

Extreme Equilibria: The Benefits of Correlation*

Kirill Rudov[†] Fedor Sandomirskiy[‡] Leeat Yariv[§]

February 15, 2025

Abstract.

We study whether a given Nash equilibrium can be improved within the set of correlated equilibria for arbitrary objectives. Our main contribution is a sharp characterization: in a generic game, a Nash equilibrium is an extreme point of the set of correlated equilibria if and only if at most two agents randomize. Consequently, any sufficiently mixed Nash equilibrium involving at least three randomizing agents can always be improved by correlating actions or switching to a less random equilibrium, regardless of the underlying objective. We show that even if one focuses on objectives that depend on payoffs, excess randomness in equilibrium implies improvability. We extend our analysis to symmetric games, incomplete information games, and coarse correlated equilibria, revealing a fundamental tension between the randomness in Nash equilibria and their optimality.

Keywords: Extreme Equilibria, Correlated Equilibria

JEL codes: C60, C72, D02, D60

*We thank S. Nageeb Ali, Benjamin Brooks, Emir Kamenica, Navin Kartik, Jacopo Peregó, Doron Ravid, and Éva Tardos for helpful comments and suggestions. We also appreciate the input of seminar audiences at Columbia University, Penn State University, University of Chicago, Higher School of Economics, ASSA 2025 Annual Meeting, and the Stony Brook Game Theory Festival. We gratefully acknowledge financial support from the National Science Foundation through grant SES-1949381.

[†]University of California at Berkeley; krudov@berkeley.edu

[‡]Princeton University; fsandomi@princeton.edu

[§]Princeton University, CEPR, and NBER; lyariv@princeton.edu

1 Introduction

Correlated equilibria are a powerful generalization of Nash equilibrium, offering both computational and strategic advantages. Unlike Nash equilibria, which can be notoriously difficult to compute in general games, correlated equilibria are computationally simple. They also naturally arise in settings where agents can communicate or collude, making them a plausible solution concept for many applications. Moreover, they are mechanism-implementable: a mediator can credibly induce agents to follow correlated strategies without requiring external enforcement. Given these advantages, a fundamental question arises: when can Nash equilibria be strictly improved through correlation?

This paper examines conditions under which correlated equilibria outperform Nash equilibria, revealing a key insight: whenever a Nash equilibrium involves significant randomization, it is improvable, irrespective of the objective. This observation has broad implications, from auction design to voting mechanisms to firm competition. In many strategic settings, decision-makers rely on mixed strategies, whether bidders randomizing in auctions, firms setting stochastic pricing policies, or voters deciding probabilistically whether to turn out. Our results show that such randomness frequently leaves room for improvement via correlation.

In Section 3, we establish that whenever a Nash equilibrium involves more than two agents using mixed strategies, it is improvable. The set of correlated equilibria is convex, meaning that any generic objective function is maximized at an extreme point of this set. Thus, unless a Nash equilibrium is an extreme point itself, it can be improved irrespective of the objective. As we show, extremality is a stringent requirement.

To build intuition, consider a game where n agents each choose between two actions. Such a setup arises in various strategic settings: voters deciding whether to cast a costly ballot, firms choosing whether to engage in costly R&D, or individuals in a social network deciding whether to adopt a new technology. A correlated equilibrium is a probability distribution over action profiles, where each agent receives a recommendation and finds it optimal to follow. The feasibility of such an equilibrium is governed by $2n$ linear constraints: one per agent per action. A foundational linear programming result (Winkler, 1988) states that any extreme point of this set must be supported on at most $2n + 1$ action profiles. Now, consider a Nash equilibrium where all agents use mixed strategies, implying that every action profile is played with positive probability. Since this yields 2^n active action profiles, the equilibrium can only be an extreme point if $2^n \leq 1 + 2n$. This condition holds only for $n = 1, 2$, meaning that for any binary-action game with three or more agents who all mix, Nash equilibria are necessarily non-extreme, and hence, improvable.

In Section 4, we extend our analysis to objectives based on payoffs, rather than dis-

tributions over action profiles. We show that even when focusing solely on utilitarian welfare—the sum of agents’ payoffs—mixed-strategy Nash equilibria remain generically improvable whenever more than two agents participate.

The intuition follows from convexity principles. The set of achievable payoffs under correlated equilibria forms a linear projection of the correlated equilibrium polytope onto a lower-dimensional space. The extreme points of a projected set are contained within the projection of the extreme points of the original set. Given that utilitarian welfare is almost always maximized at a unique extreme point, it follows that Nash equilibria with excess randomness leave room for welfare improvements via correlation.

Symmetric Games and Symmetric Improvements Symmetric games are prevalent in economic and strategic settings, appearing in auctions, voting models, and competitive markets, where identical agents face identical strategic choices. In such games, symmetric equilibria, with all agents adopting the same strategy, often emerge as focal solutions.

In Section 5, we show that in symmetric games with three or more agents, any symmetric mixed equilibrium can be improved, even when restricting attention to symmetric correlated equilibria. This result implies that even when coordination is limited to symmetric recommendations—such as a common policy guideline for firms or a uniform turnout mobilization effort in elections—improvements remain possible. Only pure-strategy equilibria resist such improvements.

Our results extend to games where agents share identical strategic roles, even if full symmetry does not hold. Specifically, we consider games invariant to a *transitive* subset of agent permutations, where each agent can be mapped onto another via some permutation.¹ Even in these broader settings, purity remains a necessary condition for non-improvability whenever more than two agents participate.

Furthermore, we establish that extreme symmetric correlated equilibria in symmetric games can be represented via a simple procedure. When the number of participants is large, these equilibria can be reformulated in terms of a small set of payoff-irrelevant states, each generating conditionally independent and identically distributed signals that guide individual decisions.

Implications In Section 6, we explore implications for incomplete information games and alternative notions of correlation. Our results connect to the general framework of Bergemann and Morris (2019), where agents hold private information under common payoff uncertainty. We show that a Bayesian Nash equilibrium is an extreme point of

¹This generalization traces back to Nash (1950) and has been studied in social choice contexts by Bartholdi, Hann-Caruthers, Josyula, Tamuz, and Yariv (2021).

the Bayesian correlated equilibrium set if and only if it is pure, reinforcing the necessity of determinism for non-improvability in broader strategic settings.

We also examine a weaker notion of coarse correlated equilibrium, where agents commit ex-ante to following recommendations generated by an external coordinating device. This framework captures real-world scenarios like firms entering binding collusive agreements (McAfee and McMillan, 1992) or consumers subscribing to algorithmic recommender systems. The set of coarse correlated equilibria strictly contains the set of correlated equilibria, making extremality within this set an even stronger constraint. Indeed, we show that, generically, a Nash equilibrium is an extreme point of the coarse correlated equilibrium set if and only if it is pure.

Summing up Our findings provide a systematic framework for understanding when correlation can enhance strategic interactions. While Nash equilibria serve as a foundational concept in game theory, our results show that they often fall short of optimality when agents randomize. This observation has direct implications for market design, regulatory policies, and algorithmic decision-making in multi-agent systems. Whether in competitive markets, voting mechanisms, or algorithmically mediated platforms, correlation provides a natural tool for improving strategic outcomes. By demonstrating the ubiquity of improvable Nash equilibria, this paper underscores the practical and theoretical importance of correlated strategies in economic and strategic decision-making.

1.1 Related literature

First introduced by Aumann (1974) and Aumann (1987), correlated equilibria have received substantial attention in the literature. They are computationally simpler than Nash equilibria; see Papadimitriou and Roughgarden (2008). Correlated equilibria also offer a reduced-form way of capturing pre-play communication without explicitly modeling the communication phase; see Forges (2020). Similarly, they result from a variety of learning heuristics; see Foster and Vohra (1997), Fudenberg and Levine (1999), and Hart and Mas-Colell (2000).

While Aumann (1974) presents a 2-agent game of chicken where correlation enhances utilitarian welfare, Ashlagi, Monderer, and Tennenholtz (2008) explore the potential magnitude of this improvement, primarily focusing on 2×2 games and analyzing the resulting welfare gains from correlation; see also Bradonjic, Ercal, Meyerson, and Roytman (2014). Peeters and Potters (1999), Hendrickx, Peeters, and Potters (2002), and Calvó-Armengol (2006) also study the structure of the correlated equilibrium set of 2×2 games. However, general conditions under which correlation improves a given equilibrium have remained elusive, a gap this paper seeks to address.

Several studies examine 2-agent games. Cripps (1995), Evangelista and Raghavan (1996), and Canovas, Hansen, and Jaumard (1999) demonstrate that in 2-agent games, all Nash equilibria are extreme within the correlated equilibrium set. As we illustrate, 2-agent games are not representative of the general case. In particular, with more than two agents, Nash equilibria are not necessarily extreme points of the correlated equilibrium set, and the degree of randomness in equilibrium profiles serves as a sufficient indicator of their extremality.²

Correlated equilibria have been analyzed in a variety of specific applications, including Cournot competition (Gérard-Varet and Moulin, 1978), abatement games (Moulin, Ray, and Gupta, 2014), quadratic games (Dokka, Moulin, Ray, and SenGupta, 2023), various auctions (Lopomo, Marx, and Sun, 2011; Feldman, Lucier, and Nisan, 2016; Agranov and Yariv, 2018; Pavlov, 2023; Ahunbay and Bichler, 2024), and voting (Gerardi and Yariv, 2007). Their empirical relevance has been seen in various experimental settings: for example, in voting contexts (Goeree and Yariv, 2011), in bargaining (Agranov and Tergiman, 2014), in auctions (Agranov and Yariv, 2018), and in symmetric bimatrix games (Georgalos, Ray, and SenGupta, 2020; Friedman, Rabanal, Rud, and Zhao, 2022).³

Technically, our results contribute to the growing literature on extreme-point methods in economic theory; see, e.g., Manelli and Vincent (2007), Kleiner, Moldovanu, and Strack (2021), Arieli, Babichenko, Smorodinsky, and Yamashita (2023), Yang and Zentefis (2024), Nikzad (2022), Kleiner, Moldovanu, Strack, and Whitmeyer (2024), and Lahr and Niemyer (2024). Results in Section 5 rely on recent developments in the study of finite exchangeable distributions and approximation results of dependent samples with independent ones. See Diaconis and Freedman (1980), Arratia and Tavaré (1994), Arratia, Barbour, and Tavaré (2003), and Stam (1978).

2 Model

Consider a game $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$, where $N = \{1, \dots, n\}$ is a finite set of agents, A_i is a finite set of actions of agent i , the function $u_i: A \rightarrow \mathbb{R}$ represents the utility of each agent i , and $A = \prod_{i \in N} A_i$ is the set of action profiles.

Consider the set $\text{CE}(\Gamma)$ of correlated equilibria. It consists of all probability distribu-

²One thread of literature considers games where Nash equilibria cannot be improved via correlation. Neyman (1997) shows this is the case for games with a continuum of actions and a concave potential. Ui (2008) extends this result to games that may not admit a potential, but where each agent’s payoff is concave in her own actions. See also Einy, Haimanko, and Lagziel (2022), Wu (2008), and Jann and Schottmüller (2015) for specific applications where there is a unique correlated equilibrium, implying non-improvability of the unique Nash equilibrium.

³Recent work uses simulations to inspect the correlated outcomes generated by certain learning algorithms, such as those minimizing regret. For example, in the auction context, see Kolumbus and Nisan (2022).

tions $\mu \in \Delta(A)$ such that for all $i \in N$ and for all distinct $a_i, a'_i \in A_i$ we have

$$\sum_{a_{-i} \in A_{-i}} \mu(a_i, a_{-i}) u_i(a_i, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} \mu(a_i, a_{-i}) u_i(a'_i, a_{-i}). \quad (1)$$

One can interpret μ as the distribution of actions recommended by a mediator, ensuring that each agent i finds it optimal to follow the recommended action a_i , as captured by the incentive constraint (1).

Correlated equilibria μ that are product distributions— $\mu = \mu_1 \times \dots \times \mu_n$ with $\mu_i \in \Delta(S_i)$ —form the set of Nash equilibria $\text{NASH}(\Gamma)$. The set $\text{NASH}(\Gamma)$ is non-empty (Nash, 1950) and thus $\text{CE}(\Gamma)$ is non-empty as well. Since $\text{CE}(\Gamma)$ is cut from $\Delta(S)$ by a finite number of incentive constraints (6), it is a non-empty convex polytope and can be described as the convex hull of its extreme points.⁴

We call a Nash equilibrium **extreme** if it is an extreme point of the set of correlated equilibria $\text{CE}(\Gamma)$. Our main goal is to provide conditions under which Nash equilibria are extreme. As we show below, when Nash equilibria are not extreme, the agents or the social planner would potentially have opportunities to improve outcomes via channels of communication, the use of intermediaries, or the design of mechanisms.

To make general structural insights about extreme Nash equilibria possible, we need to rule out trivial examples, such as degenerate games where an agent is indifferent across all actions. We exclude such examples using the classical notion of regularity introduced by Harsanyi (1973b). In essence, a Nash equilibrium of a game is **regular** if it remains stable under small perturbations of payoffs.⁵

Arguably, only regular equilibria are relevant for economic modeling, which makes regularity a standard assumption (Van Damme, 1991).⁶ Furthermore, games with all regular equilibria are prevalent. We term any game from an open everywhere dense set of games with the complement having zero Lebesgue measure a **generic game**. As shown by Harsanyi (1973b), games where all equilibria are regular are generic. In particular, a small perturbation of a regular game preserves regularity. For example, if payoffs u are drawn at random according to some density over $\mathbb{R}^{N \times A}$, the resulting game will be generic with probability one.

⁴A point x of a convex set X is called extreme if it cannot be represented as a non-trivial convex combination of other points from X , i.e., $x = \alpha y + (1 - \alpha) y'$ with $\alpha \in (0, 1)$ and $y, y' \in X$ can only hold if $y = y' = x$.

⁵A Nash equilibrium ν of a game Γ is regular if the incentive constraints outside the support of ν are not active, and the conditions of the implicit function theorem are satisfied on the support of ν . This ensures that the equilibrium weights are smooth functions of payoffs in a small neighborhood of ν . A regular equilibrium places zero weight on weakly dominated actions.

⁶For example, regularity of all equilibria is assumed in classical results on the oddness of Nash equilibria (Harsanyi, 1973b; Wilson, 1971) or purification (Harsanyi, 1973a; Govindan, Reny, and Robson, 2003).




FIGURE 1: A non-degenerate objective (left panel) and a degenerate one (right panel)

Extremality and improvability Extreme points play an important role in economic analysis due to Bauer’s Maximum Principle (see, e.g., [Border, 2006](#), Theorem 7.69). It asserts that any linear or convex objective maximized over a convex compact set attains its maximum at an extreme point; moreover, if the objective is strictly convex, the optimum is unique. In particular, a social planner maximizing *any* linear or convex objective $W: \Delta(S) \rightarrow \mathbb{R}$ over the set of correlated equilibria can restrict her attention to extreme points of $CE(\Gamma)$. Thus, non-extremality offers a conservative perspective on the improvability of an equilibrium: it can be improved for *all* objectives in a large class. We now formalize this idea.

First, consider maximizing a linear objective W —for instance, utilitarian welfare—over a polytope X that represents the set of feasible outcomes, such as the set of correlated equilibria. Crucially, linearity here pertains to linearity in probabilities, not in the actions themselves. In particular, common objectives—such as expected welfare, revenue, or the likelihood of a given action profile—are linear.

An objective is **non-degenerate** if it attains a unique optimum over X .⁷ The left panel of Figure 1 illustrates a non-degenerate objective, its maximum attained at a unique extreme point of X . The right panel depicts a knife-edge case, where the level hyperplanes of the objective are parallel to a face of X , and thus the optimum is not unique.

Although the optima of a degenerate objective may include non-extreme points, a small perturbation of such a degenerate objective suffices to rule out all non-extreme equilibria. Indeed, consider an ε -perturbed objective $W_\varepsilon(\mu) = W(\mu) + \sum_{a \in A} \varepsilon_a \cdot \mu(a)$, where $\varepsilon \in \mathbb{R}^A$ is a vector of small shocks. If shocks ε are taken at random with any absolutely continuous distribution, none of the non-extreme equilibria is optimal for W_ε with probability one. For example, one can take ε to be uniform on the ball of radius r with an arbitrary $r > 0$. In other words, for any small ball in the space of linear objectives, a non-extreme Nash equilibrium cannot be optimal for all objectives except degenerate ones, which form a measure-zero set. Accordingly, a designer who is uncertain about the exact weights of her

⁷An objective can be degenerate with respect to one polytope and non-degenerate with respect to another.

objective will never choose a non-extreme Nash equilibrium.

More generally, consider a designer with a non-degenerate (weakly) convex objective W . Convexity of W may capture risk aversion: for example, a regulator may seek to avoid action profiles triggering bank runs, assigning increasingly high costs to equilibria that put substantial weight on such outcomes. Suppose a Nash equilibrium ν is non-extreme. Hence, ν can be represented as $\alpha\mu + (1 - \alpha)\mu'$ with $\alpha \in (0, 1)$ and distinct $\mu, \mu' \in \text{CE}(\Gamma)$. By convexity, we obtain $W(\nu) \leq \alpha W(\mu) + (1 - \alpha)W(\mu')$ and, thus, one of μ, μ' gives at least as high value to W as ν . Since W has a unique optimum, we conclude that ν cannot be optimal. Thus, a non-extreme Nash equilibrium is never optimal. The conclusion extends to quasi-convex objectives since Bauer's principle admits such a generalization [Ball \(2023\)](#).

Aumann's example and improvability Consider a version of [Aumann \(1974\)](#)'s classical example illustrating the power of correlated equilibria. There are two agents, who face the following game (each entry represents an action profile and contains the corresponding payoff pair, where the first is the row agent's payoff and the second is the column's payoff):

$$\begin{pmatrix} 0, 0 & 4, 1 \\ 1, 4 & 3, 3 \end{pmatrix} \tag{2}$$

The designer aims to maximize the social welfare $W(\mu) = \sum_a (u_1(a) + u_2(a)) \cdot \mu(a)$ over the set of all correlated equilibria μ . For this game, the incentive constraints (6) defining the set of correlated equilibria take an especially simple form: the weight of each action profile (4, 1) or (1, 4) must be at least the weight of each of (0, 0) or (3, 3).

This game has a mixed Nash equilibrium where each agent randomizes uniformly. Its utilitarian welfare level of 4 can be improved in an incentive-compatible way by reducing the weight of the (0, 0) outcome. In fact, there is a correlated equilibrium in which each of the non-zero payoff pairs is reached with probability 1/3. Its utilitarian welfare of 16/3 exceeds that of the two pure Nash equilibria (1, 4) and (4, 1) and is optimal.

Although none of the Nash equilibria are optimal for the utilitarian welfare objective, they are all extreme and thus non-improvable according to our conservative perspective. Indeed, for some objectives, these equilibria are maximizers. For instance, if the designer cares about the total cubed payoffs of agents, the pure asymmetric equilibria become optimal. If the objective is to maximize the total weight assigned to the action profiles (0, 0) and (3, 3), the mixed Nash equilibrium becomes the unique optimum.

This example underscores the observation that our notion of improvability is demanding, and suggests that Nash equilibria may perhaps rarely, if ever, be non-extreme. As we will see, this intuition is limited to 2-agent games.

3 A Characterization of Extreme Nash Equilibria

We show that whether or not a Nash equilibrium is extreme depends on the amount of randomization agents invoke in their strategies. Roughly speaking, equilibria with substantial uncertainty over the action profiles implemented cannot be extreme.

Theorem 1. *A regular Nash equilibrium is extreme if and only if at most two agents randomize.*

Theorem 1 indicates that for any social planner with a strictly convex objective, if three or more agents randomize at a Nash equilibrium, the outcome can be improved, either by introducing correlation—by allowing agents to communicate or introducing mediation—or by selecting an alternative pure or “almost pure” equilibrium when such exist.

For instance, suppose the planner cares about overall agents’ efficiency, but dislikes uncertainty over outcomes. For example, the welfare may aim at maximizing the sum of overall efficiency net of the entropy of outcomes. Theorem 1 implies that the resolution to the trade-off between payoff efficiency and uncertainty is, in many ways, detail free. Naturally, a dislike of uncertainty would push the social planner to impose limited mixing by agents. The theorem indicates that, regardless of the number of agents or the payoff structure, the social planner could always improve upon a Nash equilibrium in which more than two agents mix.

In games involving two randomizing agents, Theorem 1 asserts that any Nash equilibrium cannot be improved upon. This echoes Cripps (1995), Evangelista and Raghavan (1996), and Canovas et al. (1999), who demonstrated that any Nash equilibrium of a generic 2-agent game Γ is an extreme point of $CE(\Gamma)$.

Theorem 1 also has implications for games with a unique correlated equilibrium. For example, generic two-agent conflicting-interest (constant-sum) games have this property. As demonstrated by Viossat (2008), a small enough perturbation of any game with a unique correlated equilibrium retains the property. Thus, the set of games with a unique correlated equilibrium is an open set within the set of all games. A unique correlated equilibrium is necessarily a Nash equilibrium. Theorem 1 provides insights on the structure of such equilibria.

Corollary 1. *If Γ is a game with a unique correlated equilibrium v , then v is either a pure Nash equilibrium or entails precisely two agents mixing.*

We emphasize that this corollary requires neither genericity nor regularity assumptions. While Theorem 1 implies this result for a dense subset of games within the open set of games with a unique correlated equilibrium, we can extend the conclusion to the entire set. This follows from the upper hemicontinuity of the Nash correspondence, which,

combined with the uniqueness of the correlated equilibrium, implies the continuity of the Nash equilibrium correspondence in this specific setting.

Why are there such stringent restrictions on mixing for equilibria to be non-extreme? To glean some intuition, consider the case in which n agents each have two actions. If μ is a correlated equilibrium, it must satisfy the constraints given by (1). There are at most $2n$ such constraints: for each agent who is recommended either of the two actions, we need to ensure that she does not deviate to the other. Each additional constraint adds at most one element to the support of the extreme distributions. Therefore, the support of any extreme point of the set of correlated equilibria is bounded by $2n+1$. Suppose, now, that ν is a Nash equilibrium with the n agents mixing. Then, the support of ν contains 2^n elements. For ν to be non-extreme, it must be that $2^n \leq 2n+1$. In particular, for $n > 2$, a Nash equilibrium that involves all agents mixing necessarily has greater support than an extreme point of the set of correlated equilibria and is, therefore, non-extreme.

The formal proof of the theorem is contained in Appendix 1.1. There, we formulate a slightly stronger result that also provides a lower bound on the dimension of the face carrying a non-extreme Nash equilibrium. The dimension of this face grows exponentially with the number k of randomizing agents and is at least

$$2^{k-3} - 1.$$

The proof relies on two lemmas echoing the simple example above. The first bounds the support size of any extreme point of the set of correlated equilibria.

Lemma 1. *For any game Γ , if a distribution $\mu \in \Delta(S)$ is an extreme point of the set of correlated equilibria, then*

$$|\text{supp } \mu| \leq 1 + \sum_{i \in N} |S_i| \cdot (|S_i| - 1). \quad (3)$$

Lemma 1 suggests that in games between a large number of agents, correlated equilibria have a relatively small support. As an illustration, consider a game involving n agents, each with two available actions. The lemma indicates that an extreme correlated equilibrium puts positive weight on at most $n+1$ strategy profiles, which grows linearly in n . In contrast, the total number of strategy profiles is 2^n , which grows exponentially.

The proof of Lemma 1 generalizes the example above and relies on a basic principle in linear programming formalized by Winkler (1988). Namely, if μ is an extreme point of a set of probability measures satisfying linear constraints, and k of those constraints are binding at μ , then the support of μ cannot exceed $k+1$. In our setting, the linear constraints in question are governed by those defining correlated equilibria, and given in Equation 6.

The second lemma we utilize follows from [McKelvey and McLennan \(1997\)](#). It restricts the number of actions that agents can use in a regular Nash equilibrium.

Lemma 2 ([McKelvey and McLennan \(1997\)](#)). *Consider an n -agent game Γ and a regular Nash equilibrium $\nu = (\nu_1, \nu_2, \dots, \nu_n)$. Then, for any agent i ,*

$$|\text{supp } \nu_i| - 1 \leq \sum_{j \neq i} (|\text{supp } \nu_j| - 1). \quad (4)$$

Lemma 2 links the support size of agents' strategies, bounding their plausible variability.⁸ For example, in a regular Nash equilibrium, it cannot be that all agents but one use pure strategies. In general, a regular Nash equilibrium cannot have an agent mixing across many actions, while all others mix over a far smaller set.

Lemmas 1 and 2 indicate that a regular Nash equilibrium cannot be improved if and only if both constraints (3) and (4) hold. In the proof of Theorem 1, we use majorization techniques to show that the combination of the two constraints indeed implies limited mixing.

Our results have implications for strategic settings where Nash equilibria require substantial randomization: in such cases, outcomes can be improved if agents coordinate, either autonomously or through intermediaries. Many real-world applications exhibit unique Nash equilibria with significant mixing. Consider, for instance, a simplified version of the costly voting model of [Palfrey and Rosenthal \(1983\)](#), where two voter groups, D and R , exist with $|D| < |R|$. Voters in group D derive a utility of 1 if their preferred candidate, d , wins and 0 otherwise, while voters in group R receive utility 1 if r wins and 0 otherwise. Ties are broken randomly, and casting a vote incurs a small cost, $c > 0$. For intermediate values of c , [Palfrey and Rosenthal \(1983\)](#) show that the only Nash equilibrium entails all voters in one group randomizing between voting and abstaining. By Theorem 1, this mixed-strategy Nash equilibrium is not an extreme point, meaning there exists a correlated equilibrium that strictly improves outcomes for any non-degenerate objective.

Similar conclusions arise in Bertrand price competition, where identical firms must randomize prices over a certain range to prevent competitors from systematically undercutting them ([Vives, 1999](#)). If firms priced predictably, rivals could slightly undercut them and capture the entire market. Likewise, venture capitalists' investment decisions exhibit strategic randomness: multiple investors allocate capital across competing startups, knowing that the probability of a startup's success depends on total funding received. The unique mixed equilibrium involves investors randomizing across opportunities to prevent driving down expected returns ([Hellmann and Thiele, 2015](#)). If an investor always

⁸[Kreps \(1981\)](#) obtains a "dual" result: under the same condition on the support of ν , there exist payoffs u_i , $i = 1, \dots, n$, making ν a unique Nash equilibrium.

funded the same type of startup, competitors could adjust their strategies, maximizing returns elsewhere. Similar dynamics appear in multi-agent Blotto games, where agents allocate resources across battlefields while anticipating their opponents' allocations (Rober-son, 2006), as well as in online content moderation, where platforms adjust policies to maintain engagement while mitigating backlash (Acemoglu, Makhdoumi, Malekian, and Ozdaglar, 2022). In each case, Nash equilibria involve mixing, but our results show that strategic correlation can yield strictly better outcomes. Certainly, the line between coordination and collusion is sometimes thin. In that respect, our results provide insights for regulators into the environments that are more susceptible to collusion.

4 Extremality in Payoff Space

The preceding discussion characterizes extreme equilibria in terms of their structure within the space of distributions over actions. However, economic analysis often focuses on the payoffs associated with equilibria, rather than the distributions themselves. We now show that utilitarian welfare of an equilibrium—the sum of agents' payoffs—can generically be improved whenever there is substantial mixing.

For each $\mu \in \Delta(A)$, assign the expected payoff profile $u(\mu) \in \mathbb{R}^n$ given by $u_i(\mu) = \sum_a u_i(a)\mu(a)$. For any game Γ , let $U^{CE}(\Gamma) \subset \mathbb{R}^n$ be the set of CE payoff profiles in Γ . We say that a Nash equilibrium ν with payoff vector $u(\nu)$ is **payoff-extreme** if its associated expected payoff vector is an extreme point of the set $U^{CE}(\Gamma)$ of CE-induced payoff profiles. The set of extreme points of $U^{CE}(\Gamma)$ has some appeal: it contains Pareto efficient and utilitarian efficient payoff profiles that can be attained via a correlated equilibrium.

The set $U^{CE}(\Gamma)$ can be viewed as a linear projection of the convex polytope of correlated equilibria onto a lower-dimensional space via the mapping $\mu \rightarrow u(\mu)$. Just as the sharpest edges of a shadow must originate from the sharpest edges of the object casting it, it is well known that the extreme points of a projection of a convex set are contained within the projection of the set's extreme points. In particular, the extreme payoffs in $U^{CE}(\Gamma)$ derive from the extreme points of the correlated equilibrium polytope.

Corollary 2. *In a generic game, a Nash equilibrium with more than two agents randomizing is not payoff-extreme.*

Formally, to deduce this corollary from Theorem 1, one needs to make sure that each extreme point of $U^{CE}(\Gamma)$ originates as a projection of a unique extreme point of $CE(\Gamma)$. We verify this in Appendix 1.3 by showing that, in a generic game, all extreme correlated equilibria result in different payoff vectors. Consequently, one cannot replace the genericity of a game with the regularity of a Nash equilibrium in Corollary 2.

The converse of the corollary does not hold: the projection of an extreme correlated equilibrium might not be an extreme point in the payoff space. Consequently, pure Nash equilibria and Nash equilibria with exactly two agents randomizing are not necessarily payoff-extreme. For example, the mixed Nash equilibrium in [Aumann \(1974\)](#)'s game (2) described at the end of Section 2 is in the interior of U^{CE} .

Payoff-extremality has important implications for welfare analysis. Applying Bauer's Maximum Principle in the space of payoffs, we conclude that if a Nash equilibrium is not payoff-extreme, there exists a correlated equilibrium that yields a higher value for any non-degenerate linear objective in the space of payoffs.

As it turns out, utilitarian welfare, and similarly any weighted utilitarian welfare, are generically non-degenerate. We, therefore, have the following result, echoing Theorem 1.

Proposition 1. *In a generic game, the utilitarian welfare of any Nash equilibrium with more than two agents randomizing can be improved within the set of correlated equilibria.*

The proof shows formally that utilitarian welfare is non-degenerate in generic games. Intuitively, the set of games in which utilitarian welfare is constant on an edge of U^{CE} is non-generic: those edges would need to be perpendicular to the vector of equal weights.


5 Symmetric Games

Many games considered in the literature exhibit symmetry: an agent's payoffs depend only on her actions and the actions others choose, but not on her or others' identity. Put differently, permuting the labels of agents and, accordingly, the actions they take, makes no difference to payoffs; see, for example, [Dasgupta and Maskin \(1986\)](#). We now consider a generalized notion of symmetric games, discuss the prevalence of symmetric equilibria that can be improved via analogously symmetric correlated equilibria, and explore the structure of these beneficial correlated equilibria.

5.1 Generalized Symmetric Games

The standard definition of a symmetric game requires that any permutation of agents leaves the game unchanged. This definition of symmetry rules out a variety of games that appear symmetric. For example, consider a voting game between agents 1 through 9 placed on a network as follows:⁹

⁹This example resembles an example motivating the study of [Bartholdi et al. \(2021\)](#) on equitable voting rules.



Each group of three agents— $\{1, 2, 3\}$, $\{4, 5, 6\}$, and $\{7, 8, 9\}$ —forms a “district,” denoted by D_1 , D_2 , and D_3 , respectively. Every agent selects between two possible actions, a or b . First, each district’s choice is determined by majority rule: if at least two agents within a district vote for a , the district selects a ; otherwise, it selects b . Then, the overall decision of the electorate (E) is determined by a second majority rule applied to the districts’ choices: if at least two districts choose a , the electorate selects a ; otherwise, it selects b . Ultimately, each agent’s utility depends solely on the final electorate’s decision.

This game is symmetric in the sense that no agent is treated differently: each agent has the same “role” as any other. Nonetheless, it is not symmetric in the usual sense. For example, suppose $\{1, \dots, 5\}$ vote for a and $\{6, \dots, 9\}$ vote for b (differentiated by shading in the figure), then the ultimate electorate choice is a . However, if agents 5 and 9 were swapped, the electorate’s choice would change to b . We now generalize the notion of symmetry.¹⁰


Given a set P of permutations, we say that a game Γ and a distribution μ are symmetric with respect to P , or P -**symmetric** for short, if they are invariant under all permutations from $\pi \in P$. That is, for any $\pi \in P$, $u_i(a) = u_{\pi(i)}(a_{\pi(1)}, \dots, a_{\pi(n)})$ and $\mu(a) = \mu(a_{\pi(1)}, \dots, a_{\pi(n)})$ for all $a \in A$. Analogously, we say that a welfare function W is P -symmetric if it similarly invariant under all permutations in P . That is, for any $\pi \in P$, $W(u_1, \dots, u_n) = W(u_{\pi(1)}, \dots, u_{\pi(n)})$. Without loss of generality, we focus on permutation sets P that are subgroups of the group of all permutations S_n .¹¹

When $P = S_n$, P -symmetry of a game boils down to the common assumption of game symmetry. Some interesting cases arise when $P \neq S_n$ but is transitive, i.e., for any two agents i and j , there is a permutation $\pi \in P$ such that $\pi(i) = j$. For example, suppose agents are located on a circle: agent i placed between agents $i - 1$ and $i + 1 \pmod n$ for all $2 \leq i \leq n$, as described in the figure below for $n = 8$. Suppose agents interact with their neighbors so that $u_i(s) = f(s_{i-1}, s_i, s_{i+1})$ (with $s_0 = s_n$ and $s_{n+1} = s_1$). Then, the game is invariant under the set P of all cyclic permutations $\pi: i \rightarrow i + k \pmod n$. Similarly, the example above pertaining to the tree network is also symmetric for a transitive P . The importance of the distinction between fully symmetric and P -symmetric games has recently

¹⁰In his seminal treatise on non-cooperative games, Nash (1950) suggests a more general definition of symmetry in the spirit of the one we offer.

¹¹That is, for any permutation in P , its inverse is also in P , and the composition of any two permutations in P is also in P .

been highlighted by [Bartholdi, Hann-Caruthers, Josyula, Tamuz, and Yariv \(2021\)](#).



Let P be a subgroup of S_n and consider a P -symmetric n -agent game Γ with $n \geq 3$. Let ν be a regular symmetric mixed Nash equilibrium. Under the conditions of Theorem 1, ν cannot be an extreme point of the set of correlated equilibria and, thus, improvable. What if P -symmetry is desirable in and of itself? We handle this question next.

5.2 Improvement via Symmetric Correlated Equilibria

A designer may want to treat agents symmetrically and favor equity over any other objective. Furthermore, if the objective is P -symmetric—e.g., the social (utilitarian) welfare—its optimum can be achieved via P -symmetric correlated equilibria. For a Nash equilibrium, symmetry with respect to a transitive P is equivalent to symmetry with respect to $P = S_n$. We now study when a symmetric Nash equilibrium can be improved within the set of correlated equilibria respecting the symmetry of the game.

Theorem 2. *For P -symmetric n -agent game with transitive P , a regular symmetric Nash equilibrium is an extreme point of the set of P -symmetric correlated equilibria if and only if $n = 2$ or it is pure.*

Theorem 2 implies that any mixing in a regular symmetric equilibrium can be improved, even when maintaining equity considerations. For example, firms participating in Bertrand price competition, or venture capital investment firms are often similar to one another and the equilibria discussed at the end of Section 3 are therefore symmetric and improvable, even when insisting on symmetric play. Likewise, symmetric equilibria in war of attrition games or public good games often involve mixing and are therefore improvable.

More formally, Bauer’s principle implies that the value of a generic P -symmetric linear objective can be strictly improved. We note that the restriction to P -symmetric correlated equilibria in Theorem 2 is important. If the condition of the theorem pertained to general correlated equilibria, we would be able to improve a generic linear objective but not

necessarily those respecting the symmetry of the game.¹²

To gain some intuition for the proof, consider the case of $P = S_n$. First, regularity allows us to only consider actions that feature strictly positive weight in agents' strategies.¹³

Second, notice that any symmetric correlated equilibrium corresponds to an exchangeable distribution. If each agent places positive weight on m pure actions, each of the m actions has to generate at least as high a payoff as the remaining $m - 1$ actions, yielding $m(m - 1)$ incentive constraints. Relying, again, on [Winkler \(1988\)](#), we show that an extreme point of the set of symmetric correlated equilibria can be represented as a convex combination of at most

$$m(m - 1) + 1$$

extreme exchangeable distributions, where m is the number of actions of each agent.

Last, we show that any symmetric Nash equilibrium corresponds to a mixture of at least $\binom{n+m-1}{m-1}$ extreme exchangeable distributions. With at least three agents, for any $m > 1$, we have that

$$\binom{n+m-1}{m-1} > m(m+1) + 1$$

and the result follows. To handle P -symmetry, the proof generalizes the notion of exchangeability to P -exchangeability of a finite set of random variables, where the joint distribution is invariant to variable permutations from P .

5.3 The Structure of Beneficial Correlated Equilibria

We now focus on $P = S_n$ and show that extreme symmetric correlated equilibria take a very particular form, regardless of the underlying game's details.

Proposition 2. *Any extreme correlated equilibrium of a fully symmetric game Γ with n agents and m actions a_1, \dots, a_m can be obtained from a fixed collection of $M \leq m(m - 1) + 1$ urns each with n balls marked with actions, as follows:*

- *An urn is chosen at random according to some distribution $p \in \Delta_M$. Agents do not know which urn has been chosen.*
- *Agents $\{1, \dots, n\}$ approach the urn sequentially and draw balls without replacement.*
- *Each agent takes an action corresponding to her ball's label.*

¹²Similarly to generic games, by a generic linear objective we mean the one from an open everywhere dense set with a zero-measure complement.

¹³For this to hold without loss of generality, we only need that incentive constraints outside of the support of agents' strategies are not active.

Intuitively, as already argued, an extreme point of the set of symmetric correlated equilibria can be represented as a convex combination of at most

$$m(m-1)+1$$

extreme exchangeable distributions. Each extreme exchangeable distribution is then represented as an urn. Sampling without replacement from an urn allows us to generate an arbitrary symmetric distribution over action profiles.

The proposition provides a simple parametrization of correlated equilibria, particularly convenient when the number of actions is small. For example, if $m = 2$, we only need $M = 3$ urns. The composition of each urn is determined by a single number: the fraction of balls marked a_1 in the urn. A distribution p over three urns adds 2 more parameters. As a result, all extreme symmetric correlated equilibria form a 5-parametric family. More generally, we need $(m^2 + 1)(m - 1)$ parameters. Importantly, this number does not depend on the number of agents n , while describing an arbitrary distribution over actions requires $m^n - 1$ parameters—an exponentially growing quantity in n .

Proposition 2 implies further simplification when the number of agents is large relative to the number of available actions. In that case, the joint distributions corresponding to draws with and without replacement become close to one another.¹⁴

Formally, a distribution of $\xi_1, \xi_2, \dots, \xi_n$ is ε -close to a mixture of i.i.d. distributions if, for any $k = 1, \dots, n$, the total variation distance between the joint distribution of any k random variables $\xi_{i_1}, \dots, \xi_{i_k}$ and the closest mixture of i.i.d. distributions is at most $\varepsilon \cdot k$. The approximation results of Diaconis and Freedman (1980) imply the following corollary.

Corollary 3. *Any fully symmetric correlated equilibrium μ of a fully symmetric game Γ with n agents, each having m strategies, is ε -close to a mixture of i.i.d. distribution with*

$$\varepsilon = \frac{2m}{n}.$$

Thus, in symmetric games where the number of agents is large relative to the number of actions, the set of correlated equilibria can be approximated via the following simple procedure. Pick a set of states Θ , a prior distribution $p \in \Delta(\Theta)$, a set of signals S , and a family of distributions $\nu_\theta \in \Delta(S)$ for each $\theta \in \Theta$. The correlated profile is then generated as follows. The state θ is realized according to the distribution p , and independent and identically distributed signals $s_i \sim \nu_\theta$ are then privately released to each agent i . Agents, knowing the distributions over states and signals, best reply to their observed signals. That is, symmetric correlated equilibria are close to Bayesian Nash equilibria of a game obtained

¹⁴This phenomenon underlies the famous De Finetti theorem, which states that extreme exchangeable distributions over an infinite product space $T \times T \times \dots$ are distributions of i.i.d. T -valued random variables (Diaconis and Freedman, 1980).

after adding a payoff-irrelevant random state variable θ with conditionally independent and identically distributed private signals.

6 Implications for Incomplete Information Games and Coarse Correlated Equilibria

This section illustrates implications of our techniques to two other settings. First, we consider incomplete information games and extremality within the set of Bayesian correlated equilibria. Then, we look at a relaxation of correlated equilibria, namely coarse correlated equilibria, which result from various learning dynamics.

6.1 Bayesian Correlated Equilibria

Here, we outline the implications of our analysis for strategic environments with payoff uncertainty. Recall the general framework of [Bergemann and Morris \(2019\)](#) allowing for common payoff uncertainty and private information. The common uncertainty is represented by a state variable θ from a set of states Θ . Each agent i has private information represented by a type t_i from a finite set T_i . We denote the set of type profiles by $T = \times_{i \in N} T_i$ and the joint distribution over states and types by $\pi \in \Delta(\Theta \times T)$. Each agent i 's payoff is determined by the action profile $a \in A$, the state θ , and the type profile t . Consequently, agent i 's utility function is defined as $v_i : A \times \Theta \times T \rightarrow \mathbb{R}$. This constitutes a game of incomplete information, which we denote by $\mathcal{G} = (N, (A_i)_{i \in N}, \Theta, (T_i)_{i \in N}, \pi, (v_i)_{i \in N})$.

A joint distribution $\psi \in \Delta(A \times \Theta \times T)$ is a **Bayesian correlated equilibrium** (BCE) if its marginal on $\Theta \times T$ coincides with π and the following obedience conditions are satisfied

$$\sum_{a_{-i} \in A_{-i}, \theta \in \Theta, t_{-i} \in T_{-i}} \psi(a, \theta, t) v_i(a, \theta, t) \geq \sum_{a_{-i} \in A_{-i}, \theta \in \Theta, t_{-i} \in T_{-i}} \psi(a, \theta, t) v_i(a'_i, a_{-i}, \theta, t) \quad (5)$$

for each $i \in N$, $t_i \in T_i$, $a_i \in A_i$, and $a'_i \in A_i$. This condition ensures that no agent of any type can gain by unilaterally deviating from a recommended action a_i to an alternative action a'_i , given the posterior belief induced by the recommendation and their private type. A BCE ψ is a **Bayesian Nash equilibrium** if θ, a_1, \dots, a_n are all independent conditional on the profile realized types $t = (t_1, \dots, t_n)$. A Bayesian Nash equilibrium can be identified with a profile of functions $\sigma_i : T_i \rightarrow \Delta(A_i)$ according to which each agent i randomizes her action given her realized type. Such an equilibrium is pure if each σ_i always outputs a pure strategy.

Proposition 3. *For a generic game with incomplete information, such that at least one of the following conditions holds:*

- *non-trivial common payoff uncertainty*: $|\Theta| \geq 2$, or
- *non-trivial private information*: $|T_i| \geq 2$ for at least 3 agents,

a Bayesian Nash equilibrium is an extreme point of BCE if and only if it is pure.

The requirement that the game is generic implies, in particular, that π has full support: individual types provide noisy information about the environment. The proof appears in Appendix 1.5 and follows the same lines as that of Theorem 1. It builds on an incomplete information extension of Lemmas 1 and 2. The proof uses the genericity of π to guarantee full support. Consequently, instead of assuming that the game is generic, one can assume that a full-support π is fixed and utilities u are generic.

6.2 Coarse Correlated Equilibria

We now consider the larger set of coarse correlated equilibria and show that any randomness in a Nash equilibrium renders it improvable within the set of coarse correlated equilibria.

Correlated equilibria rely on ex-post incentive constraints, meaning that each agent finds it optimal to follow the recommended action conditional on receiving that recommendation. The notion of coarse correlated equilibria, introduced by [Moulin and Vial \(1978\)](#), relaxes this requirement by imposing only ex-ante incentive constraints. Coarse correlated equilibria capture strategic settings where agents decide whether to commit to a correlating device *before* receiving any specific recommendation. This is akin to an agent deciding whether to opt-in to a recommendation system, such as those employed by platforms like Facebook or Google for ad targeting, where the alternative is to turn off personalized recommendations. Alternatively, it speaks to agents' incentives to join collusive agreements (for example, see [McAfee and McMillan, 1992](#)). Coarse correlated equilibria also represent the set of outcomes achievable by arbitrary no-regret learning dynamics ([Hart and Mas-Colell, 2001](#)).

Formally, a distribution $\mu \in \Delta(A)$ is a **coarse correlated equilibrium** (CCE) in a game $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$ if

$$\sum_{a \in A} u_i(a) \mu(a) \geq \max_{a'_i \in A_i} \sum_{a \in A} u_i(a'_i, a_{-i}) \mu(a) \quad (6)$$

for all $i \in N$. Each agent faces a single individual rationality constraint: the expected payoff from the correlated strategy must be at least as high as the best deterministic deviation.

We now investigate the extremality of Nash equilibria within the larger set of coarse correlated equilibria. Naturally, extremality is then more demanding. As an analog of Theorem 1, we establish the following result.

Proposition 4. *In a generic game, a Nash equilibrium is an extreme point of the set of coarse correlated equilibria if and only if it is pure.*

As in Theorem 1, the genericity assumption can be replaced by the assumption of regularity of the Nash equilibrium under consideration. The proof of Proposition 4 is contained in Appendix 1.6. It mirrors that of Theorem 1.

Proposition 4 demonstrates a stark contrast between correlated and coarse correlated equilibria. While mixed Nash equilibria can be extreme points of $CE(\Gamma)$ when two agents randomize, they can never be extreme points of $CCE(\Gamma)$ in generic games. This result implies that in a generic game, any mixed regular Nash equilibrium can be improved, for any non-degenerate objective, by moving to some coarse correlated equilibrium. This is consistent with the prevalence of examples in the literature, primarily focused on 2-agent games, where Nash equilibria are improvable via coarse correlated equilibria but not correlated equilibria (e.g, Moulin and Vial, 1978; Gérard-Varet and Moulin, 1978; Moulin et al., 2014; Dokka et al., 2023).

References

- ACEMOGLU, D., A. MAKHDOUNI, A. MALEKIAN, AND A. OZDAGLAR (2022): “Too much data: Prices and inefficiencies in data markets,” *American Economic Journal: Microeconomics*, 14, 218–256.
- AGRANOV, M. AND C. TERGIMAN (2014): “Communication in multilateral bargaining,” *Journal of Public Economics*, 118, 75–85.
- AGRANOV, M. AND L. YARIV (2018): “Collusion through communication in auctions,” *Games and Economic Behavior*, 107, 93–108.
- AHUNBAY, M. Ş. AND M. BICHLER (2024): “On the Uniqueness of Bayesian Coarse Correlated Equilibria in Standard First-Price and All-Pay Auctions,” *arXiv preprint arXiv:2401.01185*.
- ARIELI, I., Y. BABICHENKO, R. SMORODINSKY, AND T. YAMASHITA (2023): “Optimal persuasion via bi-pooling,” *Theoretical Economics*, 18, 15–36.
- ARRATIA, R., A. D. BARBOUR, AND S. TAVARÉ (2003): *Logarithmic combinatorial structures: a probabilistic approach*, vol. 1, European Mathematical Society.
- ARRATIA, R. AND S. TAVARÉ (1994): “Independent process approximations for random combinatorial structures,” *Advances in mathematics*, 104, 90–154.
- ASHLAGI, I., D. MONDERER, AND M. TENNENHOLTZ (2008): “On the value of correlation,” *Journal of Artificial Intelligence Research*, 33, 575–613.
- AUMANN, R. J. (1974): “Subjectivity and correlation in randomized strategies,” *Journal of mathematical Economics*, 1, 67–96.
- (1987): “Correlated equilibrium as an expression of Bayesian rationality,” *Econometrica: Journal of the Econometric Society*, 1–18.
- BALL, I. (2023): “Bauer’s Maximum Principle for Quasiconvex Functions,” *arXiv preprint arXiv:2305.04893*.
- BARTHOLDI, L., W. HANN-CARUTHERS, M. JOSYULA, O. TAMUZ, AND L. YARIV (2021): “Equitable voting rules,” *Econometrica*, 89, 563–589.
- BERGEMANN, D. AND S. MORRIS (2019): “Information design: A unified perspective,” *Journal of Economic Literature*, 57, 44–95.
- BORDER, K. (2006): *Infinite dimensional analysis: A hitchhiker’s guide*, Springer.
- BRADONJIC, M., G. ERCAL, A. MEYERSON, AND A. ROYTMAN (2014): “The price of mediation,” *Discrete Mathematics & Theoretical Computer Science*, 16.
- CALVÓ-ARMENGOL, A. (2006): “The set of correlated equilibria of 2x2 games,” *mimeo*.
- CANOVAS, S. G., P. HANSEN, AND B. JAUMARD (1999): “Nash equilibria from the correlated equilibria viewpoint,” *International Game Theory Review*, 1, 33–44.
- CRIPPS, M. (1995): “Extreme correlated and Nash equilibria in two-person games,” .

- DASGUPTA, P. AND E. MASKIN (1986): “The existence of equilibrium in discontinuous economic games, I: Theory,” *The Review of economic studies*, 53, 1–26.
- DIACONIS, P. AND D. FREEDMAN (1980): “Finite exchangeable sequences,” *The Annals of Probability*, 745–764.
- DOKKA, T., H. MOULIN, I. RAY, AND S. SENGUPTA (2023): “Equilibrium design in an n-player quadratic game,” *Review of economic design*, 27, 419–438.
- EINY, E., O. HAIMANKO, AND D. LAGZIEL (2022): “Strong robustness to incomplete information and the uniqueness of a correlated equilibrium,” *Economic Theory*, 73, 91–119.
- EVANGELISTA, F. S. AND T. RAGHAVAN (1996): “A note on correlated equilibrium,” *International Journal of Game Theory*, 25, 35–41.
- FELDMAN, M., B. LUCIER, AND N. NISAN (2016): “Correlated and coarse equilibria of single-item auctions,” in *Web and Internet Economics: 12th International Conference, WINE 2016, Montreal, Canada, December 11-14, 2016, Proceedings 12*, Springer, 131–144.
- FORGES, F. (2020): “Correlated equilibria and communication in games,” *Complex Social and Behavioral Systems: Game Theory and Agent-Based Models*, 107–118.
- FOSTER, D. P. AND R. V. VOHRA (1997): “Calibrated learning and correlated equilibrium,” *Games and Economic Behavior*, 21, 40–55.
- FRIEDMAN, D., J. P. RABANAL, O. A. RUD, AND S. ZHAO (2022): “On the empirical relevance of correlated equilibrium,” *Journal of Economic Theory*, 205, 105531.
- FUDENBERG, D. AND D. K. LEVINE (1999): “Conditional universal consistency,” *Games and Economic Behavior*, 29, 104–130.
- GEORGALOS, K., I. RAY, AND S. SENGUPTA (2020): “Nash versus coarse correlation,” *Experimental Economics*, 23, 1178–1204.
- GÉRARD-VARET, L.-A. AND H. MOULIN (1978): “Correlation and duopoly,” *Journal of economic theory*, 19, 123–149.
- GERARDI, D. AND L. YARIV (2007): “Deliberative voting,” *Journal of Economic theory*, 134, 317–338.
- GOEREE, J. K. AND L. YARIV (2011): “An experimental study of collective deliberation,” *Econometrica*, 79, 893–921.
- GOVINDAN, S., P. J. RENY, AND A. J. ROBSON (2003): “A short proof of Harsanyi’s purification theorem,” *Games and Economic Behavior*, 45, 369–374.
- HARSANYI, J. C. (1973a): “Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points,” *International journal of game theory*, 2, 1–23.
- (1973b): “Oddness of the number of equilibrium points: a new proof,” *International Journal of Game Theory*, 2, 235–250.
- HART, S. AND A. MAS-COLELL (2000): “A simple adaptive procedure leading to correlated equilibrium,” *Econometrica*, 68, 1127–1150.

- (2001): “A general class of adaptive strategies,” *Journal of Economic Theory*, 98, 26–54.
- HELLMANN, T. AND V. THIELE (2015): “Contracting among founders,” *The Journal of Law, Economics, and Organization*, 31, 629–661.
- HENDRICKX, R., R. PEETERS, AND J. POTTERS (2002): “A relation between Nash equilibria and correlated equilibria,” *International game theory review*, 4, 405–413.
- JANN, O. AND C. SCHOTTMÜLLER (2015): “Correlated equilibria in homogeneous good Bertrand competition,” *Journal of Mathematical Economics*, 57, 31–37.
- KLEINER, A., B. MOLDOVANU, AND P. STRACK (2021): “Extreme points and majorization: Economic applications,” *Econometrica*, 89, 1557–1593.
- KLEINER, A., B. MOLDOVANU, P. STRACK, AND M. WHITMEYER (2024): “The Extreme Points of Fusions,” *arXiv preprint arXiv:2409.10779*.
- KOLUMBUS, Y. AND N. NISAN (2022): “Auctions between regret-minimizing agents,” in *Proceedings of the ACM Web Conference 2022*, 100–111.
- KREPS, V. (1981): “Finite N-person non-cooperative games with unique equilibrium points,” *International Journal of Game Theory*, 10, 125–129.
- LAHR, P. AND A. NIEMEYER (2024): “Extreme Points in Multi-Dimensional Screening,” *arXiv preprint arXiv:2412.00649*.
- LOPOMO, G., L. M. MARX, AND P. SUN (2011): “Bidder collusion at first-price auctions,” *Review of Economic Design*, 15, 177–211.
- MANELLI, A. M. AND D. R. VINCENT (2007): “Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly,” *Journal of Economic theory*, 137, 153–185.
- McAFEE, R. P. AND J. McMILLAN (1992): “Bidding rings,” *The American Economic Review*, 579–599.
- McKELVEY, R. D. AND A. McLENNAN (1997): “The maximal number of regular totally mixed Nash equilibria,” *Journal of Economic Theory*, 72, 411–425.
- MOULIN, H., I. RAY, AND S. S. GUPTA (2014): “Coarse correlated equilibria in an abatement game,” Tech. rep., Cardiff Economics Working Papers.
- MOULIN, H. AND J. P. VIAL (1978): “Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon,” *International Journal of Game Theory*, 7, 201–221.
- NASH, J. F. (1950): “Non-cooperative games,” .
- NEYMAN, A. (1997): “Correlated equilibrium and potential games,” *International Journal of Game Theory*, 26, 223–227.
- NIKZAD, A. (2022): “Constrained majorization: Applications in mechanism design,” in *Proceedings of the 23rd ACM Conference on Economics and Computation*, 330–331.

- PALFREY, T. R. AND H. ROSENTHAL (1983): "A Strategic Calculus of Voting," *Public Choice*, 41, 7–53.
- PAPADIMITRIOU, C. H. AND T. ROUGHGARDEN (2008): "Computing correlated equilibria in multi-player games," *Journal of the ACM (JACM)*, 55, 1–29.
- PAVLOV, G. (2023): "Correlated equilibria and communication equilibria in all-pay auctions," *Review of Economic Design*, 1–33.
- PEETERS, R. AND J. POTTERS (1999): "On the structure of the set of correlated equilibria in two-by-two bimatrix games," .
- ROBERSON, B. (2006): "The colonel blotto game," *Economic Theory*, 29, 1–24.
- STAM, A. J. (1978): "Distance between sampling with and without replacement," *Statistica Neerlandica*, 32, 81–91.
- UI, T. (2008): "Correlated equilibrium and concave games," *International Journal of Game Theory*, 37, 1–13.
- VAN DAMME, E. (1991): *Stability and perfection of Nash equilibria*, vol. 339, Springer.
- VIOSAT, Y. (2008): "Is having a unique equilibrium robust?" *Journal of Mathematical Economics*, 44, 1152–1160.
- VIVES, X. (1999): *Oligopoly Pricing: Old Ideas and New Tools*, MIT Press.
- WILSON, R. (1971): "Computing equilibria of n-person games," *SIAM Journal on Applied Mathematics*, 21, 80–87.
- WINKLER, G. (1988): "Extreme points of moment sets," *Mathematics of Operations Research*, 13, 581–587.
- WU, J. (2008): "Correlated equilibrium of bertrand competition," in *Internet and Network Economics: 4th International Workshop, WINE 2008, Shanghai, China, December 17-20, 2008. Proceedings 4*, Springer, 166–177.
- YANG, K. H. AND A. K. ZENTEFIS (2024): "Monotone Function Intervals: Theory and Applications," *American Economic Review*, 114, 2239–2270.

1 Appendix: Proofs

1.1 Proof of Theorem 1

We prove the following theorem, which generalizes Theorem 1.

Theorem 3. *Consider a regular Nash equilibrium with k agents randomizing their actions. This equilibrium is an extreme point of the correlated equilibrium polytope if and only if $k \leq 2$. For $k \geq 3$, the equilibrium lies in the interior of a face of dimension higher or equal to*

$$2^{k-3} - 1. \tag{7}$$

The dimension of the face containing a Nash equilibrium represents the number of linearly independent perturbations that preserve the incentive constraints—essentially, the directions in which potential improvements can occur.

The proof also contains a characterization of equilibria for small dimensions. For example, a Nash equilibrium is in the interior of a face of dimension 1 if either exactly 3 agents randomize over 2 actions each, or two agents randomize over two actions and one over three. Interestingly, no Nash equilibrium can be in the interior of a face of dimension 2. Faces of dimension 3 only contain equilibria where two agents randomize over 3 actions, and one randomizes over 2, or those where one agent randomizes over 4 actions, one over 3, and one over 2.

Our first step is proving the following extended version of Lemma 1, which bounds the support of a correlated equilibrium belonging to the interior of a face of dimension $d \geq 0$. Lemma 1 focusing on extreme correlated equilibria corresponds to $d = 0$.

Lemma 3. *For any game Γ , if a distribution $\mu \in \Delta(S)$ belongs to the interior of a face of dimension d of the set of correlated equilibria, then*

$$|\text{supp } \mu| \leq 1 + d + \sum_{i \in N} |S_i| \cdot (|S_i| - 1). \tag{8}$$

Proof of Lemma 1. The set of correlated equilibria is cut from the simplex of all probability measures $\Delta(S)$ by incentive constraints taking the form of linear inequalities (6). By Theorem 2.1 of Winkler (1988), if μ is an extreme point of a set of probability measures under linear constraints, then $|\text{supp } \mu|$ cannot exceed $k + 1$, where k is the number of constraints active at μ .

First, consider the case of $d = 0$, i.e., extreme μ . For each agent i , there are $|A_i| \cdot (|A_i| - 1)$ incentive constraints since for each recommended action $a_i \in A_i$, we need to rule out each deviation $a'_i \in A_i \setminus \{a_i\}$. Applying Winkler's theorem, we obtain (8) with $d = 0$.

Now, suppose $d > 1$. We impose additional d independent linear constraints on the set of correlated equilibria so that μ becomes an extreme point of the constrained set.

Winkler's theorem implies (8). ■

Proof of Theorem 1. We first show that if ν is a regular Nash equilibrium, where at most two agents randomize their actions, then ν is an extreme point of $\text{CE}(\Gamma)$.

First, consider pure ν . Any point mass is the extreme point of the set of probability measures. Thus ν is an extreme point of $\Delta(S)$ and thus of $\text{CE}(\Gamma) \subset \Delta(S)$.

Now, suppose that agents $i, j \in N$ randomize and the rest of agents $k \in N \setminus \{i, j\}$ play pure actions. Towards a contradiction, assume that ν is represented as a convex combination of correlated equilibria $\lambda\mu + (1 - \lambda)\mu'$ with $\lambda \in (0, 1)$ and $\mu \neq \mu'$. Let $\nu_{i,j} = \nu_i \times \nu_j$, $\mu_{i,j}$, and $\mu'_{i,j}$ be the marginals of ν , μ , and μ' on $S_i \times S_j$. We get $\nu_{i,j} = \mu_{\{i,j\}} + (1 - \lambda)\mu'_{\{i,j\}}$, where $\mu_{\{i,j\}} \neq \mu'_{\{i,j\}}$ since otherwise μ would be equal to μ' as all the agents except for i and j play pure actions. Consider the 2-agent game $\Gamma_{\{i,j\}}$ obtained from Γ by replacing all the agents $k \in N \setminus \{i, j\}$ with dummies taking actions played in ν . The distribution $\nu_{i,j}$ is a regular Nash equilibrium of $\Gamma_{\{i,j\}}$.

By the results of [Cripps \(1995\)](#), [Evangelista and Raghavan \(1996\)](#), and [Canovas et al. \(1999\)](#), any regular Nash equilibrium of a 2-agent game is an extreme point of its set of correlated equilibria. Therefore, $\nu_{i,j}$ is an extreme point of $\text{CE}(\Gamma_{\{i,j\}})$. However, we represented $\nu_{i,j}$ as a non-trivial convex combination of correlated equilibria $\mu_{\{i,j\}}$ and $\mu'_{\{i,j\}}$. This contradiction implies that a regular Nash equilibrium with two agents mixing their actions is an extreme point of $\text{CE}(\Gamma)$.

In the opposite direction, we need to prove that if a regular Nash equilibrium $\nu = \nu_1 \times \dots \times \nu_n$ of an n -agent game Γ is an extreme point of $\text{CE}(\Gamma)$, then ν at most two agents randomize. It is enough to prove this claim assuming that $\text{supp } \nu_i = S_i$ for any i since the incentive constraints are strict outside of the support, by definition of a regular Nash equilibrium; see [Harsanyi \(1973b\)](#).

Denote $m_i = |S_i| = |\text{supp } \nu_i|$ and assume without loss of generality that agents are enumerated so that $m_i \geq m_{i+1}$. Lemma 8 and 2 imply that numbers m_1, \dots, m_n satisfy the following conditions

$$\begin{cases} m_1 \cdot \dots \cdot m_n & \leq 1 + \sum_{j=1}^n m_j \cdot (m_j - 1) \\ m_i - 1 & \leq \sum_{j \neq i} (m_j - 1), \quad i = 1, \dots, n \end{cases} \quad (9)$$

By Proposition 5 formulated and proved below, these inequalities imply $m_1 = m_2$ and $m_i = 1$ for $i \geq 3$. Denote the common value of m_1 and m_2 by m . We conclude that ν is pure for $m = 1$ and two agents randomize for $m > 1$.

Finally, suppose that k agents randomize at ν and let d be the dimension of a face to the interior of which ν belongs. Now, Lemmas 8 and 2 result in the following version of

system (9)

$$\begin{cases} m_1 \cdot \dots \cdot m_n & \leq 1 + d + \sum_{j=1}^n m_j \cdot (m_j - 1) \\ m_i - 1 & \leq \sum_{j \neq i} (m_j - 1), \quad i = 1, \dots, n \end{cases} \quad (10)$$

By Proposition 5, inequalities (10) can only be satisfied if $k \leq 3 + \log_2(1 + d)$ which implies (7).

For small d , formulas (14) from Proposition 5 provide the structure of solutions. For example, the only sequences solving this system for $d = 1$ but not for $d = 0$ are of the form $(2, 2, 2, 1, \dots, 1)$ (three agents all randomizing over 2 actions) and $(3, 2, 2, 1, \dots, 1)$ (one agent randomize over 3 actions and two agents randomize over two). For $d = 2$, there are no Nash equilibria as there are no solutions with $d = 2$ that would not also satisfy the system with smaller d . For $d = 3$, we get two solutions: $(3, 3, 2, 1, \dots, 1)$ (two agents randomize over 3 actions and one randomizes over 2) and $(4, 3, 2, 1, \dots, 1)$ (one agent randomize over 4 actions, one over 3, and one over 2). ■

1.2 Analysis of solutions to system (9)

In the following proposition, we consider a class of systems that contains (9) and, additionally, allows for the slack of size of Q in the right-hand side of the first inequality

$$\prod_{i=1}^n m_i \leq \sum_{i=1}^n m_i \cdot (m_i - 1) + Q \quad \text{and} \quad (11)$$

$$m_1 - 1 \leq \sum_{i=2}^n (m_i - 1) \quad (12)$$

It will also be convenient to drop all $m_i = 1$ and only consider the part of the sequence m_1, m_2, \dots that corresponds to $m_i \geq 2$.

Proposition 5. *Let (m_1, \dots, m_n) be a non-increasing sequence of $n \geq 2$ integers such that each $m_i \geq 2$ and inequalities (11) and (12) are satisfied with some integer $Q \geq 1$. Then, the length of the sequence satisfies the following bound:*

$$n \leq 3 + \log_2 Q. \quad (13)$$

If $1 \leq Q \leq 9$, the sequence must take one of the following forms:

$$\begin{aligned}
Q \geq 1 & : (m, m) \text{ for some } m \geq 1 \\
Q \geq 2 & : (2, 2, 2), (3, 2, 2) \\
Q \geq 4 & : (3, 3, 2), (4, 3, 2) \\
Q \geq 6 & : (4, 4, 2), (5, 4, 2) \\
Q \geq 8 & : (5, 5, 2), (6, 5, 2)
\end{aligned} \tag{14}$$

We will need a lemma showing that any solution (m_1, \dots, m_n) to (11) and (12) can be used to define a new solution of a particularly simple form.

Lemma 4. *Let (m_1, \dots, m_n) be a non-increasing sequence of integers numbers $m_i \geq 2$ satisfying (11) and (12). Let $m = m_1$ and $S = \sum_{i=1}^n (m_i - 1)$. Represent $S = k \cdot (m - 1) + (r - 1)$ with $r \in \{1, 2, \dots, m_1 - 1\}$. Then*

$$\underbrace{(m, \dots, m, r)}_{k \text{ times}} \tag{15}$$

also satisfies (11) and (12).

While we keep the proof an elementary application of convexity, it admits a majorization theory interpretation. In essence, the result is deduced from the monotonicity of $\sum_{i=1}^n x_i \cdot (x_i - 1) - \prod_{i=1}^n x_i$ in the majorization order, a property known as Schur convexity.

Proof of Lemma 4. For the new sequence, both sides of (12) are identical to the original one. Thus (12) holds for the new sequence. To check (11), we relax the integrality constraint and denote by $L_{m,S}$ the set of all sequences $(x_1, x_2, \dots, x_n) \in \mathbb{R}^n$ with $x_1 = m$, $1 \leq x_i \leq m$ for $i \geq 2$, and $\sum_{i=1}^n (x_i - 1) = S$. Old and new sequence belongs to $L_{m,S}$. By the quasi-convexity of $\sum_{i=1}^n x_i \cdot (x_i - 1) - \prod_{i=1}^n x_i$ on $L_{m,S}$, the maximum of this expression over $L_{m,S}$ is attained at an extreme point. All such extreme points coincide with $(m, \dots, m, r, 1, \dots, 1)$ up to a permutation of x_2, \dots, x_n . Thus, the difference between the right-hand side and the left-hand side in (11) can only increase when the old sequence is replaced with the new one. We conclude that both (11) and (12) hold for the new sequence. \blacksquare

Proof of Proposition 5. We first prove inequality (13) on the length of a sequence in terms of Q . Consider sequences of the form (m, \dots, m, r) , where m is repeated $k \geq 3$ times. Inequality (11) reads as follows

$$m^k \cdot r \leq k \cdot m(m - 1) + r(r - 1) + Q. \tag{16}$$

Since $m^k \cdot (r - 1) \geq r(r - 1)$, we get

$$m^k \leq k \cdot m(m - 1) + Q$$

and thus

$$m^k \left(1 - \frac{k \cdot (m-1)}{m^{k-1}} \right) \leq Q.$$

The ratio $k \cdot (m-1)/m^{k-1}$ is decreasing in m and in k . Thus, it attains the maximal value of $3/4$ for $k = 3$ and $m = 2$. Hence, $m^k \leq 4Q$. The right-hand side of (16) can be bounded as follows:

$$\begin{aligned} k \cdot m(m-1) + r(r-1) + Q &\leq (k+1) \cdot m(m-1) + Q \\ &\leq \frac{k+1}{k} \cdot m^k + Q \leq \frac{k+1}{k} \cdot 5Q + Q \\ &\leq \frac{4}{3} \cdot 5Q + Q \leq 8Q. \end{aligned}$$

Now consider a sequence (m_1, \dots, m_n) with $m_i \geq 2$ satisfying (11) and (12). Let (m, \dots, m, r) with m repeated $k \geq 3$ times be the associated sequence of the form (15). Since the right-hand side of (11) for (m, \dots, m, r) is at least as big as for (m_1, \dots, m_n) , we get

$$\prod_{i=1}^n m_i \leq 8Q.$$

Since $m_i \geq 2$,

$$n \leq \log_2 8Q = 3 + \log_2 Q.$$

Thus inequality (13) is established for any (m_1, \dots, m_n) corresponding to (m, \dots, m, r) with $k \geq 3$. By (12), $k \geq 2$ and thus it remains to consider the case of $k = 2$. For $k = 2$, we get

$$\prod_{i=1}^n m_i \leq 2 \cdot m(m-1) + r(r-1) + Q.$$

The inequality $n \leq 3 + \log_2 Q$ trivially holds for $n = 3$ and thus we assume $n \geq 4$. By (12), $\sum_{i=2}^n (m_i - 1) \geq m_1 - 1$. Hence, for $n \geq 4$ we get

$$\prod_{i=1}^n m_i \geq m \cdot (m-2) \cdot 2 \cdot 2 = 4m(m-2).$$

Plugging this in, we obtain

$$4m(m-2) \leq 2 \cdot m(m-1) + r(r-1) + Q$$

or

$$m^2 - 3m - 2 \leq Q.$$

Since $m^2 - 3m - 2 = (m - 2)^2 + m - 6$, we get $m \leq 2 + \sqrt{Q + 6}$. Thus,

$$\begin{aligned} \prod_{i=1}^n m_i &\leq 2 \cdot m(m-1) + r(r-1) + Q \leq 2m^2 - 5m + 2 + Q \\ &\leq 2(m^2 - 3m - 2) + m + 6 + Q \leq 2Q + m + 6 + Q \\ &\leq 3Q + 8 + \sqrt{Q + 6} = 8Q + (8 + \sqrt{Q + 6} - 5Q) \end{aligned}$$

For $Q \geq 3$, we get $\prod_{i=1}^n m_i \leq 8Q$ and thus $n \leq 3 + \log_2 Q$. The remaining case of $Q \leq 2$ will be considered in detail below. Thus inequality (13) is proved.

We now characterize the solutions for small values of Q . As before, we start with sequences of the form (m, \dots, m, r) . Note that $m^k(r-1) < r(r-1)$ and so inequality (16) gives

$$m^k < k \cdot m \cdot (m-1) + Q = k \cdot m^2 + (Q - k \cdot m)$$

or

$$m^2(m^{k-2} - k) < Q - km.$$

The left-hand side is non-negative if $k \geq 3$ and $m \geq 3$. For such k and m , for the inequality to hold, we must have $Q - km > 0$ and thus $Q \geq 10$. We conclude that for $Q \leq 9$, i.e., Q from the statement of the proposition, we must have either $k = 2$ or $k = 3$ and $m = 2$. We now analyze these two cases separately.

We first consider the case of $k = 3$ and $m = 2$, i.e., the sequence $(2, 2, 2)$. By plugging it into (11), we conclude that this sequence is a solution for $Q \geq 2$. There are no sequences (m_1, \dots, m_n) associated to $(2, 2, 2)$ except for $(2, 2, 2)$ itself.

We now consider the case of $k = 2$. Plugging $k = 2$ into (16), we get

$$m^2(r-2) \leq r(r-1) + (Q - 2 \cdot m).$$

Suppose $r \geq 4$. The left-hand side is at least $2m^2$, while the right-hand side is at most $m^2 + Q$. Hence, $m^2 \leq Q$ and so $m \leq 4$ for $Q \leq 16$. Since $r \leq m - 1$, we conclude that there are no solutions with $r \geq 4$.

Suppose $r = 3$. We get $m^2 \leq 6 + Q - 2m$. For $Q \leq 9$, we get $m^2 \leq 15$ and so it remains to check $m \in \{2, 3\}$. None of these values are compatible with $r = 3$.

Suppose $r = 2$. Inequality (16) takes the form

$$0 \leq 2 + Q - 2m.$$

For $Q \leq 9$, it implies $m \leq 5$. Since $r = 2$, we must have $m \geq 3$. Plugging $m \in \{3, 4, 5\}$ into (16), we obtain that the following sequences are solutions: $(3, 3, 2)$ for $Q \geq 4$, $(4, 4, 2)$ for $Q \geq 6$, and $(5, 5, 2)$ for $Q \geq 8$. Now consider sequences (m_1, \dots, m_n) that can be associated

with these families of solutions. For $(3, 3, 2)$, there is only one such sequence $(3, 2, 2, 2)$ which becomes a solution for $Q \geq 12$. For $(4, 4, 2)$, we need to check $(4, 3, 3)$, $(4, 3, 2, 2)$, and $(4, 2, 2, 2, 2)$, which become solutions for $Q \geq 12$, $Q \geq 26$, and $Q \geq 44$, respectively. For $(5, 5, 2)$, we obtain the following sequences $(5, 4, 3)$, $(5, 4, 2, 2)$, $(5, 3, 3, 2)$, $(5, 3, 2, 2, 2)$, and $(5, 2, 2, 2, 2, 2)$ none of which are solutions for $Q \leq 9$.

Finally, we are left with sequences of the form (m, m) , which are solutions for any $Q \geq 1$ and $m \geq 2$. It remains to consider sequences (m_1, \dots, m_n) that are associated with (m, m) . We get

$$\prod_{i=1}^n m_i \leq 2 \cdot m(m-1) + Q$$

since the right-hand side of (11) cannot decrease when (m_1, \dots, m_n) is replaced with (m, m) .

If $m_1 = m_2$, then (m_1, \dots, m_n) coincides with (m, m) and thus we focus on the case where $m_2 = m - \delta$ with integer $\delta \geq 1$. By quasi-convexity of the product,

$$\prod_{i=2}^n m_i \geq (m - \delta)(1 + \delta)$$

and so we obtain

$$\prod_{i=1}^n m_i \geq m \cdot (m - \delta) \cdot (1 + \delta).$$

Putting the pieces together, we conclude that the following inequality holds

$$m \cdot (m - \delta) \cdot (1 + \delta) \leq 2m \cdot (m - 1) + Q.$$

Without loss of generality, $m - \delta \geq 1 + \delta$, i.e., $\delta \leq (m - 1)/2$.

Suppose $\delta \geq 4$. Then the left-hand side is at least $m \cdot \frac{m+1}{2} \cdot 5$ and thus

$$m^2 \leq 2Q.$$

Since $\delta \leq (m - 1)/2$ and $\delta \geq 4$, there are no solutions unless $Q \geq 41$.

Suppose $\delta = 3$ and so $m \geq 7$. We get

$$4m(m - 3) \leq 2m(m - 1) + Q.$$

Equivalently,

$$2m^2 - 10m \leq Q.$$

The left-hand side satisfies

$$2m^2 - 10m \geq 2(m-3)^2 - 18$$

and thus

$$m \leq 3 + \sqrt{\frac{Q+18}{2}},$$

which is incompatible with $m \geq 7$ unless $Q \geq 14$.

Suppose $\delta = 2$ and so $m \geq 5$. We obtain

$$3m(m-2) \leq 2m(m-1) + Q$$

or $(m-2)^2 \leq Q+4$ and thus

$$m \leq 2 + \sqrt{Q+4}$$

For $Q \leq 11$, we get $m \leq 5$. Thus $m = 5$. The corresponding sequences take the form $(5, 3, 3)$ and $(5, 3, 2, 2)$, which become solutions for $Q \geq 13$ and $Q \geq 30$, respectively.

Suppose $\delta = 1$ and so $m \geq 3$. The sequence (m_1, \dots, m_n) takes the form $(m, m-1, 2)$. Plugging it into (11), we get

$$2m(m-1) \leq m(m-1) + (m-1)(m-2) + 2 + Q$$

and thus $(m, m-1, 2)$ is a solution for any

$$m \leq 2 + \frac{Q}{2}.$$

For $Q \leq 9$, we get the following values $m \in \{3, 4, 5, 6\}$ starting from Q equal to 2, 4, 6, 8, respectively. ■

1.3 Proofs Pertaining to Extremality in Payoff Space

First, we formulate and prove a lemma showing that, in a generic game, all extreme correlated equilibria correspond to distinct points in the payoff space. This technical result is needed to make sure that a preimage of an extreme point in the payoff space can only originate from an extreme point in the space of action distributions and thus deduce Corollary 2 from Theorem 1. Next, we prove Proposition 1.

Lemma 5. *In a generic game, no two correlated equilibria result in the same payoff to an agent.*

Proof. We fix sets of agents $N = \{1, \dots, n\}$ and actions A_1, \dots, A_n and consider a game $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$. We show that for an open set of utilities $u \in \mathbb{R}^{N \times A}$ of full measure, any two extreme points of $\text{CE}(\Gamma)$ result in different payoffs to each agent.

Extreme points of $CE(\Gamma)$ are basic feasible solutions of the corresponding system of inequalities. Thus, any extreme point μ is uniquely determined by the set of equalities formed by active constraints. In other words, there exist subsets $S \subseteq A$ (the support of μ) and a set $IC \subseteq \cup_{i \in N} \{i\} \times A_i \times A_i$ (active incentive constraints) such that μ is a unique solution to the following system of equations

$$\begin{cases} \sum_{a_{-i}} \mu(a_i, a_{-i}) (u_i(a_i, a_{-i}) - u_i(a'_i, a_{-i})) = 0, & (i, a_i, a'_i) \in IC \\ \mu(a) = 0, & a \in A \setminus S \\ \sum_a \mu(a) = 1 \end{cases} \quad (17)$$

The uniqueness of the solution means that there are exactly $|A|$ linearly independent equations on the $|A|$ coordinates of $\mu \in \mathbb{R}^A$, i.e., without loss of generality we can assume that $|IC| = |A| - |S| - 1$ and the matrix of the system is non-degenerate. Denoting this $|A| \times |A|$ matrix by $M_{u,S,IC}$ and the vector on the right-hand side by e , we rewrite the system as $M_{u,S,IC} \cdot \mu = e$. Since the solution exists and is unique, the determinant $\det M_{u,S,IC}$ must be non-zero.

Suppose that there are two distinct extreme points μ and μ' of the set of correlated equilibria such that agent j 's expected payoff is the same: $u_j(\mu) = u_j(\mu')$. Let (S, IC) and (S', IC') be the active constraints at μ and μ' , respectively, chosen so that $\det M_{u,S,IC} \neq 0$ and $\det M^{u,S',IC'} \neq 0$. We conclude that the following system on $(\mu, \mu') \in \mathbb{R}^A \times \mathbb{R}^A$ has a unique solution:

$$\begin{cases} M^{u,S,IC} \cdot \mu + 0 \cdot \mu' = e \\ 0 \cdot \mu + M^{u,S',IC'} \cdot \mu' = e \\ u_j \cdot \mu - u_j \cdot \mu' = 0 \end{cases} \quad (18)$$

This is a system of $2|A| + 1$ equations on $2|A|$ unknowns. It can only have a unique solution if the equations are linearly dependent, i.e., the determinant of the augmented matrix is zero

$$\det \begin{pmatrix} M^{u,S,IC} & 0 & e \\ 0 & M^{u,S',IC'} & e \\ u_j & -u_j & 0 \end{pmatrix} = 0.$$

This identity can be seen as an algebraic equation on $u \in \mathbb{R}^{N \times A}$ and thus defines a closed zero-measure subset of $\mathbb{R}^{N \times A}$. Taking the complement of the union of these subsets over all distinct pairs of (S, IC) and (S', IC') and agents j , we get an open set of full measure such that for any u from this set, no distinct extreme points can result in the same utility for any agent j . ■

We now prove Proposition 1. Recall that, linear objectives in the payoff space have a particularly simple structure:

$$W_\alpha(\mu) = \sum_{i \in N} \alpha_i \sum_{s \in S} u_i(s) \mu(s),$$

where α_i are the weights assigned to each agent's utility. In other words, any linear objective is a weighted welfare function. The case where $\alpha_1 = \dots = \alpha_n = 1$ corresponds to the standard utilitarian welfare, which is the sum of the agents' utilities.

Proof of Proposition 1. To deduce this result from Corollary 2, we need to show that utilitarian welfare W is a non-degenerate objective over the polytope of CE payoff profiles $U^{CE}(\Gamma)$ for a generic game. Consider a generic game Γ with utilities u_i . For generic weights α , the objective W_α is non-degenerate. Consider a new game Γ' with $u'_i = \alpha_i u_i$. In Γ' , the utilitarian welfare W' is equal to W_α in Γ . Since W_α is non-degenerate by assumption, and the set of correlated equilibria is invariant to positive affine transformations of individual utilities, the utilitarian welfare function is non-degenerate in the generic game Γ' . ■

1.4 Proofs Pertaining to Symmetric Games

We start with some preliminary analysis necessary for the proof of Theorem 2.

Let $T = \{1, \dots, m\}$ denote a finite set. We will need the following definition generalizing the notion of an exchangeable distribution. For a subgroup P of permutations of $\{1, \dots, n\}$, a distribution μ over an n -fold product space $T^n \equiv T \times \dots \times T$ is called *P -exchangeable* if for any element $(t_1, \dots, t_n) \in T^n$ and any permutation $\pi \in P$, we have

$$\mu(t_1, \dots, t_n) = \mu(t_{\pi(1)}, \dots, t_{\pi(n)}).$$

The classical notion of exchangeability corresponds to the case when P is the set of all permutations.

The set of P -exchangeable distributions is a convex set. Furthermore, any P -symmetric correlated equilibrium is a P -exchangeable distribution.

Lemma 6. *If μ is an extreme point of the set of symmetric P -correlated equilibria with transitive P , then μ can be represented as a convex combination of at most*

$$m(m-1)+1$$

extreme P -exchangeable distributions, where m is the number of actions of each agent.

The set of exchangeable distributions is universal in that it only depends on the number of agents and actions in the game, but not on their associated payoffs.

Proof of Lemma 6. To check that a P -exchangeable distribution is a correlated equilibrium of a P -symmetric game, it suffices to check the incentive constraints for a single agent since any two agents are equivalent up to relabeling determined by a permutation within P . For a single agent, there are $m(m-1)$ incentive constraints. The result then follows from [Winkler \(1988\)](#). ■

Extreme points of P -exchangeable distributions Exchangeable distributions are often considered on an infinite product space $T^{\mathbb{Z}}$. In this case, de Finetti's Theorem translates exchangeable observations to conditionally independent observations relative to some latent variable. Extreme points are then i.i.d. distributions. For a finite product space, the set of extreme exchangeable distributions has a different structure established by [Diaconis and Freedman \(1980\)](#). We generalize their result to the case of P -exchangeability.

Lemma 7. Denote by $\delta_{(t_1, \dots, t_n)}$ the distribution that places unit mass on the profile (t_1, \dots, t_n) . Any distribution μ over T^n is an extreme point of P -exchangeable distributions if and only if it can be represented as

$$\mu = \frac{1}{|P|} \sum_{\pi \in P} \delta_{(t_{\pi(1)}, \dots, t_{\pi(n)})}. \quad (19)$$

Proof. It is immediate that any such μ is an extreme point. We show that such distributions cover the set of extreme points. It suffices to demonstrate that any P -symmetric τ is in the convex hull of such distributions. To see that, we can write $\tau = \sum_{(t_1, \dots, t_n)} \tau(t_1, \dots, t_n) \cdot \delta_{(t_1, \dots, t_n)}$ and then average both sides of this identity over all $\pi \in P$. ■

Extreme exchangeable distributions may have different supports because not all terms in the sum in (19) are distinct. For example, suppose $n = 3$, $T = \{a, b\}$, and $P = S_3$. Then, there are four extreme exchangeable distributions $\delta_{(a,a,a)}$, $\frac{1}{3}(\delta_{(a,a,b)} + \delta_{(a,b,a)} + \delta_{(b,a,a)})$, $\frac{1}{3}(\delta_{(a,b,b)} + \delta_{(b,a,b)} + \delta_{(b,b,a)})$, and $\delta_{(b,b,b)}$.

Let $\Delta_{m,n}$ be the discrete simplex of dimension m and size n , which consists of all non-negative integer vectors (k_1, \dots, k_m) such that $k_1 + \dots + k_m = n$. For each $t = (t_1, \dots, t_n)$, we can assign the frequency vector $f(t)$ in $\Delta_{m,n}$ that counts the number of times each $j \in T = \{1, \dots, m\}$ appears in t . That is,

$$f(t)_j = |\{i : t_i = j\}|.$$

Since an extreme exchangeable distribution μ is obtained via symmetrization of some $t = (t_1, \dots, t_n)$, we define a frequency vector as the frequency vector $f(t)$.

The total number of different frequency vectors is given by:¹⁵

$$|\Delta_{m,n}| = \binom{n+m-1}{m-1}. \quad (20)$$

For $P = S_n$, there is a natural bijection between extreme exchangeable distributions and $\Delta_{m,n}$ and thus equality (20) provides the total number of extreme exchangeable distributions. For a general P , the number of extreme exchangeable distributions is bounded from below by $|\Delta_{m,n}|$.

Lemma 8. *Consider a product distribution $\tau = \nu \times \dots \times \nu$ over T^n with full support. If τ is represented as a mixture of extreme P -exchangeable distributions, then at least*

$$\binom{n+m-1}{m-1}$$

of them must enter the mixture with a positive weight.

Proof of Lemma 8. Sample $t = (t_1, \dots, t_n)$ from τ . Consider the frequency vector $f(t)$. Since τ is a full-support product distribution, this random vector takes all the values in $\Delta_{m,n}$ with positive probability. Now, suppose that τ is represented as a mixture of extreme P -exchangeable distributions so that all distributions corresponding to a particular frequency vector f_0 enter the mixture with zero weight. Thus f_0 must be absent from realizations of $f(t)$. This contradiction implies that for each element of $\Delta_{m,n}$, there must be an extreme exchangeable distribution with a positive weight. Thus, the total number of such distributions is at least $|\Delta_{m,n}|$ given by (20). ■

Proof of Theorem 2. Regularity of ν allows us to drop actions that are assigned zero weight. So, we can assume, without loss of generality, that each of the m actions of each agent has positive weight. By Lemma 6, the set of P -symmetric correlated equilibria is a convex hull of at most $m(m-1) + 1$ extreme P -exchangeable distributions. On the other hand, by Lemma 8, we need at least $\binom{n+m-1}{m-1}$ extreme P -exchangeable distributions to represent ν as a mixture.

Thus, ν is not an extreme point of the set of P -symmetric equilibria if

$$\binom{n+m-1}{m-1} > m(m-1) + 1. \quad (21)$$

Consider the case of $n \geq 3$ and $m \geq 2$ and show that such Nash equilibria cannot be extreme. We need to demonstrate that (21) is satisfied for any such n and m . The left-hand

¹⁵This can be derived by what is often referred to as the Stars and Bars Lemma in combinatorics. Namely, one can think of n items to be split into m separate buckets. If the items are thought of as stars on a line, the buckets can be represented by $m-1$ bars that are interspersed among the stars and partition them into m subsets. Our derivation is then equivalent to identifying the number of choices of these $m-1$ locations out of $n+m-1$ possibilities.

side is monotone-increasing in n , and thus, it is enough to consider the case of $n = 3$. We obtain

$$\binom{m+2}{m-1} > m(m-1) + 1,$$

or, equivalently,

$$\frac{(m+2)(m+1)m}{6} > m(m-1) + 1.$$

Elementary computations show that this inequality holds for all $m \geq 2$.

For $m = 1$ and any n , the Nash equilibrium is pure and is thus an extreme point. Indeed, it maximizes the expected value of an objective that equals 1 on the corresponding action profile and 0 otherwise.

For $n = 2$, any regular Nash equilibrium is an extreme point within the set of all correlated equilibria by Theorem 1 and is thus necessarily an extreme point of the smaller set of P -symmetric correlated equilibria. ■

Proof of Proposition 2. For $P = S_n$, extreme exchangeable distributions from Lemma 7 can be described via urns as the proposition claims. Indeed, μ of the form (19) corresponds to an urn with n balls marked with elements of A and a random draw without replacement. Combining this observation with Lemma 6, we obtain the proposition's result. ■

1.5 Proofs Pertaining to Bayesian Correlated Equilibria

Proof of Proposition 3. Consider a game Γ with incomplete information and a pure Bayesian Nash equilibrium ψ that assigns a single realization $a = (a_1, \dots, a_n)$ for each realization of (θ, t) . Accordingly, it is an extreme point of the set of all distributions $\Delta(A \times \Theta \times T)$ that have marginal π on $\Theta \times T$. Since the set of BCEs is a subset of this set of distributions, ψ is an extreme BCE.

Now consider a mixed Bayesian Nash equilibrium ψ , where agent i randomizes her actions according to $\sigma_i: T_i \rightarrow \Delta(A_i)$. To show that it cannot be an extreme point of the BCE set, we find an appropriate generalization of the bounds from Lemmas 1 and 2 to the incomplete information setting.

We first extend Lemma 1. Since genericity implies that π has full support, we conclude that the support of ψ contains $|\Theta| \cdot \prod_{i \in N} \left(\sum_{t_i \in T_i} |\text{supp } \sigma_i(t_i)| \right)$ elements.

By the genericity of Γ , the incentive constraints outside of the support can be assumed to be inactive: an agent i of type t_i would be strictly worse by playing a_i outside of

$\text{supp } \sigma_i(t_i)$.¹⁶ Thus, it is enough to demonstrate that ψ is not an extreme point in an auxiliary game where agent i of type t_i only has actions from $\sigma_i(t_i)$ available. Such an agent faces $|\sigma_i(t_i)| \cdot (|\sigma_i(t_i)| - 1)$ incentive constraints in a correlated equilibrium. Additionally, ψ satisfies the requirement that its marginal on $\Theta \times T$ equals π which amounts to $|\Theta| \prod_i |T_i| - 1$ constraints. By [Winkler \(1988\)](#)'s theorem, for ψ to be an extreme point, we must have

$$|\Theta| \cdot \prod_{i \in N} \left(\sum_{t_i \in T_i} |\text{supp } \sigma_i(t_i)| \right) \leq |\Theta| \prod_{i \in N} |T_i| + \sum_{i \in N} \sum_{t_i \in T_i} |\sigma_i(t_i)| \cdot (|\sigma_i(t_i)| - 1). \quad (22)$$

Denote

$$m_i = 1 + \sum_{t_i \in T_i} (|\text{supp } \sigma_i(t_i)| - 1),$$

which can be seen as the effective number of actions over which i randomizes if we factor out randomness in t_i . By the convexity of $f(x) = x(x - 1)$, we can only increase the right-hand side of (22) by replacing $\sum_{t_i \in T_i} |\sigma_i(t_i)| \cdot (|\sigma_i(t_i)| - 1)$ with $m_i(m_i - 1)$. We obtain

$$|\Theta| \cdot \prod_{i \in N} (m_i + (|T_i| - 1)) \leq |\Theta| \cdot \prod_{i \in N} |T_i| + \sum_{i \in N} m_i(m_i - 1), \quad (23)$$

which is a necessary condition for ψ to be an extreme BCE.

We now complement bound (23) with the one showing that no m_i can be much bigger than the other. Indeed, for σ_i to be an equilibrium strategy of agent i , she must be indifferent between all the actions from $\text{supp } \sigma(t_i)$ conditional on each $t_i \in T_i$. Treating these indifferences as algebraic equations on strategies of other agents $(\sigma_j)_{j \in N \setminus \{i\}}$, we obtain $\sum_{t_i \in T_i} (|\text{supp } \sigma(t_i)| - 1) = m_i - 1$ conditions on $\sum_{j \neq i} (m_j - 1)$ variables. Since u_i can be chosen arbitrarily as a function of (a_i, t_i) , all these conditions are independent in a generic game and thus can only be satisfied if the number of conditions does not exceed the number of variables. We conclude that the following inequality must hold for all agents i :

$$m_i - 1 \leq \sum_{j \neq i} (m_j - 1), \quad (24)$$

generalizing [Lemma 2](#).

Combining (23) and (24), we obtain that for ψ to be an extreme point of BCE, the

¹⁶Indeed, if there is a tie, we can lower u_i arbitrarily outside of the support, which preserves the equilibrium. Therefore, there are no ties on a dense set of games. Its complement is a closed nowhere dense algebraic set and so has zero measure. We conclude that the set itself is open and everywhere dense. Thus, such games are generic.

following system must have an integer solution (m_1, \dots, m_n) with $m_i \geq 1$ for all i :

$$\begin{cases} |\Theta| \cdot \prod_{i \in N} (m_i + (|T_i| - 1)) & \leq |\Theta| \cdot \prod_{i \in N} |T_i| + \sum_{i=1}^n m_i \cdot (m_i - 1) \\ m_i - 1 & \leq \sum_{j \neq i} (m_j - 1), \end{cases} \quad i = 1, \dots, n \quad (25)$$

Without loss of generality, we assume that $m_1 \geq m_2 \geq \dots \geq m_n$. For $|\Theta| = 1$ and $|T_i| = 1$ for all i , this system matches the one from the proof of Theorem 1.

By a simple induction argument contained in Lemma 9,

$$\prod_{i \in N} (m_i + (|T_i| - 1)) \geq \prod_{i \in N} m_i + \prod_{i \in N} |T_i| - 1$$

and thus, we obtain a system

$$\begin{cases} \prod_{i \in N} m_i & \leq 2 + \frac{1}{|\Theta|} \sum_{i=1}^n m_i \cdot (m_i - 1) \\ m_i - 1 & \leq \sum_{j \neq i} (m_j - 1), \end{cases} \quad i = 1, \dots, n \quad (26)$$

Note that any solution to this system remains a solution if we plug in $|\Theta| = 1$. From Proposition 5, we know that all such solutions take one of three forms

$$(m_1, \dots, m_n) = \begin{cases} (m, m, 1, 1, \dots, 1) & \text{with } m \geq 1 \\ (2, 2, 2, 1, \dots, 1) \\ (3, 2, 2, 1, \dots, 1) \end{cases} \quad (27)$$

Plugging this back to (26), we conclude that for $|\Theta| \geq 2$, none of these sequences work, other than the first one with $m = 1$.

Now consider the case of $|\Theta| = 1$ and assume that there are at least three agents j such that $|T_j| \geq 2$. Plugging the candidate solutions (27) into the top inequality of (25), one gets

$$(m + |T_1| - 1)(m + |T_2| - 1) \leq |T_1| |T_2| + \frac{2}{|T_j|} \cdot m(m - 1),$$

where $|T_j| \geq 2$ is the size of the type space of one of the non-randomizing agents. By decreasing $|T_1|$ and $|T_2|$ we relax the inequality, but even for $|T_1| = |T_2| = 1$, the inequality can only hold if $m = 1$. The second and the third candidate solutions in (27) are also easily ruled out. ■

Lemma 9. *Let $n \in \mathbb{N}$ and $x_i, y_i \in \mathbb{N}$, $1 \leq i \leq n$. Then*

$$\prod_{i=1}^n (x_i + y_i - 1) \geq \prod_{i=1}^n x_i + \prod_{i=1}^n y_i - 1.$$

Proof. We proceed by induction on n . For $n = 1$, the inequality is an equality:

$$x_1 + y_1 - 1 = x_1 + y_1 - 1.$$

Assume the claim holds for some $n - 1$ and deduce it for n . Define

$$A = \prod_{i=1}^{n-1} x_i, \quad B = \prod_{i=1}^{n-1} y_i, \quad P = \prod_{i=1}^{n-1} (x_i + y_i - 1),$$

so that by the induction hypothesis,

$$P \geq A + B - 1.$$

For n , we have

$$\prod_{i=1}^n (x_i + y_i - 1) = P(x_n + y_n - 1)$$

and need to show

$$P(x_n + y_n - 1) \geq Ax_n + By_n - 1.$$

Since $P \geq A + B - 1$, it suffices to prove that

$$(A + B - 1)(x_n + y_n - 1) \geq Ax_n + By_n - 1.$$

A straightforward expansion yields

$$(A + B - 1)(x_n + y_n - 1) = Ax_n + Ay_n + Bx_n + By_n - x_n - y_n - A - B + 1.$$

Thus, it is enough to verify that

$$(A - 1)y_n + (B - 1)x_n - (x_n + y_n) - (A + B) + 2 \geq 0.$$

Noting that $x_n, y_n \geq 1$, we have

$$x_n(B - 1) \geq B - 1 \quad \text{and} \quad y_n(A - 1) \geq A - 1,$$

so that

$$x_n(B - 1) + y_n(A - 1) \geq (B - 1) + (A - 1) = A + B - 2.$$

Thus, the desired inequality follows. ■

1.6 Proofs Pertaining to Coarse Correlated Equilibria

Proof of Proposition 4. Consider a Nash equilibrium μ of a game Γ . If μ is pure, it is trivially an extreme point of $\text{CCE}(\Gamma)$, as point masses are extreme in $\Delta(A)$ and $\text{CCE}(\Gamma) \subset \Delta(A)$.

Conversely, suppose μ is an extreme point of $\text{CCE}(\Gamma)$. The polytope $\text{CCE}(\Gamma)$ is defined by n linear incentive constraints within the simplex $\Delta(A)$. Thus, by Theorem 2.1 of [Winkler](#)

(1988), any extreme point μ of $\text{CCE}(\Gamma)$ satisfies

$$|\text{supp } \mu| \leq 1 + n \tag{28}$$

Now, suppose for the sake of contradiction that μ is a mixed Nash equilibrium with $k \geq 2$ agents randomizing. By regularity, the incentive constraints of non-mixing agents are slack, and so these agents are irrelevant to the question of extremality. Hence, without loss of generality, we can restrict attention to the game involving only k randomizing agents. Since μ is a Nash equilibrium, its support must contain at least 2^k action profiles. However, from (28), we have $|\text{supp } \mu| \leq 1 + k$. This leads to the inequality $2^k \leq 1 + k$, which is false for all $k \geq 2$. Thus, μ cannot be a mixed Nash equilibrium, and we conclude that μ must be pure. ■