Mediated Subgame Perfect Equilibrium

Christian Ewerhart & Haoyuan Zeng

May 20, 2026

Introduction: Motivation

Correlated equilibrium (Aumann, 1974, 1987) has become a cornerstone of noncooperative game theory. It assumes that a mediator sends private messages to the players. Compared to NE, the set of equilibrium outcomes may expand.

Question: How can one generalize correlated equilibrium to repeated games?

Problem: Even under perfect monitoring, private messages create incomplete information, which causes substantial ramifications:

- Sequential rationality is not required (Sorin, 1992; Renault and Tomala, 2011).
- With sequential rationality, the revelation principle may break down (Sugaya and Wolitzky, 2021).
- The set of implementable payoffs can often only be bounded.

This paper: Develops a simple solution concept based on transparent mediation.

Introduction: Benefits of transparent mediation

By transparent mediation, we mean that all private messages and internal records held by the mediator are disclosed at the end of each stage.

This assumption has two main benefits:

1. Simplicity

- We can essentially rely on subgame perfect equilibrium (Selten, 1965), keeping the explicit modeling of beliefs (Kreps and Wilson, 1982) at a minimum.
- There is no need for "ultimately dispensible" assumptions such as public randomization devices and observable mixed actions.

2. Practical appeal

- Transparency 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.

Introduction: Preview of results

1. New equilibrium concept

Mediated subgame perfect equilibrium (MSPE) generalizes correlated equilibrium to infinitely repeated games with perfect monitoring.

2. Revelation principle

Any MSPE can be replaced by an outcome-equivalent canonical MSPE, where recommendations are direct and players are obedient to "expected" recommendations.

3. Effective correlated minimax

The paper introduces an effective correlated minimax value that reflects both correlation in punishments and the possibility of deviations by utility-equivalent players.[a]

4. Perfect folk theorem

Necessary and sufficient conditions mirror the classical ones, but the effective correlated minimax value replaces the classic minimax value.

[a] This minimax value is the solution to a linear program.

Introduction: Some related literature

Perfect folk theorems[a]

- Friedman (1971): NE as threat points.
- Fudenberg and Maskin (1986): perfect folk theorem for $n = 2$ players or under dimensionality assumption.
- Abreu, Dutta, and Smith (1994): dimensionality can be replaced by *NEU*.
- Wen (1994): the right lower bound outside NEU is the *effective* minimax value.

[a]Parts of this literature assume public randomization devices and observable mixed actions (cf. Fudenberg et al., 2007).

Mediation and communication

- Aumann (1974, 1987): correlated equilibrium.
- Forges (1986), Myerson (1986): communication equilibrium.
- Prokopovych and Smith (2004): subgame perfect correlated equilibrium.
- Sugaya and Wolitzky (2021): failure of the revelation principle (even with internal records).

Introduction: Plan for the talk

MSPE

Examples

Revelation principle

Effective correlated minimax value

Necessary and sufficient conditions

Discussion

Conclusion

MSPE: Setup and notation

Repeated game with mediation

- Stage game: $G = \{A_i, u_i\}_{i \in N \equiv \{1,\dots,n\}}$ with finite action sets A_i and payoff functions $u_i : A \equiv \prod_{i \in N} A_i \to \mathbb{R}$.
- For each player i, a finite (endogenous) message space M_i. Moreover, $M = \prod_{i \in N} M_i$.
- Histories of length t,

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

  where $m^\tau \in M, a^\tau \in A$, form the set H^t.
- A device is a mapping $\mu : H \to \Delta(M)$, where $H = \bigcup_{t=0}^{\infty} H^t$.

Strategies and payoffs

- An information set is a pair (h^t, m_i^t).
- A behavior strategy is a mapping $\sigma_i : I_i \to \Delta(A_i)$, where I_i is the set of player i's information sets.
- Player i's expected payoff is $U_i(\sigma) = \mathbb{E}[(1-\delta)\sum_{t=0}^{\infty} \delta^t u_i(a^t)]$, where $\sigma = (\sigma_1, \dots, \sigma_n)$ and $\delta \in (0,1)$.
- Player i forms beliefs $\beta_i : I_i \to \Delta(M_{-i})$.
- Then, player i's continuation payoff $U_i(\sigma | h^t, m_i^t)$ is well-defined at every information set (h^t, m_i^t).

MSPE: Within-stage timing

1. Mediator draws
$m^t \sim \mu(\cdot \mid h^t)$
Privately sends m_i^t

→

2. Players act
Each i chooses a_i^t
given (h^t, m_i^t)

→

3. Public record updates
(m^t, a^t) becomes public
$h^{t+1} = (h^t; m^t, a^t)$

MSPE: Definition

A mediated subgame perfect equilibrium (MSPE) is a triple (μ, β, σ), consisting of a device μ, a system of beliefs $\beta = (\beta_1, \ldots, \beta_n)$, and a profile of behavior strategies $\sigma = (\sigma_1, \ldots, \sigma_n)$, satisfying:

Bayes consistency

At any information set (h, m_i) with $\mu_i(m_i \mid h) > 0$, beliefs satisfy

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

Sequential rationality

Given the device μ and beliefs β_i, player i has no incentive to deviate at any information set (h, m_i):

$$U_i(\sigma \mid h, m_i) \geq U_i(\sigma_i', \sigma_{-i} \mid h, m_i) \quad \text{for all } \sigma_i'.$$

MSPE: Basic properties (Lemma 1)

1. **One-shot deviation principle:** sequential rationality can be checked one information set at a time.
2. **MSPE payoffs nest SPE payoffs:** any payoff implementable by SPE with public randomization (irrespective of whether mixed actions are observable or not) is also implementable by MSPE (for the same δ).
3. **Irrelevance of probability-zero recommendations:** if Bayes consistency and sequential rationality hold at regular[1] information sets, the equilibrium can be completed to an outcome-equivalent MSPE.
4. **Relationship to correlated equilibrium:** the unconditional repetition of any correlated equilibrium of the stage game is an MSPE for every $\delta \in (0, 1)$.

Note: Part (4) implies existence of MSPE.

[1] We call an information set regular if the player called up to make a decision has no evidence of any device malfunction up to this point.

Example I: Private messages can be a substitute for patience

- Consider the correlated equilibrium $\alpha^* \in \Delta(A)$ that gives weight $\frac{1}{6}$ to each non-tie outcome.
- Unconditionally repeating α^* is an MSPE for every $\delta \in (0,1)$, with payoff $(1,1)$.
- But $(1,1)$ is not implementable as SPE for $\delta < \frac{1}{4}$.

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

Intuition

To implement $(1,1)$, ties must occur with probability zero. Hence, without private recommendations, at least one player's recommendation is a pure action. But then, for either this player or the other, a short-run deviation become profitable when patience is limited.

Backup on Example I

W.l.o.g., suppose player 2 chooses **R**, while player 1 mixes

$$\alpha_1 = p\,\mathbf{P} + (1-p)\,\mathbf{S},$$

where $p \in [0,1]$. Then the current expected payoff is $(2p,\ 2(1-p))$.

Player 1 can deviate to **P** and get 2 for sure, so her one-shot gain is $2-2p$. Player 2 can deviate to **S** and get $2p$, so her one-shot gain is $\max\{0,\ 4p-2\}$. Hence, some player has a deviation gain of at least

$$\min_{p\in[0,1]} \max\{2-2p, 4p-2\} = \tfrac{2}{3}.$$

However, the largest possible continuation loss from punishment is at most 2. Therefore, deterring a one-shot deviation requires

$$(1-\delta)\cdot \tfrac{2}{3} - \delta \cdot 2 \le 0,$$

i.e. $\delta \ge \frac{1}{4}$.

Example II: Mediation can simplify the implementation of a collusive outcome

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

- Consider the correlated equilibrium $\alpha^* \in \Delta(A)$, which randomizes uniformly over $(\mathbf{T}, \mathbf{C})$, $(\mathbf{M}, \mathbf{C})$, $(\mathbf{M}, \mathbf{R})$, with expected payoffs $\left(\frac{2}{3}, -\frac{1}{3}\right)$.
- Using α^* as a threat point, implement $(\mathbf{T}, \mathbf{L})$ with payoff $(1, 1)$, as an MSPE for $\delta \geq \frac{3}{4}$.
- The unique Nash equilibrium is $(\mathbf{M}, \mathbf{R})$, with payoff $(2, 0)$, so a Nash threat cannot sustain $(1, 1)$.

Intuition

The correlated equilibrium is used here as a threat point á la Friedman.[2]

[2] Obviously, this generalizes.

Backup: An analogue to Friedman's Theorem

Theorem

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

- This is the mediated analogue of Friedman's classic Nash reversion result.
- Because the set of correlated equilibria can be strictly larger than the set of Nash equilibria, there are additional threat points even in two-player games.

Example III: Mediation expands the set of implementable payoffs

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	$4, 0, 0$	$2, 0, 0$

MSPE

The payoff profile $(1, 1, 1)$ can be implemented by an MSPE for $\delta \geq \frac{4}{5}$:

- The mediator recommends $(\mathbf{T}, \mathbf{R}, \mathbf{B_1})$ and $(\mathbf{B}, \mathbf{L}, \mathbf{B_1})$ with equal probability. This yields

$$\frac{1}{2}(4, 0, 0) + \frac{1}{2}(-2, 2, 2) = (1, 1, 1).$$

- After a deviation, the mediator recommends $(\mathbf{T}, \mathbf{L}, \mathbf{B_1})$ and $(\mathbf{T}, \mathbf{R}, \mathbf{B_2})$ with equal probability, which gives $(0, 0, \frac{1}{2})$.

SPE

$(1, 1, 1)$ cannot be implemented by SPE because player 1's independent minimax value is $v_1^{\text{ind}} = 2$.

Example IV: Mediation matters also in games violating NEU

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

MSPE vs. SPE

- Consider the correlated equilibrium $\alpha^* \in \Delta(A)$ that puts weight $\frac{1}{13}$ on (F, F, F) and weight $\frac{4}{13}$ on each of (F, S, S), (S, F, S), and (S, S, F), with expected payoff of $\frac{4}{13} \approx 0.31$
- One can show that any SPE yields a payoff of at least $w_1^{\text{ind}} = \frac{4}{9} \approx 0.44$.

Revelation principle: Canonical MSPE

Definition

An MSPE is canonical if:

1. $M_i = A_i$ for every player i;
2. whenever a recommendation m_i is sent with positive probability, player i obeys it: she plays $a_i = m_i$ with probability one.

- The mediator is then **direct**: it recommends pure actions.
- Histories look like

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

  i.e., publicly recorded are recommendations and realized actions.
- Deviations are therefore publicly detectable at the end of each stage.

Revelation principle: Statement and idea of the proof

Theorem

For any MSPE, there exists an outcome-equivalent canonical MSPE.

Idea of the proof:

Step 1: Purify player randomization
Move any mixing by players into the device. The mediator now sends (m_i^t, a_i^t), i.e., the original message together with a recommended action.

→

Step 2: Make the device direct
Send only a_i^t to player i, i.e., drop the original messages. Notably, the mediator must keep the internal record $(m^0, \ldots, m^{t-1})$, extending the model at this point!

→

Step 3: Integrate the internal record out
To make the internal record obsolete, the device constructs the internal record from the (regular) public history using conditional probabilities on past recommendations.

Backup: Why the proof is nontrivial in this setting

What is lost when we go direct?

The original mediator may know more than the players because it remembers earlier private messages. A transparent canonical mediator is *not* allowed to keep that hidden state.

How the paper repairs this

At any regular public canonical history $\hat{h}^t$, we have $\Pr(\hat{h}^t) > 0$ and the mediator computes

$$\hat{\mu}(\hat{a}^t \mid \hat{h}^t) = \frac{\Pr(\hat{h}^t, \hat{a}^t)}{\Pr(\hat{h}^t)}$$

from the latent-state distribution induced by the original equilibrium.

- This is where the paper differs from models that allow internal messages to future selves.

Minimax values: Concepts in the literature

Independent minimax

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

Opponents mix independently.

Correlated minimax

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

Opponents may coordinate their punishment.

Effective independent minimax

Players i, j have equivalent utilities $(i \sim j)$,[3] if there exist $c \in \mathbb{R}$ and $d > 0$ such that $u_i(a) = c + d\, u_j(a)$ for all $a \in A$.

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

Accounts for the fact that if some $j \sim i$, then player j may also undo the punishment.

2) E.g., cartels, criminal organizations, tightly aligned teams, or political coalitions.

Minimax values: Effective correlated minimax value

- Start from a correlated action profile $\alpha \in \Delta(A)$.
- Suppose player $j \sim i$ receives recommendation a_j and then deviates optimally.[a]
- Evaluate how much utility player i can then be forced to receive.
- Finally, choose the best correlated punishment for the opponents: minimize over α.

[a] If a_j is recommended with probability zero, the inner, conditional expectation can be defined arbitrarily.

Definition

$$w_i^{\mathrm{cor}} = \min_{\alpha \in \Delta(A)} \max_{j \sim i} \mathbb{E}_{\alpha_j}\left[\max_{a'_j \in A_j} \mathbb{E}_{\alpha}[u_i(a'_j, a_{-j}) \mid a_j]\right].$$

Backup: Why w_i^{cor} is well-defined

Rewriting the outer expectation

$$\mathbb{E}_{\alpha_j}\left[\max_{a_j'} \mathbb{E}_{\alpha}[u_i(a_j', a_{-j}) \mid a_j]\right] = \sum_{a_j \in A_j} \max_{a_j' \in A_j} \sum_{a_{-j} \in A_{-j}} \alpha(a_j, a_{-j}) u_i(a_j', a_{-j}).$$

- The right-hand side is a maximum of linear functions in α and can, therefore, be determined by solving a linear program.
- This is a major computational advantage over the independent effective minimax problem when $n \geq 3$, where product probabilities create nonlinearity.

Minimax values: How the minimax concepts compare

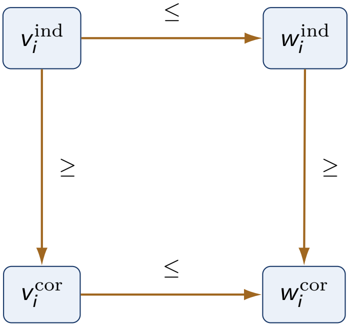

If NEU holds

There are no equivalent-utility players:

$$w_i^{\text{ind}} = v_i^{\text{ind}}, \qquad w_i^{\text{cor}} = v_i^{\text{cor}}.$$

If $n = 2$

Correlation does not matter:

$$v_i^{\text{cor}} = v_i^{\text{ind}}, \qquad w_i^{\text{cor}} = w_i^{\text{ind}}.$$

Folk theorems: Necessary conditions

Theorem

If $U_i(\sigma)$ is player i's equilibrium payoff under an MSPE, then

$$U_i(\sigma) \geq w_i^{\mathrm{cor}}.$$

For ordinary SPE, the corresponding lower bound is

$$U_i(\sigma) \geq w_i^{\mathrm{ind}}.$$

Proof intuition

- Canonicalize the equilibrium via the revelation principle.
- The initial recommendation distribution is then a correlated profile $\mu^0 \in \Delta(A)$.
- Conditional one-shot deviations by players in i's equivalence class generate the lower bound w_i^{cor}.

Folk theorem: Sufficient conditions

Theorem

Let $v \in V$ be feasible and suppose that

$$v_i > w_i^{\mathrm{cor}} \qquad \text{for all } i \in N.$$

Then there exists $\underline{\delta} < 1$ such that, for every $\delta \in (\underline{\delta}, 1)$, the repeated game has an MSPE with payoff vector v.

- This is the exact analogue of Wen's theorem, but with correlated punishments.
- The auxiliary assumption that mixed actions are observable is no longer needed.
- Transparent mediation directly encodes the relevant randomization in messages and public audit trails.

Two corollaries

If NEU holds

Since $w_i^{\text{cor}} = v_i^{\text{cor}}$, any feasible payoff with

$$v_i > v_i^{\text{cor}} \quad \forall i$$

is an MSPE payoff for sufficiently patient players.

If $n = 2$

Since $w_i^{\text{cor}} = w_i^{\text{ind}}$, the limiting set of strictly individually rational MSPE payoffs coincides with the familiar two-player SPE benchmark.[a]

[a] This is a bit tricky. If NEU holds, then $w_i^{\text{ind}} = v_i^{\text{ind}}$. If NEU fails, then w.l.o.g. $u_1 = u_2$. Then, strict individual rationality amounts to that $v_1 = v_2 > \max\{v_1^{\text{ind}}, v_2^{\text{ind}}\} = w_1^{\text{ind}} = w_2^{\text{ind}}$.

Discussion: A comment on the FM86 counterexample

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	$0,0,0$	$1,1,1$

- Fudenberg and Maskin (1986) showed that $w_i^{\text{ind}} = \frac{1}{4} > 0 = v_i^{\text{ind}}$.
- Strictly speaking, this observation does not show that the dimensionality assumption cannot be dropped from the folk theorem (because w_i^{ind} does not reflect the possiblity of public randomization).
- We show that even $w_i^{\text{cor}} = \frac{1}{4}$, i.e., mediation cannot push a player's expected payoff below $\frac{1}{4}$.

Implication

The FM86 counterexample is robust, in particular, to the introduction of public randomization.

Discussion: Weak vs. strict individual rationality

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	$0, 0, 0$	$1, -1, 1$

Revisiting Forges, Mertens, and Neyman (1986)

- The feasible payoff $v = (1, 0, 0)$ is weakly individually rational.
- The stage game satisfies NEU, so $w_i^{\mathrm{cor}} = v_i^{\mathrm{cor}} = 0$ for each player.
- Yet v is not implementable as either SPE or MSPE.

Discussion: A comment on Sugaya and Wolitzky (2021)

Purpose of the example

The example shows that the definition of an autonomous device in Forges (1986), without internal records, and Sugaya and Wolitzky (2021), with internal records, differs.

Payoffs

Rows are player 1's stage-1 actions, columns are player 2's stage-1 actions, and the three matrices correspond to player 3's stage-2 action.

E_1

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

E_2

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

E_3

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

Discussion: The nature of internal records

Two-stage environment G^7

- Players 1 and 2 move in stage 1.
- Player 3 moves in stage 2.
- Neither player 3 nor the device observes stage-1 actions.

Noncanonical mediated plan

At stage 1, the device privately sends either e_1 or e_2 to players 1 and 2, thereby selecting between *two mixed Nash equilibria with the same support but different probabilities*.

Why that matters

At stage 2, player 3 should choose E_1 after e_1 and E_2 after e_2.

- If a canonical device remembers only the public stage-1 recommendations, it may lose the distinction between e_1 and e_2.
- Then player 3 may prefer a third action E_3, breaking the equilibrium.

Why Bayesian reconstruction is needed

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

Conditional on e_1

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

Conditional on e_2

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

Aggregate

- The aggregate distribution forgets which latent message was sent.
- But player 3's optimal stage-2 action depends on that latent state.

Conclusion

- **MSPE:** combines private recommendations, sequential rationality, and ex-post transparency.
- **Revelation principle:** dynamic mediated play can be reduced to direct action recommendations.
- **Lower bound:** implementable payoffs must weakly exceed w_i^{cor}.
- **Perfect folk theorem:** any feasible payoff strictly above w_i^{cor} for every player is attainable when players are sufficiently patient.
- **Economic lesson:** mediation can make sanctions simpler/more effective, and may serve as a substitute for patience.

More speculative

What changes if ex-post observability is dropped and the mediator is allowed to keep secrets across stages? The sufficient conditions may extend, but the necessary conditions are less clear.