect Finite Horizon Folk Theorem,” *Econometrica*, 63 (2), 425–430, https://www.jstor.org/stable/2951632. 9, 19

Sorin, Sylvain (1986) “On Repeated Games with Complete Information,” *Mathematics of Operations Research*, 11 (1), 147–160, https://www.jstor.org/stable/3690060. 6, 22

——— (1992) “Repeated Games with Complete Information,” in *Handbook of Game Theory with Economic Applications*, Vol. 1, Chap. 4, pp. 71–107: Elsevier, https://www.sciencedirect.com/science/article/pii/S1574000505800074. 3, 7

von Stengel, Bernhard and Françoise Forges (2008) “Extensive-Form Correlated Equilibrium: Definition and Computational Complexity,” *Mathematics of Operations Research*, 33 (4), 1002–1022, 10.1287/moor.1080.0340. 7

Sugaya, Takuo and Alexander Wolitzky (2017) “Bounding Equilibrium Payoffs in Repeated Games with Private Monitoring,” *Theoretical Economics*, 12 (2), 691–729, http://doi.wiley.com/10.3982/TE2270. 7, 8

——— (2018) “Bounding Payoffs in Repeated Games with Private Monitoring:

N-Player Games," *Journal of Economic Theory*, 175, 58–87, https://www.sciencedirect.com/science/article/pii/S0022053118300097. 7

——— (2021) "The Revelation Principle in Multistage Games," *The Review of Economic Studies*, 88 (3), 1503–1540, https://doi.org/10.1093/restud/rdaa041. 3, 5, 8, 9, 14, 17

Tomala, Tristan (1999) "Nash Equilibria of Repeated Games with Observable Payoff Vectors," *Games and Economic Behavior*, 28 (2), 310–324, https://www.sciencedirect.com/science/article/pii/S0899825698907011. 9

——— (2009) "Perfect Communication Equilibria in Repeated Games with Imperfect Monitoring," *Games and Economic Behavior*, 67 (2), 682–694, https://linkinghub.elsevier.com/retrieve/pii/S0899825609000323. 9, 25

——— (2013) "Belief-Free Communication Equilibria in Repeated Games," *Mathematics of Operations Research*, 38 (4), 617–637, https://www.jstor.org/stable/24540874. 9

Townsend, Robert M. (1988) "Information Constrained Insurance: The Revelation Principle Extended," *Journal of Monetary Economics*, 21 (2), 411–450, https://www.sciencedirect.com/science/article/pii/0304393288900384. 8

Wen, Quan (1994) "The 'Folk Theorem' for Repeated Games with Complete Information," *Econometrica*, 62 (4), 949–954, https://www.jstor.org/stable/2951740. 2, 5, 7, 19, 21, 22, 41, 44

——— (2002) "A Folk Theorem for Repeated Sequential Games," *The Review of Economic Studies*, 69 (2), 493–512, 10.1111/1467-937X.00214. 22

Yamamoto, Yuichi (2010) "The Use of Public Randomization in Discounted Repeated Games," *International Journal of Game Theory*, 39 (3), 431–443, http://link.springer.com/10.1007/s00182-009-0219-9. 22