Mediated Subgame Perfect Equilibrium*

Christian Ewerhart† Haoyuan Zeng ‡

January 22, 2026

Abstract. This paper studies mediation in infinitely repeated games with perfect monitoring. In departure from the literature, we assume that all private messages and internal records are publicly revealed at the end of each stage. We call the resulting equilibrium concept mediated subgame perfect equilibrium (MSPE). It is shown that the revelation principle holds. We introduce an effective correlated minimax value, which can be conveniently determined as the solution of a linear program, and use it to derive necessary and sufficient conditions for the implementability of payoffs under an MSPE. These conditions are standard for two-player games with a sufficient degree of patience but are, in general, strictly more permissive. Examples illustrate the impact of effective correlated minimax profiles and the subtle role of internal records.

Keywords. Infinitely repeated games; mediation; revelation principle; perfect folk theorem; effective minimax value; correlated equilibrium; threat points

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

*) First version: December 15, 2025. 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 showed 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 comprehensive version of the folk theorem that applies to all finite stage games, i.e., even if they do not satisfy NEU.

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 growing interest in understanding the implications of assuming that a central entity provides players with additional information or recommendations for actions. By a *mediator*, we mean an autonomous device that communicates confidential messages to the players at each stage, just before decisions are made (Aumann, 1974, 1987; Forges, 1986; Myerson, 1986). It has been known for some time that mediation may have an impact on the set of implementable payoffs in dynamic games. Specifically, coordination between agents who enforce punishments may strictly lower the utility level to which an opponent can be held (Hart, 1979). 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; Renault and Tomala, 2011). However, to our knowledge, no corresponding extension of the perfect folk theorem has been established. 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 sequential equilibrium concept.

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 of Fudenberg and Maskin (1986). Our main innovation is the assumption that all private messages and internal records are *ex-post observable*, i.e., publicly revealed at the end of each stage. Moreover, we assume that the mediator is perfectly informed about all prior actions and messages, i.e., no input is required. The resulting solution concept, *mediated subgame perfect equilibrium (MSPE)*, requires Bayes consistency of beliefs and sequential rationality at all information sets. This approach turns out to be very well-behaved. In particular, the revelation principle applies. Moreover, there is an easily computable *effective correlated minimax value* that allows characterizing the set of payoffs implementable as an MSPE for sufficiently patient players, in perfect analogy to Abreu et al. (1994) and Wen (1994). We show that mediation can strictly expand the set of implementable payoffs not only in games with more than two players but also in games with impatient players. We further obtain a straightforward extension of Friedman's (1971) theorem allowing for correlated equilibria as threat points.

While our key hypothesis, 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 transparent allows us to largely rely on subgame perfect equilibrium (Selten, 1965) for intuition and to keep explicit modeling of beliefs (Kreps and Wilson, 1982) at a minimum. Indeed, all private information in our setup is temporary and related to the mediator's messages that are disclosed at the end of each stage. Second, our assumption that the autonomous device must be "publicly auditable" also has practical appeal: for instance, this principle is commonly required in public administration, law enforcement, procurement, and other hierarchical settings, where it has the potential to improve accountability, reduce corruption, and foster trust.[1] Ex-post observability may, however, also result 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* (Theorem 1) for our solution concept. The principle says that, without loss of generality, messages sent by the mediator convey recommendations for each player to choose a specific pure action, while players are obedient in the sense that they follow those recommendations (except after recommendations that are sent with zero probability given the history). Forges (1986) observed that the classic revelation principle extends with little change to multi-stage games for the Nash equilibrium concept.

[1] However, we do not examine incentives for the device to keep its promise.

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

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. That very same intuition, formalized by Sugaya and Wolitzky (2021), is the basis of our proof as well. However, the extension requires two novel arguments. First and foremost, because the canonical device in our setup may be in an informational state that is more limited than that of the original device (because it cannot keep internal records), the device computes conditional probabilities to recover the correct probability distribution for randomized recommendations.[3] Second, to deal with subgames that are reached by any (counterfactual) malfunction of the canonical device, we "reset" the stage counter after such events. Along these lines, we expand the scope of the revelation principle to reflect sequential rationality in a large class of repeated games with ex-post transparent mediation.

Next, to prepare the derivation of necessary and sufficient conditions for implementability of a payoff profile, we introduce the notion of an *effective correlated minimax* value $w_i^{\texttt{cor}}$ (for player i). To define the concept, one starts from a correlated action profile α in the stage game and determines the highest payoff that player i could possibly realize in the stage game if 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 $w_i^{\texttt{cor}}$. This concept relates to existing concepts in a natural way. In games that satisfy NEU, $w_i^{\texttt{cor}}$ coincides with player i's effective independent minimax value (Wen, 1994).[5] In games with just two players,

[3]In Forges (1988, p. 197), the canonical mediator remains *connected* to the non-canonical device and, as we explain in Section 6, the same assumption is made in the Forges (1986, Prop. 1). Similarly, in Sugaya and Wolitzky (2021), the mediator can send *confidential messages* to its future selves. Sticking with our transparency assumption, however, we will prohibit such possibilities.

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

[5]Allowing for games that violate NEU allows us to keep a comprehensive perspective. Moreover, such games arise naturally when dealing with cartels, criminal organizations, oligarchic elites, and terrorist organizations, for instance.

$w_i^{\texttt{cor}}$ coincides with player i's correlated minimax value. A major plus compared to the independent minimax value and the effective independent minimax value is the fact that $w_i^{\texttt{cor}}$ can be conveniently determined by solving a simple *linear program*.

We go on to derive necessary and sufficient conditions for the implementability of a payoff profile as an MSPE. We show that a necessary condition for the implementability of a payoff profile as an MSPE is that each player obtains at least her effective correlated minimax value (Theorem 2). Notably, the proof of this result crucially exploits the revelation principle. As for sufficient conditions, we establish a general perfect folk theorem with ex-post observable mediation (Theorem 3), where the individual rationality condition again takes the form that the expected payoff for each player strictly exceeds her effective correlated minimax value. From this general result, we obtain corollaries for games that satisfy NEU and for two-player games. These corollaries are natural analogues of results in Abreu et al. (1994) and Fudenberg and Maskin (1986, Thms. 1 & 2). The proofs of our sufficient conditions follow established lines, with one exception. Specifically, given mediation, the reference to "ultimately dispensable" assumptions, such as the observability of mixed actions, is not needed. Given that the fully rigorous treatment of such issues is quite demanding (Sorin, 1986; Fudenberg and Maskin, 1991) and dispensing with observability of mixed strategies in games that violate NEU is a delicate issue (Fudenberg et al., 2007, Sec. 4.2 and Fn. 10), the ability to avoid these complications is, in our view, an additional advantage of our approach.

Finally, we offer a discussion, covering various examples and extensions. We show that an example due to Fudenberg and Maskin (1986, Ex. 3) is robust with respect to the introduction of mediation and, hence, to the availability of a public randomization device. We develop an analogue of Friedman's (1971) folk theorem allowing for correlated threats. We show that, for a fixed discount factor, the MSPE can be more permissive than the subgame perfect equilibrium with public randomization even in the two-player case. We review an example due to Forges et al. (1986), relating to the case of weakly individually rational payoff profiles, and note that its implications also matter in our setup. And we present an example that sheds light on the role of internal messages in the canonical autonomous device constructed by Forges (1986).

1.3 Related literature

Following the Nash reversion result by Friedman (1971), 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" that has been reviewed above.[6]

The idea of *mediation* in game theory has its origins 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).[7] None of those concepts, however, assumes ex-post observable messages. The closest in spirit to the present paper is Prokopovych and Smith (2004), who defined the concept of *subgame perfect correlated equilibrium*. Like the present analysis, players are assumed to condition their choices on the history of action profiles and their current private recommendation. In contrast to our assumptions, however, the mediator and players in their model cannot condition on messages exchanged in prior periods.[8] Starting with Lehrer (1991) and

[6] Gossner and Tomala (2020) surveyed the literature. See also Sorin (1992), Mailath and Samuelson (2006), and Mertens et al. (2015).

[7] 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 Forges and Ray (2024). Solan and Vieille (2002) studied mediation in stochastic games.

[8] Interestingly, in their conclusion, Prokopovych and Smith (2004) mention the possibility of adding confidential messages to the 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.

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).[9] 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 Sugaya and Wolitzky (2017), the mediator can maintain an undisclosed state across periods and coordinate play via private recommendations without making the punishment phase common knowledge among the players.[10]

The *revelation principle* for Nash equilibrium in multi-stage games appeared first in Forges (1986). By comparison, Myerson (1986) assumed sequential rationality relative to conditional probability systems. 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 conditioned on beliefs, with posteriors determined by Bayes' rule whenever possible. A failure of the revelation principle for trembling-hand 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 setup, in which messages exchanged in prior periods are erased from the history, their proof is more straight-

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

[10]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."

forward 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 communication revelation principle may fail for sequential equilibrium, whereas it holds for conditional probability perfect Bayesian equilibrium. In their setting, however, the mediator can send confidential internal messages to its future self.[11] Makris and Renou (2023) generalized the concept of correlated Bayesian equilibrium (Bergemann and Morris, 2016) to multi-stage games.

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 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). 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). It is well-known that communication allows players to depress a deviator's payoff below the minimax value (Gossner and Tomala, 2007). Renault and Tomala (2011, Ex. 2.7 & Thm. 3.1) illustrated the fact that the correlated minimax value may be strictly lower than the independent minimax value and derived a Nash folk theorem allowing for correlation. With heterogeneous discounting, the effective minimax value is not a lower bound for subgame perfect payoffs (Chen, 2008).

[11] Interestingly, our informational setup can be replicated in Sugaya and Wolitzky (2021). For this, the mediator would send carbon copies of all private messages to its future self in the next stage, who would then disclose those messages to all players. Cf. Section 7.

Relative to this literature, the novelty of the MSPE concept is the combination of private messaging, sequential rationality, and ex-post transparency.

1.4 Overview

The remainder of this paper is structured as follows. Section 2 introduces infinitely repeated games with mediation. The revelation principle is established in Section 3. Section 4 introduces the effective correlated minimax value, while Section 5 derives necessary and sufficient conditions for MSPE implementability. An example-based 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 the *stage game*, where for each player $i \in N$, A_i is the finite set of i's actions, and $u_i : A \equiv \bigtimes_{i \in N} A_i \to \mathbb{R}$ is i's payoff function. Let M_i be a finite set of messages for player i. At any stage $t \in \{0, 1, 2, \ldots\}$, each player i receives a message $m_i^t \in M_i$ and subsequently chooses an action $a_i^t \in A_i$. A *history* of length t is a finite sequence $h^t = (m^0, a^0; \ldots; m^{t-1}, a^{t-1})$, 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 of length t will be denoted by H^t, where $h^0 = \varnothing$ is the empty history. Let $H = \bigcup_{t=0}^{\infty} H^t$ be the set of histories of any length, with typical element h. A *device* is a mapping $\mu : H \to \Delta(M)$,[12] and we

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

denote by $\mu(\cdot \mid h)$ and $\mu_i(\cdot \mid h) = \sum_{m_{-i} \in M_{-i}} \mu(\cdot, m_{-i} \mid h)$, respectively, the probability distributions over m and m_i at history h.[13] Any 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.

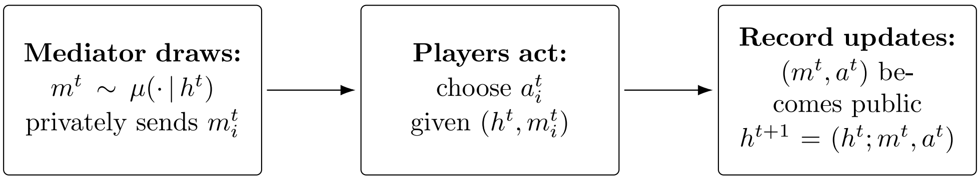

Figure 1: Within-stage timing.

The timing within each stage is illustrated in Figure 1. At the beginning of stage t, the device draws a message profile m^t according to $\mu(\cdot \mid h^t)$ and privately sends m_i^t to each player i. Players then choose actions a_i^t simultaneously and independently. Finally, (m^t, a^t) is made public and is appended to the history to form h^{t+1}.

For clarity, we repeat our main assumption:

Assumption 1. At each stage $t \in \{0, 1, \ldots\}$,

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

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

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

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 distribution over player i's action $a_i \in A_i$ at $(h, m_i) \in I_i$. Let Σ_i be

[13] As usual, the index $-i$ refers to the elements with index $j \neq i$. Thus, $M_{-i} = \bigtimes_{j \neq i} M_j$, $m_{-i} = (m_j)_{j \neq i}$, etc.

the set of player i's behavior strategies, and let $\Sigma = \bigtimes_{i\in N} \Sigma_i$. Given a device μ and a profile of behavior strategies $\sigma = (\sigma_1, \ldots, \sigma_n) \in \Sigma$, the *outcome* is the induced probability distribution over infinite paths $\{a^t\}_{t=0}^{\infty}$. Let $\delta \in (0,1)$ denote the discount factor. Then, after normalization, player i's *expected payoff* is defined as

$$U_i(\sigma) = \mathbb{E}\left[(1-\delta)\sum_{t=0}^{\infty} \delta^t u_i(a^t)\right], \tag{1}$$

where we suppress the dependence on μ for convenience.

Let $(h^t, m_i^t) \in I_i$ be an information set for some player $i \in N$. We denote by $\beta_i(\cdot \mid h^t, m_i^t) \in \Delta(M_{-i})$ player i's *belief* over other players' messages m_{-i} at (h^t, m_i^t). A *system of beliefs* β specifies a mapping $\beta_i : I_i \to \Delta(M_{-i})$ for each $i \in N$. We note that player i's belief $\beta_i(\cdot \mid h^t, m_i^t)$, the profile of behavior strategies σ, and the device μ jointly induce a probability distribution over infinite continuation paths $\{a^{t+\tau}\}_{\tau=0}^{\infty}$.[14] Player i's *continuation payoff* at (h^t, m_i^t) is then defined as

$$U_i(\sigma \mid h^t, m_i^t) = \mathbb{E}\left[(1-\delta)\sum_{\tau=0}^{\infty} \delta^\tau u_i(a^{t+\tau})\right], \tag{2}$$

where we again suppress the dependence on μ and β_i.

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

(i) (*Bayes consistency*) Players' beliefs are derived from μ via Bayes' rule whenever possible, i.e.,

$$\beta_i(m_{-i} \mid h, m_i) = \frac{\mu(m_i, m_{-i} \mid h)}{\mu_i(m_i \mid h)}$$

for any player $i \in N$, any information set $(h, m_i) \in I_i$ such that $\mu_i(m_i \mid h) > 0$, and any profile of messages $m_{-i} \in M_{-i}$.

[14] For details, see Appendix A.1.

(ii) (*Sequential Rationality*) Given the device μ and a system of beliefs β, players' choices satisfy

$$U_i(\sigma \,|\, h, m_i) \geq U_i(\sigma_i', \sigma_{-i} \,|\, h, m_i),$$

for any player $i \in N$, any information set $(h, m_i) \in I_i$, and any deviation $\sigma_i' \in \Sigma_i$.

We emphasize that the set of messages M_i for each player i is part of the description of the MSPE rather than being exogenously given.

We will say that a message m_i (a message profile m) at history $h \in H$ is *regular* if $\mu_i(m_i \,|\, h) > 0$ (if $\mu(m \,|\, h) > 0$).[15] Similarly, a history $h^t \in H^t$ is regular if $\mu(m^\tau \,|\, h^\tau) > 0$ for all $\tau \in \{0, \ldots, t-1\}$. Finally, we call an information set $(h, m_i) \in I_i$ regular if both h and m_i are regular.

Keeping the device μ fixed, MSPE is outcome-equivalent to sequential equilibrium (SE). To see this, recall our assumption that the device is not a player but merely executes moves of Nature. Therefore, SE requires Bayes consistency only at regular information sets. However, the actions taken at irregular information sets do not matter for sequential rationality at regular information sets, nor for the outcome of the game. Hence, for a fixed device, the respective sets of outcomes and payoffs coincide between SE and MSPE.

The following example illustrates Definition 1.

Example 1. Consider the game G^1 shown in Figure 2. Player 1 chooses rows, player 2 chooses columns, and player 3 chooses matrices. As long as player 1 is obedient, $\mu(\cdot \,|\, h)$ gives equal weight to $(\mathbf{T}, \mathbf{R}, \mathbf{B_1})$ and $(\mathbf{B}, \mathbf{L}, \mathbf{B_1})$, leading to the payoff profile $(1, 1, 1)$. If player 1 deviates, $\mu(\cdot \,|\, h)$ gives equal weight to $(\mathbf{T}, \mathbf{L}, \mathbf{B_1})$ and $(\mathbf{T}, \mathbf{R}, \mathbf{B_2})$,

[15] Intuitively, an irregular message m_i corresponds to a malfunction of the device that is immediately detectable for player i. In contrast, irregular message profiles need not be immediately detectable for all players, but will be publicly evident at the beginning of the next stage.

leading to the Pareto-inferior payoff profile $(0, 0, \frac{1}{2})$. For $\delta \geq \frac{4}{5}$, this describes an MSPE.[16]

B_1

	L	**R**
T	$2, 0, 1$	$4, 0, 0$
B	$-2, 2, 2$	$4, 0, 0$

B_2

	L	**R**
T	$4, 0, 0$	$-2, 0, 0$
B	$4, 0, 0$	$2, 0, 0$

Figure 2: The game G^1.

The example shows that an MSPE can be more permissive than subgame perfect equilibrium (SPE). Indeed, since player 1's independent minimax value is 2, the Pareto optimal payoff profile $(1, 1, 1)$ cannot be enforced by an SPE. A public randomization device does not help because bounding player 1's payoff to 1 requires that players 2 and 3 receive private messages. One can also check that $(1, 1, 1)$ is not a correlated equilibrium payoff of the one-shot game.

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 basic properties of our equilibrium concept.

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

(i) The one-shot deviation principle applies.

(ii) Any payoff vector implementable by an SPE with public randomization for some $\delta \in (0, 1)$, even with observable mixed actions, is implementable by some MSPE for the same discount factor.[17]

[16] A perfect folk theorem using correlated equilibrium as a threat point can be found in Section 6.
[17] As we saw in Example 1, the converse is not true in general.

(iii) *For any* (μ, β, σ) *satisfying Bayes consistency and sequential rationality at all regular information sets, there is an outcome-equivalent MSPE* $(\tilde{\mu}, \tilde{\beta}, \tilde{\sigma})$ *using the same message spaces.*

(iv) *The unconditional repetition of any correlated equilibrium of the stage game* G *is an MSPE for any* $\delta \in (0, 1)$*. In particular, an MSPE exists.*

Proof. See Appendix A.2. □

By part (i), sequential rationality can be checked by considering deviations at individual 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 public randomization, at least in terms of payoffs, can be replicated by an MSPE requires a proof, however. Indeed, while M is finite, public randomization admits a continuum of signals. For games in strategic form, this additional flexibility is without consequence for payoffs by an application of 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 over finitely many messages. In the Appendix, we offer a refined argument valid for infinitely repeated games. Part (iii) is particularly useful in applications because it obviates not only the specification of beliefs in the description of an MSPE, but also the discussion of irregularities. Since any observable deviation of the device will be public information from the subsequent stage onward, it suffices to "reset" the stage counter. This, in turn, allows defining a player's belief after a zero-probability message arbitrarily and lets her choose a best response in the stage game to the opponents' correlated action profile induced by her belief and the opponents' behavior strategies. Part (iv) is now immediate. The set of messages for a player may even be chosen as the support of the corresponding marginal of the correlated equilibrium, which directly circumvents the need to define beliefs at information sets that the mediator invokes with probability zero given the history. Recalling that the stage game admits a correlated equilibrium (Hart and Schmeidler, 1989), one obtains existence.

3 The revelation principle

In this section, we derive a revelation principle for MSPE. This result will, in particular, be useful for the derivation of necessary conditions for the folk theorem. However, given the recent findings by Sugaya and Wolitzky (2021), our version of the revelation principle might also be of independent interest.

Definition 2. An MSPE (μ, β, σ) 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 $m_i \in M_i$ such that $\mu_i(m_i \,|\, h) > 0$, the mixed action $\sigma_i(\cdot \,|\, h, m_i) \in \Delta(A_i)$ gives full weight 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 is sent with probability zero given the history h.[18]

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

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

where $\hat{a}^\tau \in A$ is the profile of actions *recommended* to the players, and $a^\tau \in A$ is the profile of actions actually *chosen* by the players at stage τ, for any $\tau \in \{0, \ldots, t -$

[18] This avoids, in particular, players choosing strictly dominated actions after a malfunction of the device.

1}. In particular, in line with the principle of transparency underlying the MSPE concept, any deviation from the recommended action profile is detectable at the end of the stage, since both the recommendations and the chosen actions become public information.

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

Proof. See Appendix A.3. □

Thus, we may always assume w.l.o.g. that, for all players $i \in N$, 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.

For mathematical convenience, the proof of Theorem 1 defines the canonical device directly from the original MSPE. We found it instructive to decompose the construction *conceptually* into three steps:

Step 1 First, to move the randomization from the players to the device, the private message m_i for player i is complemented by a *recommended* action $\hat{a}_i \in A_i$. To ensure sequential rationality and outcome equivalence, the purified device sends the message profile $(m^t, \hat{a}^t)$ at a given history with the same probability with which the combination (m^t, a^t), with *chosen* action profile $a^t = \hat{a}^t$, arises at the corresponding history in the original MSPE.

Step 2 Next, a *direct* device is constructed that sends only the recommended action $\hat{a}_i$ to each player i. To preserve the law of future recommendations, this device sends (in a strict generalization of the MSPE concept) the original message profile m^t to its future selves.[19] Obedience in the sense defined above is sequentially rational in this setup because (i) players know less, as in Forges (1986),

[19] Alternatively, the mediator keeps an "internal record" of the original messages.

and (ii) obedience remains optimal even if regular and irregular information sets are merged (because players, looking ahead, assign probability zero to any malfunction of the device).

Step 3 Finally, the internal messages are dropped. At any regular history, the *canonical* device, fully informed about the action profiles recommended and chosen at all prior stages, but ignorant of the internal messages sent by the direct device, determines the conditional distribution over those messages and uses it to replicate the randomized recommendation profile of the direct device.

Step 1 is simple but important. Indeed, without prior purification, the player does not know her own action (Aumann, 1987, p. 11), so that the information structures compared by the revelation principle, viz. abstract messages vs. actions, would not be comparable. The probability consideration in this step is a standard element of the communication revelation principle (Forges, 1986), and it is formalized in Sugaya and Wolitzky (2021, App. B).

Step 3 is needed in our framework because the MSPE does not allow the mediator to keep information undisclosed across stages. Sugaya and Wolitzky (2021, Online Appendix, pp. 81-87) established a communication revelation principle for finitely repeated multi-stage games with pure moral hazard. Their framework allows for undisclosed forms of mediation, however. In particular, the mediator's messages in later stages can be based on private information from earlier stages that is not accessible to the players. In our setting, however, the device cannot condition on private information that 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.[20]

[20]The subtle role of internal records for the revelation principle for extensive-form correlated equilibrium (Forges, 1986) is illustrated with an example in Section 6.

4 Minimax values

We first review various minimax concepts that have been proposed in the literature (Section 4.1), and then introduce our notion of effective correlated minimax value (Section 4.2).

4.1 Minimax concepts in the literature

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

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

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

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

It is clear from the definitions that $v_i^{\texttt{cor}} \leq v_i^{\texttt{ind}}$, with equality if $n = 2$. For $n \geq 3$ players, however, the ability to correlate mixed actions may allow the opponents of player i to hold i's expected payoff strictly below the independent minimax payoff, so that $v_i^{\texttt{cor}} < v_i^{\texttt{ind}}$ becomes a possibility (Hart, 1979; Renault and Tomala, 1998, 2011). This possibility is illustrated also in Example 1.

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). Further, 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

$$w_i^{\texttt{ind}} = \min_{\alpha \in \times_{k \in N} \Delta(A_k)} \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \max_{a'_j \in A_j} \mathbb{E}_{\alpha_{-j}}[u_i(a'_j, a_{-j})],$$

where $\alpha_{-j} \in \bigtimes_{k \neq j} \Delta(A_k)$ is derived from the vector α by eliminating the j-th component.

The value $w_i^{\mathtt{ind}}$ coincides with the independent minimax value $v_i^{\mathtt{ind}}$ if G satisfies the NEU condition. Indeed, the maximization over the equivalence class of player i is trivial in this case. If G does not satisfy NEU, however, then there may be some player $j \sim i$ that is able to raise her own and therefore i's utility against the joint minimax action profile α, so that, in general, only $v_i^{\mathtt{ind}} \leq w_i^{\mathtt{ind}}$ can be ascertained.[21] In fact, a well-known example by Fudenberg and Maskin (1986, Ex. 3) shows 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.[22]

4.2 Effective correlated minimax value

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

$$w_i^{\mathtt{cor}} = \min_{\alpha \in \Delta(A)} \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \mathbb{E}_{\alpha_j} \left[\max_{a'_j \in A_j} \mathbb{E}_\alpha [u_i(a'_j, a_{-j}) \big| \, a_j] \right],$$

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 realization of α_j.[23] The conditioning on a_j in the inner

[21]For players in an equivalence class, it can be assumed w.l.o.g. that their payoff functions are identical. In the case $n = 2$, this has the notable implication that $w_1^{\mathtt{ind}} = w_2^{\mathtt{ind}} = \max\{v_1^{\mathtt{ind}}, v_2^{\mathtt{ind}}\}$ (Smith, 1995, p. 427). However, no analogous representation is feasible for $n \geq 3$, as follows from Fudenberg and Maskin (1986, Ex. 3).

[22]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 showing that, *even using a public randomization device*, no player i can be held down to $v_i^{\mathtt{ind}}$. In Section 6, we show that this is indeed the case for their example (but not in general).

[23]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 $\mathbb{E}_\alpha[u_i(a'_j, a_{-j}) \big| \, a_j]$ is replaced by the unconditional expectation $\mathbb{E}_\alpha[u_i(a'_j, a_{-j})]$. In fact, any alternative convention delivers the same value for $w_i^{\mathtt{cor}}$ because the outer expectation gives zero weight to such a_j. See the proof of Lemma A.1 in the Appendix.

expectation should be intuitive because player j observes a_j before choosing a'_j. At the same time, the members of player i's equivalence class cannot coordinate their responses, so that unilateral deviations are considered.

The following lemma summarizes the discussion above and complements it by providing additional inequalities related to the effective correlated minimax value.

Lemma 2. *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 introduced above satisfy:

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

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

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

Proof. See Appendix A.4. □

In part (i), only two inequalities require a proof. The lower horizontal inequality says that player i's correlated minimax value is weakly smaller than her effective correlated minimax value. The intuition is similar to the independent case. Indeed, the presence of an additional player with equivalent utility can never make it easier to minimax a given player i, even if the minimax profile may be correlated. Next, the right-hand vertical inequality says that player i's effective independent minimax value is weakly larger than its correlated counterpart. This observation captures the intuition that correlated minimax strategies can be more effective in games with three or more players, even against an equivalence class of players. Parts (ii) and (iii) provide conditions under which the various concepts coincide. These conditions are very useful in applications. In general, however, the effective correlated minimax value may differ from the other three minimax concepts, as the following example illustrates.

F

	F	**S**
F	4, 4, 4	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 3: The game G^2.

Example 2. For the game G^2 from Figure 3, we have $v_1^{\mathtt{ind}} = v_1^{\mathtt{cor}} = 0$ and $w_1^{\mathtt{ind}} = \frac{4}{9}$. However, the effective correlated minimax value is given by $w_1^{\mathtt{cor}} = \frac{4}{13}$, and it can be implemented using the following distribution of action profiles:[24]

$$\alpha(\mathbf{F}, \mathbf{F}, \mathbf{F}) = \tfrac{1}{13},$$

$$\alpha(\mathbf{F}, \mathbf{S}, \mathbf{S}) = \alpha(\mathbf{S}, \mathbf{F}, \mathbf{S}) = \alpha(\mathbf{S}, \mathbf{S}, \mathbf{F}) = \tfrac{4}{13},$$

$$\alpha(a_1, a_2, a_3) = 0, \text{ otherwise.}$$

In less straightforward examples, the computation of the effective *correlated* minimax is tremendously simplified by the fact that the minimax problem can be rewritten as a linear program.[25] In contrast, this is not in general feasible for the *independent* minimax concepts (because individual minimax probabilities are multiplied if $n \geq 3$).

[24]For further details, see Appendix A.5.

[25]Indeed, let $\mathbf{u}_j(a_j)$ denote player i's (unconditional) maximal payoff when player $j \sim i$ is supposed to play a_j but may deviate. Further, let $\mathbf{U}_i \in \mathbb{R}$ denote the maximum to be minimized. Then, as

5 Necessary and sufficient conditions

This section first derives conditions necessary for a payoff profile to be implementable as an MSPE (Section 5.1), and then various sufficient conditions (Section 5.2).

5.1 Necessary conditions

Using the concepts introduced above, we can derive the following necessary conditions for MSPE implementability.

Theorem 2 (Necessary Conditions). *If $U_i(\sigma)$ is player i's expected payoff under an MSPE (SPE), then $U_i(\sigma) \geq w_i^{\mathtt{cor}}$ ($U_i(\sigma) \geq w_i^{\mathtt{ind}}$).*

Proof. See Appendix A.6. □

The bracketed claim, due to Wen (1994), is based on the intuition that player i's payoff cannot be depressed strictly below $w_i^{\mathtt{ind}}$ if either player i herself or some other player $j \neq i$ with equivalent utility can avert this outcome. The proof of the non-bracketed part of Theorem 2 requires an additional step. Specifically, one notes that, by the revelation principle, the mediator may w.l.o.g. be assumed to send messages in the form of pure action recommendations. Therefore, the stage payoff of the "most fortunate deviator" boils down to the effective correlated minimax value.

follows from Eq. (A.2) in the Appendix,

$$
\begin{aligned}
w_i^{\mathtt{cor}} = \min_{\alpha,\, \{\mathbf{u}_j\}_{j\sim i},\, \mathbf{U}_i} \quad & \mathbf{U}_i \\
\text{s.t.} \quad & \alpha \in \Delta(A) \\
& \mathbf{u}_j(a_j) \;\geq \sum_{a_{-j}\in A_{-j}} \alpha(a_j, a_{-j})\, u_i(a_j', a_{-j}) \quad \forall j \sim i,\ \forall a_j, a_j' \in A_j \\
& \mathbf{U}_i \;\geq \sum_{a_j \in A_j} \mathbf{u}_j(a_j) \quad \forall j \sim i.
\end{aligned}
$$

5.2 Sufficient conditions

We define, as usual, the set of *feasible* payoff profiles V as the convex hull of $\{u(a) \equiv (u_1(a), \ldots, u_n(a)) : a \in A\}$. The following result is a comprehensive folk theorem for the MSPE concept.

Theorem 3. *Let $v \in V$ be such that $v_i > w_i^{\texttt{cor}}$ ($v_i > w_i^{\texttt{ind}}$) for all $i \in N$. Then, there is some $\underline{\delta} \in (0, 1)$ such that for all $\delta \in (\underline{\delta}, 1)$, there exists an MSPE (SPE with observability of mixed actions[26]) in the infinitely repeated game in which player i's expected payoff is v_i, for all $i \in N$.*

Proof. See Appendix A.7. □

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

There is another distinction, 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) argued that 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,

[26]The MSPE concept 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.

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 bulk of the literature has preferred to make simplifying assumptions instead. Indeed, as pointed out by Fudenberg et al. (2007, Fn. 10), the sufficient conditions in Wen (1994) crucially depend on the observability of mixed actions. Allowing for ex-post transparent mediation circumvents this complication. In particular, the reference to simplifying assumptions is absent from our proof of the non-bracketed part of Theorem 3.[27]

The following version of the folk theorem with mediation is an analogue of Abreu et al. (1994, Thm. 1).

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

Proof. According to Lemma 2(ii), NEU implies $w_i^{\texttt{cor}} = v_i^{\texttt{cor}}$. The non-bracketed claim is therefore immediate from Theorem 3. The bracketed claim is Abreu et al. (1994, Thm. 1). □

Finally, we turn to two-player games. Mediation does not essentially affect the set of SPE payoffs in two-player games in the limit as $\delta \to 1$.

Corollary 2. *Suppose that $n = 2$. Then, any $v \in V$ such that $v_i > v_i^{\texttt{ind}}$ for $i \in \{1, 2\}$ is an MSPE (an approximate SPE) payoff profile in the infinitely repeated game when*

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

players are sufficiently patient.

Proof. We prove the non-bracketed case first. By Lemma 2(iii), we know that $w_i^{\texttt{cor}} = w_i^{\texttt{ind}}$, for $i \in \{1,2\}$. If G satisfies, in addition, the NEU condition, then $w_i^{\texttt{ind}} = v_i^{\texttt{ind}}$, for $i \in \{1,2\}$, by Lemma 2(ii), and the claim is immediate from Theorem 3. If G violates NEU, then we may assume w.l.o.g. that $u_1 = u_2$, so that $w_1^{\texttt{ind}} = w_2^{\texttt{ind}} = \max\{v_1^{\texttt{ind}}, v_2^{\texttt{ind}}\}$ (see Fn. 21). Since feasibility implies $v_1 = v_2$, the claim holds in this case as well. In sum, we have established the non-bracketed claim.[28] The bracketed claim is Fudenberg and Maskin (1991, Prop. 2) for $n = 2$. □

6 Discussion

This section collects additional results and clarifications. Section 6.1 revisits the full-dimensionality counterexample of Fudenberg and Maskin (1986, Ex. 3). Section 6.2 develops a folk theorem based on correlated threats. Section 6.3 illustrates that, for a fixed discount factor, private recommendations can substitute for patience. Section 6.4 revisits the example of Forges et al. (1986), related to strict individual rationality. Finally, Section 6.5 comments on the role of internal messages in the revelation principle for extensive-form correlated equilibrium.

6.1 The example of Fudenberg and Maskin (1986)

Consider the game G^3 shown in Figure 4. If two players decide to choose different pure actions, then the payoff of the third player is held down at zero. Hence, $v_i^{\texttt{ind}} = v_i^{\texttt{cor}} = 0$, for each player $i \in N$.

Lemma 3. *For the game G^3, we have $w_i^{\texttt{cor}} = \frac{1}{4}$, for each $i \in N$.*

[28] Alternatively, by Fudenberg and Maskin (1986, Thm. 1), any $v \in V$ satisfying the conditions of the corollary can be implemented as an SPE with public randomization and observable mixed actions if players are sufficiently patient. Then, the non-bracketed claim follows from Lemma 1(ii).

Proof. See Appendix A.8. □

The fact that $w_i^{\texttt{ind}} = \frac{1}{4}$ is due to Fudenberg and Maskin (1986, Ex. 3). In view of Lemma 2, the statement that $w_i^{\texttt{cor}} = \frac{1}{4}$ is stronger (but easier to check). It says that, even with 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	$1, 1, 1$	$0, 0, 0$
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^3.

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

Proof. By Lemma 1(ii), mediation is more permissive than access to public randomization. Therefore, the assertion follows from Lemma 3. □

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 profiles is of dimension strictly lower than n. Sufficient conditions have often been presented under the simplifying assumption that players have access to a public randomization device. Therefore, Corollary 3 closes a potentially important gap in the literature.

6.2 Correlated threats

As was seen in Example 1, a simple way to construct an MSPE is by using a correlated equilibrium as a threat point. Generalizing this idea leads to the following extension of Friedman's (1971) perfect folk theorem.

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

Proof. See Appendix A.9. □

Since the set of correlated equilibria can be strictly larger than the convex hull of the Nash equilibria, a correlated threat can be more effective than a Nash threat even in a two-player game. This fact is illustrated by our next example.

	L	**C**	**R**
T	$1, 1$	$1, 0$	$-3, -2$
M	$0, -3$	$-1, -1$	$2, 0$
B	$2, -1$	$0, -2$	$0, 0$

Figure 5: The game G^4.

Example 3. Consider the game G^4 in Figure 5. Suppose that, on the equilibrium path, the mediator recommends $(\mathbf{T}, \mathbf{L})$ at every stage, yielding the payoff profile $v = (1, 1)$. After any deviation by any player, the mediator switches permanently to the correlated equilibrium $\alpha^* \in \Delta(A)$ defined by

$$\alpha^*(\mathbf{T}, \mathbf{C}) = \tfrac{1}{3}, \qquad \alpha^*(\mathbf{M}, \mathbf{C}) = \tfrac{1}{3}, \qquad \alpha^*(\mathbf{M}, \mathbf{R}) = \tfrac{1}{3}.$$

The correlated equilibrium implements the payoff profile $(\frac{2}{3}, -\frac{1}{3})$. Hence, we have an MSPE for all $\delta \geq \frac{3}{4}$.

Notably, $v = (1,1)$ cannot be enforced by a Nash threat. Indeed, the stage game has the unique Nash equilibrium $(\mathbf{M}, \mathbf{R})$, with payoff profile $(2,0)$.[29]

6.3 The role of the discount factor

For $n = 2$ players, our necessary and sufficient conditions imply that the respective limiting sets, as $\delta \to 1$, of strictly individually rational payoffs implementable by an MSPE or SPE coincide. For a *fixed* discount factor, however, private messages can matter, as our next example shows.

	R	**P**	**S**
R	$0,0$	$0,2$	$2,0$
P	$2,0$	$0,0$	$0,2$
S	$0,2$	$2,0$	$0,0$

Figure 6: The game G^5.

Example 4. Consider the game G^5 in Figure 6. Let α^* denote the correlated equilibrium in which each non-tied outcome receives weight $\frac{1}{6}$. Then, the unconditional repetition of α^* is an MSPE for any $\delta \in (0,1)$, implementing the payoff profile $(1,1)$. However, if δ is too small, then $(1,1)$ is not implementable as an SPE.[30]

In the example, the MSPE outcome cannot be replicated using an SPE with public randomization. The reason is that, without private messages, at least one player can perfectly predict the other player's action, which makes it hard to deter a deviation driven by short-term considerations.

[29]Since G^4 is a two-player game and $v = (1,1)$ is strictly individually rational, v can be implemented by *some* SPE. Without observable mixed actions, however, such an SPE may be complicated.

[30]Indeed, achieving a total payoff of 2 makes it imperative to avoid any ties. Therefore, at any history on the equilibrium path, at least one player chooses a pure action. But then one of the players has a one-shot deviation gain of at least $\frac{2}{3}$. For $\delta < \frac{1}{4}$, that deviation cannot be deterred in any SPE.

6.4 Weak vs. strict individual rationality

This section revisits an example that has been used to explain why sufficient conditions for folk theorems commonly require *strict* individual rationality for all players.

Example 5. Consider the game G^6 shown in Figure 7. 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, $v_1^{\texttt{ind}} = v_2^{\texttt{ind}} = v_3^{\texttt{ind}} = 0$. However, v is not implementable as an SPE (even with public randomization).[31]

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 7: The game G^6.

Turning to the possibility of mediation, the very same arguments show that $v_i^{\texttt{cor}} = 0$ 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 unnecessary as a consequence of Lemma 2(ii) because G^6 satisfies NEU. Thus, mediation does not obviate the need for the strictness assumption.

6.5 On the revelation principle for extensive-form correlated equilibrium

A novel element of our proof of the revelation principle is the idea that the canonical mediator speculates in a Bayesian fashion about the messages that the original device has sent in prior stages. To clarify our contribution, this section elaborates on the revelation principle for extensive-form correlated equilibrium. In Forges (1986), the

[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.

use of a canonical device can be assumed w.l.o.g. because players interacting with it receive less information, and therefore have fewer ways to deviate. The original proof does not explicitly discuss the point, however, that also the device receives less information if replaced by its canonical version. In the example below, the non-canonical 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 is partly lost when the players receive direct messages containing only information about their pure choices. A canonical device *without* an internal record might then lack sufficiently precise information to make recommendations in the same way as in the non-canonical setup. The example captures the intuitive idea that multiple selves of the device are "connected." The example thereby illustrates the fact that the mediator's ability to communicate with itself (Sugaya and Wolitzky, 2021) is, *a priori*, stronger than the assumption that future selves merely have access to messages sent in prior periods (Forges, 1986).

Example 6. There are $n = 3$ players and two stages $t = 1, 2$. Players 1 and 2 make their choices at stage $t = 1$, while 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^7 are given in Figure 8.

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 $(\frac{1}{3}\mathbf{T} + \frac{2}{3}\mathbf{B},\ \frac{2}{3}\mathbf{L} + \frac{1}{3}\mathbf{R})$, which is a Nash equilibrium between the two players if player 3 chooses $\mathbf{E_1}$.

$\mathbf{E}_1$

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

$\mathbf{E}_2$

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

$\mathbf{E}_3$

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

Figure 8: The two-stage game G^7.

- Upon observing e_2, players 1 and 2 play $(\frac{2}{3}\mathbf{T} + \frac{1}{3}\mathbf{B},\ \frac{1}{3}\mathbf{L} + \frac{2}{3}\mathbf{R})$, which is a Nash equilibrium between the two players if player 3 chooses $\mathbf{E_2}$.

- At stage $t = 2$, the device recalls its earlier message (either e_1 or e_2) and informs player 3 accordingly. Player 3 chooses $\mathbf{E_1}$ if the message is e_1, and $\mathbf{E_2}$ 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. In a canonical equilibrium, players 1 and 2 are merely informed at stage $t = 1$ about their recommended pure actions. Suppose that internal messages are prohibited. Then the stage-2 device knows only which of the four pairs $(\mathbf{T}, \mathbf{L})$, $(\mathbf{T}, \mathbf{R})$, $(\mathbf{B}, \mathbf{L})$, and $(\mathbf{B}, \mathbf{R})$ it recommended at stage $t = 1$. The respective probability distributions over pairs (a_1, a_2) are shown in Table I.

By definition, a canonical device cannot give a recommendation to player 3 at stage $t = 1$, because player 3 moves only at stage $t = 2$. But if the device does not inform player 3 at all, player 3 finds it optimal to choose $\mathbf{E_3}$, breaking the equilibrium. Thus, without internal records, the canonical device needs to apply Bayes' rule to come up with the correct conditional probabilities for recommending actions to player 3.

(a) Conditional on e_1

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

(b) Conditional on e_2

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

(c) Aggregate

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

Table I: Stage-1 recommendation probabilities.

7 Concluding remarks

We have proposed a simple and tractable solution concept that naturally generalizes correlated equilibrium to the class of infinitely repeated games. MSPE aligns very well with subgame perfection, public randomization, and sequential equilibrium, satisfies an important revelation principle, suggests a natural variant of the effective minimax value, and leads to straightforward analogues of classic folk theorems. There is no need to assume that mixed actions are observable. Moreover, necessary and sufficient conditions for implementability of payoffs apply to any finite stage game.

Key takeaways include that private messaging makes minimax punishments not only more effective, as has been known before, but also simpler, sequentially rational, and less reliant on players' patience. Thus, transparent mediation can benefit players who rely on the cooperation of others. As we have also seen, this observation extends in a nontrivial way to situations in which players share equivalent utilities.[32]

What happens if the requirement of ex-post observability is dropped? Intuition suggests that our sufficient conditions generalize because, as mentioned before, the mediator can always commit to disclosing all internal records and private messages

[32] There are actually two more takeaways. First, the full-dimensionality example in Fudenberg and Maskin (1986) is robust to the introduction of a public randomization device, which closes a small but long-standing gap in the literature. Second, *internal records* kept by the mediator might play a more important role for the analysis of mediation than suggested by the existing literature.

at the beginning of the next stage. However, the implications of secrecy for necessary conditions are less clear. Since a new stage does not necessarily start a new subgame, there would be some leeway for the mediator to steer players' subjective beliefs over past and current message profiles. This might allow the mediator to punish more effectively, in particular, in games that do not satisfy NEU.

A Appendix

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

A.1 The outcome of mediated play

Given a device μ and a profile of behavior strategies σ, an infinite sequence $\{(m^t, a^t)\}_{t=0}^{\infty}$ of message and action profiles is determined as follows. Recall that the initial history is $h^0 = \varnothing$. Iteratively, for any $t \in \{0, 1, 2, \ldots\}$, draw the message profile $m^t \in M$ according to $\mu(\cdot \mid h^t)$, draw a_i^t according to $\sigma_i(\cdot \mid h^t, m_i^t)$, for each player $i \in N$, and construct the updated history $h^{t+1} = (h^t; m^t, a^t) = (m^0, a^0; \ldots; m^t, a^t)$. From the infinite sequence $\{(m^t, a^t)\}_{t=0}^{\infty}$, the *outcome* $\{a^t\}_{t=0}^{\infty}$ is obtained by omitting the message profiles.

Next, fix some player $i \in N$, an information set $(h, m_i) \in I_i$, with $h = h^t \in H^t$, and a belief $\beta_i(\cdot \mid h, m_i) \in \Delta(M_{-i})$ at (h, m_i). The infinite continuation sequence $\{(m^{t+\tau}, a^{t+\tau})\}_{\tau=0}^{\infty}$ of message and action profiles is determined as follows. First, let $m_i^t = m_i$, draw $m_{-i}^t \in M_{-i}$ according to $\beta_i(\cdot \mid h, m_i)$, and draw a_j^t according to $\sigma_j(\cdot \mid h, m_j^t)$, for any player $j \in N$. This determines the pair (m^t, a^t), and the history $h^{t+1} = (h^t; m^t, a^t)$. Iteratively, for any $\tau \in \{1, 2, \ldots\}$, draw $m^{t+\tau} \in M$ according to $\mu(\cdot \mid h^{t+\tau})$, draw $a_j^{t+\tau}$ according to $\sigma_j(\cdot \mid h^{t+\tau}, m_j^{t+\tau})$, for each $j \in N$, and construct the history $h^{t+\tau+1} = (h^{t+\tau}; m^{t+\tau}, a^{t+\tau})$. The *conditional outcome* $\{a^{t+\tau}\}_{\tau=0}^{\infty}$ is obtained again from the infinite sequence $\{(m^{t+\tau}, a^{t+\tau})\}_{\tau=0}^{\infty}$ by omitting the messages.

A.2 Proof of Lemma 1

(i) Fix the device μ, a player $i \in N$, and the profile of opponents' behavior strategies $\sigma_{-i} \in \Sigma_{-i}$. Then, at any information set $(h, m_i) \in I_i$, player i faces a discounted decision problem. By the standard one-shot deviation principle (Fudenberg and Tirole, 1991), if some deviation σ_i' yields a strictly higher continuation payoff at (h, m_i), there exists a profitable deviation that differs from σ_i only at a single information set. Thus, to verify the sequential rationality condition in Definition 1, it suffices to check one-shot deviations.

(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 profile $V(\theta^0) = (V_1(\theta^0), \ldots, V_n(\theta^0)) \in \mathbb{R}^n$. Let $\mathcal{V}^0 = \{V(\theta^0) \,|\, \theta^0 \in [0,1]\} \subseteq \mathbb{R}^n$ denote the set of payoff profiles obtained that way. Then, the SPE payoff profile $V^* = (V_1^*, \ldots, V_n^*) = \mathbb{E}[V(\theta^0)]$ is an element of the convex hull of $\mathcal{V}^0$. Thus, 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 $\sum_{\kappa=1}^{n+1} \lambda_\kappa = 1$, such that

$$V_i^* = \sum_{\kappa=1}^{n+1} \lambda_\kappa V_i(\theta_\kappa^0),$$

for all $i \in N$. Hence, we may replace the public randomization device at stage $t = 0$, without affecting expected continuation payoffs or sequential rationality at stage $t = 0$, by an autonomous device using identical sets $M_1 = \ldots = M_n = \{1, \ldots, n+1\}$ and sending the message profile $(\kappa, \ldots, \kappa) \in M$ with probability λ_κ. This defines an infinitely repeated game $\mathcal{G}^0$ in which the public randomization device has been replaced by a finite device at stage $t = 0$. Applying the same reduction iteratively for any history h^t, we can construct, for any horizon $T \geq 1$, an infinitely repeated game $\mathcal{G}^T$ in which the public randomization device has been replaced by a finite device in all stages $t \leq T$. 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 in $\mathcal{G}^\infty$. Suppose that a one-shot deviation in $\mathcal{G}^\infty$ is profitable at some stage t. Given boundedness of payoffs, the respective continuation payoffs resulting from equilibrium play and deviation in $\mathcal{G}^\infty$ are arbitrarily close to the corresponding continuation payoffs in $\mathcal{G}^T$, as $T \to \infty$. Hence, we can find a horizon $T > t$ such that the same deviation is profitable in $\mathcal{G}^T$, which is impossible. Thus, a deviation is not profitable.

(iii) For any regular history $h \in H$, let $\tilde{\mu}(\cdot \,|\, h) = \mu(\cdot \,|\, h)$. For any irregular history $h \in H$, there is a first stage at which the irregularity becomes public information. Exploiting the stationary nature of the infinitely repeated game, the device is then programmed to "reset" itself by erasing the irregular initial segment of its history. If necessary, this procedure is repeated until all inconsistencies are resolved. Formally, for any history $h^t \in H$, we define the last "reset" stage $\tau(t) \equiv \tau(h^t)$ recursively as follows. Let $\tau(0) = 0$. Further, for $t \geq 0$, let $\tau(t+1) = \tau(t)$ if appending (m^t, a^t) to the history

$$h^t \backslash h^{\tau(t)} \equiv \left(m^{\tau(t)}, a^{\tau(t)}; \ldots; m^{t-1}, a^{t-1}\right) \in H^{t-\tau(t)}.$$

creates a regular history $h^{t+1} \backslash h^{\tau(t)}$, and let $\tau(t+1) = t+1$ otherwise, in which case $h^{t+1} \backslash h^{\tau(t+1)} = \varnothing$. In the MSPE that we construct, the device, beliefs, and strategies are evaluated at the regular history $h^t \backslash h^{\tau(t)}$ rather than at the original history h^t. Formally, we define a new device $\tilde{\mu} : H \to \Delta(M)$ by $\tilde{\mu}(\cdot \,|\, h^t) = \mu(\cdot \,|\, h^t \backslash h^{\tau(t)})$, for all $h^t \in H$. Similarly, for each player $i \in N$ and any message m_i^t such that $\tilde{\mu}_i(m_i^t \,|\, h^t) > 0$, let $\tilde{\sigma}_i(\cdot \,|\, h^t, m_i^t) = \sigma_i(\cdot \,|\, h^t \backslash h^{\tau(t)}, m_i^t)$. Further, at any information set (h^t, m_i^t) with $\tilde{\mu}_i(m_i^t \,|\, h^t) = 0$, choose $\tilde{\beta}_i(\cdot \,|\, h^t, m_i^t)$ arbitrarily and let $\tilde{\sigma}_i(\cdot \,|\, h^t, m_i^t)$ place probability one on some pure best response to the correlated action profile $\alpha_{-i} \in \Delta(A_{-i})$ induced by $\tilde{\beta}_i(\cdot \mid h^t, m_i^t) \in \Delta(M_{-i})$ and the profile of behavior strategies $\tilde{\sigma}_{-i}$ in the stage game. Finally, define beliefs $\tilde{\beta}_i$ at information sets with $\tilde{\mu}_i(m_i^t \mid h^t) > 0$ by Bayes' rule from $\tilde{\mu}$. Noting that, after any zero-probability message for some player, the next stage "resets" the history to $\varnothing$ regardless of the player's action, it is now immediate from the construction that $(\tilde{\mu}, \tilde{\beta}, \tilde{\sigma})$ is an MSPE that is outcome equivalent to (μ, β, σ).

(iv) Let $\alpha^* \in \Delta(A)$ be a correlated equilibrium in the stage game G. To define an MSPE, let player i's message set be given as $M_i = \operatorname{supp} \alpha_i \subseteq A_i$, for each $i \in N$. Next, let $\mu(\cdot \mid h) = \alpha^*(\cdot)$ for any $h \in H$. For any player $i \in N$ and information set $(h, m_i) \in I_i$, define the belief $\beta_i(\cdot \mid h, m_i) \in \Delta(M_{-i})$ by Bayes' rule. Further, let $\sigma_i(\cdot \mid h, m_i) \in \Delta(A_i)$ give full weight on m_i. Then, sequential rationality for player i at the information set (h, m_i) follows directly from the optimality condition for player i in the correlated equilibrium α^*. The existence claim is now immediate. □

A.3 Proof of Theorem 1

Take an arbitrary MSPE (μ, β, σ). For any history $h^t = (m^0, a^0; \ldots; m^{t-1}, a^{t-1}) \in H^t$, define the joint distribution over message and action profiles $\widetilde{\mu}(\cdot \mid h^t) \in \Delta(M \times A)$ induced at h^t by

$$\widetilde{\mu}(m^t, a^t \mid h^t) = \mu(m^t \mid h^t) \cdot \prod_{i \in N} \sigma_i(a_i^t \mid h^t, m_i^t).$$

Further, let $\widetilde{\mu}_a(\cdot \mid h) \in \Delta(A)$ and $\widetilde{\mu}_i(\cdot \mid h) \in \Delta(M_i \times A_i)$, respectively, denote the marginals of $\widetilde{\mu}(\cdot \mid h)$ on A and $M_i \times A_i$. Let $\hat{H}^t$ be the set of canonical histories of length t, with typical element $\hat{h}^t = (\hat{a}^0, a^0; \ldots; \hat{a}^{t-1}, a^{t-1})$, and let S^t be the set of *states* of length t, with typical element $s^t = (m^0, \ldots, m^{t-1})$. Conditioning on the public action record $(a^0, \ldots, a^{t-1})$ contained in the canonical history $\hat{h}^t$, Bayesian uncertainty concerns only the latent message history s^t. Keeping this in mind, the probability that the canonical history $\hat{h}^t \in \hat{H}^t$ and the state $s^t \in S^t$ obtain jointly is given by the product

$$\mathrm{pr}(\hat{h}^t, s^t) \;=\; \prod_{\tau=0}^{t-1} \widetilde{\mu}(m^\tau, \hat{a}^\tau \,|\, h^\tau),$$

where $h^\tau = (m^0, a^0; \ldots; m^{\tau-1}, a^{\tau-1})$ is the corresponding history of length τ, built from s^τ and the record of chosen actions contained in $\hat{h}^t$. Moreover, the probability that $\hat{h}^t$ obtains is

$$\mathrm{pr}(\hat{h}^t) \;=\; \sum_{s^t \in S^t} \mathrm{pr}(\hat{h}^t, s^t),$$

while the probability that $\hat{h}^t$ and $\hat{a}^t$ obtain jointly is

$$\mathrm{pr}(\hat{h}^t, \hat{a}^t) \;=\; \sum_{s^t \in S^t} \mathrm{pr}(\hat{h}^t, s^t)\, \widetilde{\mu}_a(\hat{a}^t \,|\, h^t),$$

where h^t is derived from s^t and $\hat{h}^t$ as explained above. We specify a direct MSPE candidate $(\hat{\mu}, \hat{\beta}, \hat{\sigma})$ as follows. For any $\hat{h}^t \in \hat{H}^t$ with $\mathrm{pr}(\hat{h}^t) > 0$, and any $\hat{a}^t \in A$, let

$$\hat{\mu}(\hat{a}^t \,|\, \hat{h}^t) = \frac{\mathrm{pr}(\hat{h}^t, \hat{a}^t)}{\mathrm{pr}(\hat{h}^t)}.$$

If $\mathrm{pr}(\hat{h}^t) = 0$, define $\hat{\mu}(\cdot \,|\, \hat{h}^t)$ arbitrarily. For each $i \in N$, define beliefs $\hat{\beta}_i(\cdot \,|\, \hat{h}^t, \hat{a}_i^t)$ by Bayes' rule from $\hat{\mu}$ whenever $\hat{\mu}_i(\hat{a}_i^t \,|\, \hat{h}^t) > 0$, and choose $\hat{\beta}_i$ arbitrarily otherwise. Further, define $\hat{\sigma}_i(\cdot \,|\, \hat{h}^t, \hat{a}_i^t)$ to put full weight on $a_i = \hat{a}_i^t$. We verify sequential rationality, we may restrict attention to one-shot deviations at regular information sets. Fix a canonical information set $(\hat{h}^t, \hat{a}_i^t)$ for some player $i \in N$. Similar to the above, consider joint probabilities

$$\begin{aligned}
\mathrm{pr}(s^t, m_i^t, \hat{h}^t, \hat{a}_i^t) \;&=\; \mathrm{pr}(s^t, \hat{h}^t)\mu_i(m_i^t \,|\, h^t)\sigma_i(\hat{a}_i^t \,|\, h^t, m_i^t),\\
\mathrm{pr}(\hat{h}^t, \hat{a}_i^t) \;&=\; \sum_{s^t \in S^t} \sum_{m_i^t \in M_i} \mathrm{pr}(s^t, m_i^t, \hat{h}^t, \hat{a}_i^t).
\end{aligned}$$

Assume that the information set $(\hat{h}^t, \hat{a}_i^t)$ is regular, i.e., that $\mathrm{pr}(\hat{h}^t, \hat{a}_i^t) > 0$. Then, conditional on $(\hat{h}^t, \hat{a}_i^t)$, player i's posterior $\mathrm{pr}(\cdot \,|\, \hat{h}^t, \hat{a}_i^t) \in \Delta(S^t \times M_i)$ over the latent state s^t and the current message m_i^t is given by

$$\text{pr}(s^t, m_i^t \,|\, \hat{h}^t, \hat{a}_i^t) \;=\; \frac{\text{pr}(s^t, m_i^t, \hat{h}^t, \hat{a}_i^t)}{\text{pr}(\hat{h}^t, \hat{a}_i^t)}.$$

Then, from the definition of the continuation outcome (cf. Section A.1),

$$U_i(\hat{\sigma} \,|\, \hat{h}^t, \hat{a}_i^t) = \sum_{s^t \in S^t} \sum_{m_i^t \in M_i} \text{pr}(s^t, m_i^t \,|\, \hat{h}^t, \hat{a}_i^t)\, U_i(\sigma \,|\, h^t, m_i^t).$$

Let $\hat{\sigma}_i'$ denote the one-shot deviation that chooses $a_i' \in A_i$ at $(\hat{h}^t, \hat{a}_i^t)$ and coincides with $\hat{\sigma}_i$ thereafter. With σ_i' defined analogously at (h^t, m_i^t), we obtain

$$U_i(\hat{\sigma}_i', \hat{\sigma}_{-i} \,|\, \hat{h}^t, \hat{a}_i^t) = \sum_{s^t \in S^t} \sum_{m_i^t \in M_i} \text{pr}(s^t, m_i^t \,|\, \hat{h}^t, \hat{a}_i^t)\, U_i(\sigma_i', \sigma_{-i} \,|\, h^t, m_i^t).$$

Now, whenever $\text{pr}(s^t, m_i^t \mid \hat{h}^t, \hat{a}_i^t) > 0$, we have $\sigma_i(\hat{a}_i^t \mid h^t, m_i^t) > 0$, so choosing $\hat{a}_i^t$ does not affect the continuation payoff at (h^t, m_i^t). Sequential rationality of (μ, β, σ) implies

$$U_i(\sigma_i, \sigma_{-i} \,|\, h^t, m_i^t) \geq U_i(\sigma_i', \sigma_{-i} \,|\, h^t, m_i^t),$$

for any h^t and any m_i^t. Averaging with respect to the posterior $\text{pr}(\cdot \,|\, \hat{h}^t, \hat{a}_i^t)$, we obtain

$$U_i(\hat{\sigma}_i, \hat{\sigma}_{-i} \,|\, \hat{h}^t, \hat{a}_i^t) \geq U_i(\hat{\sigma}_i', \hat{\sigma}_{-i} \,|\, \hat{h}^t, \hat{a}_i^t).$$

Thus, $(\hat{\mu}, \hat{\beta}, \hat{\sigma})$ is indeed sequentially rational at all regular information sets. We check outcome equivalence. By definition, $\hat{\mu}(\cdot \,|\, \hat{h}^t)$ follows exactly the distribution of the recommendation $\hat{a}^t$ obtained from $\widetilde{\mu}$ once the message history s^t is integrated out. Hence the stochastic process of recommended action profiles $\{\hat{a}^t\}_{t \geq 0}$ induced by $\hat{\mu}$ coincides with the process of action profiles generated by μ and σ. Under obedience, $a^t = \hat{a}^t$ for all t on the equilibrium path, so the outcome is indeed the same as under (μ, β, σ). By Lemma 1(iii), there exists an outcome-equivalent MSPE. By construction, that MSPE is direct and reflects obedience to any regular message.

□

A.4 Proof of Lemma 2

(i) The inequality $v_i^{\texttt{ind}} \leq w_i^{\texttt{ind}}$ is immediate from the definitions. To prove $v_i^{\texttt{cor}} \leq w_i^{\texttt{cor}}$, let $\alpha \in \Delta(A)$. Then, since $\sim$ is reflexive,

$$\mathbb{E}_{\alpha_i}\left[\max_{a_i' \in A_i} \mathbb{E}_\alpha[u_i(a_i', a_{-i}) \,|\, a_i]\right] \leq \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \mathbb{E}_{\alpha_j}\left[\max_{a_j' \in A_j} \mathbb{E}_\alpha[u_i(a_j', a_{-j}) \,\big|\, a_j]\right], \tag{A.1}$$

where α_i denotes the marginal distribution of α on A_i, as before. Next, keeping $a_i \in A_i$ fixed, the conditional distribution of a_{-i} is certainly an element of $\Delta(A_{-i})$.[33] Hence,

$$\max_{a_i' \in A_i} \mathbb{E}_\alpha[u_i(a_i', a_{-i}) \,|\, a_i] \geq \min_{\tilde{\alpha}_{-i} \in \Delta(A_{-i})} \max_{a_i' \in A_i} \mathbb{E}_{\tilde{\alpha}_{-i}}[u_i(a_i', a_{-i})] = v_i^{\texttt{cor}}.$$

Taking the expectation with respect to α_i yields

$$\mathbb{E}_{\alpha_i}\left[\max_{a_i' \in A_i} \mathbb{E}_\alpha[u_i(a_i', a_{-i}) \,|\, a_i]\right] \geq v_i^{\texttt{cor}}.$$

Combining this with (A.1) and subsequently taking the minimum over all $\alpha \in \Delta(A)$ shows that, indeed, $w_i^{\texttt{cor}} \geq v_i^{\texttt{cor}}$. 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 $\alpha \in \bigtimes_{k \in N} \Delta(A_k)$ and any $a_j' \in A_j$,

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

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

$$\begin{aligned} w_i^{\texttt{cor}} &= \min_{\alpha \in \Delta(A)} \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \mathbb{E}_{\alpha_j}\left[\max_{a_j' \in A_j} \mathbb{E}_\alpha[u_i(a_j', a_{-j}) \,\big|\, a_j]\right] \\ &\leq \min_{\alpha \in \bigtimes_{k \in N} \Delta(A_k)} \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \mathbb{E}_{\alpha_j}\left[\max_{a_j' \in A_j} \mathbb{E}_{\alpha_{-j}}[u_i(a_j', a_{-j})]\right] \end{aligned}$$

[33] In line with our earlier convention (see Fn. 23), this conditional distribution is the marginal on A_{-i} if α_i assigns probability zero to a_i.

$$= \min_{\alpha \in \times_{k\in N} \Delta(A_k)} \max_{\substack{j\in N \\ \text{s.t. } j\sim i}} \max_{a'_j \in A_j} \mathbb{E}_{\alpha_{-j}}[u_i(a'_j, a_{-j})]$$
$$= w_i^{\texttt{ind}},$$

which proves the claim.

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

$$\begin{aligned} w_i^{\texttt{ind}} &= \min_{\alpha \in \times_{k\in N} \Delta(A_k)} \max_{\substack{j\in N \\ \text{s.t. } j\sim i}} \max_{a'_j \in A_j} \mathbb{E}_{\alpha_{-j}}[u_i(a'_j, a_{-j})] \\ &= \min_{\alpha \in \times_{k\in N} \Delta(A_k)} \max_{a'_i \in A_i} \mathbb{E}_{\alpha_{-i}}[u_i(a'_i, a_{-i})] \\ &= v_i^{\texttt{ind}}. \end{aligned}$$

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

$$\begin{aligned} w_i^{\texttt{cor}} &= \min_{\alpha\in\Delta(A)} \max_{\substack{j\in N \\ \text{s.t. } j\sim i}} \mathbb{E}_{\alpha_j} \left[\max_{a'_j \in A_j} \mathbb{E}_\alpha[u_i(a'_j, a_{-j}) \mid a_j] \right] \\ &= \min_{\alpha\in\Delta(A)} \mathbb{E}_{\alpha_i} \left[\max_{a'_i \in A_i} \mathbb{E}_\alpha[u_i(a'_i, a_{-i}) \mid a_i] \right]. \end{aligned}$$

Let $\alpha_{-i}^{\texttt{cor}} \in \Delta(A_{-i})$ attain $v_i^{\texttt{cor}}$, i.e., $\max_{a'_i\in A_i} \mathbb{E}_{\alpha_{-i}^{\texttt{cor}}}[u_i(a'_i, a_{-i})] = v_i^{\texttt{cor}}$. Select some $\tilde{\alpha}_i \in \Delta(A_i)$, and define the product distribution $\alpha = \tilde{\alpha}_i \otimes \alpha_{-i}^{\texttt{cor}} \in \Delta(A)$. Then, the marginal of α on A_i is $\alpha_i = \tilde{\alpha}_i$. Therefore,

$$v_i^{\texttt{cor}} = \mathbb{E}_{\tilde{\alpha}_i} \left[\max_{a'_i \in A_i} \mathbb{E}_{\alpha_{-i}^{\texttt{cor}}}[u_i(a'_i, a_{-i})] \right] = \mathbb{E}_{\alpha_i} \left[\max_{a'_i \in A_i} \mathbb{E}_\alpha[u_i(a'_i, a_{-i}) \mid a_i] \right].$$

Taking the minimum over all $\alpha \in \Delta(A)$ shows that $w_i^{\texttt{cor}} \leq v_i^{\texttt{cor}}$. Together with part (i), this yields $w_i^{\texttt{cor}} = v_i^{\texttt{cor}}$.

(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

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

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

$$w_i^{\texttt{cor}} \quad \geq \quad \min_{\alpha \in \Delta(A)} \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \max_{a'_j \in A_j} \mathbb{E}_{\alpha_{-j}}[u_i(a'_j, a_{-j})].$$

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

$$\min_{(\alpha_1, \alpha_2) \in \Delta(A_1) \times \Delta(A_2)} \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \max_{a'_j \in A_j} \mathbb{E}_{\alpha_{-j}}[u_i(a'_j, a_{-j})] \quad = \quad w_i^{\texttt{ind}}.$$

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

A.5 Details on Example 2

To determine the non-effective minimax values, suppose that player 2 chooses $a_2 = \mathbf{F}$, and that player 3 chooses $a_3 = \mathbf{S}$. Then, player 1's payoff is zero, and this is player 1's minimal feasible payoff. Hence, $v_1^{\texttt{ind}} = v_1^{\texttt{cor}} = 0$.

Next, in the effective independent minimax problem, let α_i denote the probability that player i chooses $\mathbf{F}$. By symmetry, we may assume w.l.o.g. that $\alpha_1 \geq \alpha_2 \geq \alpha_3$. If $\alpha_2 \geq \frac{1}{3}$, then player 3 has a payoff of at least $\frac{4}{9}$ from choosing $\mathbf{F}$. Otherwise, i.e., if $\alpha_2 < \frac{1}{3}$, then player 1 has a payoff of at least $\frac{4}{9}$ from choosing $\mathbf{S}$. Conversely, if $\alpha_1 = \alpha_2 = \alpha_3 = \frac{1}{3}$, then all players obtain a payoff of $\frac{4}{9}$.

Finally, let $p_{a_1 a_2 a_3}$ denote the probability of $(a_1, a_2, a_3) \in A$ under some $\alpha \in \Delta(A)$. For player 1, the maximum deviation payoff obtainable under α equals

$$\Phi_1(\alpha) = \max\{4p_{\mathbf{FFF}},\ p_{\mathbf{FSS}}\} \ + \ \max\{4p_{\mathbf{SFF}},\ p_{\mathbf{SSS}}\}\,.$$

Analogous expressions obtain for players 2 and 3 by permuting indices, so that $w_1^{\texttt{cor}} = \min_{\alpha \in \Delta(A)} \max\{\Phi_1(\alpha), \Phi_2(\alpha), \Phi_3(\alpha)\}$. Note that the objective function of the minimization problem is convex. Given symmetry, we may therefore restrict attention to distributions α satisfying $p_{\mathbf{FSS}} = p_{\mathbf{SFS}} = p_{\mathbf{SSF}}$ and $p_{\mathbf{FFS}} = p_{\mathbf{FSF}} = p_{\mathbf{SFF}}$. Hence, it suffices to minimize $\Phi_1(\alpha)$ subject to nonnegativity constraints and

$$p_{\mathbf{FFF}} + 3p_{\mathbf{FSS}} + 3p_{\mathbf{SFF}} + p_{\mathbf{SSS}} = 1.$$

It is clearly optimal to set $p_{\mathbf{FSS}} = 4p_{\mathbf{FFF}}$ and $p_{\mathbf{SSS}} = 4p_{\mathbf{SFF}}$. The problem now reads

$$\min_{\substack{p_{\mathbf{FFF}}, p_{\mathbf{SFF}} \\ \text{s.t.} 13p_{\mathbf{FFF}}+7p_{\mathbf{SFF}}=1}} 4p_{\mathbf{FFF}} + 4p_{\mathbf{SFF}} \, .$$

This yields $p_{\mathbf{FFF}} = \frac{1}{13}$ and $p_{\mathbf{SFF}} = 0$. Hence, $p_{\mathbf{FSS}} = \frac{4}{13}$ and $p_{\mathbf{SSS}} = 0$, so that the symmetric distribution stated in Example 2 is indeed optimal. Moreover, $w_1^{\texttt{cor}} = \frac{4}{13}$.

A.6 Proof of Theorem 2

We will make use of the following auxiliary result.

Lemma A.1. *Given $\alpha \in \Delta(A)$, let $\alpha_i \in \Delta(A_i)$ denote the marginal of α on A_i. Then, the mapping*

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

is continuous on $\Delta(A)$.

Proof. Given the probability distribution $\alpha \in \Delta(A)$, let $\alpha_i(a_i)$ and $\alpha(a_i, a_{-i})$, respectively, denote the marginal probability of the action $a_i \in A_i$ and the probability of the action profile $(a_i, a_{-i}) \in A$. Then,

$$\begin{aligned}
&\mathbb{E}_{\alpha_i}\left[\max_{a'_i \in A_i} \mathbb{E}_\alpha[u_i(a'_i, a_{-i}) \mid a_i]\right] \\
&= \sum_{\substack{a_i \in A_i \\ \text{s.t. } \alpha_i(a_i)>0}} \alpha_i(a_i) \cdot \max_{a'_i \in A_i} \sum_{a_{-i} \in A_{-i}} \frac{\alpha(a_i, a_{-i})}{\alpha_i(a_i)} u_i(a'_i, a_{-i}) \\
&= \sum_{a_i \in A_i} \max_{a'_i \in A_i} \sum_{a_{-i} \in A_{-i}} \alpha(a_i, a_{-i}) u_i(a'_i, a_{-i}), \qquad \text{(A.2)}
\end{aligned}$$

because $\alpha_i(a_i) = 0$ implies $\alpha(a_i, a_{-i}) = 0$ for all $a_{-i} \in A_{-i}$. Since the maxima on the RHS of equation (A.2) are continuous as functions of α, this proves the lemma. □

To prove the non-bracketed part of the theorem, we assume w.l.o.g. that $u_j = u_i$ for all players $j \sim i$. Let (μ, β, σ) be an MSPE. By Theorem 1, we may assume that (μ, β, σ) is canonical. Then, in stage $t = 0$, message profiles are drawn from A according to the probability distribution $\mu^0 \equiv \mu(\cdot \,|\, \varnothing) \in \Delta(A)$. By player i's optimality condition at her information set $(\varnothing, \hat{a}_i)$ corresponding to the recommended action $\hat{a}_i \in A_i$ at stage $t = 0$, player i's expected continuation payoff satisfies

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

where L_i denotes the infimum of player i's expected payoffs across *all* MSPE. Since an MSPE exists by Lemma 1, L_i is finite. Further, from stationarity of the game, 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 $\hat{a}_i$ according to μ_i^0 shows that i's expected payoff resulting from (μ, β, σ) satisfies

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

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

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

For any $j \in N$ such that $j \sim i$, the above inequality holds for the same limit distribution μ^0 when i is replaced by j. Hence,

$$L_i = L_j \geq \mathbb{E}_{\mu_j^0} \left[\max_{a_j' \in A_j} \mathbb{E}_{\mu^0} \left[u_i(a_j', a_{-j}) \middle|\, \hat{a}_j \right] \right].$$

It follows that

$$L_i \geq \max_{\substack{j \in N \\ \text{s.t. } j \sim i}} \mathbb{E}_{\mu_j^0} \left[\max_{a_j' \in A_j} \mathbb{E}_{\mu^0} \left[u_i(a_j', a_{-j}) \mid \hat{a}_j \right] \right].$$

Recalling that $\mu^0 \in \Delta(A)$ is just a 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.7 Proof of Theorem 3

The bracketed claim is Wen (1994, Thm. 2). Next, we prove the non-bracketed claim. We assume w.l.o.g. that $w_i^{\mathtt{cor}} = 0$ for all $i \in N$. We make use of the following auxiliary result.

Lemma A.2 (Abreu et al., 1994)**.** *Let $v = (v_1, \ldots, v_n) \in V$ such that $v_i > 0$ for all $i \in N$. Then, there exist payoff profiles $v^1, \ldots, v^n \in V$ that satisfy:*

(i) $v_i^j > 0$ *for all* $i, j \in N$*;*

(ii) $v_i > v_i^i$, *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. (Sketch) Select one representative from each equivalence class of players and apply the construction in Abreu et al. (1994, p. 942) to the feasible set for the reduced player set to obtain payoff vectors with the required strict inequalities. Subsequently, assign each player i the respective vectors of her equivalence class. For details, see Wen (1994, p. 952). □

We now prove the theorem. Take a payoff profile $v = (v_1, \ldots, v_n) \in V$, such that $v_i > 0$ for all $i \in N$. By feasibility, there is some distribution $\alpha \in \Delta(A)$ such that $u_i(\alpha) \equiv \mathbb{E}_\alpha[u_i(a)] = v_i$ for all $i \in N$. Let $\alpha_*^i \in \Delta(A)$ be an effective correlated minimax distribution against (the equivalence class of) player i, i.e., a solution to the minimization problem that defines $w_i^{\mathtt{cor}}$. Finally, for each $i \in N$, let $\alpha_{**}^i \in \Delta(A)$ be a correlated action profile that implements v^i, as identified in Lemma A.2. W.l.o.g., if $i \sim j$, we take $\alpha_*^i = \alpha_*^j$ and $v^i = v^j$, hence also $\alpha_{**}^i = \alpha_{**}^j$, so that punishment paths depend only on the deviator's equivalence class. For a given discount factor $\delta \in (0,1)$ and positive integer T, we will now characterize a canonical MSPE. By Lemma 1(iii), it suffices to specify the components of the MSPE for regular histories h. At any regular history, recommendations made by the device are necessarily regular and, hence, the obedience of players at former stages is common knowledge among the players and the device.

Phase A. If in every prior stage either all players have been obedient or at least two players have been disobedient, then $\mu(h) = \alpha$.

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 $\mu(h) = \alpha_*^i$.

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 stages ago, then $\mu(h) = \alpha_{**}^i$.

It is now standard to show (e.g., Wen, 1994, pp. 952-953) that there exists an integer $T \gg 0$ and some $\underline{\delta} \in (0,1)$ such that, for any $\delta \in (\underline{\delta}, 1)$, no player $i \in N$ has an incentive to deviate in phase A, and no player $j \in N$, whether equivalent to the last deviator i or not, has an incentive to deviate in phases B or C. Therefore, the profile above can be completed to an MSPE with expected payoff profile v. □

A.8 Proof of Lemma 3

We only prove that $w_1^{\texttt{cor}} = \frac{1}{4}$. Let $p_{a_1 a_2 a_3} \in [0,1]$ denote the probability that the triple $(a_1, a_2, a_3) \in A$ is chosen. We seek to minimize

$$\Phi(\alpha) = \max \left\{ \begin{array}{c} \max\{p_{\mathbf{FFF}}, p_{\mathbf{FSS}}\} + \max\{p_{\mathbf{SFF}}, p_{\mathbf{SSS}}\}, \\ \max\{p_{\mathbf{FFF}}, p_{\mathbf{SFS}}\} + \max\{p_{\mathbf{FSF}}, p_{\mathbf{SSS}}\}, \\ \max\{p_{\mathbf{FFF}}, p_{\mathbf{SSF}}\} + \max\{p_{\mathbf{FFS}}, p_{\mathbf{SSS}}\} \end{array} \right\},$$

subject to the probability simplex constraints. We note that $\Phi(\cdot)$ is convex and symmetric with respect to both arbitrary permutations of the set of players and a simultaneous swap of **F** and **S** in all three action spaces. Hence, we may restrict attention to solutions satisfying $p_{\mathbf{FFF}} = p_{\mathbf{SSS}}$ and

$$p_{\mathbf{FFS}} = p_{\mathbf{FSF}} = p_{\mathbf{SFF}} = p_{\mathbf{FSS}} = p_{\mathbf{SFS}} = p_{\mathbf{SSF}}.$$

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

A.9 Proof of Theorem 4

The bracketed claim is due to Friedman (1971). We prove the non-bracketed claim. Fix $v = (v_1, \ldots, v_n) \in V$ with $v_i > u_i(\alpha^*)$ for all $i \in N$. Choose $\alpha \in \Delta(A)$ such that $\mathbb{E}_\alpha[u] = v$. For any canonical history $h^t = (\hat{a}^0, a^0; \ldots; \hat{a}^{t-1}, a^{t-1})$, let

$$\mu(\cdot \,|\, h^t) = \begin{cases} \alpha, & \text{if } a^\tau = \hat{a}^\tau \text{ for all } \tau < t, \\ \alpha^*, & \text{otherwise.} \end{cases}$$

As has been explained in Section 2, it suffices to check that being obedient is sequentially rational at any regular information set $(h^t, \hat{a}_i^t)$. If h^t documents a deviation, then $\mu(\cdot \,|\, h^t) = \alpha^*$ and this remains true in future stages. Hence, a one-shot deviation affects only the current payoff. Since α^* is a correlated equilibrium, and $\mu_i(\hat{a}_i \,|\, h^t) > 0$ (from regularity), being obedient is sequentially rational in this case. If h^t documents no deviation, then $\mu(\cdot \,|\, h^t) = \alpha$. As a deviation changes the regime and $\mathbb{E}_{\alpha^*}[u_i(a)] < \mathbb{E}_\alpha[u_i(a)]$, being obedient is sequentially rational for player i also in this case, if δ is sufficiently close to one. This proves the claim. $\square$

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, 3, 5, 6, 9, 19, 25, 45

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. 2, 7

——— (1987) "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55 (1), 1–18, https://www.jstor.org/stable/1911154. 2, 7, 18

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. 8

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. 8

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

Bergemann, Dirk and Stephen Morris (2016) "Bayes correlated equilibrium and the comparison of information structures in games," *Theoretical Economics*, 11 (2), 487–522. 9

Chen, Bo (2008) "On effective minimax payoffs and unequal discounting," *Economics Letters*, 100 (1), 105–107. 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. 8

Compte, Olivier (1998) "Communication in Repeated Games with Imperfect Private Monitoring," *Econometrica*, 66 (3), 597–626, https://www.jstor.org/stable/2998576. 8

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. 2, 4, 5, 7, 8, 17, 18, 30, 31

——— (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. 7, 26

Forges, Françoise and Indrajit Ray (2024) "'Subjectivity and correlation in randomized strategies': Back to the roots," *Journal of Mathematical Economics*, 114, 103044. 7

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. 4, 6, 7, 28, 47

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, 24, 25

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, 3, 6, 9, 20, 25, 26, 27, 33

——— (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, 25, 26

Fudenberg, Drew and Jean Tirole (1991) *Game Theory*: MIT Press, https://mitpress.mit.edu/9780262061414/game-theory/. 35

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, 14, 19

Hart, Sergiu and David Schmeidler (1989) "Existence of Correlated Equilibria," *Mathematics of Operations Research*, 14 (1), 18–25, 10.1287/moor.14.1.18. 16

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. 8

Kreps, David M. and Robert Wilson (1982) "Sequential Equilibria," *Econometrica*, 50 (4), 863, https://www.jstor.org/stable/1912767?origin=crossref. 4

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. 8

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. 2, 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. 8

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. 25

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. 9, 19

——— (2011) "General Properties of Long-Run Supergames," *Dynamic Games and Applications*, 1 (2), 319–350, https://doi.org/10.1007/s13235-011-0018-3. 3, 9, 19

Rockafellar, R. Tyrrell (1970) *Convex Analysis*, 28 of Princeton Mathematical Series, Princeton, NJ: Princeton University Press. 15, 35

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, 20

Solan, Eilon and Nicolas Vieille (2002) "Correlated equilibrium in stochastic games," *Games and Economic Behavior*, 38 (2), 362–399. 7

Sorin, Sylvain (1986) "On Repeated Games with Complete Information," *Mathematics of Operations Research*, 11 (1), 147–160, https://www.jstor.org/stable/3690060. 6, 24

——— (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. 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. 8

——— (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, 9, 16, 18, 31

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

——— (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, 3, 5, 19, 23, 24, 25, 45, 46

——— (2002) "A Folk Theorem for Repeated Sequential Games," *The Review of Economic Studies*, 69 (2), 493–512, 10.1111/1467-937X.00214. 24

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. 25