

Axioms for Correlated Equilibrium

Florian Brandl[†]

Institute for Microeconomics, University of Bonn

We characterize correlated equilibrium in finite normal-form games. Interpreting correlated strategies as action recommendations, we show that correlated equilibrium is the unique solution concept that never recommends a pure-strategy dominated action, treats payoff-equivalent actions interchangeably, and respects the sure-thing principle under uncertainty about payoffs and the correlation device. A parallel characterization identifies coarse correlated equilibrium among solution concepts that recommend dominant actions whenever they exist and treat payoff-equivalent actions as strongly interchangeable.

1. Introduction

Public randomization in strategic games induces correlation among the players' strategies. Such correlated strategies enable outcomes that are infeasible when players randomize independently. Game-theoretic analysis can thus draw from a larger pool of solution concepts when describing, predicting, or explaining outcomes of games. To guide the choice of the solution concept from among this wide range of possibilities, we take an axiomatic approach: we formulate axioms that require correlated strategies to be selected coherently across games and determine which solution concepts satisfy those. We thereby obtain characterizations of correlated equilibrium and coarse correlated equilibrium. The characterizations clarify which assumptions underlie each solution concept and can guide which one to use.

In our model, the set of players is fixed, and we consider normal-form games with a finite but variable set of actions for each player. A correlated strategy is a distribution over action profiles, interpreted as a correlation device that recommends an action to each player. A solution concept assigns to each game a set of correlated strategies. Nash equilibrium, correlated equilibrium, and coarse correlated equilibrium are all examples of solution concepts. The axioms we impose require that correlated strategies are selected coherently across games and that they respect rationality.

[†]florian.brandl@uni-bonn.de

Consistency addresses payoff uncertainty. Nature draws payoffs, and each draw yields a deterministic game. If a correlated strategy is selected for every realization, consistency requires it to be selected for the lottery over realizations. Rather than model lotteries explicitly, we identify them with convex combinations of the realized games, replacing random payoffs by their expectations. In this sense, consistency applies the sure-thing principle of [Savage \(1954\)](#) to strategic games.

Consequentialism specifies how the solution concept responds when an action is duplicated. Two actions of a player are *clones* if every player’s utility is the same under either action, regardless of the other players’ actions. In a two-player game, clones correspond to rows or columns that are identical in both players’ payoff matrices. When we add a clone, consequentialism allows the probability of any action profile involving the original action to be split in a fixed ratio between that profile and the profile in which the original action is replaced by its clone; all other action-profile probabilities stay fixed. Requiring the same ratio across all affected profiles ensures that recommendations of either clone convey the same information about the other players’ actions.¹

Rationality rules out recommendations of strictly dominated actions: an action profile receives probability 0 whenever some player’s action is strictly dominated in pure strategies. We also require the solution concept to be continuous and to return a nonempty, convex set of correlated strategies in each game. Convex-valuedness keeps mixtures of plausible correlation devices plausible, so the solution concept does not exclude outcomes that arise from uncertainty about which device is used. Like consistency, this is another instance of the sure-thing principle.²

In correlated equilibrium, following the recommendation is optimal for a Bayesian expected utility maximizer. Our main result ([Theorem 1](#)) shows that correlated equilibrium is the only solution concept satisfying these axioms. Equivalently, these axioms are exactly the coherence requirements that pin down correlated equilibrium. In particular, they encode Bayesian expected-utility maximization.

The second characterization strengthens consequentialism and weakens rationality. Strong consequentialism drops the fixed-ratio requirement: when an action is cloned, probability can be split arbitrarily between the original and the clone. Clones can therefore carry different information about the other players’ strategies. Correlated equilibrium violates strong consequentialism because changing the information carried by a recommendation can make it suboptimal to follow either clone when recommended. Combined with our other axioms, this yields an impossibility. Coarse correlated equilibrium, introduced by [Moulin and Vial \(1978\)](#), requires that each player’s expected utility from always following the recommendation is at least as high as the utility from committing to a fixed action before observing the recommendation. Because this incentive constraint does not depend on how probability is split between clones, coarse cor-

¹For strategy profiles with independent randomization, consequentialism is weaker than the invariance axiom of [Kohlberg and Mertens \(1986\)](#), which also allows introducing actions that are convex combinations of existing ones.

²Section 7 discusses the behavioral implications of the axioms in detail.

related equilibrium respects strong consequentialism. It violates rationality, however, because it rules out only unconditional deviations and need not make following the recommendation optimal after it is observed. On the other hand, it is optimal to commit to dominant actions whenever they exist, so coarse correlated equilibrium satisfies the weaker form of rationality that requires dominant actions to be recommended. Our second result (Theorem 2) characterizes coarse correlated equilibrium as the unique continuous and convex-valued solution concept that satisfies consistency, strong consequentialism, and weak rationality, and that selects at least one correlated strategy for each game.

Both results contribute to the recent literature on axiomatic characterizations of solution concepts. In addition to showing which solution concepts are forced upon us when committing to a set of axioms, their necessity parts clarify which properties one commits to when using the characterized solution concept.

2. Related Work

Most axiomatic characterizations for normal-form games treat only uncorrelated strategies, i.e., solution concepts that return profiles of independent strategies. The closest work to ours is [Brandl and Brandt \(2024\)](#), who show that Nash equilibrium is the only total solution concept satisfying consistency, consequentialism, and weak rationality. Once correlation is allowed, these axioms no longer pin down a unique concept: Nash equilibrium, correlated equilibrium, and coarse correlated equilibrium all satisfy them. Their proofs rely on independent randomization and do not extend to correlated strategies. [Sandomirskiy et al. \(2025\)](#) introduce a narrow bracketing axiom: if two strategy profiles are selected in two independent games, then their product profile is selected when the games are played simultaneously.³ They show that any total solution concept satisfying narrow bracketing, anonymity, rationality, and monotonicity in expected payoffs is a refinement of Nash equilibrium.

Building on [Peleg and Tijs \(1996\)](#), [Norde et al. \(1996\)](#) characterize Nash equilibrium using utility maximization in one-player games and a consistency condition that varies the set of players. Their consistency requires that any strategy profile returned for an n -player game is also returned for the $(n - k)$ -player game obtained by fixing the strategies of k players. Because our model keeps the player set fixed and allows correlated strategies, their characterization does not apply here. Other axiomatic work on Nash equilibrium includes characterizations of pure Nash equilibrium ([Voorneveld, 2019](#)), Nash equilibrium for games with quasiconcave utility functions ([Salonen, 1992](#)), and a choice-theoretic characterization ([Crescenzi, 2026](#)).

[Brandl and Brandt \(2019\)](#) consider solution concepts for two-player zero-sum games that return a set of strategies for a single player instead of strategy profiles or correlated strategies. They show that returning all maximin strategies is the coarsest such solution concept that satisfies consistency, consequentialism, and rationality. While we consider two-player zero-sum

³As noted by [Sandomirskiy et al. \(2025\)](#), consistency and consequentialism imply narrow bracketing for uncorrelated strategies.

games in our proofs, there is no methodological connection with their work.

Epistemic game theory studies what players need to know about other players to justify Nash equilibrium. The players' knowledge is modeled using Bayesian belief hierarchies, which consist of a game and a set of types for each player; a type includes the action played by this type and a belief about the types of the other players (Harsanyi, 1967). That is, players have probabilistic beliefs about other players' types, but play a deterministic action conditional on their type. A rational player maximizes their expected payoff given their type. Aumann and Brandenburger (1995) show that for two-player games, the beliefs of every pair of types constitute a Nash equilibrium if their beliefs and rationality are mutually known. This result extends to games with more than two players if the beliefs are commonly known and admit a common prior. Barelli (2009), Hellman (2013), and Bach and Tsakas (2014) show that the results of Aumann and Brandenburger (1995) still hold under weaker common knowledge assumptions. Aumann and Drèze (2008) study the payoff a player can expect in a game with common knowledge of rationality and a common prior. They characterize rational expectations via correlated equilibria of an extended game: it is rational to expect a given payoff precisely when it is a recommendation-contingent expected payoff in some correlated equilibrium of the doubled game, where each of the player's actions is replaced by two clones.

The literature on equilibrium refinements takes equilibrium play as given and singles out equilibria with additional properties. It focuses on equilibria that are strategically stable and robust with respect to the representation of a game. The closest connection to our work is that the invariance axiom of Kohlberg and Mertens (1986)—the equilibrium selection only depends on the reduced normal form of a game—implies consequentialism. The invariance to embedding condition of Govindan and Wilson (2012) is a further strengthening that also considers variable sets of agents.

3. Preliminaries

Let U be an infinite universal set of actions, and let $\mathcal{F}(U)$ be the collection of nonempty finite subsets of U . Fix the player set $N = \{1, \dots, n\}$. For action sets $A_1, \dots, A_n \in \mathcal{F}(U)$, write $A = A_1 \times \dots \times A_n$ for the set of action profiles. An n -player normal-form game on A is a function $G: A \rightarrow \mathbb{R}^n$; player i 's payoff at $a \in A$ is $G_i(a)$. A correlated strategy is a probability distribution on A , and the set of correlated strategies is denoted by $\Delta(A)$. Interpret a correlated strategy p as a signal structure: a joint signal a is drawn from p , and player i observes a_i as a recommendation to play a_i . Throughout, *action* means a pure strategy, *strategy* a mixed strategy, and *profile* a player-indexed vector.

We say that $G: A \rightarrow \mathbb{R}^n$ is a blow-up of $G': A' \rightarrow \mathbb{R}^n$ if G is obtained from G' by replacing each action with one or more payoff-equivalent actions. That is, there is a surjection $\phi = (\phi_1, \dots, \phi_n): A \rightarrow A'$ with $\phi_i: A_i \rightarrow A'_i$ such that $G = G' \circ \phi$. Thus, actions in $\phi_i^{-1}(a'_i)$ are payoff-equivalent copies, called *clones*, of $a'_i \in A'_i$, and G is obtained from G' by replacing a'_i by $|\phi_i^{-1}(a'_i)|$ clones. A correlated strategy $p \in \Delta(A)$ induces a correlated strategy $\phi_*(p) = p \circ \phi^{-1} \in \Delta(A')$

with $\phi_*(p)(a') = \sum_{a \in \phi^{-1}(a')} p(a)$ for $a' \in A'$.⁴

A solution concept f maps every game G to a set of correlated strategies $f(G)$ on the action profiles of G . Call f *total* if $f(G) \neq \emptyset$ for each G , *continuous* if it is an upper hemi-continuous correspondence, and *convex-valued* if $f(G)$ is convex for each G .⁵

Three common solution concepts are Nash equilibrium, correlated equilibrium, and coarse correlated equilibrium. A *Nash equilibrium* is a correlated strategy $p \in \Delta(A)$ that is a product distribution of $p_i \in \Delta(A_i)$, $i \in N$, such that (p_1, \dots, p_n) is a Nash equilibrium of G in the usual sense.⁶

A *correlated equilibrium* is a correlated strategy $p \in \Delta(A)$ such that for each $i \in N$ and $a_i, b_i \in A_i$,

$$\sum_{a_{-i} \in A_{-i}} p(a_i, a_{-i}) (G_i(a_i, a_{-i}) - G_i(b_i, a_{-i})) \geq 0$$

Thus, in a correlated equilibrium, it is optimal for each player to follow their recommendation when conditioning the distribution of other players' action profiles on the recommendation. Nash equilibria are correlated equilibria where the players' recommendations are independent, so that conditioning on a player's recommendation does not change the distribution of other players' action profiles.

A *coarse correlated equilibrium* is a correlated strategy $p \in \Delta(A)$ such that for each $i \in N$ and $b_i \in A_i$,

$$\sum_{a \in A} p(a_i, a_{-i}) (G_i(a_i, a_{-i}) - G_i(b_i, a_{-i})) \geq 0$$

Each player is weakly better off always following the recommendation than committing ex ante to a fixed action. Each correlated equilibrium is a coarse correlated equilibrium since the former demands that deviations are not even profitable after observing the recommendation.

We denote by *NE*, *CE*, and *CCE* the solution concepts that return the set of all Nash equilibria, correlated equilibria, and coarse correlated equilibria, respectively, for each game.

4. The axioms

Consistency requires that any correlated strategy returned in two games with the same action sets is also returned in any convex combination of the two games.

⁴Blow-ups of games also appear in work on equilibrium refinements. The invariance axiom of [Kohlberg and Mertens \(1986\)](#) uses a more permissive notion of blowing up. There, G' is a reduced form of G if G' is obtained from G by deleting actions that are convex combinations of other actions. [Chatterji and Govindan \(2006\)](#) connect two refinements of correlated equilibrium—perfect correlated equilibrium and perfect direct correlated equilibrium (introduced by [Dhillon and Mertens \(1996\)](#)) through blow-ups: every perfect correlated equilibrium is induced by a perfect direct correlated equilibrium in a blown-up game.

⁵ f is upper hemi-continuous if for each $A \in \mathcal{F}(U)^n$, each sequence of games (G^ℓ) on A converging to G , each sequence of correlated strategies (p^ℓ) with $p^\ell \in f(G^\ell)$ converging to p , we have $p \in f(G)$.

⁶Since we work with correlated strategies throughout, it is more convenient to view a Nash equilibrium as a correlated strategy that is the product of independent strategies rather than as a strategy profile consisting of a strategy for each player.

Definition 1 (Consistency). A solution concept f satisfies consistency if for any two games G, G' on A and any $\lambda \in [0, 1]$,

$$f(G) \cap f(G') \subseteq f(\lambda G + (1 - \lambda)G').$$

Consistency constrains how a solution concept responds to payoff uncertainty. If a correlated strategy p is returned for both G and G' , then p is also returned for a lottery that chooses between G and G' . We represent such lotteries by convex combinations, replacing random payoffs by their expectations. In this sense, consistency applies the sure-thing principle of [Savage \(1954\)](#) to strategic games. [Section 7](#) analyzes the behavioral content of this axiom and the others.

In the game theory literature, consistency appears in characterizations of Nash equilibrium ([Brandl and Brandt, 2024](#); [Kalai and Kalai, 2024](#)) and of maximin strategies in two-player zero-sum games ([Brandl and Brandt, 2019](#)).⁷ Variants of consistency have been considered for single-player decision problems with uncertainty about the state of nature (see, e.g., [Chernoff, 1954](#); [Milnor, 1954](#); [Gilboa and Schmeidler, 2003](#)). Analogs of consistency feature prominently in axiomatic characterizations in social choice theory, where it relates the choices for different sets of voters to each other (see, e.g., [Smith, 1973](#); [Young, 1975](#); [Young and Levenglick, 1978](#); [Myerson, 1995](#); [Brandl et al., 2016](#); [Lackner and Skowron, 2021](#)). [Harsanyi and Selten \(1972\)](#) use a variant of consistency (their Axiom 8) to characterize an extension of Nash’s bargaining solution to bargaining under uncertainty. The characterization of the Shapley value by [Shapley \(1953\)](#) also involves an additivity axiom (which he calls law of aggregation) that is similar in spirit to consistency.

Consequentialism imposes invariance to cloning actions. When an action is replaced by two payoff-equivalent clones, probability mass can move only between action profiles that differ by which clone is used; all other action-profile probabilities remain unchanged. This invariance has two natural interpretations, which lead to different definitions of consequentialism. Consider the following two games and corresponding correlated strategies.

$$G' = \begin{pmatrix} 1, 0 & 0, 0 \\ 0, 0 & 1, 0 \end{pmatrix} \quad p' = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ 0 & 0 \end{bmatrix} \quad G = \begin{pmatrix} 1, 0 & 0, 0 \\ 1, 0 & 0, 0 \\ 0, 0 & 1, 0 \end{pmatrix} \quad p = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} \\ \frac{1}{6} & \frac{1}{6} \\ 0 & 0 \end{bmatrix} \quad \tilde{p} = \begin{bmatrix} \frac{1}{2} & \frac{1}{6} \\ 0 & \frac{1}{3} \\ 0 & 0 \end{bmatrix}$$

Here, G is obtained from G' by cloning the row player’s first action. Both p and \tilde{p} induce p' because each assigns probability $\frac{1}{2}$ to the two clones of the action profiles in the first row of G' . However, p and \tilde{p} distribute probability differently across the two clones. Under p , a recommendation of either clone induces the same distribution over columns, and therefore the same belief about the column player’s recommendation; under \tilde{p} , the induced beliefs differ. Under the weaker notion, p is returned for G exactly when p' is returned for G' . Under the stronger notion, the same equivalence must also hold for \tilde{p} .

⁷[Kalai and Kalai \(2024\)](#) call the axiom the sure thing principle instead of consistency.

Suppose G is a blow-up of G' with surjection $\phi: A \rightarrow A'$. A correlated strategy $p \in \Delta(A)$ is *clone-symmetric* if for every $i \in N$ and $a_i \in A_i$ there exists $\alpha \geq 0$ such that $p(a_i, a_{-i}) = \alpha \phi_*(p)(\phi(a_i, a_{-i}))$ for all $a_{-i} \in A_{-i}$. Equivalently, any two clones of the same action in G' either induce the same conditional distribution on A_{-i} or receive probability 0. In the example, p is clone-symmetric, whereas \tilde{p} is not.

Definition 2 ((Strong) consequentialism). A solution concept f satisfies consequentialism if for all games G, G' such that G is a blow-up of G' with surjection ϕ and each clone-symmetric $p \in \Delta(A)$,

$$p \in f(G) \text{ if and only if } \phi_*(p) \in f(G')$$

Moreover, f satisfies strong consequentialism if for all games G, G' such that G is a blow-up of G' with surjection ϕ ,

$$p \in f(G) \text{ if and only if } \phi_*(p) \in f(G')$$

Consequentialism requires that a clone-symmetric correlated strategy is returned in G if and only if its induced correlated strategy is returned in G' . Equivalently, when adding a clone a'_i of a_i of player i , (i) for any $\alpha \in [0, 1]$, the probability of any action profile involving a_i can be split among its two clones so that the first clone receives an α fraction of the probability, and (ii) the probability of any action profile not involving a_i remains the same. Strong consequentialism requires that a correlated strategy is returned in G if and only if its induced correlated strategy is returned in G' . Equivalently, (i') the probability of an action profile including a_i may be split arbitrarily among its two clones, and (ii) holds. Under strong consequentialism, the two clones may induce different distributions over the action profiles of the players $N \setminus \{i\}$. Strong consequentialism implies consequentialism, and consequentialism implies equivariance: relabeling a player's actions relabels the returned correlated strategies in the same way (see Section 7). When restricting to strategies with independent randomization, both notions of consequentialism agree since such strategies are automatically clone-symmetric.

(Strong) consequentialism appears in the characterizations of [Brandl and Brandt \(2019, 2024\)](#); [Kalai and Kalai \(2024\)](#) uses the term strategy anonymity. The invariance axioms of [Kohlberg and Mertens \(1986\)](#) and [Govindan and Wilson \(2012\)](#) strengthen consequentialism by also allowing convex combinations of existing actions. Decision theory uses closely related invariance conditions as well. For example, consequentialism corresponds to Postulate 6 (cloning a player's actions) together with Postulate 9 (cloning Nature's states, i.e., opponents' actions) in [Chernoff \(1954\)](#). Postulate 9 also appears as column duplication ([Milnor, 1954](#)) and as deletion of repetitive states ([Arrow and Hurwicz, 1972](#); [Maskin, 1979](#)). In social choice theory, [Tideman \(1987\)](#) introduced independence of clones, a parallel condition for cloning social alternatives (see also [Zavist and Tideman, 1989](#); [Brandl et al., 2016](#)).

Rationality prescribes that actions that are dominated in pure strategies are never played. Weak rationality is weaker: it requires only that a dominant action is played with positive

probability. Formally, for $a_i, a'_i \in A_i$, action a_i *dominates* a'_i if $G_i(a_i, a_{-i}) > G_i(a'_i, a_{-i})$ for each $a_{-i} \in A_{-i}$. Action a_i is *dominant* if it dominates every other action in A_i .

Definition 3 ((Weak) rationality). A solution concept f satisfies rationality if for each game G , each $p \in f(G)$, each $i \in N$, and each $a_i \in A_i$,

$$\text{if } a_i \text{ is dominated, then } p(a_i, a_{-i}) = 0 \text{ for each } a_{-i} \in A_{-i}$$

A solution concept f satisfies weak rationality if for each game G , each $p \in f(G)$, each $i \in N$, and each $a_i \in A_i$,

$$\text{if } a_i \text{ is dominant, then } p(a_i, a_{-i}) > 0 \text{ for some } a_{-i} \in A_{-i}$$

Both axioms refer only to pure-strategy dominance, so they do not constrain how players evaluate lotteries. They also impose no assumptions about beliefs about opponents' rationality. Rationality is equivalent to the strong domination axiom of [Milnor \(1954\)](#) and to Property 5 of [Maskin \(1979\)](#); it is weaker than Postulate 2 of [Chernoff \(1954\)](#).

5. Results

Correlated equilibrium satisfies consistency, consequentialism, and rationality. In a correlated equilibrium, each player maximizes expected utility given the recommendation, and these axioms capture that behavior. Among total, continuous, and convex-valued solution concepts, only *CE* satisfies all three axioms.

Theorem 1. *CE is the only total, continuous, and convex-valued solution concept that satisfies consistency, consequentialism, and rationality.*

Theorem 1 makes two claims. First, any solution concept satisfying the axioms returns every correlated equilibrium. Second, any coarsening of *CE* violates at least one axiom. For example, *NE* violates convex-valuedness and *CCE* violates rationality. The first claim forces the selection of every correlated strategy consistent with Bayesian expected utility maximization, which defines correlated equilibrium. The second claim rules out coarsenings that admit behavior that violates Bayes-rationality. Equivalently, the theorem characterizes Bayesian expected utility maximization. If consequentialism is replaced by strong consequentialism, no solution concept can satisfy the resulting list of axioms, because *CE* fails strong consequentialism.

Coarse correlated equilibrium corresponds to players who either commit to a fixed action ex ante or always follow the recommendation. Equivalently, players maximize expected utility without updating beliefs about opponents' strategies after receiving a recommendation. Under either interpretation, strong consequentialism is natural: if a player commits before observing the recommendation, or ignores the information it contains, then it does not matter how probability is split between clones. While coarse correlated equilibrium violates rationality, it

satisfies weak rationality.⁸ Committing to a dominant action ex ante is always profitable; hence any coarse correlated equilibrium plays dominant actions with probability 1. The second result characterizes *CCE* using strong consequentialism and weak rationality.

Theorem 2. *CCE is the only total, continuous, and convex-valued solution concept that satisfies consistency, strong consequentialism, and weak rationality.*

Weak rationality rules out coarsenings of *CCE*, while refinements of *CCE* (such as *NE* and *CE*) tend to violate strong consequentialism. Thus, Theorem 2 characterizes expected utility maximization when players can choose between committing to a fixed action ex ante and always following the recommendation.

6. Proof outlines

Throughout this section, fix a solution concept f that is total, continuous, and convex-valued and satisfies consistency, consequentialism, and weak rationality. The proofs of Theorem 1 and Theorem 2 extensively use decomposition: write a game as a convex combination of simpler games with the same equilibrium; then use consistency to conclude that the equilibrium is selected in the original game. Both theorems rely on two preliminary results.

- (i) f returns all pure Nash equilibria in all games (Proposition 1).
- (ii) f returns all Nash equilibria in all essentially two-player zero-sum games (Proposition 2).⁹

In an essentially two-player zero-sum game, two players play a zero-sum game; all other players are dummies, with constant utility functions and no effect on others' payoffs.

The proof of Proposition 1 has three steps. First, we show that only those action profiles can be played that survive iterated restriction to dominant actions: if a player has a dominant action, remove all their other actions, and recurse. If only a single action profile survives, it is played with probability 1. Second, any game with a quasi-strict pure Nash equilibrium is a convex combination of games for which only this action profile survives. The first part and consistency imply that quasi-strict pure Nash equilibria are played. Third, we extend the claim to all pure Nash equilibria using continuity.

The proof of Proposition 2 starts from two-player matching-pennies games. They have enough symmetries to show that the unique mixed Nash equilibrium—the uniform distribution over all

⁸Consider the two-player game G and correlated strategy p below.

$$G = \begin{pmatrix} 4, 0 & 0, 0 \\ 2, 0 & 2, 0 \\ 3, 0 & 3, 0 \end{pmatrix} \quad p = \begin{bmatrix} \frac{1}{2} & 0 \\ \frac{1}{4} & \frac{1}{4} \\ 0 & 0 \end{bmatrix}$$

The row player's second action is dominated by the third, and p is a coarse correlated equilibrium with positive probability on the second row.

⁹Proposition 2 proves this statement for games that are zero-sum after a positive affine transformation of one agent's utility function, rather than honest zero-sum games. This difference is immaterial for the argument.

action profiles—is played. Any essentially two-player zero-sum game is a convex combination of blow-ups of matching pennies games. Together, this gives that the uniform distribution over all action profiles is played in an essentially two-player zero-sum game whenever it is a Nash equilibrium. The third step heavily uses consequentialism to extend this statement to Nash equilibria that are neither uniform nor of full support.

The third part shows that any solution concept that returns all Nash equilibria in essentially two-player zero-sum games has a stronger property: it returns all correlated equilibria in all games (Appendix D). At this point, we know that f is a coarsening of CE .

To prove Theorem 1, strengthen weak rationality to rationality. For any correlated strategy that is not a correlated equilibrium of a given game, we construct a second game with two properties: the same strategy is a correlated equilibrium, and some action that is played in equilibrium is dominated in a convex combination of both games. As a coarsening of CE , f returns this strategy for the second game, and thus also for the convex combination by consistency. This contradicts rationality, and shows that f is a refinement of CE .

To prove Theorem 2, assume strong consequentialism in place of consequentialism. The argument has two parts. First, fix any coarse correlated equilibrium of any game. Strong consequentialism lets us reduce to a case in which the equilibrium has additional structure. We then express the game as a convex combination of games for which the same correlated strategy is a correlated equilibrium. Since f is a coarsening of CE , consistency implies that f also returns the strategy in the original game, so f is a coarsening of CCE . Second, to show that f is a refinement of CCE , we use the same type of construction as above, but now produce a dominant action that is never played, contradicting weak rationality.

7. Discussion

Remark 1 (Independence of the axioms). Theorem 1 and Theorem 2 cease to hold if any of totality, consistency, (strong) consequentialism, or (weak) rationality are omitted. Moreover, convex-valuedness is required for Theorem 1. We show this with a series of examples.

- (i) *Totality:* Return all correlated strategies that randomize only over pure Nash equilibria. This solution concept satisfies continuity, convex-valuedness, consistency, strong consequentialism, and rationality, but it is not total.
- (ii) *Convex-valuedness:* NE satisfies totality, continuity, consistency, consequentialism, and rationality, but it violates convex-valuedness. Note that NE also violates strong consequentialism, and thus does not prove that convex-valuedness is required for Theorem 2.
- (iii) *Consistency:* Return all correlated strategies that randomize only over action profiles that survive iterated elimination of dominated strategies. This solution concept satisfies totality, continuity, convex-valuedness, strong consequentialism, and rationality, but it violates consistency.

- (iv) *Consequentialism*: Return all correlated strategies that randomize only over action profiles for which each player’s action is optimal against uniformly randomizing opponents. This solution concept satisfies totality, continuity, convex-valuedness, consistency, and rationality, but it violates consequentialism.
- (v) *Weak rationality*: Return all correlated strategies that maximize the sum of the players’ payoffs. This solution concept satisfies totality, continuity, convex-valuedness, consistency, and strong consequentialism, but it violates weak rationality.

It is open whether continuity is needed for either result, and whether convex-valuedness is needed for Theorem 2. An intriguing open question is whether *NE* and *CE* are the only solution concepts that satisfy all axioms in Theorem 1 except for convex-valuedness.

Remark 2 (Equilibrium refinements). In line with Selten’s trembling-hand perfection (Selten, 1975), various authors have proposed refinements of correlated equilibrium based on robustness to small trembles (Myerson, 1986; Dhillon and Mertens, 1996; Luo et al., 2022; Huang et al., 2026). A starting point for future research is to use the axiomatic method to characterize refinements of correlated equilibrium. In Theorem 1, convex-valuedness and the consistency axiom force many correlated strategies into the solution sets, ruling out refinements. Thus, weakening or replacing convex-valuedness and consistency, and jointly strengthening other axioms (e.g., rationality to admissibility) is a promising direction.

Remark 3 (Behavioral content of the axioms). When interpreted as behavioral conditions, the axioms take a stance on how players react to recommendations. View the correlated strategies returned by a solution concept as those which a mediator can implement—each player follows their recommendation. Under this interpretation, the results characterize the implementable correlated strategies implied by the axioms. Since correlated equilibria are exactly the implementable outcomes when players are Bayesian expected utility maximizers, the axioms entail Bayesian expected utility maximization.

- (i) *Consistency*: Whenever a given correlated strategy is implementable for any payoff realization of a random game, it is also implementable for the lottery over games. Since we treat lotteries as convex combinations of the realizations, consistency is closely tied to expected utility theory: it replaces payoff uncertainty by expected payoffs. The same logic underlies the sure-thing principle of Savage (1954): if a decision-maker would choose the same action in every state, then they choose it when the state is uncertain. Consistency imposes this sure-thing principle on players.
- (ii) *(Strong) consequentialism*: From a mediator’s perspective, the difference between both axioms is this. Suppose the mediator can implement a correlated strategy p' for a game G' , and G is obtained from G' by adding a clone of a'_i of a_i . To implement a correlated strategy for G , draw an action profile from p' ; if i ’s action in this profile is a_i , toss a coin to decide whether i ’s recommendation is a_i or a'_i ; do not change any other recommendations. Consequentialism requires that this correlated strategy is implementable for G if

the mediator uses the same coin irrespective of the action profile of i 's opponents. Strong consequentialism demands implementability even if coins with different biases are used for different action profiles.

Adding clones is akin to enlarging the mediator's message space. Consequentialism requires that if two recommendations are payoff-equivalent and informationally equivalent, then a player follows one if and only if they follow the other. In other words, a player's choice is independent of the label of an action. This is compatible with players who update their beliefs about others' recommendations conditional on their own recommendation. By contrast, strong consequentialism demands that players follow payoff-equivalent recommendations even if those induce different beliefs. This forces players to ignore the information contained in their recommendations, and is in line with committing to an action *ex ante*. The two axioms thus differ in what they assume about a player's arsenal of behavioral strategies.

- (iii) *(Weak) rationality*: Both axioms state that players prefer higher *sure* payoffs to lower ones, and thus avoid assumptions about risk attitudes. Weak rationality allows implementability even when the mediator sometimes recommends dominated actions. If players can only choose between following the mediator's recommendation and committing to a fixed action *ex ante*, following can be optimal despite occasional dominated recommendations. When a dominant action exists, however, the only implementable recommendation is to play it. Weak rationality is thus more natural than rationality when players have to commit *ex ante*.
- (iv) *Convex-valuedness*: It states that any convex combination over implementable correlated strategies is implementable. Such a convex combination arises if the mediator covertly tosses a coin to decide from which correlated strategy the recommendations are drawn. From the players' perspective, this introduces uncertainty about the distribution of the other agents' recommendations. Convex-valuedness is thus another instance of the sure-thing principle: if a player follows their recommendation for either of two beliefs about others' recommendations, they also do so for any mixture of the two beliefs. The difference from random games is that uncertainty enters through the information about other players rather than through payoffs.

Acknowledgments

The author thanks Felix Brandt for numerous discussions on the topic. The author also acknowledges support by the DFG under the Excellence Strategy EXC-2047.

APPENDIX

A. Preliminaries

For finite sets A, B with $B \subseteq A$, $\chi_B \in \{0, 1\}^A$ denotes the indicator of B ; the standard unit vector at $a \in A$ is χ_a . For $x, y \in \mathbb{R}^A$, $\text{supp}(x) = \{a \in A: x(a) \neq 0\}$ is the support of x , and $\langle x, y \rangle = \sum_{a \in A} x(a)y(a)$ is the inner product of x and y . We denote by $\Pi(A)$ the set of permutations of A .

For $S \subseteq N$, $A \in \mathcal{F}(U)^n$, and $a \in U^N$, we write $A_S = \prod_{i \in S} A_i$ and $a_S = (a_i)_{i \in S}$, as well as $A_{-S} = A_{N \setminus S}$ and $a_{-S} = a_{N \setminus S}$. For $x \in \mathbb{R}^A$, $x(\cdot, a_{-S}) \in \mathbb{R}^{A_S}$ is the restriction of x to the coordinates in S when fixing the remaining coordinates to a_{-S} . The marginal of $p \in \Delta(A)$ with respect to S is $p_S = \sum_{a_{-S} \in A_{-S}} p(\cdot, a_{-S}) \in \Delta(A_S)$. We write $\text{uni}(B) \in \Delta(A)$ for the uniform distribution on $B \subseteq A$.

If G is a game on $A \in \mathcal{F}(U)^n$ and $\pi = (\pi_1, \dots, \pi_n)$ with $\pi_i \in \Pi(A_i)$, then $G \circ \pi$ and $p \circ \pi$ denote the relabelings obtained by applying π coordinatewise. A solution concept is equivariant if relabeling the actions of a game results in the same relabeling among the returned correlated strategies.

Definition 4 (Equivariance). A solution concept satisfies equivariance if for each game G on A and each $\pi = (\pi_1, \dots, \pi_n)$ with $\pi_i \in \Pi(A_i)$, $f(G \circ \pi) = f(G) \circ \pi$.

Consequentialism implies equivariance: $G \circ \pi$ is a blow-up of G with surjection π . We frequently apply equivariance to correlated strategies where the probabilities of all supported action profiles are the same, and the permutation maps each supported action profile to another supported action profile. This gives a new game for which the same correlated strategy is returned.

A solution concept is positively homogeneous if it is invariant under scaling the payoffs of all players by the same positive constant.

Definition 5 (Positive homogeneity). A solution concept f is positively homogeneous if for each game G and each $\alpha > 0$, $f(G) = f(\alpha G)$.

Any total, continuous, and convex-valued solution concept that satisfies consistency and consequentialism contains a positively homogeneous solution concept that inherits any of our axioms from f . For any such solution concept f , define \tilde{f} for each game G by

$$\tilde{f}(G) = \bigcap_{\alpha \geq 1} f(\alpha G)$$

and note that \tilde{f} is a refinement of f .

Lemma 1 (Homogeneous refinement). *Let f be a total, continuous, and convex-valued solution concept that satisfies consistency and consequentialism. Then, \tilde{f} is total, continuous, convex-valued, positively homogeneous, and satisfies consistency and consequentialism. Moreover, if f satisfies strong consequentialism or (weak) rationality, then \tilde{f} also satisfies that axiom.*

Proof. First, we prove that \tilde{f} is total. Let G be a game on A , and let G^0 be the game on A such that $G_i \equiv 0$ for each $i \in N$. Note that all actions of all players are clones in G^0 . Hence, $\chi_a \in f(G^0)$ for each $a \in A$ by totality and consequentialism. Convex-valuedness then implies that $f(G^0) = \Delta(A)$. Then for each $\alpha \geq 1$,

$$f(\alpha G) = f(\alpha G) \cap f(G^0) \subseteq f\left(\frac{1}{\alpha}\alpha G + \left(1 - \frac{1}{\alpha}\right)G^0\right) = f(G)$$

by consistency, and so $f(\alpha G)$ is a non-increasing family of sets. Totality and continuity imply that $f(\alpha G)$ is nonempty and closed for each α . Hence, $\tilde{f}(G)$ is nonempty.

Second, we prove that \tilde{f} is continuous. Let $(G^k)_{k \in \mathbb{N}}$ be a sequence of games on A converging to G . For each $k \in \mathbb{N}$, let $p^k \in \tilde{f}(G^k)$, and assume that $(p^k)_{k \in \mathbb{N}}$ converges to $p \in \Delta(A)$. We need to show that $p \in \tilde{f}(G)$. To this end, it suffices to show that $p \in f(\alpha G)$ for each $\alpha \geq 1$. Now, for each $\alpha \geq 1$ and each k , $p^k \in f(\alpha G^k)$ by definition of \tilde{f} . Hence, since f is continuous and the fact that αG^k converges to αG for each $\alpha \geq 1$, $p \in f(\alpha G)$. Thus, $p \in \tilde{f}(G)$.

The proof of the remaining claims of the lemma is straightforward. \square

The second lemma states that a total, positively homogeneous, and convex-valued solution concept that satisfies consistency and consequentialism is invariant under adding a constant to the utility function of a player.

Lemma 2 (Adding constants). *Let f be a total, positively homogeneous, and convex-valued solution concept that satisfies consistency and consequentialism. Let \tilde{G} be a game on A such that \tilde{G}_i is constant for each $i \in N$. Then, for each game G on A , $f(G) = f(G + \tilde{G})$.*

Proof. By the argument in the proof of Lemma 1 (where $f(G^0) = \Delta(A)$), we have $f(\tilde{G}) = \Delta(A)$. Consistency and positive homogeneity then imply that $f(G) = f(G) \cap f(\tilde{G}) \subseteq f(\frac{1}{2}G + \frac{1}{2}\tilde{G}) = f(G + \tilde{G})$, and $f(G + \tilde{G}) = f(G + \tilde{G}) \cap f(-\tilde{G}) \subseteq f(\frac{1}{2}(G + \tilde{G}) + \frac{1}{2}(-\tilde{G})) = f(G)$. \square

B. Pure Nash equilibria

We show that every total and continuous solution concept satisfying consistency, consequentialism, and weak rationality returns all pure Nash equilibria. Let NE_{pure} denote the solution concept that returns all pure strategy Nash equilibria, let NE_{fs} return all full support Nash equilibria, and let $NE_{\mathbb{Q}}$ return all Nash equilibria that assign rational-valued probabilities to all action profiles. For each game G on A ,

$$\begin{aligned} NE_{\text{pure}}(G) &= NE(G) \cap \{\chi_a : a \in A\} \\ NE_{\text{fs}}(G) &= NE(G) \cap \{p \in \Delta(A) : \text{supp}(p) = A\} \\ NE_{\mathbb{Q}}(G) &= NE(G) \cap \mathbb{Q}^A \end{aligned}$$

Weak rationality requires that dominant actions are played with positive probability. The first lemma strengthens this: an action that is dominant in a subgame containing the support of a returned correlated strategy is played with probability 1 in that correlated strategy.

Lemma 3 (Restriction to dominant actions). *Let f be a solution concept satisfying consistency, consequentialism, and weak rationality. Let G be a game on A , $p \in f(G)$, and $\tilde{A} = \tilde{A}_1 \times \cdots \times \tilde{A}_n \subseteq A$ such that $\text{supp}(p) \subseteq \tilde{A}$. Assume there are $j \in N$ and $b_j \in A_j$ such that $G_j(b_j, a_{-j}) > G_j(a_j, a_{-j})$ for each $a_j \in A_j \setminus \{b_j\}$ and each $a_{-j} \in \tilde{A}_{-j}$. Then, $\text{supp}(p) \subseteq \{b_j\} \times \tilde{A}_{-j}$.*

Proof. Assume for contradiction that $\text{supp}(p) \not\subseteq \{b_j\} \times \tilde{A}_{-j}$. We will construct a game on A where b_j is dominant and f returns p , which contradicts weak rationality. Without loss of generality, $j = 1$ and $N \setminus \{j\} = \{2, \dots, n\}$. We successively construct games $G = G^1, G^2, \dots, G^n$ on A so that b_j is dominant when restricting G^i to $A_{[i]} \times \tilde{A}_{-[i]}$ and $p \in f(G^i)$ for each $i \in [n]$. Then, player 1 does not play the dominant action b_j with probability 1 in G^n , contradicting weak rationality.

For $i = 1$, the statement holds by assumption. Now let $i > 1$ and assume the statement holds for all smaller values of i . Hence, there is a game G^{i-1} on A and $\varepsilon > 0$ such that $G_1^{i-1}(b_1, a_{-1}) \geq G_1^{i-1}(a_1, a_{-1}) + \varepsilon$ for each $a_1 \in A_1 \setminus \{b_1\}$ and each $a_{-1} \in A_{-1}$ with $a_{-[i-1]} \in \tilde{A}_{-[i-1]}$, and $p \in f(G^{i-1})$. Let $\mu = \max_{a,b \in A} G_1^{i-1}(a) - G_1^{i-1}(b)$. Let $B_i \in \mathcal{F}(U)$ be disjoint from A_i with $|B_i| > \frac{\mu}{\varepsilon} |A_i \setminus \tilde{A}_i|$, and let $\hat{A}_i = A_i \cup B_i$. Fix some $\hat{a}_i \in \tilde{A}_i$ and let \hat{G} be a game on $\hat{A} = \hat{A}_i \times A_{-i}$ so that \hat{G} is a blow-up of G^{i-1} with surjection $\hat{\phi}$ so that $\hat{\phi}_k$ is the identity for $k \neq i$, and $\hat{\phi}_i$ is the identity on A_i and maps each action in B_i to \hat{a}_i . In other words, \hat{G} is obtained from G^{i-1} by introducing $|B_i|$ clones of \hat{a}_i . Consequentialism implies that $p \in f(\hat{G})$.

Let $\hat{\Sigma}_i \subseteq \Sigma_{\hat{A}_i}$ be the set of permutations of \hat{A}_i that fix \tilde{A}_i pointwise, and for each $k \neq i$, let $\hat{\Sigma}_k \subseteq \Sigma_{A_k}$ be the set of permutations of A_k containing only the identity on A_k , and let $\hat{\Sigma} = \hat{\Sigma}_1 \times \cdots \times \hat{\Sigma}_n$. Define

$$\bar{G} = \frac{1}{|\hat{\Sigma}|} \sum_{\pi \in \hat{\Sigma}} \hat{G} \circ \pi.$$

Note that all actions in $\hat{A}_i \setminus \tilde{A}_i$ are clones of each other in \bar{G} . Moreover, by the lower bound on the size of B_i , b_1 is dominant when \bar{G} is restricted to $A_{[i]} \times \tilde{A}_{-[i]}$. Equivariance implies that $p \in f(\hat{G} \circ \pi)$ for each $\pi \in \hat{\Sigma}$, since each such π fixes all action profiles in the support of p . Hence, it follows from consistency that $p \in f(\bar{G})$. Lastly, let G^i be such that \bar{G} is a blow-up of G^i with surjection $\bar{\phi}$ such that $\bar{\phi}_k$ is the identity for each $k \neq i$, and $\bar{\phi}_i$ is the identity on A_i and maps each action in B_i to some action in $A_i \setminus \tilde{A}_i$. That is, G^i is obtained from \bar{G} by removing the actions in B_i (all of which are clones of the actions in $A_i \setminus \tilde{A}_i$). Consequentialism implies that $p \in f(G^i)$. Moreover, b_1 is dominant when G^i is restricted to $A_{[i]} \times \tilde{A}_{-[i]}$. Hence, G^i has the required properties. This proves the induction step. \square

Given any game, consider the subgame obtained by the following process: if some player has a dominant action, eliminate all of their other actions; repeat this step with the reduced game until no more eliminations are possible. We say that an action profile survives iterated restriction to dominant actions if it is not eliminated during this process. Lemma 3 shows that a solution concept satisfying consistency, consequentialism, and weak rationality returns only correlated strategies that are supported on action profiles that survive iterated restriction to dominant actions.

Lemma 4 (Iterated restriction to dominant actions). *Let f be a solution concept satisfying consistency, consequentialism, and weak rationality. Let G be a game on A and let $\tilde{A} = \tilde{A}_1 \times \cdots \times \tilde{A}_n \subseteq A$ be the action profiles that survive iterated restriction to dominant actions. Then, $f(G) \subseteq \Delta(\tilde{A})$.*

Proof. This follows from repeated application of Lemma 3. □

If only a single action profile survives iterated restriction to dominant actions, this profile is the unique Nash equilibrium of the game. Lemma 4 implies that any total solution concept satisfying consistency, consequentialism, and weak rationality returns this profile and nothing else. We show that any game with a pure quasi-strict Nash equilibrium is a convex combination of games for which only this profile survives iterated restriction to dominant actions. Together, this proves that all pure quasi-strict Nash equilibria are returned.

Lemma 5 (Quasi-strict pure Nash equilibria I). *Let G be a game on A and $\tilde{a} \in A$ such that $\chi_{\tilde{a}}$ is a pure quasi-strict Nash equilibrium of G . Then, there are games G^1, \dots, G^n on A such that \tilde{a} is the only action profile that survives iterated restriction to dominant actions in G^i for each $i \in [n]$, and G is a convex combination of G^1, \dots, G^n .*

Proof. For each $j \in N$, let $\tilde{A}^j = A_j \times \{\tilde{a}_{-j}\}$, and let \hat{G}^j be a game on A so that for each $i \in N \setminus \{j\}$, \tilde{a}_i dominates all actions in $A_i \setminus \{\tilde{a}_i\}$, and for each $i \in N$ and $a \in \bigcup_{k \in N} \tilde{A}^k$,

$$\hat{G}_i^j(a) = G_i(a)$$

This is feasible since \tilde{a} is a pure quasi-strict Nash equilibrium of G . For the same reason, \tilde{a}_j dominates all actions in $A_j \setminus \{\tilde{a}_j\}$ if \hat{G}^j is restricted to \tilde{A}^j . Note that, by construction, for each $i \in N$ and $a \in \bigcup_{k \in N} \tilde{A}^k$,

$$\sum_{j \in N} \frac{1}{n} \cdot \hat{G}_i^j(a) = G_i(a). \tag{1}$$

Fix $j \in N$, and let G^j be the game on A so that for each $i \neq j$, $G_i^j = \hat{G}_i^j$, and

$$G_j^j(a) = \begin{cases} \hat{G}_j^j(a) & \text{for } a \in \bigcup_{k \in N} \tilde{A}^k, \text{ and} \\ n \cdot G_j(a) - \sum_{i \neq j} \hat{G}_j^i(a) & \text{for } a \in A \setminus \bigcup_{k \in N} \tilde{A}^k. \end{cases}$$

Observe that (1) also holds with G^j in place of \hat{G}^j since both games agree on $\bigcup_{k \in N} \tilde{A}^k$. Moreover, for each $i \neq j$, \tilde{a}_i dominates all actions in $A_i \setminus \{\tilde{a}_i\}$, and \tilde{a}_j dominates all actions in $A_j \setminus \{\tilde{a}_j\}$ if G^j is restricted to \tilde{A}^j . Hence, \tilde{a} is the only action profile in G^j that survives iterated restriction to dominant actions. Lastly, for each $j \in N$ and each $a \in A \setminus \bigcup_{k \in N} \tilde{A}^k$,

$$\begin{aligned} \sum_{i \in N} G_j^i(a) &= G_j^j(a) + \sum_{i \neq j} G_j^i(a) \\ &= n \cdot G_j(a) - \sum_{i \neq j} \hat{G}_j^i(a) + \sum_{i \neq j} G_j^i(a) \\ &= n \cdot G_j(a) - \sum_{i \neq j} G_j^i(a) + \sum_{i \neq j} G_j^i(a) \\ &= n \cdot G_j(a). \end{aligned}$$

For $a \in \bigcup_{k \in N} \tilde{A}^k$, the same equality follows from (1). Hence, $G = \sum_{i \in N} \frac{1}{n} G^i$ as required. \square

Lemma 6 (Quasi-strict pure Nash equilibria II). *Let f be a total solution concept that satisfies consistency, consequentialism, and weak rationality. Let G be a game on A and $\tilde{a} \in A$ such that $\chi_{\tilde{a}}$ is a pure quasi-strict Nash equilibrium of G . Then, $\chi_{\tilde{a}} \in f(G)$.*

Proof. Let G^1, \dots, G^n be the games promised by Lemma 5. For each $j \in N$, \tilde{a} is the only action profile in G^j that survives the iterated restriction to dominant actions. Hence, by Lemma 4 and the assumption that f is total, $\chi_{\tilde{a}} \in f(G^j)$ for each $j \in N$. Since $G = \sum_{j \in N} \frac{1}{n} \cdot G^j$, consistency implies that $\chi_{\tilde{a}} \in f(G)$. \square

For any pure Nash equilibrium in any game G , there is a sequence of games converging to G in which that equilibrium is pure and *quasi-strict*. Hence, any continuous solution concept that returns pure quasi-strict Nash equilibria returns all pure Nash equilibria. We thus have the following immediate consequence of Lemma 6.

Proposition 1 (Pure Nash equilibria). *Let f be a total and continuous solution concept that satisfies consistency, consequentialism, and weak rationality. Then, $NE_{\text{pure}} \subseteq f$.*

C. Two-player zero-sum games

We show that every total, continuous, and convex-valued solution concept that satisfies consistency, consequentialism, and weak rationality has to return all Nash equilibria for each zero-sum game in which all but two players are dummy players. Abusing terminology, we say that a game is zero-sum if there exist positive affine transformations of the players' utility functions such that for each action profile, the sum of all players' payoffs is 0.¹⁰

Definition 6 (Zero-sum games). A game G is zero-sum if there are $\alpha \in \mathbb{R}_{>0}^n$ and $\beta \in \mathbb{R}$ such that $\sum_{i \in N} \alpha_i G_i + \beta \equiv 0$.

We say that a player i is a dummy player if i 's payoff is constant and all of i 's actions are clones. Equivalently, i 's payoff is constant and no player's payoff depends on i 's action, irrespective of the other players' actions. An essentially two-player game is one in which all but two players are dummies.

Definition 7 (Essentially two-player games). Let G be a game on A . Player $i \in N$ is a dummy player if G_i is constant on A and all actions in A_i are clones. G is an essentially two-player game if there are $j, k \in N$ such that all players in $N \setminus \{j, k\}$ are dummy players. If G is an essentially two-player game and $\gamma G_j + G_k$ is constant for some $\gamma > 0$, we say that G is (j, k, γ) -zero-sum.

We prove that for any pair of players j, k , there exists some $\gamma > 0$ such that f returns all Nash equilibria in each (j, k, γ) -zero-sum game.

¹⁰Any game that is zero-sum in this sense is strategically zero-sum as defined by [Moulin and Vial \(1978\)](#). The converse is not true.

Proposition 2 (Nash equilibria of essentially two-player zero-sum games). *Let f be a total, continuous, and convex-valued solution concept that satisfies consistency, consequentialism, and weak rationality. Then, for all $j, k \in N$, there is $\gamma > 0$ such that for each (j, k, γ) -zero-sum game G , $NE(G) \subseteq f(G)$.*

The strategy is to construct more and more essentially two-player zero-sum games for which f returns some of the Nash equilibria. By consistency, we have succeeded if for each of the games G in Proposition 2 and each Nash equilibrium p of G , G is a convex combination of games for which p is a Nash equilibrium and we have established that f returns p .

C.1. Nash equilibria of bistochastic essentially two-player zero-sum games

The first step is to find some nontrivial game for which f returns a full-support Nash equilibrium. We use the game in which two players play a matching pennies game and all other players are dummies.

Definition 8 (Matching pennies games). For $j, k \in N$ and $z \in \mathbb{R}_{\geq 0}^{\{j, k\}}$, the matching pennies game between j and k with stakes z is the game $G^{j, k, z}$ on A with $A_j = A_k = \{1, 2\}$ and $A_i = \{0\}$ for each $i \neq j, k$, where for each $a \in A$,

$$G_j(a) = \begin{cases} z_j & \text{if } a_j + a_k \equiv 1 \pmod{2}, \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad G_k(a) = \begin{cases} z_k & \text{if } a_j + a_k \equiv 0 \pmod{2}, \\ 0 & \text{otherwise,} \end{cases}$$

and $G_i \equiv 0$ for each $i \neq j, k$. If $z_j = z_k = 1$, we write $G^{j, k}$ for short.

For example, for $n = 2$,

$$G^{1, 2} = \begin{pmatrix} 0, 1 & 1, 0 \\ 1, 0 & 0, 1 \end{pmatrix}$$

The unique coarse correlated equilibrium of any matching pennies game for $n = 2$ is uniform randomization over all action profiles. We prove that this uniform distribution is returned. The proof uses that f is a refinement of NE when restricting to strategies with independent randomization.

Lemma 7 (Brandl and Brandt, 2024). *Let f be a total solution concept that satisfies consistency, consequentialism, and weak rationality. Let G be a game on A , and let $p = p_1 \otimes \cdots \otimes p_n$ with $p_i \in \Delta(A_i)$. Then, $p \in f(G)$ implies $p \in NE(G)$.*

Proof. Theorem 1 of Brandl and Brandt (2024) shows that NE is the unique total solution concept returning strategies with independent randomization that satisfies consistency, consequentialism, and weak rationality. The proof of their Lemma 1 (used in the proof of their Theorem 1) remains true for solution concepts that return correlated strategies. The proof of the part $f \subseteq NE$ of their Theorem 1 assumes that there is $p = p_1 \otimes \cdots \otimes p_n \in f(G) \setminus NE(G)$ and derives a contradiction. Their arguments remain true for solution concepts that return correlated strategies. Hence, the claim follows. \square

Matching pennies games have a useful symmetry—a permutation of the action sets under which the game is invariant. Indeed, switching the labels of player 1’s actions and player 2’s actions leaves the game unchanged. Totality, equivariance, and convex-valuedness imply that f returns a correlated strategy that is invariant under this symmetry. Moreover, the matching pennies game with $z_j = 0$ or $z_k = 0$ has two pure Nash equilibria. Combined with Proposition 1 and Lemma 7, it follows that f returns the uniform distribution in a matching pennies game for some stakes.

Lemma 8 (Matching pennies games). *Let $j, k \in N$ and let f be a total, continuous, and convex-valued solution concept that satisfies consistency, consequentialism, and weak rationality. Then, there is $z \in \mathbb{R}_{>0}^{\{j,k\}}$ such that for each $\alpha > 0$, $\text{uni}(A) \in f(G^{j,k,\alpha z})$, where $A_j = A_k = \{1, 2\}$ and $A_i = \{0\}$ for each $i \neq j, k$.*

Proof. We prove the statement for $n = 2$ and $(j, k) = (1, 2)$ for convenience. The proof of the general case is the same up to notational changes. By Lemma 1, we may assume that f is positively homogeneous.

Let $G = G^{j,k}$, and let $p \in f(G)$, which exists by totality. Let $\pi_* \in \Pi(\{1, 2\})$ be the permutation that swaps 1 and 2, and let $\pi = (\pi_*, \pi_*)$. Note that $G = G \circ \pi$. Hence, by equivariance, $p \circ \pi \in f(G \circ \pi) = f(G)$. Thus, convex-valuedness implies that $\tilde{p} = \frac{1}{2}p + \frac{1}{2}(p \circ \pi) \in f(G)$. Observe that $\tilde{p}(1, 1) = \tilde{p}(2, 2)$ and $\tilde{p}(1, 2) = \tilde{p}(2, 1)$, and so there are $s, t \in [0, 1]$ such that $\tilde{p} = p^{s,t}$, where

$$p^{s,t} = \begin{pmatrix} s & t \\ t & s \end{pmatrix}$$

If $s = t$, the statement follows since f is positively homogeneous. Thus, assume without loss of generality that $s > t$.

For $\alpha, \beta \geq 0$, let $G^{(\alpha,\beta)} = G^{j,k,(\alpha,\beta)}$.

$$G^{(\alpha,\beta)} = \begin{pmatrix} 0, \beta & \alpha, 0 \\ \alpha, 0 & 0, \beta \end{pmatrix}$$

We prove that there is $\beta \in [0, 1]$ such that $\text{uni}(A) \in f(G^{(1,\beta)})$. For each $\beta \in [0, 1]$, there are $s_\beta, t_\beta \in [0, 1]$ such that $p^{s_\beta, t_\beta} \in f(G^{(1,\beta)})$. This follows from totality, equivariance, and convex-valuedness as above. Assume for contradiction that for each $\beta \in [0, 1]$, $\text{uni}(A) = p^{\frac{1}{4}, \frac{1}{4}} \notin f(G^{(1,\beta)})$. Then, by continuity and convex-valuedness, $s_0 > t_0$. By Proposition 1, $\chi_{(1,2)}, \chi_{(2,1)} \in f(G^{(1,0)})$ since these action profiles are pure Nash equilibria. Let $\lambda \in (0, 1)$ be the (unique) such number for which $q = \lambda p^{s_0, t_0} + (1 - \lambda)\chi_{(1,2)}$ is a product distribution. That is, $q = q_1 \otimes q_2$ is the product distribution of its marginals q_1 and q_2 , and since $s_0 > t_0$, $q_2(1) < q_2(2)$. In particular, q is not a Nash equilibrium since q_1 is not a best response to q_2 . Convex-valuedness implies that $q \in f(G^{(1,0)})$. But this contradicts Lemma 7. Hence, there is $\beta^* \in [0, 1]$ such that $\text{uni}(A) \in f(G^{(1,\beta^*)})$.

If $\beta^* > 0$, the statement follows from positive homogeneity. So assume that $\beta^* = 0$ and $\text{uni}(A) \notin f(G^{(1,\beta)})$ for each $\beta \in (0, 1]$. Let $\beta \in (0, 1]$ and recall that $s_\beta > t_\beta$ by continuity. By

Proposition 1 and convex-valuedness, $p^{t_\beta, s_\beta} \in f(G^{(1,0)})$ (note the transposition of s_β and t_β). Let $\pi_1, \pi_2 \in \Pi(\{1, 2\})$ such that π_1 swaps 1 and 2 and π_2 is the identity, and let $\pi = (\pi_1, \pi_2)$. Then, $p^{s_\beta, t_\beta} \in f(G^{(1,0)} \circ \pi)$. Observe that

$$\tilde{G} = \frac{1}{2}G^{(1,\beta)} + \frac{1}{2}\left(G^{(1,0)} \circ \pi\right) = \frac{1}{2}\begin{pmatrix} 0, \beta & 1, 0 \\ 1, 0 & 0, \beta \end{pmatrix} + \frac{1}{2}\begin{pmatrix} 1, 0 & 0, 0 \\ 0, 0 & 1, 0 \end{pmatrix} = \frac{1}{2}\begin{pmatrix} 1, \beta & 1, 0 \\ 1, 0 & 1, \beta \end{pmatrix}$$

Consistency implies that $p^{s_\beta, t_\beta} \in f(\tilde{G})$. Then, it follows from Lemma 2 and positive homogeneity that $p^{s_\beta, t_\beta} \in f(G^{(0,1)})$. Note that, if p^{s_β, t_β} has a convergent subsequence that does not converge to $\text{uni}(A)$ as β converges to 0, we obtain a contradiction to Lemma 7 as above. Thus, it follows from continuity that $\text{uni}(A) \in f(G^{(0,1)})$. Consistency and positive homogeneity then imply that $\text{uni}(A) \in f(G^{(1,1)})$, which suffices by positive homogeneity. This finishes the proof. \square

A matrix is deterministic if each entry is either 0 or 1, and it is bistochastic if all its entries are nonnegative and all row and column sums are 1.

Definition 9 (Deterministic and bistochastic matrices). Let $A = A_1 \times A_2$, and let $T: A \rightarrow \mathbb{R}_{\geq 0}$. Then, T is deterministic if $T(a) \in \{0, 1\}$ for each $a \in A$, and T is bistochastic if for each $i \in \{1, 2\}$ and each $a_{-i} \in A_{-i}$, $\sum_{a_i \in A_i} T(a_i, a_{-i}) = 1$.

Starting from a matching pennies game, we construct more games where the uniform distribution is a Nash equilibrium and f returns this correlated strategy, by blowing-up, permuting actions, and taking convex combinations. The next two lemmas show that one can use these operations to reach any bistochastic matrix as the utility function of the first player up to a positive affine transformation.

Lemma 9 (Building bistochastic matrices). Let $k, m \in \mathbb{N}$ such that $k \geq 2$ and m is a multiple of k , and let $T: [m]^2 \rightarrow \mathbb{R}_{\geq 0}$ be bistochastic. Let $\tilde{T}: [m]^2 \rightarrow \mathbb{R}_{\geq 0}$ be obtained from a deterministic bistochastic matrix on $[k]^2$ by replacing, for each $i \in \{1, 2\}$, each $a_i \in [k]$ by m/k clones. Then, there are $\lambda \in \Delta(\Pi([m])^2)$, $\alpha > 0$, and $\beta \in \mathbb{R}$ such that $\alpha T + \beta = \sum_{\pi \in \Pi([m])^2} \lambda_\pi \tilde{T} \circ \pi$.

Proof. First, consider the case that T is deterministic (i.e., a permutation matrix). Without loss of generality, T is the identity matrix, and \tilde{T} is obtained by replacing each 1 in the $k \times k$ identity matrix by a block of 1's of size $m/k \times m/k$. Let $\tilde{\lambda} \in \mathbb{R}_{\geq 0}^{\Pi([m])^2}$ be a vector that assigns a weight to each pair of permutations of $[m]$ such that $\tilde{\lambda}_\pi = 1$ if $\text{supp}(T) \subseteq \text{supp}(\tilde{T} \circ \pi)$ and $\tilde{\lambda}_\pi = 0$ otherwise. That is, $\tilde{\lambda}_\pi = 1$ if and only if $\pi = (\pi_1, \pi_1 \circ \pi_2)$ for some $\pi_1, \pi_2 \in \Pi([m])$ such that π_2 fixes each of the sets $\{1, \dots, m/k\}, \dots, \{m - m/k + 1, \dots, m\}$. Note that $|\tilde{\lambda}| = m!((m/k)!)^k$, and let $\lambda = \tilde{\lambda}/|\tilde{\lambda}|$. Let $\hat{T} = \sum_{\pi} \lambda_\pi \tilde{T} \circ \pi$. For each $a \in [m]^2$, $\hat{T}(a) = 1$ if $T(a) = 1$, and $\hat{T}(a) = \frac{m/k-1}{m-1}$ if $T(a) = 0$. Thus, letting $\alpha = 1 - \frac{m/k-1}{m-1}$ and $\beta = \frac{m/k-1}{m-1}$ yields the required expression. The case when T is not necessarily deterministic follows from the fact that every bistochastic matrix is a convex combination of permutation matrices, and that convex combinations of affine transformations of a matrix are again affine transformations of that matrix. \square

The case of Lemma 9 where $k = m$ and β is required to be 0 (and thus $\alpha = 1$) is the Birkhoff-von Neumann theorem. Thus, Lemma 9 can be viewed as a variant of that theorem, where the basic matrices are blow-ups of smaller permutation matrices rather than honest permutation matrices.

Using Lemma 9, we show that for any bistochastic matrix T , there exists an essentially two-player zero-sum game in which the utility function of one non-dummy player is T and the uniform distribution over all action profiles is a Nash equilibrium and is returned by f .

Lemma 10 (Uniform equilibria of essentially two-player zero-sum games I). *Let f be a total, continuous, and convex-valued solution concept that satisfies consistency, consequentialism, and weak rationality. Let $m, n \in \mathbb{N}$ and $j, k \in N$. Then, there is $\gamma > 0$ such that for each bistochastic $T: [m]^{\{j,k\}} \rightarrow \mathbb{R}_{\geq 0}$, there exists a game G on $A = [m]^{\{j,k\}} \times \{0\}^{N \setminus \{j,k\}}$ such that for each $a_{-j,k} \in A_{-j,k}$, $G_j(\cdot, a_{-j,k}) = T$, $G_k(\cdot, a_{-j,k}) = -\gamma T$, for each $i \neq j, k$, $G_i \equiv 0$, and $\text{uni}(A) \in f(G) \cap NE(G)$.*

Proof. We prove the statement assuming $N = \{j, k\}$. The proof of the general case, i.e., with $n - 2$ dummy players, is the same, albeit more notationally heavy. Replacing m by $2m$ and using consequentialism, we may assume that 2 divides m . Moreover, by Lemma 1, we may assume that f is positively homogeneous.

By Lemma 8, there is $z \in \mathbb{R}_{>0}^{\{j,k\}}$ such that $\text{uni}(A) \in f(G^{j,k,\alpha z})$ for each $\alpha > 0$. It follows that there is $\gamma > 0$ such that for $z = (1, \gamma)$, $\text{uni}(A) \in f(G^{j,k,z})$. We write G^z instead of $G^{j,k,z}$ for short. Let \tilde{G}^z be the game on A obtained from G^z by replacing each action in A_j and A_k by $m/2$ clones. Consequentialism implies that $\text{uni}(A) \in f(\tilde{G}^z)$. Since $G_j^z: A \rightarrow \mathbb{R}$ is deterministic and bistochastic, it follows from Lemma 9 that there are $\alpha > 0$, $\beta \in \mathbb{R}$, and $\lambda \in \Delta(\Pi([m])^{\{j,k\}})$ such that

$$\alpha T + \beta = \sum_{\pi \in \Pi([m])^{\{j,k\}}} \lambda_\pi (\tilde{G}_j^z \circ \pi).$$

Let $\bar{G} = \sum_{\pi \in \Pi([m])^{\{j,k\}}} \lambda_\pi (\tilde{G}^z \circ \pi)$. Equivariance and consistency imply that $\text{uni}(A) \in f(\bar{G})$. Letting G^β be the game with $G_j^\beta \equiv \beta$ and $G_k^\beta \equiv -\gamma\beta$, and letting $G = \frac{1}{\alpha}(\bar{G} - G^\beta)$, we have that $G_j = T$ and $G_k = -\gamma T$. Moreover, it follows from Lemma 2 and positive homogeneity that $\text{uni}(A) \in f(G)$, and since T is bistochastic, $\text{uni}(A) \in NE(G)$. \square

We prove a weaker version of Proposition 2: any total, continuous, and convex-valued solution concept that satisfies consistency, consequentialism, and weak rationality returns the uniform distribution over all action profiles in any essentially two-player zero-sum game in which this correlated strategy is a Nash equilibrium, assuming the payoffs of the two active players have a fixed ratio.

Lemma 11 (Uniform equilibria of essentially two-player zero-sum games II). *Let f be a total, continuous, and convex-valued solution concept that satisfies consistency, consequentialism, and weak rationality. Let $j, k \in N$, let $\gamma > 0$ be as promised by Lemma 10, and let G be a (j, k, γ) -zero-sum game on A with $\text{uni}(A) \in NE(G)$. Then, $\text{uni}(A) \in f(G)$.*

Proof. To simplify notation, we assume that $N = \{j, k\}$; the proof extends straightforwardly to the general situation. We may assume that f is positively homogeneous by Lemma 1 and that $|A_j| = |A_k|$ by consequentialism. Moreover, we may assume that $G_j(a) \geq 0$ for each $a \in A$ and that $\gamma G_j + G_k \equiv 0$ by Lemma 2.

Observe that all row sums of G_j are equal and all its column sums are equal since $\text{uni}(A) \in NE(G)$. That is, $\sum_{a_k \in A_k} G_j(a_j, a_k)$ is independent of $a_j \in A_j$, and $\sum_{a_j \in A_j} G_k(a_j, a_k) = -\gamma \sum_{a_j \in A_j} G_j(a_j, a_k)$ is constant on A_k . Hence, G_j is a nonnegative multiple of a bistochastic matrix. If $G_j \equiv 0$, then also $G_k \equiv 0$, so that all actions in A_j are clones and all actions in A_k are clones. Thus, $\text{uni}(A) \in f(G)$ follows from consequentialism. If $G_j \not\equiv 0$, by positive homogeneity, we may assume that G_j is bistochastic. It follows from Lemma 10 that there is a G' on A with $G'_j = G_j$, $G'_k = -\gamma G'_j$, and $\text{uni}(A) \in f(G') \cap NE(G')$. But then $G = G'$, and $\text{uni}(A) \in f(G)$ follows. \square

C.2. Reduction to uniform full support Nash equilibria

Lemma 11 shows that the statement Proposition 2 holds for the uniform distribution over all action profiles. We extend it to arbitrary Nash equilibria in two steps: first, to rational-valued Nash equilibria that do not necessarily have full support; second, to all Nash equilibria.

Lemma 12 (Reduction to uniform Nash equilibria). *Let f be a solution concept that satisfies consistency, consequentialism, and weak rationality, let $j, k \in N$, and let $\gamma > 0$. Assume that for each (j, k, γ) -zero-sum game G on A with $\text{uni}(A) \in NE(G)$, $\text{uni}(A) \in f(G)$. Then, for each (j, k, γ) -zero-sum game G on A , $NE_{\mathbb{Q}}(G) \subseteq f(G)$.*

Proof. Assume that $N = \{j, k\}$; the general case with $n - 2$ dummy players follows similarly. By Lemma 1, we may assume that f is positively homogeneous.

First, observe that for any (j, k, γ) -zero-sum game G and any $p \in NE_{\mathbb{Q}}(G) \cap NE_{\text{fs}}(G)$, it follows directly from consequentialism that $p \in f(G)$. Indeed, by replacing each action a_i of each player i by a number of clones proportional to $p_i(a_i)$ and using consequentialism, one can reduce to the case where p is the uniform distribution, which is covered by assumption. The remainder of the proof removes the assumption that p has full support.

Let G be a (j, k, γ) -zero-sum game on A and $p \in NE_{\mathbb{Q}}(G)$. Recall that p_j and p_k are the players' marginals, so that $p = p_j \otimes p_k$. Let $\tilde{N} = \{i \in N : \text{supp}(p_i) \subsetneq A_i\}$ be the set of players whose strategy does not have full support, and let $\tilde{n} = |\tilde{N}|$. We prove that $p \in f(G)$ by induction (with at most two steps). The base case $\tilde{n} = 0$ holds by assumption. Now assume that $\tilde{n} > 0$ and the statement holds for all smaller values of \tilde{n} . Assume without loss of generality that $j \in \tilde{N}$.

Step 1 ($G_j(a_j, p_k) = 0$ for each $a_j \in A_j$). We first assume $G_j(\cdot, p_k) \equiv 0$, i.e., player j 's expected payoff given p is 0 for each action. Let $b_j \in \text{supp}(p_j)$, and let \tilde{G} be the game on \tilde{A} , where $\tilde{A}_j = A_j \cup \bigcup_{a_j \in A_j \setminus \text{supp}(p_j)} \{\tilde{a}_j, b_j^{a_j}\}$ for distinct $\tilde{a}_j, b_j^{a_j} \in U \setminus A_j$, $\tilde{A}_k = A_k$, and for each $i \in \{j, k\}$

and $c \in A$,

$$\tilde{G}_i(c) = \begin{cases} G_i(c) & \text{if } c_j \in A_j \\ G_i(a_j, c_{-j}) & \text{if } c_j = \tilde{a}_j \text{ for some } a_j \in A_j \setminus \text{supp}(p_j) \\ 2G_i(b_j, c_{-j}) - G_i(a_j, c_{-j}) & \text{if } c_j = b_j^{a_j} \text{ for some } a_j \in A_j \setminus \text{supp}(p_j) \end{cases}$$

That is, \tilde{G} is obtained from G by adding, for each action $a_j \in A_j \setminus \text{supp}(p_j)$, a clone \tilde{a}_j of a_j and a new action $b_j^{a_j}$ that corresponds to a combination of b_j and a_j . Note that \tilde{G} is a (j, k, γ) -zero-sum game since G is.

Let $\tilde{p}_j \in \Delta(\tilde{A}_j) \cap \mathbb{Q}^{\tilde{A}_j}$ such that $\text{supp}(\tilde{p}_j) = \tilde{A}_j$, for each $a_j \in \text{supp}(p_j) \setminus \{b_j\}$, $\tilde{p}_j(a_j) = p_j(a_j)$, and for each $a_j \in A_j \setminus \text{supp}(p_j)$, $\tilde{p}_j(a_j) + \tilde{p}_j(\tilde{a}_j) = \tilde{p}_j(b_j^{a_j})$. Let $\tilde{p} = \tilde{p}_j \otimes p_k$ be the product distribution of \tilde{p}_j and p_k , and observe that $\tilde{p} \in NE_{\mathbb{Q}}(\tilde{G})$. Thus, by the induction hypothesis $\tilde{p} \in f(\tilde{G})$. Let $\hat{p}_j \in \Delta(\tilde{A}_j)$ such that $\hat{p}_j(c_j) = \tilde{p}_j(c_j)$ for each $c_j \in \text{supp}(p_j) \cup \{b_j^{a_j} : a_j \in A_j \setminus \text{supp}(p_j)\}$, and $\hat{p}_j(\tilde{a}_j) = \tilde{p}_j(a_j) + \tilde{p}_j(\tilde{a}_j)$ for each $a_j \in A_j \setminus \text{supp}(p_j)$. In other words, \hat{p}_j is obtained from \tilde{p}_j by shifting the probability on a_j to \tilde{a}_j . Let $\hat{p} = \hat{p}_j \otimes p_k \in \Delta(\tilde{A})$. Consequentialism implies that $\hat{p} \in f(\tilde{G})$.

Let $\pi_j \in \Pi(\tilde{A}_j)$ be the permutation that swaps \tilde{a}_j and $b_j^{a_j}$ for each $a_j \in A_j \setminus \text{supp}(p_j)$, let π_k be the identity on \tilde{A}_k , and let $\pi = (\pi_j, \pi_k)$. Since $\hat{p}_j(b_j^{a_j}) = \hat{p}_j(\tilde{a}_j)$ for each $a_j \in A_j \setminus \text{supp}(p_j)$, $\hat{p} = \hat{p} \circ \pi$, and so by equivariance, $\hat{p} \in f(\tilde{G} \circ \pi)$. Let $\hat{G} = \frac{1}{2}\tilde{G} + \frac{1}{2}\tilde{G} \circ \pi$. Consistency implies that $\hat{p} \in f(\hat{G})$. Observe that, by construction of \tilde{G}_j , \hat{G} is obtained from G by adding two clones of b_j for each $a_j \in A_j \setminus \text{supp}(p_j)$. Since $\hat{p}_j(a_j) = p_j(a_j)$ for each $a_j \in A_j \setminus \{b_j\}$, it follows from consequentialism that $p \in f(G)$.

Step 2 ($G_j(a_j, p_{-j}) \neq 0$ for some $a_j \in A_j$). For each $a_j \in A_j$, let $\alpha_{a_j} = G_j(a_j, p_k)$ be player j 's expected payoff for a_j against p_k . Note that $\alpha_{a_j} \geq \alpha_{b_j}$ for each $a_j \in \text{supp}(p_j)$ and $b_j \in A_j$. Let \tilde{G} be the game on A such that for each $a_j \in A_j$, $\tilde{G}_j(a_j, \cdot) = -\frac{1}{\gamma}\tilde{G}_k(a_j, \cdot) \equiv \alpha_{a_j}$. Note that \tilde{G} is a (j, k, γ) -zero-sum game. Moreover, all actions in $B_j = \{a_j \in A_j : \alpha_{a_j} \geq \alpha_{b_j} \text{ for each } b_j \in A_j\}$ are clones, and all actions in A_k are clones. Consequentialism and Proposition 1 thus imply that $p \in f(\tilde{G})$.

Let $\bar{G} = G - \tilde{G}$. Then, \bar{G} is (j, k, γ) -zero-sum, $p \in NE_{\mathbb{Q}}(\bar{G})$, and for each $a_j \in A_j$, $\bar{G}_j(a_j, p_k) = \alpha_{a_j} - \alpha_{a_j} = 0$. Hence, by Step 1, $p \in f(\bar{G})$. Consistency and positive homogeneity imply that $p \in f(G)$. This completes the proof. \square

Lemma 13 (Reduction to rational-valued Nash equilibria). *Let f be a continuous and convex-valued solution concept, let $j, k \in N$, and let $\gamma > 0$. Assume that for each (j, k, γ) -zero-sum game G on A , $NE_{\mathbb{Q}}(G) \subseteq f(G)$. Then, for each (j, k, γ) -zero-sum game G on A , $NE(G) \subseteq f(G)$.*

Proof. Assume that $N = \{j, k\}$; the proof straightforwardly extends to the general case with $n - 2$ dummy players. For each essentially two-player zero-sum game G on A , $NE_{\{j, k\}}(G) = \{(p_j, p_k) : p_j \otimes p_k \in NE(G)\}$ is a convex polytope in $\Delta(A_j) \times \Delta(A_k)$. Now fix a (j, k, γ) -zero-sum

game G on A , and let $p \in NE(G)$. Since $f(G)$ is convex, it suffices to consider the case that $p = p_j \otimes p_k$ is the product distribution of p_j and p_k , where (p_j, p_k) is a vertex of $NE_{\{j,k\}}(G)$.

Let $(G^\ell)_{\ell \in \mathbb{N}}$ and $(p^\ell)_{\ell \in \mathbb{N}}$ be sequences of (j, k, γ) -zero-sum games and correlated strategies on A such that the following holds: for each $\ell \in \mathbb{N}$, $p^\ell \in NE(G^\ell)$, $G_i^\ell \in \mathbb{Q}^A$ for each $i \in N$, $(G_i^\ell)_{\ell \in \mathbb{N}}$ converges to G_i for each $i \in N$, and $(p^\ell)_{\ell \in \mathbb{N}}$ converges to p .¹¹ For each $\ell \in \mathbb{N}$, the convex polytope $NE_{\{j,k\}}(G^\ell)$ has rational-valued vertices. Thus, by assumption, $p^\ell \in NE_{\mathbb{Q}}(G^\ell) \subseteq f(G^\ell)$, and so $p \in f(G)$ by continuity. \square

Proof of Proposition 2. By Lemma 11, for any $j, k \in N$, there is $\gamma > 0$ such that for each (j, k, γ) -zero-sum game G on A with $\text{uni}(A) \in NE(G)$, $\text{uni}(A) \in f(G)$. Then, by Lemma 12 and Lemma 13, for any (j, k, γ) -zero-sum game G on A , $NE(G) \subseteq f(G)$. This completes the proof. \square

D. Reduction to zero-sum games

Let f be a total, continuous, and convex-valued solution concept that satisfies consistency, consequentialism, and rationality. We prove that if f returns all Nash equilibria of essentially two-player zero-sum games, then $f = CE$.

Let $A = A_1 \times \dots \times A_n$ and $B \subset A$. A path of length k in B is a sequence $a^0, \dots, a^k \in B$ such that for each $\ell \in [k]$, $a^{\ell-1}$ and a^ℓ differ in exactly one coordinate. For $a, b \in B$, the distance between a and b in B is the length of a shortest path from a to b in B , or infinity if there is no such path; a and b are adjacent if they are at distance 1, i.e., if they differ in exactly one coordinate. We say that B is connected if there is a path between any two elements of B .

Lemma 14 (Basic decompositions). *Let $T: A \rightarrow \mathbb{R}$ and $p \in \Delta(A)$ such that $\langle p, T \rangle = 0$. If $\text{supp}(p)$ is connected, then there exist $T^1, \dots, T^k: A \rightarrow \mathbb{R}$ such that $\sum_{\ell \in [k]} T^\ell = T$, and for each $\ell \in [k]$, $\langle p, T^\ell \rangle = 0$ and the support of T^ℓ consists of one element of A or two adjacent elements of $\text{supp}(p)$.*

Proof. It suffices to prove the statement for the case $\text{supp}(T) \subseteq \text{supp}(p)$ since $\langle p, T' \rangle = 0$ for any $T': A \rightarrow \mathbb{R}$ that is supported on $A \setminus \text{supp}(p)$. Let $B = \text{supp}(p)$, let $B^+ = \{a \in B: T(a) > 0\}$, and let $B^- = \{a \in B: T(a) < 0\}$.

¹¹Theorem 2 of Shapley and Snow (1950) states that if G is a two-player zero-sum game with $N = \{j, k\}$ and $p \in NE(G)$, then (p_j, p_k) is a vertex of $NE_{\{j,k\}}(G)$ if and only if there exists a nonsingular square submatrix \hat{G}_j of G_j defined by $\hat{A} = \hat{A}_j \times \hat{A}_k \subseteq A_j \times A_k$ such that $p_j^\top = \frac{\mathbf{1}^\top \hat{G}_j^{-1}}{(\mathbf{1}, \hat{G}_j^{-1} \mathbf{1})}$, $p_k = \frac{\hat{G}_j^{-1} \mathbf{1}}{(\mathbf{1}, \hat{G}_j^{-1} \mathbf{1})}$, and $v = \frac{1}{(\mathbf{1}, \hat{G}_j^{-1} \mathbf{1})}$, where v is the value of G . Moreover, \hat{A} contains the support of p , and $\hat{A}_j \subseteq \{a_j \in A_j: G_j(a_j, p_k) = v\}$, and similarly for \hat{A}_k . Hence, (p_j, p_k) is also a vertex of $NE_{\{j,k\}}(G')$, where G' is obtained from G by decreasing j 's payoff for actions in $A_j \setminus \hat{A}_j$ and decreasing k 's payoff for actions in $A_k \setminus \hat{A}_k$, while keeping the game zero sum. Then, for any $p' \in NE(G')$, $\text{supp}(p') \subseteq \hat{A}$, and so nonsingularity of \hat{G}_j implies that p is the unique Nash equilibrium of G' . Perturbing the payoffs of G' so that all payoffs are rational-valued and using continuity of NE , we obtain a game close to G for which a correlated strategy close to p is the unique Nash equilibrium.

Claim 1. Let $a \in B^+$ and $b \in B^-$ minimize the distance in B between elements of B^+ and B^- , and let m be that distance. If $m > 1$, there are $v, w: A \rightarrow \mathbb{R}$ such that $T = v + w$, $\langle p, v \rangle = 0$ and $\langle p, w \rangle = 0$, the distance between $\text{supp}(v_+)$ and $\text{supp}(v_-)$ is smaller than m , the distance between $\text{supp}(w_+)$ and $\text{supp}(w_-)$ is smaller than m , $|\text{supp}(v)|, |\text{supp}(w)| \leq |\text{supp}(T)|$, and $\text{supp}(v), \text{supp}(w) \subseteq B$.

Proof of Claim 1. Let a^0, \dots, a^m be a path from a to b in B . Note that $T(a^\ell) = 0$ for each $\ell \in [m-1]$ since there is no path from B^+ to B^- of length at most $m-1$ in B . Recall that χ_a is the standard unit vector at $a \in A$, and define $v = T(a^0)\chi_{a^0} - T(a^0)\frac{p(a^0)}{p(a^1)}\chi_{a^1}$ and $w = T - v$. Then, $\langle p, v \rangle = 0$, and so $\langle p, w \rangle = 0$. Moreover, the shortest path from $\text{supp}(v_+)$ to $\text{supp}(v_-)$ in B has length 1 since a^0 and a^1 are adjacent in B , and a^1, \dots, a^m is a path from $\text{supp}(w_+)$ to $\text{supp}(w_-)$ in B of length $m-1$. Lastly, $T = v + w$, $|\text{supp}(v)|, |\text{supp}(w)| \leq |\text{supp}(T)|$, and $\text{supp}(v), \text{supp}(w) \subseteq B$ since a^0, \dots, a^m is a path in B and T is supported in B . \square

Now we prove the statement of the lemma by induction on $|\text{supp}(T)|$. The base case $|\text{supp}(T)| = 0$ is trivial. Observe that $|\text{supp}(T)| = 1$ is not possible since $\text{supp}(T_+), \text{supp}(T_-) \neq \emptyset$ whenever $T \neq 0$, $\langle p, T \rangle = 0$, and $\text{supp}(T) \subseteq B$. Now consider the case that $|\text{supp}(T)| \geq 2$, and assume the statement holds for all smaller support sizes. We run a second induction on the distance between $\text{supp}(T_+)$ and $\text{supp}(T_-)$. Denote this distance by m .

If $m = 1$, there are $a \in \text{supp}(T_+)$ and $b \in \text{supp}(T_-)$ that are adjacent in B . Consider first the case that $p(a)T(a) \leq p(b)T(b)$. Let $v = T(a)\chi_a - T(a)\frac{p(a)}{p(b)}\chi_b$ and $w = T - v$. Then, $\langle p, v \rangle = 0$, $\langle p, w \rangle = \langle p, T \rangle - \langle p, v \rangle = 0$, $|\text{supp}(v)| = 2$, and $|\text{supp}(w)| < |\text{supp}(T)|$. The statement holds trivially for v , and it holds for w by the hypothesis of the induction on the support size. Hence, it also holds for $T = v + w$. The case $p(a)T(a) \geq p(b)T(b)$ is analogous.

If $m > 1$, let v, w be as obtained from Claim 1. Then, the statement holds for v and w by the hypothesis of the induction on m . Hence, it also holds for $T = v + w$. This completes both inductions and thus the proof. \square

Lemma 15 (From Nash equilibria to correlated equilibria). *Let f be a total, continuous, and convex-valued solution concept that satisfies consistency, consequentialism, and weak rationality. Assume that for all $j, k \in N$, there is $\gamma_{j,k} > 0$ such that $NE(G) \subseteq f(G)$ for each $(j, k, \gamma_{j,k})$ -zero-sum game G . Let G be a game on A and $j \in N$ such that for each $i \neq j$, $G_i \equiv 0$. Then, $CE(G) \subseteq f(G)$.*

Proof. By Lemma 1, we may assume that f is positively homogeneous. We establish a special case of the lemma first.

Claim 1. Let $j \in N$, $G'_j: A \rightarrow \mathbb{R}$, $b_j \in A_j$, and $p \in \Delta(A)$ such that for each $a_j \neq b_j$, $p(a_j, \cdot) \equiv 0$, $\langle p(b_j, \cdot), G'_j(a_j, \cdot) \rangle \leq 0$, and $G'_j(b_j, \cdot) \equiv 0$. Then, there exists a game G on A such that $G_j = G'_j$, for each $i \in N$, $G_i(b_j, \cdot) \equiv 0$, and $p \in f(G)$.

Note that $p \in CE(G)$.

Proof of Claim 1. We proceed in multiple steps.

Step 1. First, assume that p is connected and $\langle p(b_j, \cdot), G'_j(a_j, \cdot) \rangle = 0$ for each $a_j \in A_j$. For each $a_j \neq b_j$, let $\hat{T}^{a_j, 1}, \dots, \hat{T}^{a_j, m_{a_j}} : A_{-j} \rightarrow \mathbb{R}$ be as obtained from Lemma 14 applied to $G'_j(a_j, \cdot)$ and $p(b_j, \cdot)$. For $\ell \in [m_{a_j}]$, let $G^{a_j, \ell}$ be the game on A with $G_j^{a_j, \ell}(a_j, \cdot) = \hat{T}^{a_j, \ell}$, for each $c_j \neq a_j$, $G_j^{a_j, \ell}(c_j, \cdot) \equiv 0$, and for $i \neq j$, define $G_i^{a_j, \ell}$ as follows.

- If $\text{supp}(\hat{T}^{a_j, \ell}) = \{a_{-j}\}$ for some $a_{-j} \in \text{supp}(G'_j(a_j, \cdot))$, let $k \neq j$ be arbitrary, let $G_k^{a_j, \ell} = -\gamma_{j,k} G_j^{a_j, \ell}$, and for $i \neq j, k$, let $G_i^{a_j, \ell} \equiv 0$. Note that $a_{-j} \notin \text{supp}(p(b_j, \cdot))$ since $\langle p(b_j, \cdot), \hat{T}^{a_j, \ell} \rangle = 0$. For each $c_{-j} \in A_{-j} \setminus \{a_{-j}\}$ and each $c_j \neq b_j$, $G_j^{a_j, \ell}(b_j, c_{-j}) = G_j^{a_j, \ell}(c_j, c_{-j}) = 0$, and so $\chi_{(b_j, c_{-j})}$ is a pure Nash equilibrium of $G^{a_j, \ell}$ for each such c_{-j} . Thus, p is in the convex hull of the pure Nash equilibria of $G^{a_j, \ell}$, and so by Proposition 1 and convex-valuedness, $p \in f(G^{a_j, \ell})$.
- If $\text{supp}(\hat{T}^{a_j, \ell}) = \{a_{-j}, b_{-j}\}$ for adjacent $a_{-j}, b_{-j} \in \text{supp}(p(b_j, \cdot))$, let $k \neq j$ such that $a_{-j,k} = b_{-j,k}$, let $G_k^{a_j, \ell} = -\gamma_{j,k} G_j^{a_j, \ell}$, and for $i \neq j, k$, let $G_i^{a_j, \ell} \equiv 0$. Note that $G^{a_j, \ell}$ is a $(j, k, \gamma_{j,k})$ -zero-sum game. Let $\hat{p} = \frac{1}{p(b_j, a_{-j}) + p(b_j, b_{-j})} (p(b_j, a_{-j}) \chi_{(b_j, a_{-j})} + p(b_j, b_{-j}) \chi_{(b_j, b_{-j})}) \in \Delta(A)$. Since $\langle p(b_j, \cdot), G_j^{a_j, \ell}(a_j, \cdot) \rangle = 0$ and $G_k^{a_j, \ell}(b_j, \cdot) \equiv 0$, it follows that $\hat{p} \in NE(G^{a_j, \ell})$. Moreover, for each $c_{-j} \in A_{-j} \setminus \{a_{-j}, b_{-j}\}$, $\chi_{(b_j, c_{-j})}$ is a pure Nash equilibrium of $G^{a_j, \ell}$ since $G_j^{a_j, \ell}(b_j, c_{-j}) = G_j^{a_j, \ell}(c_j, c_{-j}) = 0$ for each $c_j \in A_j$, and so $\chi_{(b_j, c_{-j})} \in f(G^{a_j, \ell})$ by Proposition 1. Hence, $p \in f(G^{a_j, \ell})$ by convex-valuedness.

Let $G = \sum_{a_j \neq b_j} \sum_{\ell \in [m_{a_j}]} G^{a_j, \ell}$. By consistency and positive homogeneity, $p \in f(G)$, and G has the properties stated in Claim 1 by construction.

Step 2. Second, we generalize to $\langle p(b_j, \cdot), G'_j(a_j, \cdot) \rangle \leq 0$ for each $a_j \neq b_j$. Let $\alpha \in \mathbb{R}^{A_j}$ such that for each $a_j \in A_j$, $\alpha_{a_j} = \langle p(b_j, \cdot), G'_j(a_j, \cdot) \rangle \leq 0$. Let \hat{G} be the game on A such that for each $a_j \in A_j$, $\hat{G}_j(a_j, \cdot) \equiv \alpha_{a_j}$, and for each $i \neq j$, let $\hat{G}_i \equiv 0$. Note that for each $a_j \in A_j$, $\langle p(b_j, \cdot), G'_j(a_j, \cdot) - \hat{G}_j(a_j, \cdot) \rangle = 0$. By Step 1 applied to $G'_j - \hat{G}_j$, there is \tilde{G} with the properties stated in Claim 1. Moreover, in \hat{G} , for each $i \neq j$, all actions in A_i are clones, and for each $a_j \neq b_j$, $\hat{G}_j(a_j, \cdot) \leq \hat{G}_j(b_j, \cdot)$. Hence, for each $a_{-j} \in A_{-j}$, (b_j, a_{-j}) is a pure Nash equilibrium of \hat{G} , and so $p \in f(\hat{G})$ by Proposition 1 and convex-valuedness. Consistency and positive homogeneity then imply that $p \in f(\tilde{G} + \hat{G})$, and so $\tilde{G} + \hat{G}$ is as required by Claim 1.

Step 3. Third, we remove the assumption that p is connected. Observe that, by consistency and positive homogeneity, it suffices to prove Claim 1 under the assumption that there is $c_j \in A_j$ such that $G'_j(a_j, \cdot) \equiv 0$ for each $a_j \neq b_j, c_j$.

We first prove this statement if $A_j = \{b_j, c_j\}$. By continuity, we may assume that $\langle p(b_j, \cdot), G'_j(c_j, \cdot) \rangle < 0$ with a strict inequality. Hence, there is $a_{-j} \in A_{-j}$ such that $G'_j(b_j, a_{-j}) > G'_j(c_j, a_{-j})$, and so there is $\hat{p} \in CE(G)$ such that $\text{supp}(\hat{p}) = \{b_j\} \times A_{-j}$ (e.g., any \hat{p} with probability sufficiently close to 1 on a_{-j} works). For each $\varepsilon > 0$, $(1 - \varepsilon)p + \varepsilon\hat{p}$ is connected, and so $(1 - \varepsilon)p + \varepsilon\hat{p} \in f(G)$ by Step 2. Then, since f is continuous and thus closed-valued, $p \in f(G)$.

Now consider the case $|A_j| > 2$. Let \tilde{G}'_j and \tilde{p} be the restrictions of G'_j and p to $\{b_j, c_j\} \times A_{-j}$. Let \tilde{G} be the game on $\{b_j, c_j\} \times A_{-j}$ obtained by applying Claim 1 to \tilde{G}'_j and \tilde{p} . Then, if G is the game on A obtained from \tilde{G} by replacing b_j by $|A_j| - 1$ clones, G has all required properties by consequentialism. This finishes the proof of Claim 1. \square

We use Claim 1 to prove the lemma. Let G be a game on A and $j \in N$ such that for each $i \neq j$, $G_i \equiv 0$, and let $p \in CE(G)$. For each $b_j \in A_j$ with $p(b_j, \cdot) \neq 0$, let $p^{b_j} \in \Delta(A)$ such that $p^{b_j}(b_j, \cdot) = \frac{p(b_j, \cdot)}{|p(b_j, \cdot)|}$. Thus, $p^{b_j}(b_j, \cdot)$ is the correlated strategy of agents other than j if j receives the signal b_j .

Fix $b_j \in A_j$ with $p(b_j, \cdot) \neq 0$. We decompose G into three types of games.

- First, let $G'_j: A \rightarrow \mathbb{R}$ such that for each $a_j \in A_j$, $G'_j(a_j, \cdot) = G_j(a_j, \cdot) - G_j(b_j, \cdot)$. Let G^{b_j} be the game on A obtained by applying Claim 1 to j , G'_j , b_j , and p^{b_j} . The hypotheses hold since p is a correlated equilibrium of G . By the claim, $G_j^{b_j} = G'_j$ and $p^{b_j} \in f(G^{b_j})$.
- Second, for each $k \neq j$, let $\bar{G}^{b_j, k}$ be the game on A such that $\bar{G}_k^{b_j, k} = -G_k^{b_j}$, and for each $i \neq k$, $\bar{G}_i^{b_j, k} \equiv 0$. Then, for each $a_{-j} \in A_{-j}$, $\chi_{(b_j, a_{-j})}$ is a pure Nash equilibrium of $\bar{G}^{b_j, k}$ since $\bar{G}_k^{b_j, k}(b_j, \cdot) \equiv 0$. Hence, by Proposition 1 and convex-valuedness, $p^{b_j} \in f(\bar{G}^{b_j, k})$.
- Third, let \tilde{G}^{b_j} be the game on A such that for each $a_j \in A_j$, $\tilde{G}_j^{b_j}(a_j, \cdot) \equiv G_j(b_j, \cdot)$, and for each $i \neq j$, $\tilde{G}_i^{b_j} \equiv 0$. Note that each action profile is a pure Nash equilibrium of \tilde{G}^{b_j} . Hence, by Proposition 1 and convex-valuedness, $f(\tilde{G}^{b_j}) = \Delta(A)$. In particular, $p^{b_j} \in f(\tilde{G}^{b_j})$.

Observe that $G = G^{b_j} + \sum_{k \neq j} \bar{G}^{b_j, k} + \tilde{G}^{b_j}$. Consistency and positive homogeneity imply that $p^{b_j} \in f(G)$.

Since $p^{b_j} \in f(G)$ for each $b_j \in A_j$ with $p(b_j, \cdot) \neq 0$, convex-valuedness implies that $p \in f(G)$ as required. \square

It is straightforward to remove the assumption that $G_i \equiv 0$ for all $i \neq j$ from Lemma 15 using consistency.

E. $f = CE$

We prove Theorem 1. The first lemma shows that CE satisfies all of the axioms.

Lemma 16 (CE satisfies the axioms). *CE satisfies totality, continuity, convex-valuedness, consistency, consequentialism, and rationality.*

Proof. CE is total since it is a coarsening of NE . For every game G on A , $CE(G)$ is defined by linear constraints, and the constraints depend linearly on G . Thus, CE is continuous and convex-valued, and it satisfies consistency. CE satisfies consequentialism since adding clones only introduces redundant constraints. Lastly, CE satisfies rationality since if a_i is dominated by b_i , player i prefers b_i to a_i when recommended to play a_i . \square

Lemma 17 (Containment in CE). *Let f be a total, continuous, and convex-valued solution concept that satisfies consistency, consequentialism, and rationality. Assume that for all $j, k \in N$, there is $\gamma_{j,k} > 0$ such that $NE(G) \subseteq f(G)$ for each $(j, k, \gamma_{j,k})$ -zero-sum game G . Then, $f \subseteq CE$.*

Proof. Assume for contradiction that there is a game G on A such that $f(G) \setminus CE(G) \neq \emptyset$, and let $p \in f(G) \setminus CE(G)$. Then, there are $j \in N$, $b_j, c_j \in A_j$, and $\epsilon > 0$ such that

$$\langle p(c_j, \cdot), G_j(c_j, \cdot) - G_j(b_j, \cdot) + \epsilon \mathbf{1} \rangle < 0 \quad (2)$$

That is, j can increase its payoff by more than ϵ by playing b_j instead of c_j when receiving the signal c_j . By introducing a clone of b_j and using consequentialism, we may assume that $p(b_j, \cdot) \equiv 0$. Let

$$\begin{aligned} A_j^- &= \{a_j \in A_j : \langle p(a_j, \cdot), G_j(c_j, \cdot) - G_j(b_j, \cdot) + \epsilon \mathbf{1} \rangle < 0\}, \text{ and} \\ A_j^+ &= \{a_j \in A_j : \langle p(a_j, \cdot), G_j(c_j, \cdot) - G_j(b_j, \cdot) + \epsilon \mathbf{1} \rangle \geq 0\} \end{aligned}$$

That is, A_j^- is the set of actions a_j of j so that if the signal is a_j , then the payoff for b_j is higher than that of c_j by more than ϵ . Note that $b_j \in A_j^+$ (since $p(b_j, \cdot) \equiv 0$), and that $c_j \in A_j^-$ by (2). Let \hat{G} be the game on A such that for each $a_j \in A_j^-$, $\hat{G}_j(a_j, \cdot) \equiv 0$, for each $a_j \in A_j^+$, $\hat{G}_j(a_j, \cdot) = G_j(c_j, \cdot) - G_j(b_j, \cdot) + \epsilon \mathbf{1}$, and for each $i \neq j$, $\hat{G}_i \equiv 0$. Observe that A_j^- and A_j^+ are sets of clones in \hat{G} , respectively. Then, $p \in CE(\hat{G})$ by definition of A_j^- and A_j^+ , and $p \in f(\hat{G})$ by Lemma 15.

To conclude, let $\bar{G} = \frac{1}{2}G + \frac{1}{2}\hat{G}$. Consistency implies that $p \in f(\bar{G})$. Note that b_j dominates c_j in \bar{G} . Since $p(c_j, \cdot) \not\equiv 0$ by (2), this contradicts rationality.¹² \square

Proof of Theorem 1. By Lemma 16, CE satisfies the axioms. It remains to show that CE is the only such solution concept.

By Proposition 2, for all $j, k \in N$, there is $\gamma_{j,k} > 0$ such that $NE(G) \subseteq f(G)$ for each $(j, k, \gamma_{j,k})$ -zero-sum game G . Then Lemma 17 implies that $f \subseteq CE$. To show that $CE \subseteq f$, let G be a game on A , and let $p \in CE(G)$. For each $j \in N$, let G^j be the game on A with $G^j_j = nG_j$, and for each $i \neq j$, $G^j_i \equiv 0$. Note that $G = \frac{1}{n} \sum_{j \in N} G^j$, and for each $j \in N$, $p \in CE(G^j)$. Lemma 15 implies that $p \in f(G^j)$ for each $j \in N$. Thus, $p \in f(G)$ by consistency, concluding the proof. \square

F. Characterization of CCE

Lemma 18 (CCE satisfies the axioms). *CCE satisfies totality, continuity, convex-valuedness, consistency, strong consequentialism, and weak rationality.*

¹²Note that this is the only step in the proof of Theorem 1 for which the additional strength of rationality over weak rationality is needed.

Proof. Totality, continuity, convex-valuedness, and consistency follow from arguments similar to those used in Lemma 16 to show that CE has these properties. CCE satisfies strong consequentialism since for any game G , whether a correlated strategy satisfies the constraints defining $CCE(G)$ does not depend on how probability is distributed over clones. Lastly, CCE satisfies weak rationality since if b_i is a dominant action, player i 's expected utility of committing to b_i is higher than that of always following their recommendation for any correlated strategy. \square

Lemma 19 (Containment of CCE). *Let f be a total, continuous, and convex-valued solution concept that satisfies consistency, strong consequentialism, and weak rationality. Then, $CCE \subseteq f$.*

Proof. By Lemma 1, we may assume that f is positively homogeneous. Several times during the proof, we use that $CE(G) \subseteq f(G)$ for any game G by Lemma 15 and Proposition 2. Let G be a game on A , and let $p \in CCE(G)$. We prove that $p \in f(G)$. By consistency and positive homogeneity, we may assume that there is $i \in N$ such that $G_j \equiv 0$ for all $j \neq i$.

First, consider the case that $G_i \in \mathbb{Q}^A$. Let p be a vertex of the convex polytope $CCE(G)$, and note that $p \in \Delta_{\mathbb{Q}}(A)$. Strong consequentialism allows reducing to a simpler situation.

Claim 1. We may assume that there is $B \in \mathcal{F}(U)$ such that for each $i \in N$, $B \subseteq A_i$, and $p = \text{uni}(D)$, where $D = \{(b, \dots, b) : b \in B\}$ is the diagonal in B^n .

Proof of Claim 1. Choose $k \in \mathbb{N}$ such that $kp(a) \in \mathbb{N}_0$ for each $a \in A$. Let $B \in \mathcal{F}(U)$ be disjoint from $\bigcup_{j \in N} A_j$ with $|B| = k$. Since $\sum_{a \in A} kp(a) = k$, we can fix a map $\sigma : B \rightarrow A$ such that for each $a \in A$,

$$|\sigma^{-1}(a)| = kp(a)$$

For each $j \in N$, let $\tilde{A}_j = A_j \cup B$ and $\tilde{A} = \tilde{A}_1 \times \dots \times \tilde{A}_n$. Let $\phi : \tilde{A} \rightarrow A$ be the surjection such that for each $j \in N$, $\phi_j(a_j) = a_j$ for each $a_j \in A_j$ and $\phi_j(b) = \sigma(b)_j$ for each $b \in B$. Let \tilde{G} be the game on \tilde{A} given by $\tilde{G} = G \circ \phi$, so that \tilde{G} is a blow-up of G with surjection ϕ . Intuitively, for each action $a \in A$, we add $kp(a)$ clones of a_j for each j .

Let $\tilde{p} = \text{uni}(D)$. For any $a \in A$,

$$\phi_*(\tilde{p})(a) = \sum_{\tilde{a} \in \phi^{-1}(a)} \tilde{p}(\tilde{a}) = \frac{1}{|B|} |\{b \in B : \sigma(b) = a\}| = \frac{kp(a)}{k} = p(a)$$

and so $\phi_*(\tilde{p}) = p$. Since CCE and f both satisfy strong consequentialism, $p \in CCE(G)$ if and only if $\tilde{p} \in CCE(\tilde{G})$, and $p \in f(G)$ if and only if $\tilde{p} \in f(\tilde{G})$.

The triple $(\tilde{G}, \tilde{A}, \tilde{p})$ has the properties claimed in Claim 1. Relabeling this triple to (G, A, p) finishes the proof. \square

If $|B| = 1$, then $p \in CE(G) \subseteq f(G)$ and we are done. For the rest of the proof, assume that $|B| \geq 2$. We show that G can be written as a convex combination of three types of games. For each such game, f returns p as a consequence of $CE \subseteq f$ and strong consequentialism. Consistency then implies that $p \in f(G)$.

For $b^* \in B$, let $\tilde{A} = (\{b^*\} \cup (A_i \setminus B)) \times A_{-i}$, and let \tilde{G}^* be the game on \tilde{A} such that

- (i) $\tilde{G}_i^*(b^*, b, \dots, b) = G_i(b, \dots, b)$ for each $b \in B$,
- (ii) $\tilde{G}_i^*(b^*, a_{-i}) = 0$ for each $a_{-i} \in A_{-i} \setminus D_{-i}$,
- (iii) $\tilde{G}_i^*(a) = G_i(a)$ for each $a \in A$ with $a_i \in A_i \setminus B$, and
- (iv) $\tilde{G}_j^* \equiv 0$ for each $j \neq i$.

Let $\tilde{p} = \text{uni}(\tilde{D})$, where $\tilde{D} = \{b^*\} \times D_{-i}$. Observe that i 's payoff for \tilde{p} in \tilde{G}^* equals i 's payoff for p in G . Since $p \in CCE(G)$, it follows that $\tilde{p} \in CE(\tilde{G}^*) \subseteq f(\tilde{G}^*)$. Let $\tilde{G} = \tilde{G}^* \circ \phi$ be the blow-up of \tilde{G}^* with surjection $\phi: A \rightarrow \tilde{A}$, where $\phi_i^{-1}(b^*) = B$, ϕ_i is the identity on $A_i \setminus B$, and for each $j \neq i$, ϕ_j is the identity on A_j . That is, \tilde{G} is obtained from \tilde{G}^* by replacing b^* by $|B|$ clones. Strong consequentialism and $\tilde{p} \in f(\tilde{G}^*)$ imply that $p \in f(\tilde{G})$. Note that \tilde{G}_i agrees with G_i on D and on $(A_i \setminus B) \times A_{-i}$.

Let $c, c' \in B$ be distinct, and let \tilde{G}^c be the game on $\tilde{A}^c = (\{c, c'\} \cup (A_i \setminus B)) \times A_{-i}$ such that

- (i) $\tilde{G}_i^c(c, b, \dots, b) = G_i(c, b, \dots, b) - G_i(b, \dots, b)$ for each $b \in B$,
- (ii) $\tilde{G}_i^c(c, a_{-i}) = 0$ for each $a_{-i} \in A_{-i} \setminus D_{-i}$,
- (iii) $\tilde{G}_i^c(a) = 0$ for each $a \in \tilde{A}^c$ with $a_i \in \{c'\} \cup (A_i \setminus B)$, and
- (iv) $\tilde{G}_j^c \equiv 0$ for each $j \neq i$.

Let $\tilde{p}^c = \text{uni}(\{(c, \dots, c)\} \cup \{(c', b, \dots, b) : b \in B \setminus \{c\}\}) \in \Delta(\tilde{A}^c)$. Since $p \in CCE(G)$, we have $\sum_{b \in B} p(b, \dots, b) (G_i(b, \dots, b) - G_i(c, b, \dots, b)) \geq 0$. That is, player i cannot profitably deviate to c before observing their recommendation in G . Hence, player i cannot profitably deviate to c when their recommendation is c' in \tilde{G}^c . Moreover, $\tilde{G}_i^c(a_i, c, \dots, c) = 0$ for each $a_i \in \tilde{A}_i^c$, and so player i cannot profitably deviate to any other action when their recommendation is c . Hence, $\tilde{p}^c \in CE(\tilde{G}^c) \subseteq f(\tilde{G}^c)$. Let $G^c = \tilde{G}^c \circ \phi$ be the blow-up of \tilde{G}^c with surjection $\phi: A \rightarrow \tilde{A}^c$, where $\phi_i^{-1}(c') = B \setminus \{c\}$, ϕ_i is the identity on $(A_i \setminus B) \cup \{c\}$, and for each $j \neq i$, ϕ_j is the identity on A_j . Strong consequentialism implies that $p \in f(G^c)$.

Lastly, let \hat{G} be the game on A with

- (i) $\hat{G}_i(a) = 0$ for each $a \in A$ with $a_i \in B$ and $a_{-i} \in D_{-i}$,
- (ii) $\hat{G}_i(a) = G_i(a)$ for each $a \in A$ with $a_i \in B$ and $a_{-i} \notin D_{-i}$,
- (iii) $\hat{G}_i(a) = 0$ for each $a \in A$ with $a_i \in A_i \setminus B$, and
- (iv) $\hat{G}_j \equiv 0$ for each $j \neq i$.

Since $\hat{G}_i(a) = 0$ for each $a \in A$ with $a_{-i} \in D_{-i}$, it follows that $p \in CE(\hat{G}) \subseteq f(\hat{G})$.

Observe that $G = \tilde{G} + \hat{G} + \sum_{c \in B} G^c$. Thus, consistency and positive homogeneity imply that $p \in f(G)$. This shows that $f(G)$ contains all vertices of $CCE(G)$, and thus that $CCE(G) \subseteq f(G)$ by convex-valuedness.

Second, consider the general case $G_i \in \mathbb{R}^A$, and let p be a vertex of $CCE(G)$. By convex-valuedness, it suffices to prove that $p \in f(G)$. There exist sequences of games $(G^t)_{t \in \mathbb{N}}$ on A and $(p^t)_{t \in \mathbb{N}} \subseteq \Delta_{\mathbb{Q}}(A)$ such that, for each $t \in \mathbb{N}$, $G_i^t \in \mathbb{Q}^A$, $G_j^t \equiv 0$ for all $j \neq i$, and $p^t \in CCE(G^t)$ and $G_i^t \rightarrow G_i$ and $p^t \rightarrow p$.¹³ Then it follows from the first part that $p^t \in f(G^t)$ for each $t \in \mathbb{N}$, and thus $p \in f(G)$ by continuity. This finishes the proof. \square

Lemma 20 (Containment in CCE). *Let f be a total, continuous, and convex-valued solution concept that satisfies consistency, strong consequentialism, and weak rationality. Then, $f \subseteq CCE$.*

Proof. Assume for contradiction that $f \not\subseteq CCE$. Then, there exists a game G on A , $p \in f(G)$, $i \in N$, $a_i^* \in A_i$, and $\varepsilon > 0$ such that

$$\sum_{a \in A} p(a) (G_i(a) - G_i(a_i^*, a_{-i})) < -\varepsilon \quad (3)$$

That is, i can increase its payoff by more than ε by deviating to a_i^* before observing its signal. We use strong consequentialism to construct a game \tilde{G} on \tilde{A} and a correlated strategy $\tilde{p} \in f(\tilde{G}) \setminus CCE(\tilde{G})$ such that for each $a_{-i} \in \tilde{A}_{-i}$, there is at most one action $a_i \in \tilde{A}_i$ with $\tilde{p}(a_i, a_{-i}) > 0$. This step is similar to the proof of Claim 1 in Lemma 19.

Let $B \in \mathcal{F}(U)$ be disjoint from $\bigcup_{j \in N} A_j$ and $|B| = |A|$. Fix a bijection $\sigma: B \rightarrow A$. Let $\tilde{A} = (A_1 \cup B) \times \cdots \times (A_n \cup B)$, and let $\tilde{G} = G \circ \phi$ be the blow up of G with surjection $\phi: \tilde{A} \rightarrow A$ such that for each $j \in N$ and $a_j \in A_j$, $\phi_j^{-1}(a_j) = \{a_j\} \cup \{b \in B: a_j = \sigma(b)\}$. That is, for each action a_j of each player j , we add $|A|/|A_j| = |A_{-j}|$ clones of a_j labeled $\sigma^{-1}(a_j, a_{-j})$, where a_{-j} ranges over A_{-j} . Let $\tilde{p} \in \Delta(\tilde{A})$ such that for each $b \in B$, $\tilde{p}(b, \dots, b) = p(\sigma(b))$. Hence, \tilde{p} is supported on the diagonal $D = \{(b, \dots, b): b \in B\}$ of B^n . By construction, $\phi_*(\tilde{p}) = p$, and so $\tilde{p} \in f(\tilde{G})$ by strong consequentialism. By (3),

$$\sum_{b \in B} \tilde{p}(b, \dots, b) \left(\tilde{G}_i(b, \dots, b) - \tilde{G}_i(a_i^*, b, \dots, b) \right) < -\varepsilon \quad (4)$$

Let $C > \max_{a, a' \in A} (G_i(a) - G_i(a'))$ be a number larger than the maximal payoff difference between any two action profiles of i . Let \hat{G} be the game on \tilde{A} such that

- (i) $\hat{G}_i(b, \dots, b) = \tilde{G}_i(a_i^*, b, \dots, b) - \tilde{G}_i(b, \dots, b) - \varepsilon$ for each $b \in B$,
- (ii) $\hat{G}_i(a_i^*, \cdot) \equiv 0$,
- (iii) $\hat{G}_i(a) = -C$ for each $a \in \tilde{A} \setminus D$ with $a_i \neq a_i^*$, and

¹³This follows from a standard continuity property of solutions to linear programs and the fact that $CCE(G)$ is the set of solutions to a system of linear constraints. For $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$, let $P = \{x \in \mathbb{R}^n: Ax \leq b\}$ and let x be a vertex. Then some set of constraints $J \subseteq [m]$ with $|J| = n$ is tight at x and A_J is nonsingular, hence $x = A_J^{-1}b_J$. Let $\delta = \min_{k \notin J} (b_k - a_k^\top x) > 0$. Approximate (A, b) by rationals $(A^t, b^t) \rightarrow (A, b)$ so that for all large t , A_J^t stays nonsingular and the perturbation is small compared to δ . Set $x^t = (A_J^t)^{-1}b_J^t$. Then $x^t \rightarrow x$ and, for t large, the constraints in J are tight at x^t while all others remain slack; thus x^t is a basic feasible solution of $P^t = \{x: A^t x \leq b^t\}$, which has only rational-valued constraints.

(iv) $\hat{G}_j \equiv 0$ for each $j \neq i$.

Then, $\tilde{p} \in CCE(\hat{G}) \subseteq f(\hat{G})$ by (4) and Lemma 19.

Consistency implies that $\tilde{p} \in f(\frac{1}{2}\tilde{G} + \frac{1}{2}\hat{G})$. Note that a_i^* is a dominant action for i in $\frac{1}{2}\tilde{G} + \frac{1}{2}\hat{G}$, and $\tilde{p}(a_i^*, \cdot) \equiv 0$. This contradicts weak rationality. \square

Theorem 2 is immediate from Lemma 19 and Lemma 20.

References

- K. J. Arrow and L. Hurwicz. An optimality criterion of decision-making under ignorance. In C. F. Carter and J. L. Ford, editors, *Uncertainty and expectations in economics: essays in honour of G.L.S. Shackle*, pages 1–11. Basil Blackwell, 1972.
- R. J. Aumann and A. Brandenburger. Epistemic conditions for Nash equilibrium. *Econometrica*, 63(5):1161–1180, 1995.
- R. J. Aumann and J. H. Drèze. Rational expectations in games. *American Economic Review*, 98(1):72–86, 2008.
- C. W. Bach and E. Tsakas. Pairwise epistemic conditions for Nash equilibrium. *Games and Economic Behavior*, 85:48–59, 2014.
- P. Barelli. Consistency of beliefs and epistemic conditions for nash and correlated equilibria. *Games and Economic Behavior*, 67(2):363–375, 2009.
- F. Brandl and F. Brandt. Justifying optimal play via consistency. *Theoretical Economics*, 14(4):1185–1201, 2019.
- F. Brandl and F. Brandt. An axiomatic characterization of Nash equilibrium. *Theoretical Economics*, 19(4):1473–1504, 2024.
- F. Brandl, F. Brandt, and H. G. Seedig. Consistent probabilistic social choice. *Econometrica*, 84(5):1839–1880, 2016.
- S. Chatterji and S. Govindan. Message spaces for perfect correlated equilibria. *Economic Theory*, 28(2):475–479, 2006.
- H. Chernoff. Rational selection of decision functions. *Econometrica*, 22(4):422–443, 1954.
- M. Crescenzi. A choice-based axiomatization of Nash equilibrium. Technical report, arXiv.org, 2026.
- A. Dhillon and J.-F. Mertens. Perfect correlated equilibrium. *Journal of Economic Theory*, 68(2):279–302, 1996.
- I. Gilboa and D. Schmeidler. A derivation of expected utility maximization in the context of a game. *Games and Economic Behavior*, 44(1):172–182, 2003.
- S. Govindan and R. Wilson. Axiomatic equilibrium selection for generic two-player games. *Econometrica*, 80(4):1639–1699, 2012.
- J. C. Harsanyi. Games with incomplete information played by “Bayesian” players, part I. *Management Science*, 50(12):1804–1817, 1967.

- J. C. Harsanyi and R. Selten. A generalized Nash solution for two-person bargaining games with incomplete information. *Management Science*, 18(5):80–106, 1972.
- Z. Hellman. Weakly rational expectations. *Journal of Mathematical Economics*, 49(6):496–500, 2013.
- W. Huang, J. J. Kline, and P. Man. Correlated perfect equilibrium. 2026.
- A. T. Kalai and E. Kalai. Beyond dominance and Nash: Ranking equilibria by critical mass. *Games and Economic Behavior*, 144:378–394, 2024.
- E. Kohlberg and J.-F. Mertens. On the strategic stability of equilibria. *Econometrica*, 54:1003–1037, 1986.
- M. Lackner and P. Skowron. Consistent approval-based multi-winner rules. *Journal of Economic Theory*, 192:105173, 2021.
- X. Luo, Y. Qiao, and Y. Sun. A revelation principle for correlated equilibrium under trembling-hand perfection. *Journal of Economic Theory*, 200:105396, 2022.
- E. Maskin. Decision-making under ignorance with implications for social choice. *Theory and Decision*, 11(3):319–337, 1979.
- J. Milnor. Games against nature. In *Decision Processes*, chapter 4, pages 49–59. Wiley, 1954.
- H. Moulin and J.-P. Vial. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3–4):201–221, 1978.
- R. B. Myerson. Acceptable and predominant correlated equilibria. *International Journal of Game Theory*, 15(3):133–154, 1986.
- R. B. Myerson. Axiomatic derivation of scoring rules without the ordering assumption. *Social Choice and Welfare*, 12(1):59–74, 1995.
- H. Norde, J. Potters, H. Reijnierse, and D. Vermeulen. Equilibrium selection and consistency. *Games and Economic Behavior*, 12(2):219–225, 1996.
- B. Peleg and S. H. Tijs. The consistency principle for games in strategic form. *International Journal of Game Theory*, 25(1):13–34, 1996.
- H. Salonen. An axiomatic analysis of the Nash equilibrium concept. *Theory and Decision*, 33:177–189, 1992.
- F. Sandomirskiy, P. H. Sung, O. Tamuz, and B. Wincelberg. Narrow bracketing and risk in games. 2025. Working paper.

- L. J. Savage. *The Foundations of Statistics*. Wiley Publications in Statistics. Wiley and Sons, 1954.
- R. Selten. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1):25–55, 1975.
- L. S. Shapley. A value for n-person games. *Annals of Math Studies*, 28:307–317, 1953.
- L. S. Shapley and R. N. Snow. Basic solutions of discrete games. *Annals of Mathematics Studies*, 24, 1950.
- J. H. Smith. Aggregation of preferences with variable electorate. *Econometrica*, 41(6):1027–1041, 1973.
- T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.
- M. Voorneveld. An axiomatization of the Nash equilibrium concept. *Games and Economic Behavior*, 117:316–321, 2019.
- H. P. Young. Social choice scoring functions. *SIAM Journal on Applied Mathematics*, 28(4):824–838, 1975.
- H. P. Young and A. B. Levenglick. A consistent extension of Condorcet’s election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978.
- T. M. Zavist and T. N. Tideman. Complete independence of clones in the ranked pairs rule. *Social Choice and Welfare*, 6(2):167–173, 1989.